

Security & Stability Assessment of SAC126 Recommendation 1

DNSSEC DS Record Automation via CDS/CDNSKEY

23 June 2026

Interisle Consulting Group, LLC

Table of Contents

Executive Summary	3
Scope and approach	3
Conclusion	3
Findings structure	4
Areas of residual concern	4
Use of this assessment	4
1 Introduction	5
2 Methodology	6
3 Mechanism Description	8
3.1 Actors and trust relationships	8
3.2 The ongoing-maintenance data flow (RFC 7344)	8
3.3 The bootstrapping data flow (RFC 9615)	9
3.4 The going-insecure data flow (RFC 8078 section 4)	9
3.5 Deployment topology	10
3.6 Validation and pre-checks performed by the parent	10
3.7 State transitions covered by the mechanism	10
3.8 Out-of-scope variations	11
4 Threat Model and Security Pillar Findings	12
4.1 Threat model overview	12
4.2 Branch S1 — Authentication of DS update requests	12
4.3 Branch S2 — Authentication of bootstrapping	13
4.4 Branch S3 — Interaction with EPP registration locks	14
4.5 Branch S4 — Cross-nameserver consistency	15
4.6 Branch S5 — Multi-channel collision behavior	16
4.7 Summary of security pillar findings	16
5 Stability Pillar Findings	18
5.1 Branch T1 — Standards conformance	18
5.2 Branch T2 — Validation pre-checks and resolution coherence	19
5.3 Branch T3 — TTL handling and rollback timing	20
5.4 Branch T4 — Scanning load and cadence	21
5.5 Branch T5 — Multi-signer setups per RFC 8901	22
5.6 Summary of stability pillar findings	23
6 Cross-Cutting Concerns	24
6.1 Concern X1 — The going-insecure path	24
6.2 Concern X2 — Reporting and observability	26
6.3 Composite finding for cross-cutting concerns	28
7 Findings and Recommendations	29
7.1 Hard requirements	29
7.2 Disclosure requirements	30
7.3 Recommended practices	31
7.4 Composite assessment	32

7.5 Treatment of subsequent operator submissions	33
Appendix A: Mapping of Findings to SAC126 Sections.....	34
A.1 Hard requirements	34
A.2 Disclosure requirements	36
A.3 Recommended practices	37
A.4 Reverse index by SAC126 section.....	37
A.5 Findings without direct SAC126 source	39

Executive Summary

This report presents a security and stability assessment of the mechanism identified in Recommendation 1 of [SAC126](#), *DNSSEC DS Record Automation* (SSAC, 12 August 2024). Recommendation 1¹ holds that registries and registrars wishing to automate DS record management for third-party DNSSEC operations should use the CDS/CDNSKEY mechanism specified in RFCs [7344](#), [8078](#), and [9615](#). The assessment is conducted within the analytical frame established by the [Registry Services Evaluation Policy](#) (RSEP).

Scope and approach

The mechanism under assessment is the publication, by a Child DNS operator, of CDS and CDNSKEY records at the apex of a child zone, and the corresponding ingest, validation, and DS publication actions performed by the parent registry or its sponsoring registrar. Three operational flows are within scope:

1. Ongoing DS maintenance under an existing chain of trust (RFC 7344)
2. Initial DS provisioning where no chain of trust yet exists (RFC 9615)
3. The going-insecure transition (RFC 8078 section 4)

It should be noted that push-based variants, out-of-band APIs, and the CSYNC mechanism are excluded, consistent with the bounds of Recommendation 1 itself.

The assessment is structured as two parallel pillars derived directly from RSEP section 1.2 and 1.3, decomposed into ten branches and two cross-cutting concerns. Each branch is evaluated through documentation review, structured threat modeling, test scenario analysis, and reference to operational evidence from existing deployments. SAC126 Appendix B is treated as primary input where its analytical work pre-articulates the design space the operator must navigate.

Conclusion

The CDS/CDNSKEY mechanism, when implemented in conformance with the seven hard requirements set out in Section 7.1 and accompanied by the corresponding disclosure documentation in Section 7.2, does not raise significant security or stability issues within the meaning of RSEP section 1.2 and 1.3. The mechanism rests on three standards-track RFCs whose conformance is checkable against design documentation, and its operational properties — when subject to the requirements identified — produce neither unauthorized mutation of registry data nor conditions adverse to the operation of validating resolvers. A positive preliminary determination by ICANN under RSEP section 2.2 is supported by the technical evidence developed in this report.

¹ SAC126 Recommendation 1: If a registry or a registrar wishes to implement DS automation for third-party DNSSEC operations, the current recommended interoperable mechanism is CDS/CDNSKEY (RFCs 7344, 9615). This mechanism has limitations as described in this document but has been deployed and there is operational experience in using it.

Findings structure

The assessment's findings are consolidated in Section 7 of the body of this report and organized into three categories reflecting their normative weight. Seven *hard requirements* (HR-1 through HR-7) are preconditions for a positive finding; their absence produces a negative result under one or both RSEP pillars. Seven *disclosure requirements* (DR-1 through DR-7) must be documented in the operator's design materials, with the substance of each choice left to operator judgment. Four *recommended practices* (RP-1 through RP-4) strengthen the deployment posture but are not required for a positive finding under the present standards landscape. Together, the eighteen items form a complete checklist against which any deployment of the mechanism can be evaluated.

Areas of residual concern

The assessment identifies three areas where the mechanism's properties depend materially on deployment-time choices that the published standards do not fully resolve.

1. The *multi-signer case per [RFC 8901](#) model 2*, where each participating provider holds its own keys and a unilateral CDS publication by one provider could, under naive implementation, expel the others from the delegation.
2. The *reporting and observability gap*: the cited RFCs do not specify reporting mechanisms, leaving registrants and Child DNS operators dependent on operator-specific channels to detect events affecting their delegations.
3. The *bootstrapping precondition* that authoritative nameserver hostnames lie in DNSSEC-signed zones — a structural property of RFC 9615's authentication construction that limits the mechanism's applicability to delegations whose nameservers meet the precondition.

Each area is addressed by specific items in Section 7 but warrants explicit attention in any submission and in subsequent ICANN evaluation.

Use of this assessment

The findings in Section 7 are mechanism-specific rather than operator-specific and apply to any deployment of the CDS/CDNSKEY mechanism for third-party DNSSEC operations. Operators submitting RSEP requests can engage each item directly: confirming the hard requirements, documenting the disclosure items, and addressing the recommended practices either by adoption or by reasoned acknowledgement. ICANN staff conducting a preliminary determination under RSEP section 2.2 can use the same set as the technical evidence base. A submission that addresses each item explicitly provides, by construction, the basis for a positive preliminary determination; a submission that omits items or treats them without specificity leaves residual ambiguity that may justify further engagement or referral to the Registry Services Technical Evaluation Panel under RSEP section 2.4.

The assessment supports SAC126's central observation that automated DS maintenance is a natural and desirable extension of the broader trend toward automation in DNSSEC operations, and that the CDS/CDNSKEY mechanism — limitations notwithstanding — is the appropriate technical foundation for that automation in the present standards landscape.

1 Introduction

This assessment evaluates the security and stability properties of the mechanism identified in Recommendation 1 of SAC126, *Towards DNSSEC DS Record Automation* (SSAC, 12 August 2024). The recommendation states that registries and registrars wishing to implement DS automation for third-party DNSSEC operations should use the CDS/CDNSKEY mechanism specified in RFCs 7344 and 9615, and acknowledges that while the mechanism has known limitations, it is the only interoperable approach with meaningful deployment experience.

The assessment is conducted in accordance with the Registry Services Evaluation Policy (RSEP), which governs ICANN's process for evaluating new or modified registry services. RSEP defines security and stability as discrete, narrowly scoped concepts. Section 1.2 defines an effect on security as either the unauthorized disclosure, alteration, insertion, or destruction of Registry Data, or the unauthorized access to or disclosure of information or resources on the Internet by systems operating in accordance with applicable standards. Section 1.3 defines an effect on stability as either non-compliance with authoritative published standards (such as IETF standards-track or Best Current Practice RFCs), or the creation of a condition that adversely affects the throughput, response time, consistency, or coherence of responses to Internet servers or end systems operating per applicable standards. These two definitions form the analytical frame for the entire assessment. Findings that do not map to one of these four conditions — unauthorized mutation, unauthorized access, standards non-conformance, or measurable performance/coherence degradation — are out of scope.

The mechanism under evaluation is the publication, by a Child DNS operator, of CDS and/or CDNSKEY resource records at the apex of a child zone, and the corresponding ingest, validation, and DS publication actions performed by the parent registry or its sponsoring registrar.

The following RFCs are relevant in this context:

- RFC 7344 (Automating DNSSEC Delegation Trust Maintenance) specifies the mechanism for ongoing DS maintenance under an existing chain of trust.
- RFC 9615 (Automatic DNSSEC Bootstrapping Using Authenticated Signals from the Zone's Operator) specifies an in-band authenticated DNSSEC bootstrapping protocol for the case where no chain of trust yet exists.
- RFC 8078 (Managing DS Records from the Parent via CDS/CDNSKEY) provides operational guidance for the parent side, including the explicit signal for transitioning a delegation to insecure status.

SAC126 Recommendation 1 contemplates these three RFCs as a single composed mechanism, and this assessment treats them accordingly.

The assessment is bound to the case identified in SAC126 itself: the domain's DNS service is operated by a third party that is not the sponsoring registrar. Where the registrar also serves as the DNS operator, DS provisioning reduces to an internal coordination problem that does not invoke the CDS/CDNSKEY mechanism. The case where the registrar is the signer is therefore noted but not assessed. Likewise, push-based and out-of-band API approaches discussed in SAC126 section 4.2.2 fall outside Recommendation 1 and outside this scope.

2 Methodology

The assessment proceeds along two parallel pillars derived directly from RSEP sections 1.2 and 1.3. The security pillar examines the mechanism's resistance to unauthorized insertion, alteration, or destruction of DS records — which are Registry Data within the meaning of RSEP section 1.2(A) — and unauthorized disclosure or access via standards-conformant systems within the meaning of RSEP section 1.2(B). The stability pillar examines, first, whether the mechanism as deployed conforms to the cited standards-track RFCs, and second, whether its operational characteristics introduce conditions that adversely affect throughput, response time, consistency, or coherence of DNS responses.

Each pillar is decomposed into discrete branches representing concrete attack surfaces or operational characteristics. Within the security pillar, these are

- authentication of DS update requests under an existing chain of trust (RFC 7344);
- authentication of initial DS provisioning where no chain of trust exists (RFC 9615);
- the interaction of automation with EPP registration locks;
- defense against single-nameserver compromise via cross-nameserver consistency requirements; and
- collision behavior when both registry and registrar perform automation against the same delegation.

Within the stability pillar, the branches are:

- standards conformance against the MUST and SHOULD clauses of the cited RFCs;
- the parent-side validation pre-checks required by RFC 7344 section 4.1 to prevent publication of breaking DS sets;
- TTL handling and the rollback timing window;
- scanning load and cadence under pull-based operation; and
- the handling of multi-signer setups per RFC 8901 model 2, for which no standardized CDS/CDNSKEY processing currently exists.

Two cross-cutting concerns — the integrity of the going-insecure path (RFC 8078 section 4) and the absence of standardized reporting for failed or rejected updates — apply to both pillars and are treated separately.

For each branch, the assessment applies four methodological instruments:

1. Documentation review: the operator's design specifications, DNSSEC Practice Statement (DPS), EPP extension definitions, change-management runbooks, and any registrant-facing documentation.

