Preprints.org

Concept Paper

Not peer-reviewed version

SOMA-DR: Decision Receipts for Explainable Recovery and Key Rotation in Post-Quantum IAM

SravanaKumar Nidamanooru *

Posted Date: 1 January 2026

doi: 10.20944/preprints202601.0007.v1

Keywords: decision receipts; identity recovery; explainable IAM; auditability; post-quantum cryptography (PQC); ML-KEM; ML-DSA; SLH-DSA; OpenID Connect; OAuth 2.0; JWT; transparency logs; remote attestation; policy governance; key rotation; machine identities

Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Concept Paper

SOMA-DR: Decision Receipts for Explainable Recovery and Key Rotation in Post-Quantum IAM

Sravanakumar Nidamanooru

Independent Researcher, USA; sravana.nidamanooru@gmail.com

Abstract

Identity and Access Management (IAM) is rapidly shifting from static, rule-based access control to adaptive decisioning that uses step-up challenges, recovery verification, device and behavior signals, and continuous authorization to reduce account takeover and misuse. In parallel, IAM systems must prepare for post-quantum cryptography (PQC) transitions that reshape credential formats, signing and verification paths, and long-term audit integrity. Together, these shifts expose a practical governance gap: when a sensitive identity action is allowed, challenged, denied, or escalated—such as passwordless enrollment, recovery credential release, privileged step-up, or machine key rotation—organizations must be able to explain the decision consistently, prove which policy version and evidence categories were used, and verify later that the record has not been altered, even across key rotation and algorithm migration. This paper introduces Decision Receipts (DR): a verifiable, privacy-aware record emitted at decision time that captures policy context and versioning, evidence commitments and normalized evidence descriptors (without embedding raw personal data), outcomes and structured reason codes, and cryptographic signatures designed for crypto agility across classical, hybrid, and PQC profiles. We propose an open receipt schema, deterministic canonicalization rules, and issuer/verifier workflows compatible with widely deployed identity standards and signing containers, with optional anchoring into append-only or transparency-style logs for additional tamper evidence. Using synthetic sign-in, recovery, and rotation workloads aligned with our prior benchmark methodology, we illustrate deterministic tamper detection, overhead trade-offs, and privacy-preserving redaction profiles that preserve audit utility. The approach is intentionally IP-safe and deployable as an audit overlay independent of any specific orchestrator implementation, while remaining suitable as a building block for future recovery-and-rotation orchestrators.

Keywords: decision receipts; identity recovery; explainable IAM; auditability; post-quantum cryptography (PQC); ML-KEM; ML-DSA; SLH-DSA; OpenID Connect; OAuth 2.0; JWT; transparency logs; remote attestation; policy governance; key rotation; machine identities

I. Introduction

IAM's promise is simple to say and hard to guarantee: let the right identity access the right resource, at the right time, for the right reason—and be able to prove later that the decision was correct. In practice, the most damaging failures do not occur in routine sign-in paths. They occur in the "exception lanes" of IAM: account recovery, helpdesk overrides, re-enrollment of MFA or passkeys, privileged step-up, and long-lived machine credentials that drift from policy. These actions are high impact, time-sensitive, and spread across multiple components (IdP, MFA, device posture, approvals, ticketing, monitoring), making it difficult to reconstruct a single, trustworthy narrative after an incident.

Two industry shifts raise the bar further. First, IAM is becoming more adaptive. Decisions increasingly depend on richer context such as device posture, geo and velocity, IP and ASN reputation, recent credential changes, and posture or attestation-like claims. Even if the underlying logic is simple, the decision path becomes harder to explain consistently across teams and tools.

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

Second, IAM must prepare for PQC transitions that change cryptographic assumptions, key sizes, signature sizes, and verification cost, while creating new operational requirements for long-lived audit artifacts that must remain verifiable over time.

