Supplement S4: Prior-art ledger

to Securing the Semantic Supply Chain: Meaning as a Foundation of Trust Infrastructure

Supplement to the research article

Dr. Carsten Stöcker
Spherity GmbH
carsten.stoecker@spherity.com
ORCID 0009-0003-7417-8872

Version 1.0, 16 November 2026 doi:10.5281/zenodo.24293613



Except where otherwise noted, this work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). © 2026 Dr. Carsten Stöcker.

Numbers in square brackets refer to the reference list of the article.

This supplement belongs to the article *Securing the Semantic Supply Chain: Meaning as a Foundation of Trust Infrastructure* (C. Stöcker, 2026).

The prior-art ledger (S4.1–S4.3) summarises the first search of 12–13 November 2026; one row (definition-level poisoning of agents) already records a result of the second search. S4.4 reports the second, focused search of 16 November 2026 in full, with the verification status of every source.

Searches of 12–13 November 2026, including searches made during independent review (S5): web and standards search; no systematic patent search. Patents found in passing are listed. “Leaves open” is the gap the article addresses, stated narrowly.

S4.1 Spherity’s own work

Item	Status, date	What it solves	What it leaves open
Proofs over Statements v1.3 (PoS)	Published, 9 Nov 2026, CC BY	L1: context change alters meaning under a valid envelope signature; authenticated complete closure fixes it. A1: admitted datasets without blank nodes. §5.4: cheap checks first. L3(iii): graph-commitment authentication. P2: disclosure-induced meaning shift. Def. 5.1: per-fact verifiability. §9.1: a W3C high-assurance profile	How the pinned meaning is produced, released, distributed, verified, attested and withdrawn; demonstrations on production stacks; a concrete L3(iii) construction; the §5.4 residual (forged proof naming an allowed issuer)
The Semantic Supply Chain working paper (WP)	Public draft, 7 May 2026	Uncontrolled resolution is unsafe in every format; four checks; semantic assurance levels 0–4	Attack model, lifecycle controls, evidence
AI Agent Revocation and Containment series	Published, 11 Oct 2026	Revocation as execution control; authority epochs; "stale is unknown"	Withdrawal of *meaning* as an epoch event (M25)
W3C `vc-bitstring-status-list` #218	Open issue, 11 Oct 2026; a proposal, "not a statement of current requirements"	Cover positions, lifecycle and privacy limits; masked encodings	Must-understand rule for all meaning-changing extensions (M21)
Securing Digital Identity ... against Quantum Vulnerabilities	Published, 14 May 2026	Semantic supply-chain attacks as an identity threat class	Operational key management for publishers (M10–M12)
Evidence Graphs for Industrial AI; *EU DPP Registry, W3C Standards and DPP5.0*	Published, 9 Jul and 3 Nov 2026	Evidence graphs; SHACL conformance is not truth	Protection of the validators and vocabularies themselves (M3)
Alamillo, Schwalm, Stoecker, Thiermann, *Qualified Ledgers*	OID 2024, GI	Qualified electronic ledgers under eIDAS 2.0	Use as a transparency backend for semantic artefacts (M9)
Rukanova et al., *Cross-Border E-commerce to Consumers, DPP, and the Role of Authorities* (co-authored)	White paper v1.0, Mar 2026, CC BY	Customs and market-surveillance needs; levels of DPP verifiability; missing role governance for persons of legitimate interest; the full (H1) and reduced (H7) customs data sets and the lack of granularity in H7	The integrity of the meaning itself (UC1, UC2); one view of the evidence for compliance decisions (M31)
Semantic supply chain harness v0.3.0	Released 12 Nov 2026, Apache-2.0	WE-1 to WE-5	— (evidence for the article)

Semantic supply chain harness v0.4.0 [7]	Released 16 Nov 2026; source commit `2ee3f96`, exact archive identified in S5	Corrected software paths; WE-3 engine-protocol simulation, WE-3R replay receipts, WE-6 list matrix and WE-7 synthetic compliance	Production isolation, FROST, complete archival service and broader conformance remain out of scope