2. Structured threat modeling using STRIDE² applied to the CDS ingest pipeline, with each identified threat mapped explicitly to one or both clauses of RSEP section 1.2.
3. Concrete test scenario execution, covering at minimum the cases SAC126 identifies as operationally non-trivial: algorithm rollover, transitioning to insecure, multi-signer key changes, registration lock interaction, conflicting CDS and CDNSKEY records, and inconsistent records across nameserver instances.
4. Operational evidence drawn from existing deployments — the ccTLD registries operating CDS scanning today, regional Internet registries (e.g., RIPE NCC) for reverse zones, and large operators (e.g., Cloudflare and DNSimple) publishing CDS/CDNSKEY on the child side — to ground theoretical findings in observed behavior at scale.

Each finding is recorded against three attributes: which RSEP clause it engages, what the mechanism does in response, and whether the response is sufficient. A finding of insufficiency triggers either a remediation recommendation, a documented residual risk to be accepted by the operator, or a referral to standards work in progress where the gap is industry-wide rather than operator-specific. SAC126 Appendix B is treated as primary input for several branches because it pre-articulates the design space the operator must navigate; this assessment does not duplicate that analysis but evaluates the operator's specific resolution of each open question.

The assessment deliverable is a written report structured to support an RSEP preliminary determination by ICANN. It states findings in terms of the section 1.2 and section 1.3 definitions, identifies the standards conformance posture against each cited RFC, and provides the technical evidence base needed for ICANN to determine whether the proposed service raises significant security or stability issues within the meaning of the policy.

² STRIDE (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege) is a threat model for identifying computer security threats.

3 Mechanism Description

This section describes the CDS/CDNSKEY mechanism as composed by RFCs 7344, 8078, and 9615, in the form it must be assessed. The description is normative for the remainder of this report; subsequent findings reference the data flows, actors, and decision points defined here.

3.1 Actors and trust relationships

Five actors participate in the mechanism.

- The **registrant** holds the domain registration and bears ultimate authority over its DNSSEC configuration but takes no direct technical action under the mechanism.
- The **Child DNS operator** controls the authoritative contents of the child zone, including its DNSKEY, CDS, and CDNSKEY records, and signs the zone using one or more private keys.
- The **registrar** is the registrant's contractual interface to the registry under the RRR (Registry-Registrar-Registrant) model and may, depending on deployment choice, perform the automation function or merely relay results received from the registry.
- The **registry** controls the parent zone and is the sole entity that publishes DS records into that zone, regardless of which party initiated the change.
- The **resolver population** is not an actor in the provisioning sense but is the system whose correct operation defines the stability pillar of this assessment; resolvers must continue to validate signatures throughout any DS state transition.

The pre-existing trust relationships are asymmetric. The registry trusts the registrar through an authenticated EPP channel established under their contractual relationship. The registrar trusts the registrant through an authenticated customer interface. Neither the registry nor the registrar has any pre-existing trust relationship with the Child DNS operator when that operator is a third party, and the absence of that relationship is the central problem the mechanism solves. The mechanism does not introduce an out-of-band credential exchange between Child DNS operator and the parent; instead, it derives trust either from the child's existing DNSSEC chain of trust (RFC 7344) or from a constructed authentication path rooted in the parent zone's own signed delegation data (RFC 9615).

3.2 The ongoing-maintenance data flow (RFC 7344)

When a child zone is already securely delegated — i.e., a valid DS RRset exists in the parent zone and the corresponding DNSKEY signs the child — the mechanism for changing the DS RRset is as follows.

- The Child DNS operator publishes one or more CDS records, one or more CDNSKEY records, or both, at the apex of the child zone. These records are signed by the existing zone-signing keys, which are themselves authenticated by the existing DS RRset in the parent.
- The parent (registry or registrar, depending on deployment) periodically queries the child zone for these records, typically by direct DNS query to the authoritative nameservers listed in the delegation. On receipt, the parent validates the DNSSEC signatures over the CDS or CDNSKEY RRset using the existing chain of trust. If validation succeeds, and if the resulting DS RRset

would not break DNS resolution of the child zone (per the safety check required by RFC 7344 section 4.1), the parent publishes the new DS RRset.

Three properties of this flow are central to the assessment.

1. The authentication of the change request is intrinsic: a third party cannot forge a CDS or CDNSKEY record without access to the child's existing DNSKEY private keys, which are held by the legitimate Child DNS operator.
2. The trigger for change is the parent's poll, not an action by the Child DNS operator; this introduces latency between publication and effect.
3. The parent's validation pre-check is a defense in depth against the legitimate operator publishing a self-injurious update — for example, a CDS RRset referencing a key that does not appear in the published DNSKEY RRset.

3.3 The bootstrapping data flow (RFC 9615)

When a child zone is signed but not yet securely delegated — no DS RRset exists in the parent zone — the chain of trust that authenticates the ongoing-maintenance flow does not yet exist, and a different authentication method is required. RFC 9615 specifies that the Child DNS operator publishes a Signaling Zone under a well-known name derived from the child zone and from each of its authoritative nameserver hostnames. The Signaling Zone contains CDS and CDNSKEY records for the child and is signed under the nameserver's own DNSSEC key, which is itself authenticated through the chain of trust of the nameserver hostname's parent zone.

Authentication therefore derives from a property the parent can independently verify: the entity publishing the bootstrapping signal is the same entity operating the nameservers listed in the delegation. The construction relies on the nameserver hostnames being themselves DNSSEC-signed and validatable. Where this precondition holds, the parent can establish trust in the initial DS RRset without prior credential exchange with the Child DNS operator. Where it does not hold — for example, where the nameserver hostnames live in an unsigned zone — bootstrapping falls back to either an out-of-band method or the registrant-centric process described in SAC126 section 4.1, which is outside the scope of Recommendation 1.

3.4 The going-insecure data flow (RFC 8078 section 4)

A third flow exists for transitioning a delegation from secure to insecure. Per RFC 8078 section 4, the Child DNS operator publishes a specially formed CDS RRset whose contents indicate the explicit removal of all DS records, i.e., a single CDS record with algorithm 0. On detection, the parent removes the DS RRset from the parent zone, ending DNSSEC validation for the delegation.

This flow is operationally significant for two reasons. It is the only standardized path for a registrant to leave DNSSEC under the mechanism without manual intervention, and it is also a destructive operation that, if accepted under fraudulent input, removes the cryptographic protection the mechanism otherwise provides. The going-insecure flow therefore appears in both the security pillar (as a potential attack target) and the stability pillar (as a required functional capability) of the assessment.

3.5 Deployment topology

The mechanism admits two deployment topologies under the RRR model.

- In the **registry-led topology** where the registry performs the periodic scan, validates the records, and publishes the resulting DS RRset directly into the parent zone. The registrar is informed of the change after the fact, typically via an EPP poll message per [RFC 8590](#) or an equivalent mechanism, so that the registrar's own state remains consistent with the registry.
- In the **registrar-led topology**, the registrar performs the scan, validates the records, and submits the resulting DS data to the registry through the standard EPP DNSSEC extension per [RFC 5910](#). The registry treats the submission as it would any other registrar-initiated DS change, applying its standard acceptance and publication logic.

The two topologies have materially different security and stability profiles even though the underlying protocol mechanisms are identical. The registry-led topology centralizes the scanning workload and produces a single source of truth for DS state but bypasses the registrar's view of the registration. The registrar-led topology preserves the registrar's role as primary mutator but distributes the scanning workload across registrars of varying technical capability and creates the possibility of inconsistent behavior across registrars within the same TLD. Where both topologies operate concurrently against the same delegation — registry and registrar each running independent scanners — collision behavior becomes a distinct concern, addressed in Branch S5 of the assessment.

3.6 Validation and pre-checks performed by the parent

Before publishing any DS RRset derived from a CDS or CDNSKEY signal, the parent is required by RFC 7344 section 4.1 to perform validation that the resulting state will not break resolution. At minimum this entails verifying that for each DS record to be published, a corresponding DNSKEY record exists in the child zone with matching algorithm and key tag, and that the child zone's DNSKEY RRset is signed by at least one key referenced by the new DS RRset. A parent that publishes a DS RRset without performing this check can be induced — by a misconfigured but legitimate Child DNS operator, or in some adversarial scenarios — to produce a delegation whose DNSSEC validation fails, rendering the child zone unresolvable for validating clients.

Beyond the minimum, several deployed implementations apply additional checks: rejection of CDS records using digest types deprecated under [RFC 9904](#), [RFC 9905](#) and [RFC 9906](#); rejection of DNSKEY records using prohibited signing algorithms; consistency checks between CDS and CDNSKEY records when both are published; and consistency checks across the authoritative nameserver set, per [RFC 9975](#) (*Clarifications on CDS/CDNSKEY and CSYNC Consistency*), which updates RFCs 7344 and [7477](#) to mandate cross-nameserver consistency checks. The cross-nameserver consistency check is now a published standards-track requirement; the remaining checks are not uniformly mandated by the cited RFCs and represent implementation latitude that the assessment will examine on a per-operator basis.

3.7 State transitions covered by the mechanism

The mechanism, taken as a whole, supports three state transitions for a delegation.

- **Insecure to secure** is the bootstrapping case per RFC 9615.

- **Secure to secure with key change** is the ongoing-maintenance case per RFC 7344, including key rollovers and algorithm rollovers.
- **Secure to insecure** is the going-insecure case per RFC 8078 section 4.

Each of the three transitions presents a distinct authentication problem, a distinct set of validation requirements, and a distinct failure mode if executed incorrectly. The remaining sections of this assessment treat each transition independently where its security or stability properties differ materially from the others, and jointly where they do not.

3.8 Out-of-scope variations

Three variations of the mechanism are recognized in the literature but excluded from this assessment.

- CSYNC records (RFC 7477), which provide an analogous mechanism for delegation NS and glue records, are not part of Recommendation 1 and are not assessed.
- Push-based and notification-triggered variants in which the Child DNS operator actively notifies the parent — including the Generalized DNS Notifications mechanism standardized in [RFC 9859](#) (September 2025), which post-dates SAC126 — are likewise outside Recommendation 1.
- Out-of-band APIs offered by individual operators, while operationally significant, are not the standardized mechanism cited by Recommendation 1 and are not assessed here. Separately, [draft-ietf-dnsop-ds-automation](#) (Operational Recommendations for DNSSEC Delegation Signer (DS) Automation) is in-progress IETF work that covers, with normative SHOULDs, much of the same operational territory as SAC126 Appendix B and this assessment; the draft is intended as Best Current Practice and was in IESG Evaluation at the time of writing. Although in-scope material from this draft is not yet authoritative, it is referenced throughout where its content materially informs the analysis. It is expected that this IETF draft will be published in a matter of months; SSAC may wish to publish a document to update SAC126 Appendix B following its publication. Also see Section 4.4 of this report for further analysis.

4 Threat Model and Security Pillar Findings

This section establishes the threat model used throughout the security pillar of the assessment and applies it to each of the five security branches identified in the methodology. The threat model is grounded in the STRIDE framework — spoofing, tampering, repudiation, information disclosure, denial of service, elevation of privilege — and each identified threat is mapped explicitly to one or both clauses of RSEP section 1.2 to maintain analytical discipline.

4.1 Threat model overview

The asset under protection is the DS RRset published in the parent zone for any given child delegation. This RRset is Registry Data within the meaning of RSEP section 1.2(A), and any unauthorized insertion, alteration, or destruction of it constitutes a security effect under the policy. A secondary asset is the integrity of the chain of trust that connects the parent's DS publication decision to the legitimate intent of the registrant; compromise of this chain may not directly mutate registry data but may cause registry data to be mutated in ways the registrant did not authorize, which is operationally indistinguishable from direct unauthorized mutation.

Three adversary classes are considered:

1. The **external network adversary** can observe and tamper with DNS traffic in transit between the parent and the child zone's nameservers but does not control the nameservers themselves.
2. The **nameserver-position adversary** controls one or more — but not all — of the authoritative nameservers listed in the child's delegation, whether through compromise, misconfiguration, or a malicious provider in a multi-signer setup.
3. The **full child operator adversary** controls all nameservers and DNSSEC keys for the child zone but is not the legitimate registrant; this corresponds to scenarios such as a complete operator compromise, an unauthorized provider transition, or an inside-attacker scenario at the Child DNS operator.