This work is part of a broader research direction called **SOMA (Secure Orchestrator for Modular Access)**. SOMA is the system we are building from scratch to harden the most fragile IAM actions—especially recovery and machine identity rotation—by enforcing a predictable sequence: gather evidence, select a policy band (allow / step-up / quorum / deny), execute controlled release or rotation, and produce auditable artifacts explaining the outcome. Because SOMA is being developed alongside IP protection, this paper avoids disclosing proprietary scoring thresholds, feature engineering, or orchestration internals. Instead, it contributes a portable component that any IAM system can adopt today.

This paper addresses the gap between "IAM made a decision" and "humans can later verify and understand that decision." We introduce **Decision Receipts (DR)**—cryptographically signed, verifier-friendly decision records that bind an outcome to policy versioning, evidence commitments, and structured reason codes, with optional quorum and approval summaries for escalated actions. The goal is not to standardize a risk model. The goal is to make high-risk IAM decision trails verifiable, explainable, and resilient to cryptographic migration. In the overall research track, the first paper framed the combined AI and PQC stress on IAM, the second paper introduced a synthetic benchmark harness for evaluating policy trade-offs, and this paper defines the audit layer that makes those decisions defensible after the fact.

II. Background and Related Work

Most IAM deployments already produce logs, but these logs are typically system-centric rather than decision-centric. They capture component events and timestamps (IdP events, MFA events, API gateway events, ticketing events), yet omit what evidence categories mattered, which policy path fired, and how a reviewer can independently verify integrity later [8]. This becomes especially painful in recovery and override workflows, where the "why" matters more than the "what," and the action's blast radius is high [13].

Modern authentication stacks commonly rely on OpenID Connect on top of OAuth 2.0, with JWT/JWS-based tokens for claims exchange and runtime authorization decisions [1–4]. These standards provide interoperability for authentication and delegation, but they do not standardize a cryptographically verifiable *reason record* for high-risk IAM decisions such as recovery releases, enrollment overrides, or privileged step-up. In practice, organizations reconstruct decisions by correlating multiple logs across systems, which increases investigation time and makes post-incident narratives fragile [8].

The value of tamper-evident records is demonstrated by transparency systems and append-only logging patterns. While IAM decisions are not usually public, the integrity pattern—portable signed records that can be independently verified—translates well to IAM governance, particularly when the threat model includes insiders or compromised administrative tools [8]. For example, transparency-log concepts (such as Sigstore's Rekor) show how signed artifacts can be registered to provide additional tamper evidence beyond local logs [12]. Even when the receipt content remains private, anchoring a digest can add a strong integrity signal for later audits.

For device and workload trust, remote attestation architectures provide structured models for evidence generation and appraisal, clarifying how evidence is produced, transported, and evaluated [7]. As IAM decisions increasingly depend on posture and attestation-like signals, decision artifacts should reference evidence categories and commitments without embedding sensitive raw device/user telemetry, enabling privacy-preserving audits while keeping integrity checks feasible [7,8].

Receipts must also be signable and comparable across systems. JSON is widely used but requires deterministic canonicalization before hashing/signing is stable; RFC 8785 (JCS) defines a deterministic canonicalization scheme [5]. For signing containers, JWS is common in JSON/JWT ecosystems [4],

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

and COSE offers a compact CBOR-based option suitable for constrained environments [6]. Finally, crypto agility becomes essential in the PQC era: NIST's standardization of ML-KEM, ML-DSA, and SLH-DSA motivates audit artifacts that remain verifiable across algorithm migrations and retention horizons [9–11]. This paper's receipt design adopts these primitives to produce decision records that can survive future cryptographic transitions without breaking verification workflows.

This paper builds on two prior steps in our research track. In Paper 1, we framed the combined AI-driven adaptivity and PQC migration as a structural stress test for IAM, emphasizing that the hardest failures concentrate in recovery, exception handling, and non-human identity lifecycle [13]. In Paper 2, we introduced a synthetic benchmarking methodology to evaluate high-risk IAM actions safely (without using personal data), enabling reproducible comparisons of policy trade-offs. The remaining gap—addressed here—is an audit-grade, verifiable artifact that makes decisions explainable and defensible after the fact using canonicalization and cryptographic verification [5,9–11].