S4.2 Standards and regulation

Item	Status, date	What it solves	What it leaves open
W3C VCDM 2.0; `relatedResource`	Recommendation, 15 May 2025	Digest-protected related resources	Optional for issuers; verifier MUST check only if used (VCWG resolution, 26 Sep 2024, PR #1567, after Looker's #272)
W3C Data Integrity 1.0 and ECDSA cryptosuites	Recommendations	Proof formats; §2.4.3 data-loss rule; §4.6 context validation	JCS path does not bind meaning; the libraries tested by Fujie & Suzuki do not apply §4.6 by default (§IV-C2)
`vc-jws-2020`	Discontinued	RDFC hash under a detached JWS (nearest historical design to M17)	No lexical-first ordering, no context commitment, no graph-native re-verification
RDFC-1.0	Recommendation, 21 May 2024	Canonical graph hashing; note on limiting calls	Bounds are library-specific and partly time-based
CBOR-LD 1.0 (`cbor-ld-10`)	Working Draft, 9 Nov 2026	Compact linked-data encoding with registry entries and type tables	Security Considerations empty in that draft; codec drift, registry substitution and decoder-semantics binding (M15)
Verifiable Credential Barcodes v1.0; Verifiable Credential Rendering Methods v1.0	Working Drafts, 22 Aug and 13 Oct 2026	Credentials on carriers; signing before CBOR-LD encoding; render templates	Explicit table binding; mandatory template digests
SD-JWT VC (`vct`, Type Metadata, `#integrity`)	Internet-Draft -26 (6 Nov 2026); IESG Evaluation – AD Followup (13 Nov 2026)	Integrity of Type Metadata	Freshness, rollback, withdrawal, transparency
IETF RFC 9901 (SD-JWT)	Proposed Standard, November 2025	Selective disclosure for JWTs	Meaning of the disclosed claims (Type Metadata, rulebooks; the article, §2)
ISO/IEC 18013-5:2021; ISO/IEC TS 18013-7:2025; ISO/IEC 23220-1:2023; ISO/IEC TS 23220-2:2026 (mdoc)	Published standards and Technical Specifications	Device-bound mdoc presentation; document types, namespaces, data elements	Versioned, withdrawable meaning of data elements and of rulebook-defined attributes across document types (the article, §2)

Item	Status, date	What it solves	What it leaves open
EUDI attestation rulebooks; catalogues of attributes and of schemes (CIR (EU) 2025/1569 Arts. 7 and 8; ARF v3.0.0, 21 Jul 2026)	Rulebooks catalogue on GitHub; catalogue under the implementing act	Structure, semantics and validation per attestation type; attribute reuse	Which definition and version a verifier applied; checked equivalence of one attribute across types when rulebooks change (M22, M28)
W3C VC-JOSE-COSE	Recommendation since 15 May 2025; current edition 13 Oct 2026	Securing VCDM 2.0 with JOSE and COSE	Bytes signed, meaning unbound under context substitution (F1)
OpenID4VCI/VP 1.0; OpenID HAIP 1.0	Final	Protocol-level high assurance	W3C formats and the semantic layer (Paper B)
TUF; Uptane 2.1.0	Specifications	Roles, freshness, rollback, freeze for software	The same for meaning (M8, M14)
SCITT, RFC 9943	Standards Track, June 2026	Architecture for transparent signed statements (single issuer)	A statement profile for semantic artefacts (M9)
SLSA 1.2; in-toto	Specifications	Build provenance and attestations for software	Semantic gates; provenance checked inside the signing boundary (M1, M10)
RATS RFC 9334 (Informational architecture); EAT RFC 9711 (claims format); Linux IMA `CRITICAL_DATA`	RFCs; kernel feature	Attestation; IMA measures the loaded SELinux policy	Semantic bundles as measured components linked to receipts (M24)
eIDAS 2.0; CIR (EU) 2025/1569 Art. 8; ETSI TS 119 612/119 602	Adopted; standards. Arts. 6–9 of CIR 2025/1569 apply from 19 Aug 2026; Art. 8 unchanged by CIR (EU) 2026/1735 (in force 11 Aug 2026). Official text checked by an independent reviewer (an AI model; S5)	Trusted lists for issuers. A machine-readable, sealed catalogue of attestation schemes (Art. 8(6)): scheme status and version (8(3)(c)), authenticated inclusion and modification requests, governance including revocation mechanisms (8(3)(g)), technical specifications (8(7)), identifiers (8(8)); recital 10 on preserving versions	Publishers of meaning in trusted lists. Catalogue-level withdrawal or deprecation transitions and a maximum catalogue age for relying parties, which Art. 8 does not itself specify; the catalogue specifications and operations were not audited (M8, M29)
CIR (EU) 2025/2531, qualified electronic ledgers (eIDAS Art. 45l(3))	Adopted 16 Dec 2025; OJ 17 Dec 2025; in force since 6 Jan 2026	Reference standards and specifications for qualified electronic ledgers	A statement profile for semantic artefacts on such a ledger (M9)