Each branch below considers which adversary classes are within scope for that branch and how the mechanism resists or fails to resist them.

The trust assumptions under which the mechanism operates must be stated explicitly, because the security findings depend on them. The parent zone is assumed to be correctly signed and its keys correctly managed by the registry. The resolver population is assumed to perform DNSSEC validation in accordance with [RFC 4035](#) and [RFC 6840](#). The Child DNS operator's private keys are assumed to be held only by the legitimate operator; key compromise scenarios are not within the mechanism's scope.

4.2 Branch S1 — Authentication of DS update requests

This branch addresses RSEP section 1.2(A) — unauthorized alteration or insertion of registry data — for the secure-to-secure transition per RFC 7344. The question is whether the mechanism, as deployed, prevents an adversary from causing the parent to publish a DS RRset that the legitimate Child DNS operator did not authorize.

The mechanism's authentication construction for this transition relies on the existing chain of trust. The CDS or CDNSKEY RRset published at the child apex is signed by the child's existing zone-signing keys, and those signatures are validated by the parent against the existing DS RRset before any change is applied. An **external network adversary** therefore cannot inject a forged CDS RRset into the parent's view of the child zone without also being able to forge a DNSSEC signature, which by construction requires possession of the child's private signing keys. Against the external network adversary, this branch yields a positive finding under RSEP section 1.2: the mechanism resists the relevant tampering threat to the standard expected of DNSSEC-protected data.

The **nameserver-position adversary** presents a more subtle case. An adversary controlling one of several authoritative nameservers can serve a CDS or CDNSKEY RRset that differs from what the legitimate operator publishes through the other nameservers, provided the adversary has access to a key authorized by the existing DS RRset. In a single-operator setup with a single signing key holder, this requires key compromise, which falls outside the mechanism's threat model. In a multi-signer setup per RFC 8901 model 2, however, each operator holds its own signing keys, and any one of them can sign a CDS or CDNSKEY RRset that authenticates against the existing DS RRset — because the existing DS RRset, by design, references all participating operators' keys. A malicious or compromised operator in such a setup can therefore push a DS update that removes the other operators' keys, effectively performing a hostile takeover of the delegation. The mechanism does not address this case in its base form, and the cross-nameserver consistency check examined under Branch S4 is the principal mitigation. The finding for this branch under multi-signer conditions is therefore conditional on the deployment of Branch S4 mitigations and is recorded as a residual risk that must be accepted, mitigated, or excluded by deployment policy.

The **full child operator adversary** is by construction able to mutate the DS RRset through the mechanism, because the mechanism's authentication test cannot distinguish a legitimate Child DNS operator from an adversary with full control of the legitimate operator's signing infrastructure. This is not a failure of the mechanism but a property of any delegation-of-authority model; once the registrant has delegated DNSSEC operation to an operator, the operator's authority over the DS RRset is, by design, indistinguishable from the registrant's. The mechanism's response to this scenario is the registrar lock and registry lock interaction examined under Branch S3, and the manual override path preserved by RFC 8078 section 4 and examined under cross-cutting Concern X1. The finding for this branch against the full child operator adversary is that the mechanism behaves as designed; the security property is provided by the surrounding controls, not by RFC 7344 itself.

4.3 Branch S2 — Authentication of bootstrapping

This branch addresses the same RSEP section 1.2(A) concern but for the insecure-to-secure transition per RFC 9615. The question is sharper because no chain of trust yet connects the parent to the child, and the construction by which RFC 9615 establishes one for the first time is the entire substance of the assessment for this branch.

RFC 9615 derives authentication from the chain of trust covering the authoritative nameserver hostnames. Concretely, for each nameserver listed in the child's delegation, the Child DNS operator publishes the bootstrapping signal under a name derived from both the child zone and the nameserver hostname, in a Signaling Zone served by that nameserver and signed by keys whose DS records are anchored in the nameserver hostname's parent zone. The parent validates the signal against this chain

of trust before publishing the new DS RRset for the child. Against the **external network adversary**, the construction provides cryptographic authentication equivalent to that of any DNSSEC-validated query, and the branch yields a positive finding for that adversary class.

The **nameserver-position adversary** is more interesting in the bootstrapping context than in the maintenance context. Because RFC 9975 requires cross-nameserver consistency across the full authoritative set for CDS, CDNSKEY, and CSYNC operations — an adversary controlling fewer than all nameservers cannot induce the parent to publish a DS RRset of their choosing. The all-nameservers requirement is a deliberate defense against this adversary class and yields a positive finding when the parent enforces it. Implementations that relax the requirement — for example, by accepting the signal from any one nameserver — degrade the security posture of this branch and constitute an implementation defect with security consequences under RSEP section 1.2(A).

The bootstrapping construction has a precondition that warrants explicit examination: the nameserver hostnames themselves must lie in DNSSEC-signed zones whose chain of trust is anchored in the root. Where this precondition fails — most commonly because the nameserver hostnames are in an unsigned TLD or an unsigned parent zone — RFC 9615 cannot establish authentication, and the mechanism falls back to either out-of-band methods or the registrant-centric process. The fallback is not within the scope of Recommendation 1 and is not assessed here, but the assessment must record the precondition as a deployment constraint: bootstrapping under Recommendation 1 is feasible only for delegations whose nameserver hostnames are themselves under DNSSEC.

The **full child operator adversary** in the bootstrapping context is again constructive rather than adversarial in the protocol sense; the mechanism authenticates the operator that controls the nameservers, and if that operator is in fact the adversary, the registrant has a separate problem that the bootstrapping mechanism cannot solve. The relevant defense is the registrant's choice of nameservers in the first place, which is outside the mechanism's scope.

4.4 Branch S3 — Interaction with EPP registration locks

This branch addresses a different category of unauthorized mutation: changes that are properly authenticated by the cryptographic mechanism but that the registrant has explicitly indicated should not be applied. RSEP section 1.2(A) covers any unauthorized alteration of registry data, and a registrant who has set a server-side update lock has communicated through the registration system that the registry data should not be mutated, regardless of cryptographic authentication of the request.

SAC126 Appendix B.3 provides a detailed analysis of the lock interaction question and concludes that automated DS maintenance should be suspended when a *serverUpdateProhibited* lock is in place but should proceed when only a *clientUpdateProhibited* lock is in place. The assessment adopts this distinction. However, the IETF has reached a different conclusion on this point. Draft-ietf-dnsop-ds-automation section 6.1 Recommendation 2, with detailed rationale in section 6.2.1, holds that when DS automation is performed by the registry, automation **MUST NOT** be suspended on the basis of a registry update lock (*serverUpdateProhibited*) alone (that is, some other signal is needed for automation to be suspended). The draft's analysis rests on a narrower reading of EPP semantics: RFC 5731 section 2.3 specifies that *serverUpdateProhibited* demands only that the registrar's "[r]equests to update the object [...] **MUST** be rejected", and the same section explicitly permits an EPP server (registry) to override status values set by an EPP client (registrar) subject to local server policies. Therefore, registry-side

automation is not formally constrained by the lock, and the security model gain from suspension is illusory because the lock does not protect against actions performed by the registry itself in any case. The draft's position would not in itself produce a section 1.2 finding, provided that the registry provides a mechanism for registrars to opt out of registry-led DS automation, for their entire portfolio, and optionally for individual domain names (e.g. via an EPP extension). For the avoidance of doubt, if a registrar opts a domain out of DS automation, only DS information provided by the registrar via the SRS should be published. SSAC may wish to publish a document to update SAC126 Appendix B in the light of the draft's analysis following its publication as an RFC.

The *clientUpdateProhibited* case is more nuanced. As SAC126 observes, this lock is removable by the registrar without out-of-band verification and therefore does not constitute a meaningful defense against actions taken by the registrar or the registry themselves. Treating it as a barrier to DS automation would prevent legitimate ongoing maintenance — algorithm rollovers in particular — without providing a corresponding security benefit and would force registrants to choose between operational maintainability and a lock that has unrelated protective purposes. The assessment finds that proceeding with automation under *clientUpdateProhibited* is consistent with both RSEP section 1.2 (because the lock does not represent authorization withdrawal in this case) and with operational reality, provided the parent continues to perform the validation pre-checks described in Section 3.6.

A separate concern under this branch is the preservation of manual recovery paths. Where authentication credentials at the Child DNS operator are destroyed, lost, or compromised — and the operator can no longer publish authenticated CDS or CDNSKEY records — the registrant must retain the ability to mutate the DS RRset through the registrar's manual interface. SAC126 explicitly warns against disabling manual DS management interfaces, and the assessment treats the preservation of this interface as a precondition for a positive finding under this branch. A deployment that automates DS maintenance but removes the manual interface fails under both RSEP section 1.2 — because the registrant cannot withdraw authorization from a compromised operator — and under RSEP section 1.3 — because the registrant cannot recover from a stuck-state condition that affects resolution coherence. It may be noted that [Section 8 of the Registrar Accreditation Agreement](#) obliges registrars to provide a manual interface at all times, which addresses this concern.

4.5 Branch S4 — Cross-nameserver consistency

This branch addresses the **nameserver-position adversary** class identified under Branch S1 and the related concern of negligent multi-signer operation. The question is whether the parent, when collecting CDS or CDNSKEY records from the child, queries all authoritative nameservers and verifies consistency before acting, or queries any single nameserver and acts on whatever response it receives.

RFCs 7344 and 9615 do not, in their base specifications, require multi-nameserver consistency checks. This gap, referenced in SAC126 section 7 and Appendix B.5, has since been closed by [RFC 9975](#) (*Clarifications on CDS/CDNSKEY and CSYNC Consistency*), which updates RFCs 7344 and 7477 to mandate that parent-side entities ensure consistency across all the child's authoritative nameservers before acting on CDS, CDNSKEY, or CSYNC records. RFC 9975 is now the authoritative published source for the substance of HR-2 in this assessment. The assessment finds that the absence of cross-nameserver consistency checks introduces a security exposure under RSEP section 1.2(A) in two scenarios: nameserver compromise where the adversary controls a subset of the authoritative set, and multi-signer setups where one participating operator publishes CDS or CDNSKEY records inconsistent

with the others. In both scenarios, a parent that queries any single nameserver may receive and act upon a signal that does not represent the legitimate intent of the delegation as a whole.

A deployment that implements cross-nameserver consistency checks per RFC 9975 — collecting the signal from each authoritative nameserver and aborting the operation on any discrepancy — yields a positive finding under this branch. A deployment that does not implement such checks is non-conformant with RFC 9975 and is therefore a section 1.3(A) finding under RSEP, in addition to constituting a residual security risk under section 1.2(A). Operators must either implement the RFC 9975 consistency requirement or document a specific deployment policy that justifies departing from it, such as restricting CDS automation to single-nameserver delegations (which is not generally a desirable operational posture).

4.6 Branch S5 — Multi-channel collision behavior

This branch addresses the case where both the registry and the registrar have deployed CDS automation against the same delegation, and the question is whether their concurrent operation introduces conditions that violate RSEP section 1.2 or section 1.3. The branch sits at the boundary between security and stability because the relevant failure modes include both unintended DS state changes (security) and inconsistency between registry and registrar views of DS state (stability).

The base case is benign. If both the registry and the registrar correctly read the same CDS or CDNSKEY records, validate them under the same chain of trust, and produce the same DS RRset, the resulting submissions are idempotent: whichever submission the registry processes second produces no observable state change. The collision becomes consequential only when the two parties see different views of the child's signal — for example, because their scans occur at different points within the propagation window of a child-side update, or because they query different nameservers in a setup with not-yet-converged state. In such cases, the parent zone's DS RRset may flap between two valid states until the child's state stabilizes, and the next scan converges both observers on the same signal.