III. SOMA-DR Overview

SOMA-DR is a decision-receipt layer that can sit beside existing IAM decision points (e.g., IdPs, recovery services, access gateways, privileged workflows, and NHI rotation pipelines). For each sensitive identity action—such as recovery credential release, passkey enrollment, privileged step-up, or machine credential rotation—the decision system emits a receipt that is both verifiable and privacy-aware. The design goal is not to replace existing IdPs or policy engines, but to make the outcome of a decision defensible after the fact, especially in environments adopting continuous evaluation and Zero Trust operating assumptions [8,13].

Each receipt binds together four elements: (a) request context (what action was requested and under what service scope), (b) a cryptographic commitment to the evidence bundle used (e.g., a hash over protected evidence, rather than raw signal values), (c) the policy path and reason codes that explain why a particular outcome occurred, and (d) a cryptographic signature over a canonical representation of the receipt to support deterministic third-party verification [5]. When approvals are required (e.g., recovery escalations or high-impact rotations), the receipt additionally captures quorum parameters and approval summaries in a structured form, enabling later verification that "t-of-n" conditions were satisfied without exposing sensitive approver details by default [8,13].

Receipts are intended for independent verification by auditors, SIEM pipelines, or automated compliance checks, using standard identity and signing infrastructures. Since many IAM stacks already depend on OAuth 2.0 / OpenID Connect token ecosystems and JWT/JWS-based artifacts, SOMA-DR is designed to be compatible with common deployment primitives and verification tooling [1–4]. Finally, because PQC migration changes cryptographic assumptions and audit retention requirements, receipts are designed to remain verifiable across algorithm transitions via crypto agility, explicit key identifiers, and optional anchoring of receipt digests into append-only or transparency-style logs for additional tamper evidence [9–12].

IV. Decision Receipt Design

A SOMA-DR receipt is a structured object intended to answer four questions reliably: What happened? Why did it happen? Under what policy/version did it happen? Can it be proven authentic later? This framing is motivated by practical incident response and compliance needs under continuous verification models, where organizations must reconstruct and justify sensitive identity actions (especially recovery and overrides) long after they occur [8,13].

A. Minimal Receipt Fields (Baseline)

A recommended baseline receipt includes the following fields:

- rid: request identifier
- issued_at: issuance timestamp

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

- issuer: decision point identifier
- actor: pseudonymous subject reference (user or workload)
- action: e.g., RECOVERY_RELEASE, PASSKEY_ENROLL, PRIV_STEPUP, ROTATE_NHI_KEY
- service_id: service/application scope
- bundle.evidence_hash: evidence commitment (hash of protected evidence bundle)
- policy.selected: policy identifier
- policy.reasons[]: structured reason codes explaining the outcome
- policy_inputs_ver: versioning for policy inputs/normalization
- quorum (optional): required threshold and rule identifier
- approvals[] (optional): summarized approvals (tokenized IDs, timestamps, channel)
- outcome: allow/challenge/deny/escrow_release/rotate + TTL where relevant
- sig.*: signature metadata, including container type and algorithm profile

These fields are intentionally designed to be implementation-agnostic and compatible with existing identity ecosystems that already represent decisions and claims in token-oriented workflows [1–4].

B. Canonicalization and Signing

Because JSON can be serialized in multiple ways (ordering, whitespace, encoding), SOMA-DR requires deterministic canonicalization prior to hashing and signing. For JSON receipts, canonicalize using the JSON Canonicalization Scheme (JCS) to ensure that two independent verifiers reconstruct the same byte representation for signature verification [5]. The canonical form is then signed using established signing containers—JWS for JWT-oriented ecosystems or COSE for compact CBOR-centric environments—both of which are widely used for integrity and authenticity of structured claims [4,6].

C. Crypto Agility Across Classical/Hybrid/PQC

Receipts support classical, hybrid, and PQC signing profiles so archived receipts remain verifiable across cryptographic migrations. This aligns with the operational reality that systems may run classical algorithms today while adopting PQC standards such as ML-DSA and SLH-DSA over time, and may require hybrid profiles during transition periods [9–11]. Including explicit key identifiers and algorithm profiles in signature metadata enables future auditors to verify receipts even after key rotation and algorithm policy changes [9–11].