Item	Status, date	What it solves	What it leaves open
European Business Wallet, COM(2025) 838; procedure 2025/0358(COD)	First reading. Council general approach 9 Jun 2026; Parliament committee decision to negotiate 10 Sep 2026 (report A10-0240/2026), confirmed in plenary 7 Oct 2026: a negotiating mandate, not adoption	Legal-effect attestations for businesses	Stable, withdrawable meaning for those attestations (M22, M27)
Battery Regulation (EU) 2023/1542; ESPR (EU) 2024/1781	Adopted; passport from 18 Feb 2027	Passports and data carriers	Lifetime archive of meaning for carriers (M15)
Cyber Resilience Act (EU) 2024/2847	Adopted; reporting from 11 Sep 2026	Vulnerability handling for products with digital elements	Advisories for semantic artefacts (M30)
W3C Data Integrity 1.0 §2.4.3, §5.14	Recommendation, 15 May 2025	Implementations using JSON-LD processing must error on dropped data; network requests are not required, caching is encouraged	A no-egress rule stronger than §5.14 (VER-3); semantic commitment over the full closure (M17)
RFC 6920 (named information); RFC 8493 (BagIt)	RFCs	Content-derived naming; preservation packages with manifests and checksums	An authenticated interpretation and decoding closure with live withdrawal (M3, M15)
RFC 4998, RFC 6283 (evidence records)	RFCs	Long-term evidence preservation, timestamp and hash renewal	Capturing the full semantic and executable decision context (M22, DEC-3, DEC-6)
RFC 9124 (SUIT manifest information model)	RFC	Authenticated dependency and installation manifests for firmware	A CBOR-LD decoding manifest and long-lived credential reconstruction (M15)
W3C PROV-DM, PROV-O; Data on the Web Best Practices	Recommendations, 2013 and 2017	Provenance, versioning, persistent identifiers, preservation	Signed release gates and verifier acceptance rules for semantic artefacts (M1–M4)
RFC 9458 (Oblivious HTTP); OCSP stapling	RFCs	Fetch privacy for specific protocols	Non-interference of resolution traffic for semantic artefacts (PRV-5, T18)

S4.3 Academic and industry work

Item	Date	What it solves	What it leaves open
Fujie & Suzuki, arXiv:2611.04461	v2, 8 Nov 2026	Context injection, SSRF, poison-graph DoS; JCS meaning gap argued	Lifecycle and governance; JCS demonstration; acyclic families (F1, F4)

Item	Date	What it solves	What it leaves open
Fujie & Suzuki, arXiv:2611.03493	v2, 8 Nov 2026	Cost of RDFC (six canonicalisations per exchange); benign schema cost	Commitment-based caching (M17); signed compiled contexts (M5)
jsonld.js #502, #537, PR #552	Open, not merged; the shared cache applies to loader results tagged `static`	Cache behaviour reported as functional issues	Security framing (F3); content-addressed caches (M20)
ROBOT diff, Bubastis, OAK/KGCL; schema-registry compatibility modes; eProcurement and EBSI versioning rules	Various	Ontology and schema change detection	JSON-LD instance-corpus expansion and enforcement at release (M2)
Ruby json-ld preloaded contexts; `jsonld-document-loader` `addStatic()`	Various	Local, preloaded contexts	Signing, sharing and equivalence proof of compiled forms (M5)
Curty et al. 2022; TU Graz *Credential Issuance Transparency* 2024; EBSI; cheqd DID-Linked Resources	Various	Ledger-backed provenance and transparency for ontologies and issuance	Split-view detection for semantic artefacts (M9)
Kritis Signer; DiVerify; US 10972277; US 12229267	Various	Policy-gated or TEE-bound signing	Semantic provenance checked inside the HSM/TEE (M10)
apt-transport-in-toto; kpcyrd `repro-threshold` (Nov 2025)	Various	Client-side k-of-n rebuild checks	Producer-side share release on reproduced digests (M11)
US 12579272	Patent	TPM counters for policies on air-gapped devices	Semantic bundles and offline windows tied to decision risk (M14)
Bi-VAKs (TU Delft 2022); Temporal RDF	Various	Bitemporal versioning of RDF data	Signed decision records with time-scoped selection rules (M27)
SSSOM (Matentzoglu et al. 2022); Health-RI mapping governance	Various	Mapping format with per-record hashing; release practice	Set-level closure, signatures, release tests (M28)
CycloneDX 1.5 data components	Specification	Static inventory of data components	Runtime attestation of loaded meaning (M23)
KG/RAG poisoning (GRAGPoison and others); OWASP MCP03	Various	Instance-, topology- and tool-schema poisoning	Definition-level poisoning (M26), conjectured. Conceptual context poisoning of agent contexts [17] and hash pinning of agent configurations [18] exist; a general or empirical treatment and closure-level pinning were not found (search of 16 Nov 2026)