The assessment finds that collision-induced flapping does not in itself violate RSEP section 1.2 — because each flap is, by construction, an authenticated and validated change — and the stability impact under RSEP section 1.3 is bounded: each interim state is a state that would individually pass the parent's validation pre-checks, so DNSSEC resolution remains functional throughout. The duration of the flap is bounded by the convergence time of the child's authoritative nameserver replication, typically within a single resource record TTL.

The mitigation, where collision behavior is judged unacceptable, is for the registry and the registrar to agree by deployment policy that only one party will perform CDS automation for any given delegation. SAC126 section 4.4 and Appendix B.2 discuss this choice in operational terms; the assessment finds that either choice is consistent with RSEP provided it is consistently applied and clearly documented in the operator's design. A deployment in which both parties run automation without coordination is recorded with a finding of acceptable but suboptimal stability posture, with a recommendation to formalize the coordination model by means of the opt-out mechanism described in Section 4.4 above.

4.7 Summary of security pillar findings

The five security branches yield the following composite assessment.

- Against the **external network adversary**, the mechanism — when deployed in conformance with RFCs 7344 and 9615 and with the parent-side validation pre-checks of RFC 7344 section 4.1 — provides authentication of DS changes to a standard equivalent to that of DNSSEC-validated DNS data and yields positive findings under RSEP section 1.2 for both the maintenance and bootstrapping transitions.
- Against the **nameserver-position adversary**, positive findings are conditional on the deployment of cross-nameserver consistency checks under Branch S4; in the absence of such checks, a residual security risk exists that is most acute in multi-signer setups per RFC 8901 model 2.
- Against the **full child operator adversary**, the mechanism behaves as designed and the relevant security properties are provided by the surrounding controls examined under Branch S3. The lock interaction findings under Branch S3 require that the deployment may proceed under both *clientUpdateProhibited* and *serverUpdateProhibited* and must preserve a manual recovery path. The collision findings under Branch S5 yield no negative finding but recommend coordination between registry and registrar to bound interim-state duration.

Three findings are recorded as conditional or residual and warrant explicit documentation in any RSEP submission: the cross-nameserver consistency exposure (S4), the multi-signer takeover scenario (S1 under multi-signer deployment), and the precondition that nameserver hostnames be DNSSEC-signed for bootstrapping (S2). None of the three rises, in the assessment's view, to the level of a significant security issue under RSEP section 1.2 when accompanied by appropriate deployment-time mitigations, but all three should be disclosed and addressed in the operator's design documentation.

5 Stability Pillar Findings

This section assesses the mechanism against RSEP section 1.3, which defines a stability effect as either non-compliance with authoritative published standards or the creation of a condition that adversely affects the throughput, response time, consistency, or coherence of responses to Internet servers or end systems operating per applicable standards. The two clauses are addressed in turn: Branch T1 covers standards conformance under section 1.3(A), and Branches T2 through T5 cover the operational stability properties under section 1.3(B). The five branches are not orthogonal — a standards conformance finding may have operational consequences, and vice versa — but the assessment maintains the distinction because the two clauses produce different remediation pathways under RSEP.

5.1 Branch T1 — Standards conformance

Recommendation 1 cites three RFCs: RFC 7344, RFC 8078, and RFC 9615. All three are on the IETF standards track and meet the RSEP section 1.3(A) test of being authoritative and published by a well-established, recognized and authoritative standards body. A deployment that conforms to the MUST and SHOULD clauses of these three RFCs satisfies the standards-conformance prong of the stability definition for the substantive scope of the mechanism.

The assessment must, however, examine three categories of conformance questions that the citation alone does not resolve. The first is the handling of MUST clauses whose plain-text reading admits more than one implementation: for example, RFC 7344 section 3 requires that the parent verify that the new DS RRset would not break delegation but does not specify the exact set of checks that constitute that verification. An implementation that performs only the minimum checks described in Section 3.6 of this assessment is conformant, as is one that performs additional checks; an implementation that performs no verification at all is not conformant and constitutes a section 1.3(A) finding. The assessment treats the minimum-check posture as the conformance floor and any additional checks as defensive depth.

The second category is the handling of SHOULD clauses. SHOULD-level requirements admit deviation where the implementer has documented reasons. For RFC 7344 and RFC 9615, the relevant SHOULD clauses concern the frequency of polling, the handling of failed validations, and the treatment of records returned with inconsistent contents from different nameservers. SAC126 Appendix B treats most of these as design choices to be documented by the operator. The assessment adopts this posture: a deployment that deviates from a SHOULD with documented and operationally sound reasoning is conformant under section 1.3(A), and the documentation requirement is itself the assessment criterion. Undocumented deviation is recorded as a finding of insufficiency.

The third category is the relationship to standards work that has progressed since SAC126 was published. SAC126 referenced two IETF drafts as in-progress: Generalized DNS Notifications and CDS/CDNSKEY consistency. Both have since been published as standards-track RFCs. Generalized DNS Notifications is now [RFC 9859](#); CDS/CDNSKEY consistency is now [RFC 9975](#) (*Clarifications on CDS/CDNSKEY and CSYNC Consistency*), which updates RFCs 7344 and 7477. RFC 9859 specifies a notification-triggered variant that is outside Recommendation 1's scope (see Section 3.8); its publication does not extend the conformance scope of this assessment but does make the standard available as a practical option for operators wishing to improve the response-time properties of their CDS/CDNSKEY deployment. RFC 9975 is squarely within the conformance scope of this assessment: it updates RFC 7344 — one of the three RFCs Recommendation 1 cites — and adds a normative MUST-

level requirement that parent-side entities ensure cross-nameserver consistency before acting on CDS or CDNSKEY records. A deployment of the CDS/CDNSKEY mechanism that does not enforce cross-nameserver consistency is now non-conformant with a published standards-track RFC that updates a cited RFC and therefore produces a section 1.3(A) finding under RSEP.

A third piece of in-progress IETF work, draft-ietf-dnsop-ds-automation, did not exist as a working group document when SAC126 was published but has since become the IETF's direct operational follow-on to SAC126 Appendix B. The draft is intended as Best Current Practice, was in IESG Evaluation at the time of this assessment and explicitly derives its core issues from SAC126 section 4.4. Its recommendations span acceptance checks, reporting and transparency, registration locks, and multiple submitting parties — substantially the same operational territory addressed by Branches T2, T3, S3, S5, and Concerns X1 and X2 of this assessment. On publication, the draft will become an authoritative published standard under section 1.3(A) and may serve as a conformance benchmark for operators implementing CDS/CDNSKEY automation. Where its substance materially informs an assessment branch, it is referenced in the relevant section.

Three operational standards from the IETF REGEXT Working Group are referenced in SAC126 Appendix B and bear on conformance: the EPP TTL extension, the EPP Change Poll extension, and various error reporting drafts. Two of the three were drafts at the time SAC126 was published. The EPP TTL extension has since been published as [RFC 9803](#) (June 2025) and is now an authoritative published standard. The EPP Change Poll extension (RFC 8590) was already published. Both are addressed under their respective branches of this assessment: RFC 8590 conformance is treated as required where the deployment uses EPP for registry-to-registrar notification of automation-driven DS changes; RFC 9803 is treated as the now-available standard for asymmetric TTL handling under Branch T3. The remaining error reporting drafts in REGEXT are not yet published as RFCs and are treated as operational considerations under the cross-cutting reporting Concern X2.

The assessment finds that the standards-conformance posture under section 1.3(A) is straightforward: the cited RFCs are authoritative, their conformance criteria are checkable against design documentation, and the deployment's posture against both MUST and SHOULD clauses can be definitively recorded. Conditional findings under this branch are limited to undocumented SHOULD deviations and to the minimum-validation question under RFC 7344 section 3, both of which are resolvable through the operator's own design documentation.

5.2 Branch T2 — Validation pre-checks and resolution coherence

This branch addresses RSEP section 1.3(B) — the prevention of conditions that adversely affect the coherence of DNS responses — through the mechanism's validation pre-checks. The question is whether the parent, when ingesting a CDS or CDNSKEY signal, performs sufficient verification that the resulting DS RRset will not produce a delegation that fails DNSSEC validation at resolvers operating per RFC 4035 and RFC 6840.

The minimum validation set has been identified in Section 3.6 of this assessment: for each DS record to be published, a corresponding DNSKEY must exist in the child zone with matching algorithm and key tag, and the child's DNSKEY RRset must be signed by at least one key referenced by the new DS RRset. A deployment performing this minimum set will not, under any signal accepted from the child, produce a parent-side DS RRset whose corresponding DNSKEY is unreachable or unauthorized in the child. This is the floor required to avoid creating a coherence-breaking condition under section 1.3(B), and the

assessment treats it as a hard requirement; a positive finding under this branch is conditional on the deployment performing this minimum set.

Three additional checks merit explicit examination because their omission has been observed in practice to cause stability incidents. The first is the digest type check: a CDS RRset using a digest type that is no longer recommended under RFC 9904 — currently SHA-1 (deprecated by [RFC 9905](#)) and GOST R 34.11-94 (deprecated by [RFC 9906](#)), marked MUST NOT — should be rejected, because resolvers conformant to current DNSSEC algorithm guidance may treat the DS RRset as insecure. A deployment that accepts SHA-1 or GOST R 34.11-94 CDS records produces a coherence risk under section 1.3(B) for the validating resolver population, even though the CDS record is technically valid under RFC 7344. The assessment records the digest type check as a recommended addition to the minimum validation set.

The second is the algorithm check: a DNSKEY using a signing algorithm marked MUST NOT IMPLEMENT under RFC 9904 should not be the basis for a published DS record, for the same coherence reason. The third is the consistency check between CDS and CDNSKEY when both are published, addressed in SAC126 Appendix B.5: a deployment that accepts a CDS record without verifying that a corresponding CDNSKEY (when published) is consistent with it admits a class of misconfiguration in which the published DS RRset references a key that the child does not in fact serve.

The assessment finds that a positive finding under T2 requires the minimum validation set as a hard floor, and considers the three additional checks (digest type, algorithm, CDS/CDNSKEY consistency) as strongly recommended for any deployment seeking to maintain a robust coherence posture. The validation pre-check requirement is the principal defense against the legitimate-but-misconfigured Child DNS operator scenario, which is the most common cause of self-inflicted DNSSEC outages observed in operational practice.

5.3 Branch T3 — TTL handling and rollback timing

This branch addresses section 1.3(B) through the lens of operational recovery: when a DS publication produces an unintended outcome — whether through legitimate Child DNS operator error, an as-yet-undetected attack, or an implementation defect — how quickly can the parent zone return to a known-good state, and what coherence and response-time properties does the recovery window have?

The dominant variable is the TTL of the DS RRset in the parent zone. A long TTL — in the range historically common for parent-side delegation data, hours to a day — means that a misconfigured DS RRset, once published, will be cached by validating resolvers for the full TTL before any correction can take effect. During that window, validation failures at the child are observable to end users as resolution failures, even after the parent has corrected the published DS RRset. A short TTL — in the 5-to-15-minute range examined in operational deployments cited in SAC126 Appendix B.1 — bounds the resolution-failure window to that shorter period at the cost of additional query load against the parent's authoritative nameservers.

SAC126 Appendix B.1 reports that real-world deployments at .goog with permanent DS TTLs of 15 minutes, and experimental work at ISC with 5-minute TTLs, observe negligible impact on the parent's authoritative nameserver load. The assessment treats this as operational evidence supporting the feasibility of short DS TTLs for the bounded purpose of post-update rollback windows. The recommended posture, consistent with SAC126's analysis, is asymmetric TTL handling: a long normal TTL for steady-state operation, reduced to a short TTL at the time of any DS update and restored to the normal value after a duration sufficient to ensure the previous DS RRset has expired from caches.