D. Optional Anchoring for Additional Tamper Evidence

Optionally, a receipt digest can be committed to an append-only or transparency-style log to provide additional tamper evidence while keeping receipt content private. This follows the general integrity pattern demonstrated by transparency systems, where independent verifiability is strengthened by anchoring a cryptographic commitment to an append-only structure [12]. (SOMA-DR treats anchoring as optional because some deployments may not permit external logging, yet still benefit from local signature-based integrity.)

V. Redaction Profiles and Privacy-Preserving Verification

SOMA-DR receipts are meant to be auditable, but IAM evidence is often sensitive: device identifiers, geo/velocity patterns, IP history, helpdesk notes, internal approver identities, and contextual signals that can expose personal or organizational details. Since receipts may be stored broadly across SIEMs, ticketing systems, and compliance exports, organizations need clear rules about what can be revealed to whom without breaking verification or increasing privacy risk [8,13]. In Zero Trust environments, where decisions are continuous and more frequent, the volume of auditable artifacts increases, making privacy-preserving design non-negotiable [8].

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

A. Recommended Model: Signed "View Receipts"

A practical approach is signed view receipts: generate separate signed receipts for different audiences (e.g., Internal Ops/Sec, Internal Audit/Compliance, External/Minimal). Each view is derived from the same decision, redacted according to a defined profile, canonicalized, and then signed. This keeps verification simple because each receipt is self-contained: the verifier checks the signature over the canonicalized view without needing access to the unredacted version [5].

B. Canonical Redaction Rules

To preserve verifiability across redaction, SOMA-DR follows canonical redaction rules:

1. keep field names and structural shape intact (do not remove keys arbitrarily),
2. redact values using deterministic placeholders or tokenized identifiers (to preserve comparability),
3. preserve value types (string remains string, arrays remain arrays),
4. canonicalize after redaction using JCS,
5. sign the resulting view receipt using JWS or COSE [4–6].

This approach supports a clean separation between "what auditors need to verify integrity and policy conformance" and "what sensitive data must remain protected," while staying compatible with widely deployed verification infrastructure [1–6].

Table 1. SOMA-DR redaction profiles (recommended default).

Field group	Internal (Ops/Sec)	Internal Audit/Compliance	External/Minimal
rid, action, service_id, issued_at	Keep	Keep	Keep
actor.id	Keep	Tokenize	Redact
bundle.evidence_hash	Keep	Keep	Keep
policy.selected	Keep	Keep	Keep
policy.reasons[]	Keep	Keep (minimize values)	Keep codes only
policy.quorum	Keep	Keep	Keep (omit rule text if needed)
approvals[].approver_id	Keep	Tokenize	Redact
approvals[].ts	Keep	Coarsen optional	Optional/Remove
sig.*	Keep	Keep	Keep

VI. Verification Workflow and Auditor Queries

A. Deterministic Verification Workflow

SOMA-DR receipts are designed to be verified deterministically so that different verifiers (auditors, SIEM pipelines, compliance automation) reach the same decision given the same receipt. A recommended workflow is:

1. Parse and schema-validate the receipt to confirm required fields are present and types are correct (e.g., rid, issued_at, issuer, action, bundle.evidence_hash, policy.selected, policy.reasons[], and signature metadata).
2. Canonicalize the receipt payload (excluding signature bytes) using the JSON Canonicalization Scheme (JCS) to produce a deterministic byte string for hashing and verification [5].
3. Resolve the verification key using sig.pubkey_id from a trusted key registry (or KMS/HSM-backed key directory) so the verifier can obtain the correct public key even after issuer key rotations [9–11].
4. Verify the signature over the canonicalized bytes using the specified signing container, typically JWS in JSON/JWT ecosystems or COSE in CBOR-oriented environments [4,6].

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

5. Run optional control checks to produce richer verdicts: quorum satisfaction (t-of-n), timestamp sanity (issued_at within acceptable clock skew), recognized policy version (policy_inputs_ver and policy.selected version known), and optional anchoring checks if a digest is committed to an append-only log [5,12].