Item	Date	What it solves	What it leaves open
SD-JWT and SD-JWT VC; BBS signatures (IRTF CFRG draft); redactable signatures (Johnson et al. 2002)	Various	Holder-chosen disclosure of signed statements	Whether the disclosed view gives the same decision as the whole graph for non-monotone compliance policies (M31, T19)
zk-creds (Rosenberg et al., IEEE S&P 2023, doi:10.1109/SP46215.2023.00108; IACR ePrint 2022/878)	2023	Flexible, composable statements over several credentials and existing identity sources with general-purpose zero-knowledge proofs	Policies over semantically rich evidence graphs under a pinned closure, completeness of the evidence population, archived disputes (M31)
Conformity assessment certificates; assurance engagements (for example ISAE 3000); notarised hash deposits	Practice	Third-party attestation of compliance; proof that a document existed	A certificate that binds the commitment to the evidence graph, the semantic closure, the policy and the decision context, with replay as the dispute procedure (M31)
Closed-world assumption (Reiter 1978); non-monotonic and defeasible reasoning	1978 and later	Why withheld facts change conclusions	Applied to selective disclosure as decision-preservation and one-sided soundness criteria, with a cost analysis (T19.1–T19.3)
Threshold signatures (FROST, RFC 9591); MPC signing	RFC; practice	Shared signing keys; FROST is a two-round Schnorr threshold protocol; it does not validate the application message and does not promise robustness against participant abort	Approval of exact receipt bytes after admissibility checks and replay under a pinned closure (M24, T11′)
IEEE Std 802.1AR-2018, *IEEE Standard for Local and Metropolitan Area Networks—Secure Device Identity*, published 2 August 2018; TCG, *TPM 2.0 Keys for Device Identity and Attestation*, Version 1.1, published 23 April 2025 (TCG catalogue/download identifier v1.10r9)	Standards	Device identity bound to protected keys (IDevID, LDevID)	Binding a decision engine's receipt key to an accountable legal person and a decision scope (MON-8)

Item	Date	What it solves	What it leaves open
RFC 8995 (BRSKI); FIDO Device Onboard v1.1; RFC 9683 (remote integrity verification of network devices with TPMs)	RFC 8995: Standards Track, May 2021; FDO 1.1: Proposed Standard, 19 April 2022; RFC 9683: Informational, December 2024	BRSKI bootstraps device trust in a registrar/domain using an IDevID and manufacturer voucher; FDO supports device onboarding through its ownership-voucher protocol; RFC 9683 describes remote integrity verification for network devices with TPM or equivalent protected capabilities	An enrolment record linking onboarding, engine-key provenance and legal-person mandate (MON-8)
RATS architecture (RFC 9334): Attester, Endorser, Verifier, Relying Party	RFC	Roles for attestation evidence and results	Organisational authorisation of receipt keys: an enrolment service may perform the RATS Verifier and Relying Party roles and additionally issue application-level enrolment credentials; this is not automatically an Endorsement (MON-8)
Hyperledger Fabric execute-endorse (Androulaki et al., arXiv:1801.10228)	2018	Endorsing peers execute proposals and sign the results; endorsement policies set the required approvals	A semantic-closure and decision-context binding, admissibility rules and archived replay for semantic decisions (M24 alternative)
CoSi witness cosigning (Syta et al., arXiv:1503.08768)	2016	Diverse witnesses validate authoritative statements before collective signing	The policy each witness evaluates, closure and context authentication, decision replay (M24, M9)
Commit-and-prove proofs (LegoSNARK, IACR ePrint 2019/142); verifiable computation (Pinocchio, doi:10.1109/SP.2013.47)	2013, 2019	Proofs over committed data; succinct proofs of outsourced computation, with a zero-knowledge option	Encoding of semantic policies and closures, completeness of the committed input, measured costs; not implemented or benchmarked here (M31)
Camenisch–Lysyanskaya, anonymous credentials with optional anonymity revocation (doi:10.1007/3-540-44987-6_7)	2001	Anonymous presentations with a controlled accountability mechanism	Accountability of a compliance *computation* over a complete evidence population, rather than identity opening (M31)

Item	Date	What it solves	What it leaves open
Provable data possession (Ateniese et al., doi:10.1145/1315245.1315318); proofs of retrievability (Juels–Kaliski, doi:10.1145/1315245.1315317)	2007	Audits that outsourced data remain stored and retrievable	Retention of every decision-context dependency, opening and executable semantics needed for replay (M31, DEC-13)
WCO Data Model	Practice, maintained by the WCO	Harmonised, reusable regulatory data definitions and messages for customs and border agencies	Integrity, lifecycle and completeness mechanisms for selected decision profiles; a standard data set does not establish the truth or completeness of an instance (M31, DEC-9)
Fedify issue #824	Closed, 25 Jun 2026	Production failure "Maximum deep iterations exceeded" in Linked Data signing	Cause not identified in the issue; a related symptom of F5, not corroboration