The assessment makes three observations specific to the section 1.3(B) frame.

1. The asymmetric posture introduces no coherence risk at the resolver because each cached state is self-consistent; the cost is a brief period of additional query load following each DS change.
2. The posture requires a mechanism for adjusting the DS RRset's TTL on demand. RFC 9803 provides the standardized EPP-based mechanism for this purpose, allowing registrars to convey TTL changes to the registry through a published standards-track extension; alternatively, the registry may use an internal mechanism. Absent either, the registry must choose a single TTL that balances steady-state load against rollback latency.
3. The deployment of permanent short TTLs — as practiced by .goog — is a stronger section 1.3(B) posture than asymmetric TTLs because it bounds the rollback window at all times, not only following a deliberate update; the operational evidence suggests that the additional query load is acceptable in practice.

The assessment finds no negative section 1.3(B) effect from the use of short DS TTLs in the ranges examined by operational practice. The branch yields a positive finding for any deployment that documents its TTL handling posture and provides a rollback path with a bounded latency under operational conditions.

5.4 Branch T4 — Scanning load and cadence

This branch addresses section 1.3(B) under the throughput and response-time clauses. The mechanism, in its pull-based form, requires the parent — registry or registrar — to query each child zone under management for CDS and CDNSKEY records on a recurring cadence, because the mechanism provides no other signal that a child wishes to update its DS RRset. The question is whether the resulting query load creates a condition adverse to throughput or response time at either the parent's scanning infrastructure or the children's authoritative nameservers.

SAC126 Appendix B and the SSAC's referenced operational survey indicate that daily scanning of zones on the order of several hundred million domains is currently feasible on a small number of scanning machines, and that the resulting query load is small relative to the children's normal query volume. The throughput dimension under section 1.3(B) is therefore not a finding of concern at present scale: the mechanism, deployed as described in current operational practice, does not produce a throughput-degrading condition at either end.

The response-time dimension is more nuanced. The cadence of scanning sets a lower bound on the latency between a Child DNS operator publishing a CDS or CDNSKEY signal and the parent acting upon it. A daily cadence implies a worst-case latency of approximately one day, with a mean of approximately twelve hours. For ongoing maintenance — algorithm rollovers, key rotations, scheduled changes — this latency is operationally acceptable because the Child DNS operator can plan around it; for emergency response — recovering from a key compromise, withdrawing a misconfigured DS RRset — the latency is operationally significant. The assessment finds that a daily-cadence pull-based deployment is consistent with section 1.3(B) for ongoing maintenance but introduces a latency property for emergency response that operators should disclose to registrants and Child DNS operators.

The hybrid pull-with-notification approach, contemplated as in-progress draft work at the time SAC126 was published, has since been standardized as RFC 9859. The mechanism allows the Child DNS operator to signal the parent that a CDS or CDNSKEY scan is desired, substantially reducing the

response-time latency relative to a pure-polling deployment. RFC 9859 introduces a new DSYNC resource record type for notification endpoint discovery, with the initial registry entries explicitly covering parental NS and DS record updates including DNSSEC bootstrapping. RFC 9859 specifies a different mechanism than the one in scope under Recommendation 1 — the recommendation cites RFCs 7344, 8078, and 9615 — and so its publication does not change the conformance frame of this assessment. It is, however, directly available to operators wishing to improve their section 1.3(B) response-time posture. A deployment that pairs CDS/CDNSKEY automation under Recommendation 1 with RFC 9859 notifications to trigger immediate scanning addresses the emergency-response latency disclosed under DR-2 without requiring any departure from Recommendation 1's substantive scope. This assessment treats the availability of RFC 9859 as a positive development that operators are recommended to consider, while noting that its adoption is not required under Recommendation 1, and for gTLD registry operators, would be considered a new Registry Service, with the corresponding requirement for the operator to submit an RSEP request.

A specific operational concern under this branch is the registry-side scanning configuration relative to the registrar-side: where both perform automation, the combined query load against children is doubled relative to a single-party deployment. SAC126 Appendix B.2 treats this as an operational matter rather than a stability finding, and the assessment concurs. The combined load remains small in absolute terms and does not create a section 1.3(B) condition; the recommendation to coordinate between registry and registrar made under Branch S5 of the security pillar applies here for the secondary purpose of avoiding redundant work.

The assessment finds that the scanning load and cadence properties of the mechanism, as deployed in current operational practice, do not create conditions adverse to throughput or response time within the meaning of section 1.3(B), with the noted disclosure requirement around emergency response latency.

5.5 Branch T5 — Multi-signer setups per RFC 8901

This branch addresses the most substantive open question in the section 1.3(B) analysis: the handling of CDS and CDNSKEY records in delegations operated by multiple DNS providers under RFC 8901 model 2, where each provider holds its own signing keys. SAC126 section 7 and Appendix B.5 explicitly note that no standards-defined process exists for handling CDS or CDNSKEY records in this configuration, and that practical deployment experience is currently limited.

The structural problem is straightforward to state. In RFC 8901 model 2, the delegation's DS RRset contains entries for keys held by each participating provider, so that a response signed by any provider's keys can be validated through the chain of trust. A CDS or CDNSKEY signal published by one provider, in isolation, represents only that provider's view of the desired DS RRset; if applied directly by the parent, it would replace the multi-provider DS RRset with one that references only the publishing provider's keys, causing responses signed by the other providers to fail validation. The mechanism's base form has no construct for distinguishing *this is my contribution to the joint DS RRset* from *this is the complete intended DS RRset*, and the consequence under naive implementation is that any provider's CDS publication can effectively expel the others from the delegation.

SAC126 Appendix B.5 recommends a conservative posture: each participating provider should publish, at its own apex, the union of all participating providers' CDS and CDNSKEY entries, and the parent should require consistency across the responses received from all authoritative nameservers before acting. Where this posture is in place, a single provider acting unilaterally cannot succeed in expelling

the others, because the inconsistency between its publication and the others' will be detected and the operation aborted. Where the posture is not in place — either because the providers have not coordinated their CDS publications or because the parent does not enforce cross-nameserver consistency — the mechanism creates a section 1.3(B) coherence risk that is not present in single-provider deployments.

The assessment makes three findings under this branch.

1. The absence of standardized handling for multi-signer CDS/CDNSKEY processing is a known gap, acknowledged by SAC126 itself, that has been substantially closed by the publication of RFC 9975. A deployment supporting multi-signer delegations without enforcing the cross-nameserver consistency requirements of RFC 9975 is now a section 1.3(A) non-conformance finding, not merely a scope limitation. RFC 9975 does not, however, fully specify multi-signer operation under RFC 8901 model 2 — questions of CDS RRset composition (whether each participating provider publishes the union of all participating providers' entries or only its own) remain operator-specific.
2. Deployments that admit multi-signer delegations under CDS automation without enforcing cross-nameserver consistency present a section 1.3(B) coherence risk that should be disclosed in the operator's design documentation and either mitigated through the consistency posture described above or excluded by deployment policy.
3. Implementation of the cross-nameserver consistency check per RFC 9975 is now mandatory under that published standards-track RFC; doing so resolves both this branch's concern and the related Branch S4 concern under the security pillar.

5.6 Summary of stability pillar findings

The five stability branches yield the following composite assessment. Under section 1.3(A), the mechanism rests on three standards-track RFCs whose conformance is checkable against design documentation; the conformance posture is generally strong, with conditional findings limited to undocumented SHOULD deviations and to the minimum-validation question per RFC 7344 section 3. Under section 1.3(B), the mechanism's validation pre-checks under T2 are the principal defense against coherence-breaking conditions and constitute the assessment's hard floor; TTL handling under T3 is a positive finding when documented; scanning load under T4 is a positive finding at present operational scale, with a disclosure requirement around emergency response latency; and multi-signer handling under T5 is the principal residual finding, requiring either active mitigation through cross-nameserver consistency checks or exclusion by deployment policy.

Consistent with the security pillar findings in Section 4, three operational practices emerge as common requirements across branches: enforcement of the minimum validation pre-check set, cross-nameserver consistency checks for both CDS and CDNSKEY records, and documentation of the operator's choices on TTL handling, scanning cadence, and multi-signer policy. None of the assessed branches yields a finding that the mechanism, properly deployed, creates a significant stability issue under RSEP section 1.3. Three branches yield conditional findings — T1 (undocumented SHOULD deviation), T2 (minimum validation set as hard floor), and T5 (multi-signer handling) — each of which is resolvable through deployment-time controls and design documentation.

6 Cross-Cutting Concerns

This section addresses two concerns that do not fit cleanly within either the security or stability pillar but bear on both. The first is the integrity of the going-insecure transition per RFC 8078 section 4, which is simultaneously a potential attack surface (security) and a required functional capability whose failure produces a stuck-state condition (stability). The second is the absence of standardized reporting for failed or rejected DS automation events, which affects both the registrant's ability to detect unauthorized changes (security) and the operator's ability to maintain coherent operation in the face of error conditions (stability). Because these concerns engage both pillars simultaneously, they are treated separately rather than apportioned between Sections 4 and 5.

6.1 Concern X1 — The going-insecure path

The mechanism's going-insecure transition, specified in RFC 8078 section 4, allows a Child DNS operator to signal that a delegation should cease to be secured by DNSSEC. The signal takes the form of a CDS RRset containing a single record with algorithm 0, which the parent interprets as an instruction to remove the existing DS RRset entirely. On detection and validation of this signal, the parent removes the DS RRset, after which the delegation operates as an insecure delegation in the global DNS.

This transition is operationally necessary. A registrant migrating between DNS providers may need to disable DNSSEC during the transition window if the destination provider does not support signed transfers, and SAC126 section 2.2 lists this scenario explicitly as a routine reason for going insecure. A registrant whose Child DNS operator has lost or destroyed the keys authorized by the existing DS RRset has no other path to recover ongoing operation of their domain through the automated mechanism, because no new authenticated CDS or CDNSKEY signal can be produced. A registrant who has decided, for reasons of operational simplicity or risk preference, to operate without DNSSEC must be able to act on that decision through the same channels by which they enabled DNSSEC in the first place. The going-insecure transition is therefore a load-bearing component of the mechanism, not an optional adjunct.

Under RSEP section 1.2, the going-insecure transition presents a focused security question: does the mechanism prevent an unauthorized party from removing the DS RRset, and thereby downgrading the delegation from secure to insecure? The authentication construction is identical to that of an ordinary DS update per RFC 7344 — the algorithm-zero CDS record is itself signed by the existing chain of trust — so an external network adversary cannot forge such a signal without access to the child's existing keys. The nameserver-position adversary in a multi-signer setup can produce the signal if their key is referenced by the existing DS RRset, which is the same exposure identified under Branch S1, and the same cross-nameserver consistency mitigation under Branch S4 applies. The full child operator adversary can, by construction, perform the transition; this is a property of any delegation-of-authority model and is not specific to the going-insecure path.

The assessment finds that the going-insecure transition's authentication posture under RSEP section 1.2 is equivalent to that of any other DS update, and that the security findings made under Branches S1, S2, and S4 apply unchanged. There is no additional security exposure introduced specifically by the going-insecure path beyond what is already addressed in those branches.

Under RSEP section 1.3, however, the going-insecure transition presents a distinct stability concern: the consequences of the parent failing to honor a legitimate going-insecure signal. If the parent cannot or will not remove the DS RRset on receipt of an authenticated algorithm-zero signal, the registrant is left with a delegation that continues to require successful DNSSEC validation, but whose Child DNS operator has indicated — through the standardized mechanism — that DNSSEC operation is no longer intended. The delegation may rapidly become inoperable as the Child DNS operator ceases to maintain the signing infrastructure. This is a coherence-breaking condition under section 1.3(B): responses to validating resolvers will fail, and the registrant has no automated path to recover.