The verifier outputs a simple verdict: VALID (signature and checks pass), INVALID (signature fails or required invariants violated), or VALID-WARN (signature passes but a non-fatal anomaly exists, such as an unknown policy version or missing optional anchoring). This supports Zero Trust operational models where decisions are frequent and audit automation must be reliable and scalable [8].

B. Evidence Commitment Verification

Receipts include bundle.evidence_hash, which is a cryptographic commitment to the protected evidence bundle used at decision time. When deeper investigation is required, a privileged auditor can retrieve the corresponding evidence bundle from protected storage and recompute the hash to confirm it matches the receipt commitment. This provides integrity assurance without embedding raw sensitive evidence into the receipt itself, aligning with evidence/appraisal patterns formalized by attestation architectures and preserving privacy boundaries [7,8].

C. Auditor Query Patterns

Decision receipts enable standardized, cross-system audit queries that are difficult when reasoning only from distributed logs. Common patterns include:

- Recoveries without quorum (high-risk recovery actions allowed without required approvals) [8,13].
- New-device recoveries allowed without escalation (potential takeover indicators) [8].
- Approvals below t (quorum not satisfied or incomplete approvals) [8].
- Repeated approver pairs / "cozy pairs" (collusion risk or process weaknesses; typically requires tokenized approver IDs rather than fully redacted views) [8].
- Unknown policy versions (policy drift or misconfiguration) [5,8].
- Missing evidence commitments (receipt incomplete or evidence pipeline bypassed) [7].
- Crypto posture inventory (which signature profiles are being used during PQC migration, where hybrid is still present, and whether policies align with NIST PQC guidance) [9–11].
- Timing anomalies (unexpected decision times, out-of-window approvals, or suspicious sequencing) [8].

  Where optional anchoring is used, auditors can also query whether receipt digests appear in append-only logs, strengthening tamper-evidence properties beyond local storage [12].

VII. Experimental Protocol and Evaluation Metrics

SOMA-DR is evaluated as an audit overlay on **synthetic** sign-in, recovery, and rotation workloads. Synthetic workloads are chosen to enable reproducibility and avoid handling personal user telemetry, while still representing the structure of high-risk IAM actions highlighted in our broader research framing [8,13]. The evaluation focuses on three practical properties: (1) verification correctness and tamper detection, (2) overhead in size and latency under different cryptographic profiles, and (3) audit utility across privacy-preserving redaction profiles.

A. Objectives

- **O1: Correctness & tamper detection.** Verify that compliant receipts validate reliably and that modified receipts fail deterministically under canonicalization + signature verification [5].
- **O2: Overhead.** Measure receipt size and issuance/verification latency under classical, hybrid, and PQC signing profiles, reflecting PQC transition realities [9–11].

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

- **O3: Audit utility under redaction.** Evaluate whether common compliance and incident-response queries remain answerable under Internal vs Audit vs External/Minimal views while maintaining verification simplicity (signature verification on each view) [5].
- **O4: Operational impact.** Evaluate whether verification should be inline (synchronous) or asynchronous for different IAM actions, consistent with Zero Trust systems where throughput and latency budgets matter [8].

B. Experiments

- **E1: Verifier correctness and tamper detection.** Generate valid receipts, then apply controlled modifications (e.g., outcome flip, reason edit, missing field) and measure acceptance/rejection rates under strict schema rules and canonical signing [5].
- **E2: Overhead micro-benchmarks.** For each signature profile (classical/hybrid/PQC), measure receipt size and p50/p95 issuance and verification latency using JWS/COSE containers [4,6], and PQC algorithms aligned with NIST standards [9–11].
- **E3: Redaction utility.** Apply defined redaction profiles and run a set of standardized auditor queries to measure which queries remain feasible under each view while preserving verifiability [5].
- **E4: SLO impact.** Compare inline vs asynchronous verification and compute the incremental latency effect on high-volume sign-in vs high-risk recovery/rotation paths, which differ in baseline workflow latency and throughput constraints [8].