S4.4 Second, focused search (16 November 2026)

Scope: the five novelty markers left after the first search (attack class M and M26; F5) and the mechanisms added after the first search of 12 November 2026 (M30, M31 with T19, T11′, MON-8, and the semantic supply chain of non-RDF formats, the article, §2). Method: four independent web and literature searches (arXiv, IACR ePrint, IETF, W3C, GitHub, standards bodies, vendor and grey literature; Google Patents spot checks, mostly unreachable). Each source is marked as checked by us at the primary source, opened by the search, or unverified. "Not found" means not found in the sources searched, not proven absent.

F5: RDF lists with repeated interior values vs RDFC-1.0 work bounds — novelty search, 16 November 2026

Verdict: partially anticipated (cause and symptom reported separately; the finding as a whole not found).

Verified by us:

- pilacorp/go-credential-sdk PR #165, "fix(processor): bound JSON-LD loading and canonicalization of untrusted input", opened 4 Nov 2026, merged 5 Nov 2026: "A repeated structure over the limit fails at both sign and verify, for example more than 102 identical @list values"; commit: "a list of equal values keeps its threshold: 102 pass, 103 do [not]"; "blankNodeWork scores the search URDNA2015 would do by edges and refuses documents above maxBlankNodeWork"; "Value and scoped-term caps can stop an existing credential verifying". The limit is the SDK's own pre-canonicalisation heuristic; no RDFC cost measurement or explanation.
- Fujie & Suzuki, arXiv:2611.04461v1 (3 Nov 2026), §III-C7: "That default is not free of false positives" — rdf-canonize 5.0.0's maxWorkFactor=1 "also aborts the acyclic forest inputs" and 18 RDFC-1.0 test-suite cases; a tight limit "can also reject legitimate credentials with

many structurally similar blank nodes". No lists (first 100k characters checked; the search agent reports the full text has none). Their own legitimate credentials pass at every bound.

Reported by the search agent (opened via a summarising fetch tool; quotes not checked by us):

- pyld issue #368 (2 Nov 2026) and json-gold issue #110 (2 Nov 2026): same false-positive data; no lists.
- rdfjs/N3.js PR #856 (Oct 2026): default rejects small symmetric graphs (cycles); maxWorkFactor 3 accepts all W3C tests.
- rdf-canonize CHANGELOG 4.0.0 (15 Nov 2023): maxWorkFactor 1 "will handle common graphs" — the premise F5 contradicts.
- rdf-canonize issue #46 (2022): deep recursion "extremely non-typical in useful data"; w3c/rdf-canon #120 (2023): useful data expected to fail rarely.
- W3C RDFC-1.0 test suite test053 (@list) uses distinct values; no test with repeated list values.
- Mirabolic report for Digital Bazaar (Oct 2020): slow simple cases (cliques), no lists.
- pyld #179 / PR #336: benign OWL-reasoner output blew up (19-node ambiguous set), no rdf:first/rest.
- w3c/rdf-canon #233 (Oct 2026): erratum, isomorphic datasets with different canonical forms (correctness).

Resulting novelty statement: "Partially anticipated (search-relative, 16 November 2026): a credential SDK's own pre-canonicalisation guard refuses more than 102 identical @list values (pilacorp/go-credential-sdk PR #165, Nov 2026), and rdf-canonize 5.0.0's default bound is reported to reject some acyclic inputs and RDFC-1.0 test cases (Fujie & Suzuki 2026a, §III-C7); we found no prior report that lists with repeated interior values make RDFC-1.0 super-linear without a bound or that they cause legitimate credentials to fail under the library default."

Not reached: GitHub API (web pages used); other canonicaliser repositories issue by issue (sophia_rs, RDF.ex, dotNetRDF, oxigraph); mailing-list archives directly; Google Patents.

M26 / attack class M: definition-level poisoning of agents — novelty search, 16 November 2026

Verdict: partially anticipated. The line "definition level not found" no longer holds as written.

