Supplement S5: Review and reproducibility record

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.

Contents

Method and the seven rounds of independent review 3

S5.1 Iterations 3

S5.2 First independent review (12 November 2026): findings and disposition 4

S5.3 What the model checker is, and is not 6

S5.4 Second independent review (13 November 2026): findings and disposition 6

S5.5 Third independent review (13 November 2026): findings and disposition 10

S5.6 Fourth independent review (13 November 2026): the harness code 13

S5.7 Fifth independent review (14 November 2026): the release candidate and the paper 16

S5.8 Sixth and seventh reviews (14 and 15 November 2026): the release gate 19

S5.9 Harness release record 22

S5.10 A discrepancy in the v0.3.0 release record 22

S5.11 Pre-publication review (16 November 2026): the article, the supplements and the Executive Brief **22**

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 formal model, its proofs and the model checking are given in S3. The model checker (`model/mc.mjs`, 66 cases, strict `--check` with non-vacuity checks and the declared safe-but-unavailable `MC6.offline.deferEffect` exception) and the experiments are released with harness v0.4.0. Bounded model checking covers the instances checked, not every size; the general statements are the theorems.

The reviews recorded here examined the model, the harness and the manuscript from which the article derives: a claims note in successive versions (v0.9.x). Section references (§) in this record refer to the sections of that manuscript, which the article reorganises; its appendices correspond to the supplements (Appendix A to S2, B to S4, C to S1, D to S3, D.0 to this supplement).

Method and the seven rounds of independent review

All reviews recorded here were carried out by AI models; they are not peer review. A pre-publication review of the article, the supplements and the Executive Brief followed (S5.11).

The model was built in seven iterations. After the third, an independent reviewer (a Claude subagent) re-ran every experiment and model check and reported twenty findings. After the fourth, a second independent reviewer (OpenAI GPT-6 Astra) worked through the complete material in eight passes. It ran the harness and model checker in its own sandbox, mutated the model, wrote its own probes, checked every standards and legal statement against primary sources, and reported thirty-three findings. Each of these reviews included three critical findings. A third review by the same reviewer examined the material added after the author's decisions (T11′, T19, M31) and reported twenty-eight findings, one critical. A fourth review by the same reviewer examined the code of harness v0.4.0, reproduced every result, and reported thirty-three findings, seven critical, with an integrated patch and regression tests. A fifth review examined the release candidate and the paper together, against primary sources, and reported twenty-three findings, two critical, with a patch and exact text replacements. We reproduced the decisive findings before accepting them, and accepted all findings of these five reviews; the dispositions follow. Two release-gate rounds came after them (S5.8).

S5.1 Iterations

Iteration	Model	Defect found	Consequence
v1	Signature + dataset	No separation of RDF meaning from IRI meaning; no time, state or attestation	Meaning layers, verifier state
v2	Layers, state, time, network adversary	Register reset treated apart from arbitrary extend	Device adversary classes

Iteration	Model	Defect found	Consequence
v3	Taint appraisal, register kinds	(first review) reboot; in-application decisions; MC-4 assumed verdicts; version counters allowed un-withdrawal; loose windows	v4
v4	Attested engine with PCR-bound key; metadata counters; skew; mutants	(second review) PCR-only key callable by the host; closure missed embedded contexts; freshness bound on the wrong object; overloaded digest term; redundant guards reported as necessary; check incomplete; privacy overclaims	v5
v5	Isolated engine with exclusive key use, or a threshold key with replaying co-signers; processing-aware closure; typed digests; status-anchored freshness; OBJ-11; T11′, T17–T19; 51 model-checking cases	(third review) V6 guard assumed execution knowledge; T19 used a view count as a lower bound; completeness, admissibility, scope, lifecycle, archive and transcript privacy unspecified	v5.1
v5.1	T11′ as replay-validated outcomes under threshold EUF-CMA; T19 rebuilt (information, cost, coverage); MC-6; 65 cases	(fourth review, code) MC-6 never reached the withdrawal guard; T11 did not separate exclusive use from attribution	v5.2
v5.2	T11 (ii-a)/(ii-b); T11′ receipt v2 and participant enrolment; MC-6 online-before-withdrawal schedule; 66 cases	(fifth review) T11's attribution conclusion lacked actual-use, relevant-time and compromise premises; the enrolment chain was stated as one sequence	T11 restated within v5.2 (governance-attribution corollary); model checker unchanged; the manuscript