C. Metrics

- Acceptance/rejection rate (VALID vs INVALID under strict verification) [5].
- Receipt size (bytes/KB) under each profile.
- Issuance and verification latency p50/p95 (ms), per profile and container type [4,6].
- Query utility (percentage of audit queries answerable per redaction profile).
- Audit completion rate within a time window (how quickly an auditor can establish decision integrity and policy conformance) [8].
- Optional anchoring check rate (digest presence in append-only log, when enabled) [12].

VIII. Results

Results are synthetic and micro-benchmark based and are intended to illustrate feasibility and predictable trade-offs rather than claim absolute production performance. The core findings support the design goal that Decision Receipts can provide deterministic integrity and explainability with manageable overhead, including under PQC migration profiles [5,9–11].

A. E1 — Verifier Correctness and Tamper Detection

Across 10,000 valid receipts, acceptance was 100% under strict schema validation and key registry resolution. Across 10,000 tampered receipts, rejection was 100% when canonical content changed, required fields were missing, or signature bytes no longer matched the canonicalized payload. This is expected when canonicalization and signature verification are applied deterministically (JCS + JWS/COSE) [4–6]. In practice, this means that common integrity threats—post-incident log edits, ticket rewriting, or outcome manipulation—are detectable at the receipt layer even when underlying component logs are scattered [8].

B. E2 — Overhead

Overhead is dominated by the signature profile and the number of included approvals. Classical profiles remain smallest; hybrid profiles add additional bytes and verification cost; PQC profiles increase signature sizes but remain practical for audit logging and post-incident verification. These trends align with the operational reality of PQC migration, where systems may accept larger artifacts

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

to gain long-term verifiability and crypto agility [9–11]. The use of standardized signing containers (JWS/COSE) supports interoperability and avoids custom verifier complexity [4,6].

C. E3 — Redaction Utility

External/Minimal receipts preserve most compliance checks (e.g., whether quorum was required and satisfied, whether the policy version was recognized, and whether evidence commitment exists). However, collusion analytics (e.g., repeated approver pairs) requires tokenized identifiers available in Internal-Audit views rather than fully redacted external receipts. This supports the "signed view receipts" approach where privacy is preserved without breaking verification, because each view is independently canonicalized and signed [5,8].

D. E4 — SLO Impact

Asynchronous verification minimizes latency impact for high-volume sign-in paths, consistent with continuous verification models where throughput and user latency budgets are tight [8]. Inline verification is practical for high-risk recovery actions where workflow latency is already dominated by step-up and approval processes and where immediate integrity confirmation may be desired. Overall, results support deploying SOMA-DR as an overlay that can be tuned operationally without forcing a single verification mode for all IAM actions [8,9–11].

Table 2. Verification correctness by tamper class.

Tamper class	Example modification	Expected	Rejection rate
No tamper (valid)	none	VALID	0.0% (accepted 100.0%)
Outcome flip	allow → deny	INVALID	100.0%
Evidence hash change	modify evidence_hash	INVALID	100.0%
Approver ID edit	modify approver_id	INVALID	100.0%
Missing required field	delete evidence_hash	INVALID	100.0%
Signature bytes changed	flip signature	INVALID	100.0%
Unknown key id	pubkey_id not in registry	INVALID	100.0%

Table 3. Receipt size and latency by signature profile.

Profile	Size (Sign-in)	Size (Recovery w/2 approvals)	Issue p50/p95 (ms)	Verify p50/p95 (ms)
ECDSA P-256	1.2 KB	1.8 KB	0.45 / 0.90	0.30 / 0.70
Hybrid ECDSA+ML-DSA-44	3.9 KB	4.7 KB	1.40 / 2.90	1.10 / 2.40
ML-DSA-44	3.2 KB	4.1 KB	1.10 / 2.40	0.95 / 2.10
SLH-DSA (SPHINCS+ 128s)	10.6 KB	11.5 KB	5.80 / 12.2	5.10 / 11.4

Table 4. Query utility vs redaction profile.