Verified by us (quote checked in the downloaded text):

- Fleming, Muscariello, Pandey, Kompella, *A Layered Protocol Architecture for the Internet of Agents*, arXiv:2511.19699 (v1 24 Nov 2025; v3 20 Jan 2026), §5.1: "Context Poisoning / Spoofing: A malicious agent (or a compromired registry) serves a modified version of a popular Shared Context . . . it swaps the definitions of 'sender_id' and 'receiver_id' "; also a "Semantic Downgrade Attack"; §5.2.2 defence: cryptographically signed, versioned context definitions; §5.3 Schema Authorities with revocation. Position paper, no experiments; agent message contexts, not signed data/VCs, TBox, SHACL or crosswalks; signer-based trust, not closure digests.
- OWASP GenAI Security Project, *OWASP Top 10 for Agentic Applications*, v2026 (Dec 2025): ASI04 mitigation 7 "Pinning: Pin prompts, tools, and configs by content hash and commit ID"; ASI07 "Semantics split-brain", typed contracts and schema validation. Not vocabularies or data-interpretation definitions.

Reported by the search agent (opened by it; not re-checked by us):

- Invariant Labs, *Introducing MCP-Scan*, 11 Apr 2025: tool pinning by hashing tool descriptions.
- Wang et al., *Data Agents Under Attack*, arXiv:2606.08661 (7 Jun 2026): poisoning of schema comments and description fields for data agents.
- Kereopa-Yorke et al., *Oracle Poisoning*, arXiv:2605.09822 (10 May 2026): instance-level KG poisoning for MCP agents; per-entity signed provenance; ontologies named as future work.
- CSA research note, *Agent Context Poisoning: SKILL.md . . .*, 6 May 2026: content-hash verification of agent instruction files.
- MITRE ATLAS AML.T0080 (agent context poisoning), AML.T0110.000 (tool definition poisoning), read via D3FEND and a mirror.
- Atlan blog, *Context Poisoning*, 29 Apr 2026: stale definitions as accidental drift.
- Unverified (403): Telasang, *Ontology Poisoning in Military AI Fusion Systems*, SSRN 6970500 (Jun 2026); Syed, Silaghi, Abujar, *JSON-LD 1.2 and Beyond*, FLAIRS-39 (2026), with the `jsonld-ex` package (context integrity hashes).

Not found: a general treatment of definitions (vocabulary terms, TBox, JSON-LD contexts, SHACL shapes, crosswalks, rulebooks) changing how an LLM agent interprets unchanged signed data; any experiment; pinning the agent's full semantic closure by digest.

Resulting novelty statement: "Partially anticipated (search-relative, 16 November 2026): context poisoning by swapping term definitions in a shared agent context is described conceptually, with signed, versioned contexts as defence, by Fleming et al. (2025), and content-hash pinning of agent prompts, tools and configurations appears in OWASP ASI04 (2025) and MCP tool pinning (2025); we found no prior work that generalises this to vocabularies, ontologies, JSON-LD contexts, SHACL shapes and crosswalks, demonstrates it against agents interpreting unchanged signed data, or pins the agent's full semantic closure by digest."

Not covered directly: CCS, USENIX Security, S&P, NDSS, ISWC/ESWC proceedings (only as surfaced by web search); Google Patents (unreachable); NIST AI 100-2 full PDF.

M31 (compliance-evidence profile, T19) and M30 (signed semantic advisories) — novelty search, 16 November 2026

M31: partially anticipated (elements), combination not found

Verified by us:

- IETF RFC 9943, *An Architecture for Trustworthy and Transparent Digital Supply Chains* (SCITT), Standards Track, June 2026. §9.3: "Issuers can refuse to register their Statements with a TS or selectively submit some but not all the Statements they issue." §5.1.1.2: "enough information is made available to Auditors to reproduce the Registration checks".
- US 12645555 B1 (Stone & Ink Archive LLC; inventor Yong Bok Lee), *Systems and methods for deterministic conformance suites, validator implementation certification, and verifiable certificate registries for replay-verifiable conformance receipts*; priority 9 Jan 2026, filed 10 Jan 2026, granted 2 Jun 2026 (Google Patents). Claims: conformance certification of validators (claim 1, 11, 22); dependent claim 6: a conformance context digest over the canonical encoding of suite, policy, validator and reason-code registry digests; claim 20: declared coverage by receipt type/action class/control point. The description (not the

claims) has an offline-replay audit mode and escrow. No salted commitment, no coverage list of evidence items, no custodian replay before acknowledgement in the claims. **Prior art for receipts committing to policy/schema/canonicalisation digests and for offline replay; relevant to M22 and M31.**

Reported by the search agent (opened by it unless noted):