A second stability concern arises from the interaction between the going-insecure path and the parent's routine validation pre-checks. Under RFC 7344 section 4.1 and the minimum validation set examined under Branch T2, the parent verifies that the resulting DS RRset will not break resolution. For an algorithm-zero CDS signal, the resulting DS RRset is empty, and the routine pre-check is structurally inapplicable: there are no DS records to validate against the child's DNSKEY contents. A naive implementation that applies the validation pre-check to algorithm-zero signals without recognizing their special semantics may incorrectly reject the signal, producing exactly the failure mode described above. The assessment treats correct handling of the algorithm-zero special case as a hard requirement: a deployment that fails to honor authenticated algorithm-zero CDS signals fails under section 1.3(B) for the same reason it would fail under section 1.2 if it failed to apply other authenticated changes.

A third concern bears on the interaction between going-insecure and registration locks. The *clientUpdateProhibited* and *serverUpdateProhibited* cases under Branch S3 were found not to require suspension of routine maintenance because these locks, given the parallel ability for the registrar to opt out of registry-led DS automation, do not represent meaningful authorization withdrawal; the same analysis applies to going-insecure transitions under that lock.

A fourth concern, closely related, is the preservation of the manual going-insecure path. Where the automated path is unavailable — the lock is in place, the Child DNS operator's keys are destroyed, the operator is uncooperative in a multi-provider transition, or the registrant has affirmatively chosen to act outside the automation — the registrant must retain the ability to remove the DS RRset through the registrar's manual interface. SAC126 explicitly warns against disabling manual DS management interfaces, and Branch S3 of this assessment treats the preservation of the manual interface as a precondition for a positive finding. The same precondition applies to going-insecure: a deployment that automates DS maintenance but removes the manual going-insecure path fails both pillars, because it produces stuck-state conditions that the registrant cannot resolve through any available channel.

The assessment finds that the going-insecure path is correctly handled in any deployment that meets four requirements. The deployment must recognize and honor authenticated algorithm-zero CDS signals according to RFC 8078 section 4. The deployment must apply the same cross-nameserver consistency checks to going-insecure signals as to other automated DS changes, addressing the multi-signer concern identified under Branches S1 and S4. The deployment must honor a registrar opt-out from DS automation in the going-insecure context, on the same basis as for ongoing maintenance under Branch S3. The deployment must preserve a manual going-insecure path through the registrar's standard interfaces. A deployment meeting these four requirements yields a positive finding under both RSEP pillars for the going-insecure transition; a deployment failing any of the four presents a finding that should be disclosed and resolved before the mechanism is approved for deployment.

6.2 Concern X2 — Reporting and observability

The cited RFCs do not specify a normative reporting mechanism for events arising during automated DS maintenance. RFC 7344 and RFC 9615 describe what the parent does; they do not require the parent to communicate the outcome to the Child DNS operator, the registrant, the registrar (in the registry-led topology), or any other party. SAC126 Appendix B.4 discusses this gap in operational terms, observes that early reporting of rejected updates would help Child DNS operators handle conditions in a timely manner, and proposes a categorization of report-worthy events. The proposal is not yet a published standard; it represents the SSAC's view of operational best practice, which is the appropriate input for this branch but does not have the standing of a section 1.3(A) conformance benchmark.

The reporting gap engages both RSEP pillars in distinct ways. Under section 1.2, the absence of reporting affects the registrant's and the Child DNS operator's ability to detect that an unauthorized event has occurred. If a CDS automation system silently rejects a legitimate signal — perhaps because the validation pre-checks fail, perhaps because of a registration lock the registrant did not know was in place — the legitimate party may not be aware that their action did not take effect. If the silent rejection is the result of a successful attack on the underlying signaling infrastructure, the absence of reporting prolongs the window in which the attack remains undetected. Conversely, if a successful but unintended event has occurred — a misconfigured Child DNS operator publishes a CDS RRset that the parent accepts — the registrant may not be aware that their delegation has changed until end users encounter resolution failures. In both cases, the section 1.2 concern is not that the mechanism produces unauthorized events, but that its observability properties limit the registrant's ability to detect events the registrant did not intend.

Under section 1.3(B), the absence of reporting affects the operator's ability to maintain coherent operation through error conditions. SAC126 Appendix B.4 lists four categories of report-worthy events:

1. successful provisioning (whether manual, automated bootstrapping, or first-time automation after a manual change),
2. successful removal (whether manual or via the algorithm-zero path),
3. pending updates that cannot be applied (because validation fails or a lock prevents application), and
4. detected drift between the published DS RRset and the child's current DNSSEC state (where no update has been requested but the existing DS RRset has become incompatible with the child's keys).

Each of these categories represents a state in which operator action — by the registrant, the Child DNS operator, the registrar, or the registry itself — may be required to restore coherent operation. Without reporting, such action cannot be taken in a timely manner, and the duration of any incoherent state is bounded only by the time until end users notice and report failures, which is a poor operational property.

The drift-detection case under category four merits separate examination. A delegation may operate correctly for an extended period after its DS RRset is published, only to enter a failure state when the Child DNS operator changes the child's DNSKEY contents in a way that is no longer compatible with the existing DS RRset. SAC126 section B.4 cites the example of an algorithm rollover in which the old algorithm is removed from the child while the DS RRset still references a key with that algorithm. The mechanism's normal flow does not detect this condition because no CDS or CDNSKEY signal is being

processed at the time of failure; the failure is the absence of a signal that should have been published. A deployment that performs occasional drift checks — re-validating the published DS RRset against the child's current contents and reporting discrepancies — addresses a class of stability incident that no other branch of this assessment captures. The assessment treats drift-detection reporting as a recommended enhancement under section 1.3(B), not a hard requirement.

The reporting question has a multi-recipient dimension that the assessment must engage. SAC126 Appendix B.4 identifies three categories of recipient:

1. the Child DNS operator (via the report channel proposed in RFC 9859, or via email derived from SOA contact data),
2. the registrar (in the registry-led topology, via EPP under RFC 8590 or an equivalent), and
3. the registrant or technical contact (via email).

The recipients have different operational roles — the Child DNS operator can correct a misconfigured signal, the registrar can adjust the registration's lock state, the registrant can authorize out-of-band action — and the appropriate recipient depends on the event being reported. A deployment that reports only to one recipient category, or that conflates the categories, may fail to engage the party best positioned to act on the report.

The IETF draft-ietf-dnsop-ds-automation section 5 articulates the same reporting territory in normative SHOULD form. Specifically, the draft recommends: notification of relevant points of contact for DS updates and DS deactivation (Recommendation 5.1.1); notifying the Child DNS operator and technical contact first for error conditions, with the registrant notified only if the problem persists for a prolonged period such as three days (Recommendation 5.1.2); use of [RFC 9567](#) report queries to the agent domain per RFC 9859 section 4 as the standardized channel to the Child DNS operator (Recommendation 5.1.3); RFC 8590 EPP notification from registry to registrar (Recommendation 5.1.4); and registrant access to the current DS configuration through the customer portal, with DS update history optionally provided (Recommendation 5.1.5). On publication of the draft as a BCP RFC, these recommendations will become authoritative best practice and a deployment that omits any of them should disclose the omission with reasoning. The disclosure-required and recommended-practice findings under this branch (DR-5, RP-1) are consistent with these draft recommendations.

The assessment treats reporting as a finding of recommended practice rather than required conformance, with three distinct positions:

1. RFC 8590 conformance is required where the deployment uses the Extensible Provisioning Protocol (EPP) for registry-to-registrar notification of automation-driven DS changes, because RFC 8590 is itself an authoritative published standard and its applicability to this case is direct. A registry that performs DS automation but does not notify the sponsoring registrar through RFC 8590 or an equivalent mechanism produces a registry-registrar state inconsistency that is itself a section 1.3(B) coherence concern at the protocol layer between the two parties.
2. Reporting to the Child DNS operator and the registrant for the four event categories identified in SAC126 Appendix B.4 is recommended practice; its absence does not constitute a section 1.3 finding under the present standards landscape but is a disclosure requirement in the operator's design documentation.

3. Drift-detection reporting is recommended practice under section 1.3(B) but is not a hard requirement; deployments that omit it should disclose the omission and the resulting bound on detection latency for drift-induced failures.

A specific observation under this branch concerns the relationship between reporting and the scanning cadence examined under Branch T4. The pull-based scanning model produces a natural reporting cadence: each scan produces a state observation that can be compared to the previously published DS RRset and to the previously observed child contents, and changes can be reported on detection. A deployment that scans for CDS or CDNSKEY records but does not retain or use the resulting state for reporting purposes is leaving available signal unused. The assessment recommends that any deployment performing pull-based scanning extend its scanning logic to include drift detection and reporting on the same cadence, at marginal additional cost.

6.3 Composite finding for cross-cutting concerns

The two cross-cutting concerns yield a composite finding that bears on the assessment's overall conclusion. Concern X1 (going-insecure) is a hard requirement: a deployment that fails to handle the algorithm-zero path correctly, fails to honor a registrar opt-out in that context, fails to extend cross-nameserver consistency to the algorithm-zero signal, or fails to preserve a manual going-insecure path produces conditions that the assessment treats as failing under both RSEP pillars. The four requirements set out in Section 6.1 are accordingly elevated to the same status as the minimum validation pre-check set under Branch T2 and the lock interaction requirements under Branch S3.

Concern X2 (reporting) is a graduated finding. The RFC 8590 conformance requirement for registry-to-registrar notification is hard and immediate; the recommendations for reporting to the Child DNS operator and the registrant are best practice with a disclosure requirement; drift-detection reporting is a recommended enhancement whose absence requires explicit acknowledgement. The graduation reflects the present state of standardization: where authoritative published standards exist (RFC 8590), the assessment treats their application as required; where best practice has been articulated by the SSAC but not yet standardized, the assessment treats it as disclosure-required; and where the practice represents an extension beyond articulated best practice, the assessment treats it as recommended.

The cross-cutting concerns do not change the assessment's overall posture: the mechanism, properly deployed, does not raise significant security or stability issues under RSEP. They do, however, expand the set of deployment-time requirements that must be present for the assessment to yield a positive finding. The full set of such requirements is consolidated and presented in Section 7.

7 Findings and Recommendations

This section consolidates the findings developed across Sections 4 through 6 into a form suitable for submission alongside an RSEP request, or for use by ICANN staff in conducting a preliminary determination under RSEP section 2.2. The findings are organized into three categories reflecting their normative weight: hard requirements whose absence produces a negative finding under one or both RSEP pillars; disclosure requirements whose substance must be documented in the operator's design but whose specific resolution is left to operator judgment; and recommended practices whose adoption strengthens the deployment posture but whose absence does not, in itself, produce a negative finding under the present standards landscape.

7.1 Hard requirements

The following requirements emerged across multiple branches as preconditions for a positive finding. Each is stated in normative terms, with reference to the RSEP clause it engages and to the assessment branches in which it was developed.