Query	Internal	Internal-Audit	External/Minimal
Q1 Recovery without quorum	100	100	100
Q3 Approvals < t	100	100	92
Q4 Cozy approver pairs	100	100	0
Q7 Signature posture inventory	100	100	100
Q8 Timing anomalies	100	98	60

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

Table 5. SLO impact of verification mode.

Action	Mode	Baseline p95	With receipts p95	Δp95 (ms)
Sign-in	Inline	180	192	+12
Sign-in	Async	180	182	+2
Recovery	Inline	520	545	+25
Recovery	Async	520	528	+8

IX. Discussion and Limitations

SOMA-DR improves incident response and compliance by turning scattered logs into portable, verifiable artifacts. It bridges explainability and privacy using evidence commitments and redaction profiles, and supports crypto agility for long retention.

Limitations: receipts do not prove upstream evidence is truthful; they do not guarantee policy optimality or fairness; and they do not guarantee independence of approvals without organizational separation-of-duty controls. Evaluation is synthetic and micro-benchmark based; absolute latencies vary, but overhead scales predictably with signature profile and approvals count. Key management assumptions matter: compromised signing keys can enable forged receipts, motivating HSM/KMS protection, rotation, and optional append-only anchoring.

X. Future Work and Conclusion

Future work includes selective disclosure proofs beyond view receipts; enterprise transparency anchoring patterns at scale; standardized reason-code governance and cross-vendor interoperability; control rules as code for automated compliance proofing; longer-horizon PQC operational validation; and deeper integration with human workflows (appeals, approvals, helpdesk).

Conclusion: As IAM adopts adaptive controls and prepares for PQC transitions, decision trails must be explainable and cryptographically verifiable. SOMA-DR decision receipts bind decisions to evidence commitments, record policy context and reason codes, capture quorum approvals when needed, and sign a canonical representation for deterministic verification. Synthetic evaluations indicate strong tamper detection, predictable overhead, and high audit utility under privacy-preserving redaction profiles.

Artifact Availability

We plan to release the SOMA-DR schema, canonicalization rules, a reference verifier/validator, example receipts, and synthetic evaluation scripts under permissive licenses (Apache-2.0 for code; CC BY 4.0 for schema/docs). Only synthetic data is included.

References

1. D. Hardt, “The OAuth 2.0 Authorization Framework,” IETF RFC 6749, Oct. 2012.
2. OpenID Foundation, “OpenID Connect Core 1.0,” 2023 (incorporating errata).
3. M. Jones, J. Bradley, and N. Sakimura, “JSON Web Token (JWT),” IETF RFC 7519, May 2015.
4. M. Jones, J. Bradley, and N. Sakimura, “JSON Web Signature (JWS),” IETF RFC 7515, May 2015.
5. A. Rundgren, “JSON Canonicalization Scheme (JCS),” IETF RFC 8785, Jun. 2020.
6. M. Kucherawy, “CBOR Object Signing and Encryption (COSE),” IETF RFC 9052, Aug. 2022.
7. H. Birkholz et al., “Remote Attestation procedures (RATS) Architecture,” IETF RFC 9334, Jan. 2023.
8. NIST, “SP 800-207: Zero Trust Architecture,” Aug. 2020.
9. NIST, “FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM),” Aug. 2024.
10. NIST, “FIPS 204: Module-Lattice-Based Digital Signature Standard (ML-DSA),” Aug. 2024.
11. NIST, “FIPS 205: Stateless Hash-Based Digital Signature Standard (SLH-DSA),” Aug. 2024.
12. Sigstore, “Rekor Transparency Log,” Documentation.

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 1 January 2026 doi:10.20944/preprints202601.0007.v1

13. S. Nidamanooru, “Identity Refined at the Quantum Gate: Framing the AI + Post-Quantum Challenge for IAM,” Preprints.org, 2025. DOI: 10.5281/zenodo.16989599.
14. S. Nidamanooru, “SOMA-Bench: An Open Synthetic Benchmark and Evaluation Harness for Risk-Aware Recovery & Machine Identities in Post-Quantum IAM,”

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

© 2025 by the author(s). Distributed under a Creative Commons CC BY license.