S5.2 First independent review (12 November 2026): findings and disposition

The reviewer re-ran all experiments and model checks, wrote eleven probe scripts, and reported 20 findings. We reproduced the critical ones before accepting them (S3.5, X7–X11).

#	Severity	Finding (short)	Disposition
1	critical	Reboot or TPM reset erases the log; taint fails; swtpm's control channel resets PCR 14	Accepted. T11 rewritten (attested engine, V4); MC-3 models reboot; swtpm reset noted as a software-TPM artefact
2	critical	A compromised application can decide in its own code with unmeasured semantics	Accepted. **Central correction of M24**: the decision engine must be inside the measured boundary
3	critical	M15 manifest not transitive; JCS inner proof does not bind meaning; MC-4 assumed its verdicts	Accepted. T10 and DST-9 corrected; MC-4 rebuilt as a case analysis over drift kinds and proof kinds; reproduced (X7, X8)
4	major	Stage 1 quadratic (base58 of an unbounded `proofValue`); "linear impossible with JCS" false	Accepted. T3 premise "bounded pre-authentication decoding"; radix ordering makes JCS O(n); reproduced (X11)
5	major	No key-to-issuer binding; `verifyFromQuads` without allow-list	Accepted. T2 premise; VER-1 sharpened; reproduced (X10)
6	major	T9.2 source list incomplete; "iff" wrong	Accepted. Sufficiency only; design changed to a certified closure (M5)
7	major	Replay undoes a withdrawal; MC-2 adversary could not replay	Accepted. Metadata counter, permanent withdrawal; MC-1/MC-2 v4
8	major	T7: k witnesses with one honest is not enough	Accepted. Quorum intersection
9	major	T13 false for closed-world verifiers	Accepted. Open-world premise; trade-off stated
10	major	T14 receipt fields insufficient; OBJ-6 inconsistent with L4	Accepted. Decision context defined; OBJ-6 redefined
11	major	I-JSON is more than duplicate names	Accepted. VER-14 and $\mathbb{J}$ redefined
12	minor/major	Use binding (VER-11) missing from T2; base IRI and options	Accepted
13	minor	F5 mechanism imprecise; bound semantics differ between library versions	Accepted. T9.3 restated; κ records library and bound semantics
14	minor	MC windows not tight; skew not modelled	Accepted. Exact windows, skew variant, mutants
15	minor	T4(b) non-interactive premise; T5 fail-closed premise; single online status key; OBJ-1 no theorem	Accepted
16	minor	Taint's operational cost; duplicate decision events; PCR 14 used by shim	Accepted (moot in V4; noted for real platforms)
17	minor	T15 overclaims	Accepted: "removes backdating"
18	minor	Only successes cached; ECDSA malleability	Accepted. Negative cache; normalised key
19	minor	Script and number errors (+17 bytes, not +16; X4 is not WE-2's configuration; multi-target bound)	Accepted
20	minor	T1.1 "iff"	Accepted: "if"

S5.3 What the model checker is, and is not

`mc.mjs` explores small instances exhaustively and prints a shortest counterexample for every violation. A "holds" covers the instance checked, not every size; the general statements are the theorems.

v5 adds four kinds of check:

- **Non-vacuity:** every "holds" case must reach an acceptance, use or commit, except `MC6.offline.deferEffect`, which deliberately has no acceptance and demonstrates safety at the cost of availability.
- **REDUNDANT rows:** single-guard removals that still hold in the instance. They show what the instance *cannot* demonstrate; they do not mean the guard is unnecessary.
- **Tightness mutants:** variants that must be violated.
- **A strict check:** every case ID present exactly once, with the expected verdict.