- **HR-1: Minimum validation pre-check set.** Before publishing any DS RRset derived from a CDS or CDNSKEY signal, the parent must verify that for each DS record to be published, a corresponding DNSKEY record exists in the child zone with matching algorithm and key tag, and that the child's DNSKEY RRset is signed by at least one key referenced by the new DS RRset. This requirement derives from RFC 7344 section 4.1 and is the floor for both RSEP section 1.2 (preventing unauthorized publication of incoherent state) and section 1.3(B) (preventing coherence-breaking conditions at validating resolvers). Developed under Branches T2, S1.
- **HR-2: Cross-nameserver consistency for CDS and CDNSKEY signals.** The parent must collect CDS and CDNSKEY records from each authoritative nameserver listed in the child's delegation and verify consistency across the responses before acting. Where any nameserver returns a response inconsistent with the others, the operation must be aborted and retried at a later time. This requirement engages RSEP section 1.2(A) (defense against the nameserver-position adversary, including in multi-signer setups), section 1.3(A) (conformance with RFC 9975, which updates RFC 7344 to make this a normative MUST), and section 1.3(B) (preventing coherence breakage in multi-provider deployments). Developed under Branches S1, S2, S4, T5; codifies the consistency requirements of RFC 9975. Without this requirement, the multi-signer takeover scenario identified in S1 and the multi-provider coherence risk identified in T5 are unmitigated, and the deployment is non-conformant with RFC 9975.
- **HR-3: Providing a registrar opt-out.** The parent must provide a mechanism for registrars to opt out of automated DS maintenance — including ongoing maintenance, bootstrapping, and the going-insecure transition — for all domains under its management, and optionally for specific domains via an EPP extension. This requirement engages RSEP section 1.2(A): the opt-out represents the registrar or registrant's expressed intent that DS automation should not take place. Developed under Branches S3, X1.
- **HR-4: Permitting clientUpdateProhibited.** Automated DS maintenance must not be suspended on the basis of a *clientUpdateProhibited* lock alone. The lock does not represent meaningful authorization withdrawal because the registrar — the party performing or relaying the automation — can remove the lock without out-of-band verification. Suspending

automation under this lock would prevent legitimate ongoing maintenance without providing any security benefit and would force registrants into manual recovery for routine algorithm rollovers. Developed under Branches S3, X1; consistent with SAC126 Appendix B.3.

- **HR-5: Preservation of manual interfaces.** The deployment must preserve a manual interface for DS record management at the registrar, including the ability to set, modify, remove, and entirely withdraw the DS RRset for a delegation. The manual interface is the recovery path for scenarios in which the automated mechanism is unavailable: destroyed or compromised authentication credentials at the Child DNS operator, registry locks set out-of-band, and uncooperative providers in multi-provider transitions. A deployment that automates DS maintenance but removes the manual interface fails under both RSEP pillars. Developed under Branches S3, X1; explicit in SAC126. It is noted that the Registrar Accreditation Agreement (RAA) requires that all registrars provide a manual interface for DNSSEC configuration.
- **HR-6: Correct handling of the going-insecure transition.** The deployment must recognize and honor authenticated CDS records with algorithm 0 according to RFC 8078 section 4, and must apply the same authentication, validation, and consistency requirements to such signals as to other automated DS changes. The validation pre-check requirement (HR-1) is structurally inapplicable to the empty resulting DS RRset and must not be applied as a basis for rejecting the signal. The cross-nameserver consistency requirement (HR-2) and the lock requirements (HR-3, HR-4) apply unchanged. Developed under Branch X1.
- **HR-7: RFC 8590 conformance for registry-to-registrar notification.** Where the deployment uses EPP for registry-to-registrar communication and the registry performs DS automation, the registry must notify the sponsoring registrar of automation-driven DS changes through the RFC 8590 Change Poll Extension or an equivalent mechanism. The absence of such notification produces a registry-registrar state inconsistency at the EPP protocol layer, which is itself a section 1.3(B) coherence concern within the registration system. Developed under Concern X2.

7.2 Disclosure requirements

The following items must be documented in the operator's design materials, including the DNSSEC Practice Statement. The substance of each choice is left to operator judgment; the requirement is that the choice be made deliberately, recorded, and made available to relevant parties.

- **DR-1: Validation pre-check posture beyond the minimum.** The minimum validation set (HR-1) is the hard floor; deployments routinely apply additional checks, including digest type rejection per RFC 9904, signing algorithm rejection per RFC 9904, and CDS/CDNSKEY consistency checks where both record types are published. The operator must document which additional checks are applied and the rationale for any that are omitted. Developed under Branch T2.
- **DR-2: Scanning cadence and the resulting response-time bound.** The pull-based scanning cadence sets the worst-case latency between a Child DNS operator publishing a signal and the parent acting upon it. The operator must document the cadence in use and disclose the resulting latency bound to registrants and Child DNS operators, with particular attention to the implications for emergency response such as recovery from key compromise. Developed under Branch T4.
- **DR-3: TTL handling posture.** The operator must document the DS RRset TTL applied in steady state, any TTL reduction applied around DS update events, and the duration of any reduced-TTL

window. The documentation must be sufficient to allow registrants and Child DNS operators to predict the rollback latency available in the event of a misconfigured DS publication. Developed under Branch T3.

- **DR-4: Multi-signer policy.** The operator must document the deployment's posture toward multi-signer setups per RFC 8901. Acceptable postures include requiring cross-nameserver consistency under HR-2 and admitting multi-signer delegations on that basis; excluding multi-signer delegations from automation by deployment policy; or accepting them with documented residual risk and explicit registrant disclosure. Developed under Branches T5, S1.
- **DR-5: Reporting mechanism and recipient set.** The operator must document the events that trigger reporting, the recipients to whom reports are sent (Child DNS operator, registrar, registrant, technical contact), the channels used (email derived from SOA contact data, EPP poll messages), and the back-off behavior applied to repeated conditions. The minimum content of this documentation derives from SAC126 Appendix B.4. Developed under Concern X2.
- **DR-6: Bootstrapping precondition disclosure.** Where the operator supports automated bootstrapping per RFC 9615, the documentation must disclose the precondition that the child's authoritative nameserver hostnames lie in DNSSEC-signed zones whose chain of trust is anchored in the root and must describe the operator's behavior for delegations that do not meet this precondition. Developed under Branch S2.
- **DR-7: Notification of deployment.** If an operator wishes to deploy DS automation, it should provide registrars sufficient notice of its intention to do so, so that they can submit requests to be opted out.

7.3 Recommended practices

The following items are not required for a positive finding under the present standards landscape but are recommended as enhancements to the deployment's security or stability posture. Their absence does not produce a negative finding but limits the deployment's resilience to conditions not yet addressed by published standards.

- **RP-1: Drift-detection reporting.** The deployment is recommended to perform periodic re-validation of published DS RRsets against the child's current DNSSEC state, and to report detected drift to the appropriate recipients per DR-5. The recommendation addresses a class of stability incident — child-side changes that render the existing DS RRset invalid without any CDS or CDNSKEY signal being published — that no other requirement of this assessment captures. Developed under Concern X2.
- **RP-2: Adoption of RFC 9859 (Generalized DNS Notifications).** Generalized DNS Notifications has been published as RFC 9859, which provides a notification-triggered alternative to pure-polling DS automation; its adoption alongside CDS/CDNSKEY automation under Recommendation 1 substantially improves the response-time properties analyzed under Branch T4 and is recommended where the operational benefit warrants the implementation effort. RFC 9859 specifies a different mechanism than the one within Recommendation 1's scope; its adoption is a complement to, not a substitute for, conformant deployment of RFCs 7344, 8078, and 9615. *Note:* the CDS/CDNSKEY consistency requirement, formerly tracked

alongside Generalized DNS Notifications in SAC126 and in earlier versions of this assessment, has been published as RFC 9975 and is now reflected as a hard requirement in HR-2.

- **RP-3: Permanent short DS TTLs.** The asymmetric TTL posture documented under DR-3 — long TTL in steady state, short TTL after updates — is consistent with operational practice and meets the assessment's stability requirements. Operational evidence cited in SAC126 Appendix B.1 indicates that permanent short TTLs in the 5-to-15-minute range produce negligible additional load on registry infrastructure. The deployment is recommended to consider permanent short TTLs as a stronger posture than asymmetric TTLs, because the rollback window is bounded at all times rather than only following deliberate updates.
- **RP-4: Adoption of RFC 9803 (EPP TTL extension) and [draft-ietf-regext-rdap-ttl-extension](#).** The IETF REGEXT Working Group's EPP TTL extension, which SAC126 referenced as in-progress draft work, has since been published as RFC 9803 on the standards track. Where the asymmetric TTL posture under DR-3 is in use, the deployment is recommended to support RFC 9803 to allow registrars to convey TTL changes to the registry through standard channels rather than registry-internal mechanisms. RFC 9803 is not part of Recommendation 1's mechanism scope, so its adoption is not a conformance requirement under section 1.3(A); it is, however, the standardized path for implementing TTL adjustments at the registry-registrar interface and is recommended for deployments seeking that interoperability. Correspondingly, the RDAP TTL extension defined in [draft-ietf-regext-rdap-ttl-extension](#) (a standards track document which is currently in the RFC Editor queue as of writing) provides access to configured TTL values for parties (such as DNS operators and registrants) without access to EPP.

7.4 Composite assessment

The mechanism described in Recommendation 1 of SAC126 — namely, the CDS/CDNSKEY mechanism specified in RFCs 7344, 8078, and 9615, deployed for third-party DNSSEC operations — does not, when implemented in conformance with the seven hard requirements set out in Section 7.1 and with the disclosure requirements of Section 7.2, raise significant security or stability issues within the meaning of RSEP sections 1.2 and 1.3. The mechanism rests on three standards-track RFCs whose conformance posture is checkable, and its operational properties — when subject to the hard requirements above — produce neither unauthorized mutation of registry data nor conditions that adversely affect throughput, response time, consistency, or coherence of DNS responses at validating resolvers.

The assessment identifies three areas where the mechanism's properties depend materially on deployment-time choices that the published standards do not fully resolve.

1. The multi-signer case per RFC 8901 model 2, where the cross-nameserver consistency requirement (HR-2) is the principal mitigation for both the security exposure identified in S1 and the coherence concern identified in T5.
2. The reporting and observability gap, where the absence of standardized normative requirements is partially offset by RFC 8590 (HR-7) for the registry-registrar interface but remains an open area for the broader recipient set (DR-5).
3. The bootstrapping precondition that nameserver hostnames be DNSSEC-signed (DR-6), which is a structural property of RFC 9615's authentication construction rather than a defect, but which limits the mechanism's applicability and must be disclosed as a deployment constraint.

None of these three areas, in this assessment's view, rises to the level of a significant security or stability issue under RSEP when accompanied by the documented mitigations. All three should, however, be addressed explicitly in the operator's submission and in any subsequent ICANN evaluation, to ensure that the residual risks are knowingly accepted by the parties involved and disclosed to the broader community of registrants, Child DNS operators, and validating resolver operators.

The assessment concludes that the proposed deployment of the CDS/CDNSKEY mechanism, subject to the hard requirements of Section 7.1, the disclosure requirements of Section 7.2, and adoption of the recommended practices of Section 7.3 to the extent practical, is consistent with the operator's obligations under the Registry Agreement and with the security and stability mandate of the Registry Services Evaluation Policy. A positive preliminary determination by ICANN under RSEP section 2.2 is supported by the technical evidence developed in this assessment.

7.5 Treatment of subsequent operator submissions

A note bearing on the use of this assessment for subsequent RSEP submissions: the findings in this section are mechanism-specific rather than operator-specific. The seven hard requirements, seven disclosure requirements, and four recommended practices apply to any deployment of the CDS/CDNSKEY mechanism for third-party DNSSEC operations, regardless of the registry, registrar, or technical implementation in question. Operators submitting RSEP requests for CDS automation services should be able to engage each item directly, with the operator's specific design choices and operational policies addressing the disclosure requirements and the recommended practices as applicable to their deployment.

Where an operator's submission addresses each item explicitly — confirming the hard requirements, documenting the disclosure items, and engaging the recommended practices either by adoption or by reasoned acknowledgement — ICANN's preliminary determination under RSEP section 2.2 has, by construction, the technical evidence base required to conclude that the proposed service does not raise significant security or stability issues. Where a submission omits items or treats them with insufficient specificity, the residual ambiguity is itself a basis for either further engagement during the preliminary determination period or referral to the Registry Services Technical Evaluation Panel under RSEP section 2.4 for the formal 45-day review.