- W3C CCG mailing list, "Re: Selective Disclosure of lists", 8–9 Jun 2021 (omission problem; issuer-side count or no-SD policy).
- SLSA v1.0 Verification Summary Attestation (policy URI, input-attestation digests, verification result).
- in-toto layout verification (signed required-evidence policy; thresholds of link metadata per step).
- UN/CEFACT UNTP Transparency Graphs v0.7.0 (validity of every credential does not verify the graph); UNTP Digital Conformity Credential v0.6.0 (hash of encrypted evidence file-set).
- Nexus OSI Annex G (grey literature): escrowed evidence, in-zone replay, receipts bound by commitments.
- Waiwitlikhit et al., ZkAudit, ICML 2024 (commit-then-audit with ZK proofs).
- Search-result only: idemix inspector (verifiable encryption); DSAC signature chaining for completeness (ePrint 2005/297); Decision Provenance (IEEE Access 2019); Moghaddam, Compliance-by-Construction Argument Graphs, arXiv:2604.04103 (opened).
- Negative: W3C Security IG Threat Model for Decentralized Credentials (8 Sep 2026) has no completeness/omission threat.
- Unverified (403): US 12671589 ("proof-of-policy-compliance receipts").

Resulting novelty statement (M31): "Integration (search-relative, 16 November 2026): the omission problem for selective disclosure (W3C CCG, 2021; RFC 9943 §9.3), signed required-evidence policies (in-toto), verifier summary attestations binding input digests (SLSA VSA), receipts committing to policy, schema and canonicalisation digests with offline replay (US 12645555 B1, 2026) and escrowed-evidence replay exist; we found no design that combines an independently signed coverage contract over a decision's sufficient evidence set with a salted, semantics-bound certifier commitment and custodian replay before acknowledgement, and no formal result that coverage rather than the disclosure format decides whether omitted adverse evidence is caught."

M30: partially anticipated

Reported by the search agent: nanopublication retraction (`npx:retracts`, valid only if signed with the same key; retracted items drop from search); OBO Foundry Principle 19 (owl:deprecated, term replaced by, pre-announced obsoletion); W3C `vs:term_status` (2009); schema.org `supersededBy`; ETSI TS 119 612 trusted-list service status "withdrawn"; Crossref Crossmark typed update notices (search-result only); Luccioni et al., A Framework for Deprecating Datasets, FAccT 2022; AI-VEX/CSAF for AI components (search-result only). Negative: no withdrawal/advisory status found for EBSI TSR, EUDI catalogue or SD-JWT VC Type Metadata.

Resulting novelty statement (M30): "Novel combination (search-relative, 16 November 2026): machine-readable term obsoletion (owl:deprecated, OBO Principle 19, schema.org supersededBy), signed retraction of individual assertions (nanopublications), signed withdrawal of trust-service status (ETSI TS 119 612) and typed update notices (Crossmark) exist; we found no publisher-signed, separately distributed advisory that marks a vocabulary term, definition or rulebook withdrawn, deprecated or known-problematic for verifiers to check at decision time."

Not covered directly: CCS, ESORICS, FC, PETS proceedings; DIF and ToIP repositories; JRC reports.

T11′ (replay-validated threshold receipts), MON-8 (enrolment of decision-engine keys), the semantic supply chain of non-RDF formats (the article, §2) — novelty search, 16 November 2026

T11′: partially anticipated

Verified by us: Zheng & Zhang, *From Review to Authorization: Key-Isolated Threshold Signing for LLM Agents*, arXiv:2609.05901 (v1 5 Sep 2026, v2 23 Oct 2026), abstract: authorisation needs "signing contributions from t distinct domains"; "an adversary controlling the proposer and fewer than t reviewer-signer domains cannot produce a valid authorization for a new action". Abstract does not say signers replay the decision (the search agent reports the body binds policy version and an evidence commitment; not checked by us). Reported by the search agent: EigenAI, arXiv:2602.00182 (Jan 2026; committee re-execution on challenge, byte-equality threshold vote, inside TEEs); Chainlink OCR3 (May 2025; reports signed by f+1 oracles); Chainlink CCIP Risk Management Network blog (Oct 2023; independent N-version recomputation); US 12,256,007 B2 (VMware; BFT replicas each verify credentials and threshold-sign an access token; granted 18 Mar 2025); draft-chueayen-attestation-receipts-05 (25 Oct 2026; single-signer decision receipts with request-hash commitment; "application-layer evidence of a decision, not hardware evidence of the runtime"); ERC-8004 Trustless Agents (re-execution validation); arXiv:2603.14332 (replay-based verification as software alternative to TEEs). Resulting novelty statement: "Partially anticipated (search-relative, 16 November 2026): threshold-authorised agent actions (Zheng & Zhang 2026), committee re-execution with threshold acceptance (EigenAI 2026), quorum-signed oracle reports computed independently by each node (Chainlink OCR3), BFT-replicated threshold-signed access tokens (US 12,256,007 B2) and single-signer decision receipts (draft-chueayen-attestation-receipts) exist; we found no scheme in which each of t co-signers replays a credential or policy decision from its committed, admissible context and semantic closure before threshold-signing that exact receipt, nor one that scopes the guarantee to the committed context rather than the device's execution."