S5.4 Second independent review (13 November 2026): findings and disposition

The reviewer received the self-contained review pack (SHA-256 `9a8e0fd3...`) and the public harness v0.3.0. It worked in a sandbox and ran:

- the harness: 20 pass, 1 skip without swtpm;
- the model checker: 32 as expected;
- an isolated guard-mutation study;
- a nested-context probe and a cache probe;
- the F5 matrix on two library stacks.

We reproduced R2-02 (X12), R2-15 (X13) and R2-01 (in the model: MC3.V4.D1) before accepting them. Against official sources, we checked the Federal Register notice (91 FR 32985), the EU customs reform (Regulation (EU) 2026/2108, adopted 16 September 2026) and the AI Omnibus dates.

#	Sev.	Finding (short)	Disposition	Integrated in
R2-01	critical	PCR-only key policy: root can call the engine key	**Accepted, reproduced in model** (V4 fails in 2 steps); isolated engine with exclusive key use required	T11; MC-3 V5; MON-2, MON-5; M24
R2-02	critical	Closure misses contexts embedded in the document	**Accepted, reproduced** (X12)	S3.1.2 dep(); T1.1; VER-16; PRD-3; M17
R2-03	critical	Per-credential HMAC IRIs link presentations	**Accepted**; claim narrowed to cross-credential separation	PRV-1; M18; §4.4
R2-04	major	Freshness bound on the wrong object	**Accepted, model-checked** (MC-5)	T5; DST-11; M8

#	Sev.	Finding (short)	Disposition	Integrated in
R2-05	major	"Closure digest" names three things	**Accepted**	S3.1.3 typed digests; M3, M22, M24
R2-06	major	"Every safeguard necessary" false	**Accepted, reproduced** (REDUNDANT rows)	S3.4; §0, §2.5
R2-07	major	MC-4 assumes I-RDFC	**Accepted**	MC-4 relabelled as conditional
R2-08	major	EUF-CMA proves origin, not honest computation	**Accepted**	T2 (a)/(b)
R2-09	major	OBJ-1 unscoped	**Accepted**	OBJ-1; T17; GOV-7
R2-10	major	Metadata vs target monotonicity; byte vs graph integrity	**Accepted**	T6; OBJ-2 (a)/(b); DST-3
R2-11	major	Reject-all satisfies all safety objectives	**Accepted**	OBJ-11; T4; PRF-6; non-vacuity in MC
R2-12	major	Fork exclusion ≠ conflict detection	**Accepted**	T7; DST-13
R2-13	major	Decoder semantics not bound	**Accepted, case analysis extended** (MC-4 full manifest)	T10; DST-12
R2-14	major	Admission certificate unspecified	**Accepted**; trusted-certifier assumption stated until a checker exists	T9.2; PRD-10; VER-4
R2-15	major	Cache reuse ignores processing context	**Accepted, reproduced** (X13)	PRF-2, PRF-5; M17
R2-16	major	Stage-1 "light" mode can authorise semantic actions	**Accepted**; light mode removed: no semantic authorisation before stage 2	M17
R2-17	major	Authentic subset can hide a restriction	**Accepted**	VER-8; M21
R2-18	major	Timestamps certify existence, not event time; M27 used `created`	**Accepted**	T15; DEC-2; M27
R2-19	major	Operation counts ≠ determinism; archive scope	**Accepted**	T8; T14; DEC-6
R2-20	major	No-egress still leaks through the fetcher	**Accepted**	OBJ-9; T18; PRV-5; DST-10
R2-21	major	Key commitment ≠ auditable assignment	**Accepted**; M13 marked incomplete (E1)	M13; PRV-2
R2-22	major	Module and ownership conflicts; MUST in Dep	**Accepted**	S2 profile column
R2-23	major	I-JSON numeric rule misstated	**Accepted**	§2.1 $\mathbb{J}$; VER-14
R2-24	major	Corpus results promoted to universal guarantees	**Accepted**	PRD-2, PRD-4, PRD-5; M2, M4, M5
R2-25	major	E4 of predecessors attached to corrected designs	**Accepted**	§5 evidence matrix; §7 register