Operators should also be aware that the IETF draft-ietf-dnsop-ds-automation, once published as a Best Current Practice RFC, will become the operational reference for many of the recommendations in this assessment. Submissions made during the interim period between this assessment and the draft's publication should expect ICANN evaluators to read both documents together; the draft's recommendations may be cited as authoritative in subsequent determinations once they reach RFC status.

Appendix A: Mapping of Findings to SAC126 Sections

This appendix provides a cross-reference between the findings developed in this assessment and the corresponding source material in SAC126. The mapping serves three purposes: it allows reviewers to verify that the assessment's analytical lineage is grounded in the SSAC's documented analysis; it identifies findings that extend beyond SAC126's scope and rely on additional reasoning developed in this assessment; and it surfaces the SAC126 sections that should be read alongside any operator submission as primary supporting material.

The mapping is organized in four tables. Table A.1 maps the seven hard requirements (HR-1 through HR-7) to their SAC126 sources and to any extensions made by this assessment. Table A.2 does the same for the seven disclosure requirements (DR-1 through DR-7). Table A.3 covers the four recommended practices (RP-1 through RP-4). Table A.4 provides a reverse index, organized by SAC126 section, identifying which findings draw on each part of the source document.

A.1 Hard requirements

Finding	SAC126 source	Extension by this assessment
HR-1 Minimum validation pre-check set	Section B.1 (validity checks and safety measures); references RFC 7344 section 4.1 directly	Treats the minimum check set as a hard floor under RSEP sections 1.2 and 1.3(B); SAC126 frames it as recommended operational practice rather than as a normative requirement
HR-2 Cross-nameserver consistency	Section B.5 (consistency considerations) and section 7 (future and related work, multi-signer item); references RFC 9975 (<i>Clarifications on CDS/CDNSKEY and CSYNC Consistency</i>), titled at SAC126's publication as <i>draft-ietf-dnsop-cds-consistency</i> , <i>Consistency for CDS/CDNSKEY and CSYNC is Mandatory</i>	Promotes cross-nameserver consistency from recommended posture to a hard requirement under RSEP, now backed by the authoritative published standard RFC 9975, which updates RFC 7344
HR-3 Providing a registrar opt-out	Section B.3 (registration locks); explicit in SAC126's analysis of the registry lock	Adopts SAC126's analysis directly, with the addition that the requirement extends to the going-insecure transition (HR-6 boundary case)
HR-4 Permitting clientUpdateProhibited	Section B.3 (registration locks); explicit in SAC126's analysis of the registrar lock	Adopts SAC126's analysis directly; the assessment's RSEP framing reinforces the conclusion by reference to section 1.2's authorization criterion

Finding	SAC126 source	Extension by this assessment
HR-5 Preservation of manual interfaces	Section 4.4 (Disabling manual DS management interfaces is therefore strongly discouraged); section B.2 (multiple submitting parties)	Promotes strongly discouraged to a hard requirement on the basis that the manual path is the only available recovery channel for credential loss and provider-transition scenarios
HR-6 Going-insecure transition handling	Section 2.2 (key changes, listing going insecure as a routine case); section B.4 case 2(b); references RFC 8078 section 4	Develops the integrated requirement that combines RFC 8078 section 4 conformance with the lock requirements (HR-3, HR-4) and the consistency requirement (HR-2); SAC126 treats these as separate operational considerations
HR-7 RFC 8590 conformance for registry-to-registrar notification	Section B.4 (the registrar should always stay informed of any DS changes, e.g., via EPP (RFC 8590))	Promotes the SAC126 Recommendation to a hard requirement on the basis that RFC 8590 is itself an authoritative published standard whose application to this case is direct, and that the absence of notification produces a section 1.3(B) coherence concern within the registration system

A.2 Disclosure requirements

Finding	SAC126 source	Extension by this assessment
DR-1 Validation pre-check posture beyond the minimum	Section B.1 (validity checks); section B.5 (consistency considerations); references RFC 9904 for digest type and algorithm guidance	Establishes the documentation requirement as the operator's response to the choice space SAC126 articulates
DR-2 Scanning cadence and response-time bound	Section 4.2.1 (pull-based); section 4.3 (discussion of scanning inefficiency and delays); cites Thomassen, Scalability of Scanning of CDS/CDNSKEY Records (ICANN 77)	Adds the explicit emergency-response framing to SAC126's general uncertainty and delays analysis
DR-3 TTL handling posture	Section B.1 (TTL reduction analysis); cites the .goog and ISC operational evidence	Adopts SAC126's analysis directly; the documentation requirement formalizes what SAC126 frames as design choice
DR-4 Multi-signer policy	Section 4.4 (multi-signer paragraph); section 7 (future and related work, multi-signer); section B.5 (Practical experience with multi-signer setups is currently limited)	Adopts SAC126's framing that no standards-defined process exists, and translates this into a per-operator disclosure requirement
DR-5 Reporting mechanism and recipient set	Section B.4 (reporting); the four event categories in DR-5 derive directly from SAC126's case enumeration	Adopts SAC126's enumeration as the minimum content of the disclosure; the disclosure requirement structure is added by this assessment
DR-6 Bootstrapping precondition disclosure	Section 3.1 (initial DS provisioning); section 4.2 (automatic methods, authentication problem); references RFC 9615	Extends SAC126's discussion of authentication into an explicit precondition statement; the precondition is implicit in SAC126's treatment of bootstrapping but is not stated as a deployment constraint
DR-7 Notification of deployment	Section 4.4 (operational considerations — "Should DS automation involve the registrar or the registry, or both?" and the collision question where both registry and registrar automate); Appendix B.2 (registry/registrar coordination)	SAC126 raises registry/registrar coordination as an operational question but does not address how an operator should transition into automation; this assessment adds the requirement that an operator give registrars advance notice of an intended deployment and an opportunity to opt out, formalized through the opt-out mechanism described in Section 4.4 of this assessment (developed under Branches S3 and S5)

A.3 Recommended practices

Finding	SAC126 source	Extension by this assessment
RP-1 Drift-detection reporting	Section B.4 case 4 (the Child zone is no longer compatible with the existing DS record set)	Frames SAC126's case 4 as a recommended enhancement rather than as a required reporting category, on the basis that drift detection requires extension of the scanning logic beyond what RFC 7344 contemplates
RP-2 Adoption of RFC 9859 (Generalized DNS Notifications)	Section 4.2.3 (hybrid pull/push, Generalized DNS Notifications); section 7 (future and related work)	SAC126 referenced this work as an in-progress draft; the work has since been published as RFC 9859 and the recommendation is updated accordingly. The CDS/CDNSKEY consistency draft, formerly tracked alongside Generalized DNS Notifications in SAC126 and in earlier versions of this assessment, has been published as RFC 9975 and is now reflected in HR-2.
RP-3 Permanent short DS TTLs	Section B.1 (TTL discussion, .goog and ISC evidence)	Extends SAC126's asymmetric-TTL recommendation by noting that the operational evidence supports permanent short TTLs as a stronger posture; SAC126 does not draw this conclusion explicitly
RP-4 Operational coordination between registry and registrar	Section 4.2.1 (race condition discussion); section B.2 (multiple submitting parties)	Adopts SAC126's analysis; the recommendation is consistent with SAC126's own conclusion that the task is best performed by only one of these parties

A.4 Reverse index by SAC126 section

The following table identifies, for each substantive section of SAC126, the assessment findings that draw on it. Sections of SAC126 not listed (preface, table of contents, introduction, terminology, acknowledgements) are not source material for specific findings but provide the context within which the findings are framed.

SAC126 section	Title	Findings drawing on this section
Section 2.1	Initial provisioning (DNSSEC bootstrapping)	DR-6; mechanism description in Section 3.3 of this assessment

SAC126 section	Title	Findings drawing on this section
Section 2.2	Key changes	HR-6; mechanism description in Section 3.7 of this assessment
Section 2.3	Roles and responsibilities	Mechanism description in Section 3.1 of this assessment
Section 3.1	Initial DS provisioning	S2; DR-6
Section 3.2	DS updates	S1; T1
Section 4.1	Registrant-centric	Out-of-scope variations in Section 3.8; comparative basis for Section 4–Section 5 findings
Section 4.2	Automatic methods (preamble)	T1 (standards landscape); the authentication-problem framing in Section 4 of this assessment
Section 4.2.1	Pull-based	T2; T4; DR-2; mechanism description in Section 3.5 of this assessment
Section 4.2.2	Push-based	Out-of-scope variations in Section 3.8
Section 4.2.3	Hybrid pull/push	RP-2; T4 (response-time analysis)
Section 4.3	Discussion	T1; T4; framing of overall conclusion in Section 7.4 of this assessment
Section 4.4	Operational considerations for DS automation	HR-5; DR-4; entry point to Appendix B material
Section 5	Findings (SAC126 Findings 1–4)	Framing of overall conclusion in Section 7.4 of this assessment
Section 6	Recommendations (Recommendation 1)	Scope statement in Section 1 and Section 3 of this assessment
Section 7	Future and related work	HR-2; DR-4; RP-2 (multi-signer item, registration locks item)
Section B.1	Validity checks and safety measures	HR-1; DR-1; DR-3; RP-3

SAC126 section	Title	Findings drawing on this section
Section B.2	Multiple submitting parties	HR-5; RP-4; S5 (collision analysis)
Section B.3	Registration locks	HR-3; HR-4; HR-6 (lock interaction component); S3
Section B.4	Reporting	HR-7; DR-5; RP-1; X2
Section B.5	Consistency considerations	HR-2; DR-1; T5 (multi-signer coherence); X1 (consistency for going-insecure)
Appendix C	Steps registrants have to do to secure DNSSEC delegation	Comparative basis for the assessment's framing in Section 4.1 (registrant-centric, out of scope) and Section 1 (motivation for automation)

A.5 Findings without direct SAC126 source

Three findings in this assessment do not have a direct correlate in SAC126 and rest on reasoning developed within this assessment, drawing on RSEP definitions, the cited RFCs, and operational evidence not specifically discussed by SAC126. Each is identified here for transparency, with its basis stated explicitly.

The promotion of cross-nameserver consistency (HR-2) from a recommended practice in SAC126 section B.5 to a hard requirement in this assessment rests on the analysis in Section 4.5 (Branch S4) and Section 5.5 (Branch T5) of this report, which together establish that the multi-signer takeover scenario and the multi-provider coherence risk are unmitigated without consistency enforcement. SAC126 references what was then an in-progress IETF draft on consistency but does not formally advance the requirement to the level of a hard precondition; this assessment does so on the basis of the RSEP sections 1.2(A) and 1.3(B) implications. That draft has since been published as RFC 9975 (*Clarifications on CDS/CDNSKEY and CSYNC Consistency*), which updates RFCs 7344 and 7477. HR-2 of this assessment now codifies a published standards-track requirement, and the residual analytical extension is limited to the explicit promotion of the requirement to a hard precondition under RSEP rather than the underlying technical content.

The promotion of manual interface preservation (HR-5) from strongly discouraged to disable in SAC126 section 4.4 to a hard requirement in this assessment rests on Section 6.1 (Concern X1) of this report, which extends SAC126's analysis to identify the manual interface as the recovery path for both credential loss and provider-transition scenarios. The promotion is consistent with SAC126's own analysis but adopts a more normative posture on the basis that the absence of a manual path produces stuck-state conditions failing under both RSEP pillars.

The integrated treatment of the going-insecure transition (HR-6) draws on multiple SAC126 sections (section 2.2, section B.4, RFC 8078 section 4) but the specific composition — combining standards conformance, lock interaction, consistency requirements, and the special-case handling of the validation pre-check — is developed in Section 6.1 of this assessment and is not stated as a single integrated requirement in SAC126.

These three extensions are identified to allow reviewers to engage them directly and, where appropriate, to challenge or refine the reasoning. The assessment's overall conclusion does not rest on these three extensions in isolation; each is supported by the broader analysis in the body of the report, and each is consistent with — though more normatively framed than — the corresponding SAC126 material.