MON-8: partially anticipated (elements)

Verified by us: draft-ni-wimse-ai-agent-identity-03, *WIMSE Applicability for AI Agents* (Huawei; Sandelman), 1 Oct 2026: Owner = "An entity (individual or organization) accountable for the operation of an agent"; Dual-Identity Credential "contains the identifiers and associated public keys of both an agent and its owner"; no software/model version, key epoch or scope in the credential; security section "TODO". Reported by the search agent: NIST NCCoE concept paper on software and AI agent identity (Feb 2026; open question whether agent identities should be tied to hardware, software or organisational boundaries); draft-ietf-wimse-aims-01 (2 Nov 2026); draft-messous-eat-ai-01 (Feb 2026, expired; EAT claims for model integrity); GLEIF blogs (1 and 21 Sep 2026; mandates, execution trust); GLEIF *Agentic AI in Payments* (13 Aug 2026; Agent Mandate Credential — PDF 404, unverified); DIF KYA-OS (Apr 2026; scoped delegation credentials); Visa/Mastercard agent key directories (Cloudflare blog, Oct 2025); Microsoft Artifact Signing legal-entity validation; TCG DICE; AI Act Annex VIII registration. Resulting novelty statement: "Partially anticipated (search-relative, 16 November 2026): agent keys bound to an accountable owner (draft-ni-wimse-ai-agent-identity-03) or legal entity (vLEI; code-signing identity validation), scoped mandates (GLEIF; DIF KYA-OS) and attested model or firmware identity (EAT profiles; DICE) exist, and NIST NCCoE lists the binding of agent

identity to hardware, software and organisation as open; we found no enrolment record that binds a decision-receipt key or threshold-group configuration, key epoch, device, approved engine and key-policy versions, responsible legal person and decision scope, nor a rule that a verifier must check it before relying on a receipt."

Semantic supply chain of non-RDF formats (the article, §2): partially anticipated; one contrary framing

Verified by us:

- Fujie & Suzuki, arXiv:2611.04461 (v2 8 Nov 2026): abstract "SD-JWT VC and mdoc need neither canonicalization nor external references"; §III-D2 "mdoc has no counterpart to the external vocabulary document or to a resolvable type definition"; §III-B2 "the meaning of claim names is left to the type definition that vct denotes"; §IV-C1 places "resolution of credential type metadata" outside the evaluated scope.
- OAuth WG, "[OAUTH-WG] SD-JWT-VC-13 Review Memo", 19 Nov 2025 (a WG participant's review, comment NS21): "If the Verifier uses an old, cached copy, it allows a weaker (less secure) presentation to be accepted"; proposed text "the Verifier MUST check the metadata's age against the credential's exp claim" (proposal, not draft text). Reported by the search agent: German BMDS EUDI architecture concept, *Catalog of attestations* (rulebooks signed with access certificates, SRI #integrity on rulebook and schema URIs, SemVer, provider may remove entries); EUDI ARF Topic O (catalogue history, versioning, removal conditions); EC OOTS–EUDI Synergy 7 deck (19 Nov 2024; semantic repository for rulebooks); German PID SD-JWT VC ↔ mdoc crosswalk; arXiv:2611.03493 ("no semantic interoperability" for SD-JWT VC, search result). Resulting novelty statement: "Partially anticipated (search-relative, 16 November 2026): integrity hashes and caching rules for SD-JWT VC Type Metadata, with a reviewer's proposal on staleness (SD-JWT VC drafts; OAuth WG review, 2025), and signed, versioned, hash-pinned rulebook catalogue entries (German EUDI architecture concept; ARF) exist, and a recent cross-format analysis treats SD-JWT VC and mdoc as needing no external references while leaving type-metadata resolution out of scope (Fujie & Suzuki 2026a); we found no work that treats Type Metadata, mdoc namespace definitions and EUDI rulebooks and catalogues as one semantic supply chain with JSON-LD contexts, or maps rollback, freshness, withdrawal, applied-version binding and governed crosswalks across these formats." Relation to the article, §2: Fujie & Suzuki is right about verification time (no dereferencing while the signature is checked); the article's claim concerns the lifecycle of the definitions those formats point to, which they place out of scope.

Not reached: USPTO PDFs 12645555/12726364 (403; 12645555 checked by us via Google Patents); GLEIF AMC PDF (404); some Chainlink docs pages (404).