#	Sev.	Finding (short)	Disposition	Integrated in
R2-26	major	Legal and standards attribution errors	**Accepted**; primary-source ledger applied	§1.2, §3, §6, S4, PRD-7, PRV-3, VER-3
R2-27	major	Off-graph IRIs and digests are still personal data	**Accepted**	PRV-4
R2-28	major	PQ countersignature contract; hash bounds need an idealised model	**Accepted**	M12, M15; KEY-5; DEC-3; T10
R2-29	major	t-of-n $\neq$ t honest rebuilds	**Accepted**	M11 (t $-$ f honest); KEY-3
R2-30	minor	F5 explanation scope; "not exploitable"	**Accepted**	T9.3; §8 of the paper
R2-31	minor	Regression check validates a prefix only	**Accepted, fixed and self-tested**	MC v5
R2-32	minor	+23 bytes is one configuration	**Accepted**	PRF-3; §0 takeaway
R2-33	minor	"Per-fact" overstates whole-dataset re-verification	**Accepted**	T2.1; M17; DEC-3

Refinements we made to the reviewer's proposals.

- **Novelty for new requirements:** the reviewer gave none. We assigned:
  - ★ to PRD-10, DST-12 and VER-16;
  - ◐ to the others with a named basis;
  - ◯ to VER-15, which is in Data Integrity already.
- **Profile labels:** We normalised them to Core, Core (if ...), Opt: module, Ref, Dep.
- **M17 light mode:** removed rather than restricted.
- **F5:** we kept the reviewer's wording that rejection can be an availability problem, and we report the independently reproduced matrix as E4 rather than E5. The reproduction ran the same libraries in a sandbox; it is not a field trial.

Primary-source ledger (applied).

Item	Result	Applied wording
Battery passport scope (Reg. (EU) 2023/1542 Art. 77)	LMT batteries, industrial batteries above 2 kWh, EV batteries (Parliament source); date 18 Feb 2027 retained with Art. 77(1) citation	UC1, §6.1
EU customs reform	Adopted 16 Sep 2026; Regulation (EU) 2026/2108; entered into force 20 September 2026; phased implementation (**reproduced**)	§6.1
AI Act high-risk dates after the AI Omnibus	Annex III from 2 Dec 2027; Annex I from 2 Aug 2028 (**reproduced**)	§6.1
CRA	Reporting from 11 Sep 2026; main obligations from 11 Dec 2027	§6.1, MON-3
Federal Register 2026-10979	91 FR 32985; announced 2 Jun 2026; test from 2 Jul 2026 for one year (**reproduced**)	UC2; verification marker removed
HS 2028	WCO announced acceptance on 21 Jan 2026; in force 1 Jan 2028	UC2

Item	Result	Applied wording
SD-JWT VC	draft -26 (6 Nov 2026), IESG Evaluation – AD Followup (13 Nov 2026)	S4
BSI AIS 31 (2024 reference v3.0)	DRG.2, DRG.3, DRG.4, DRT.1, PTG.2, PTG.3, NTG.1	M13
W3C	VCDM 2.0, DI 1.0 and ECDSA cryptosuites: Recommendations 15 May 2025; CBOR-LD 1.0 (`cbor-ld-10`) WD 9 Nov 2026 with empty Security Considerations in that version; *VC Rendering Methods* WD 13 Oct 2026	S4
RFC 9334	Informational architecture	S4
DPP Registry package	Dated 30 Oct 2026; contains JSON-LD, RDF/SKOS and SHACL (inspected by the reviewer)	§1.2
Community Group	Proposed 11 Nov 2026 by Phil Archer; S2 is our proposed baseline, not adopted	S2, §6.2
#218	Open proposal, explicitly "not a statement of current requirements"	PRV-3, M21
jsonld.js PR #552	Open, not merged; shared cache needs a loader result tagged `static`	F3
Fedify #824	Closed; comments do not identify lists	F5 (symptom only)
EBW negotiating mandate (searched after the author's decisions)	OEIL procedure file 2025/0358(COD): committee decision to open interinstitutional negotiations (10 Sep 2026; report A10-0240/2026, tabled 23 Sep 2026), announced in plenary 5 Oct 2026, confirmed 7 Oct 2026. Council general approach 9 Jun 2026 (Council press release; page not retrievable from the sandbox)	§1.2, §6.1, S4; verification marker removed
Qualified-ledger implementing act	CIR (EU) 2025/2531, reference standards and specifications for qualified electronic ledgers (eIDAS Art. 45l(3)): adopted 16 Dec 2025, OJ 17 Dec 2025, applies from 6 Jan 2026 (EFTA EEA-Lex record)	§4.3 (M9), §6.1, S4; verification marker removed; superseded by the round-3 ledger (in force since 6 Jan 2026)
CIR (EU) 2025/1569 Art. 8	Read in a reproduction of the text and EUDI issue #662; EUR-Lex was not reachable from the sandbox. Registration requests include status and version (8(3)(c)); the catalogue is machine-readable and sealed for integrity and authenticity (8(6)); unique identifiers (8(8)); no withdrawal or deprecation provisions found. Consolidated 11 Aug 2026 after CIR (EU) 2026/1735, whose changes were not checked	§6.1, M29, S4; claim narrowed; superseded by the round-3 ledger, which checked the Official Journal

Prior art added (S4):

- RFC 6920 (named information);
- RFC 8493 (BagIt);
- RFC 4998 and RFC 6283 (evidence records);
- W3C PROV-DM and PROV-O;
- W3C Data on the Web Best Practices;
- RFC 9124 (SUIT manifest information model).

These narrow the novelty of M3, M4, M15, M22 and DEC-3 to the specific semantic profile and its integration.

S5.5 Third independent review (13 November 2026): findings and disposition

Scope and method.

- **Input.** The reviewer received the focused round-3 pack (SHA-256 `0496bf51...`) and the full paper v0.9.3 (`4b0d8cde...`).
- **Runs.** It ran the model checker (51 as expected) and three mutations of MC-3. It also wrote 13 standalone bounded extensions on the premises of replay co-signing, a synthetic test of certificate privacy, and the F5 reproduction (24 of 24 classifications reproduced).
- **Sources.** It checked the EU acts against the Official Journal (Publications Office XHTML).
- **Findings.** 28: one critical, nineteen major, eight minor.

Our reproductions before acceptance:

- **R3-02.** With the co-signer's guard changed to outcome replay, the execution-identity property fails, and the outcome property holds. Now in the checker as MC3.V6x.D1 and MC3.V6.D1.
- **R3-05, R3-06.** The reviewer's 13 extensions, rebuilt as MC-6 in our checker; all 13 verdicts agree.
- **R3-12.** The two-candidate dictionary recovers G from the public context digest without the salt. Binding h instead closes it (`round3/r3-12-digest-dictionary.mjs`).
- **R3-24, R3-25.** Checked in the `rdf-canonize` 5.0.0 source:
  - the reference blank node is `_:a` and all others `_:z`;
  - the README says `maxDeepIterations` defaults to `Infinity`, but the code defaults to `-1` and derives the budget.
- **R3-01.** The counterexamples are elementary and correct. For "deny iff e is present", one membership test decides exact preservation.

EU facts. EUR-Lex and the Publications Office were again not reachable from our sandbox. We adopt the reviewer's Official Journal reading and record that we did not retrieve the texts ourselves.

All 28 findings are accepted.

#	Sev.	Finding (short)	Disposition	Integrated in
R3-01	critical	View count used as an exponential lower bound; "someone must see all of G"	**Accepted.** T19 rewritten. We keep the black-box oracle bound ($N - 1$ evaluations) as the honest form of the optionality cost	T19.1–T19.3; paper §0, §2.5, §4.7

#	Sev.	Finding (short)	Disposition	Integrated in
R3-02	major	V6 guard assumed knowledge of the device's execution	**Accepted, reproduced.** Outcome property, not execution; V6x added	T11′; MC-3; M24
R3-03	major	No threshold-EUF-CMA game, corruption model or key-generation premises	**Accepted**	T11′; MON-7
R3-04	major	Approval not tied to the exact signed bytes	**Accepted**	T11′ message m; DEC-1; MON-6
R3-05	major	Authentic is not admissible (closure, status, time chosen by the device)	**Accepted, reproduced** (MC-6)	T11′ step 2; MON-6
R3-06	major	"Provisional" decisions could cause effects	**Accepted, reproduced** (MC-6)	DEC-7; PRF-6; M24
R3-07	major	One view ≠ equality with complete evidence; a certificate is not a view	**Accepted**	T19 setting; T19.3
R3-08	major	Exact, grant and deny soundness conflated	**Accepted**	T19.1
R3-09	major	Completeness premise without a coverage contract	**Accepted**	T19.5; DEC-9
R3-10	major	"Linear assurance" overclaim	**Accepted**	T19.3; paper §0, §4.7
R3-11	major	One commitment does not make every party use the same one	**Accepted**	T19.4; DEC-15
R3-12	major	Public deterministic context digest leaks private input	**Accepted, reproduced**	T19; PRV-6; M31
R3-13	major	No scope check at acceptance	**Accepted**	DEC-11
R3-14	major	No source validation before certification	**Accepted**	DEC-10
R3-15	major	Re-certification semantics undefined; policy-only changes missed	**Accepted**	DEC-12
R3-16	major	Archive promise is not a verified deposit	**Accepted**	DEC-13
R3-17	major	DEC-8 mixed owners and undefined terms	**Accepted**: DEC-8 replaced; DEC-9 to DEC-15 split by role; KEY-6 extended	S2
R3-18	major	"Minimise at standardisation time" overclaims	**Accepted with refinement** (below)	§4.7; DEC-8; VER-8; PRV-6
R3-19	major	Prior art: execute-endorse, witness cosigning, commit-and-prove	**Accepted**; novelty narrowed to integration	S4; §7
R3-20	major	"ZK not practical" unsupported	**Accepted**: "not implemented or benchmarked here"	§4.7; T19
R3-21	minor	2^(n−m) assumes independent statements	**Accepted**	T19.1(e)
R3-22	minor	Timestamp field inside the object it timestamps	**Accepted**: signed core, external time evidence	M31; DEC-10
R3-23	minor	Dispute outcomes over-attributed	**Accepted**	T19.4; DEC-14
R3-24	minor	F5 explanation: `_:a/_:z`; `[0,1,0,1,0]`	**Accepted, checked in source**	F5 report

#	Sev.	Finding (short)	Disposition	Integrated in
R3-25	minor	Budget semantics; README default inconsistent	**Accepted, checked in source**	F5 report
R3-26	minor	Impact and channel overclaimed	**Accepted**	F5 report
R3-27	minor	Planned tests described in the present tense	**Accepted**	§4.5, §5, §7
R3-28	minor	Art. 8 comparison omits modification and revocation governance	**Accepted**	§6.1; S4; M29

Refinements we made.

- **R3-18 and the author's position.** The reviewer recommends allowing selective disclosure whenever sufficiency is established. We agree, and keep the author's design point in a precise form: *the holder has no choice over policy-relevant evidence.* The profile fixes the sufficient evidence for each purpose, and fields outside it may be withheld. "Full disclosure of standardised key claims" therefore means full disclosure of the *sufficient* set, not of every attribute.
- **R3-01.** We keep a cost statement about optionality, but only in its valid form: for black-box policies, certifying exact preservation over N admissible views needs $N - 1$ evaluations. Structure can reduce this, and incomplete observation cannot be repaired by computation (T19.2).
- **R3-17.** New rows are marked ◐ or ◯ with a named basis. None is marked ★ new, because the review's prior art covers their elements.

Primary-source ledger, round 3 (the reviewer's Official Journal check; texts not retrieved by us):

Item	Result	Applied wording
EBW mandate	Confirmed: committee vote and decision to negotiate 10 Sep 2026; report A10-0240/2026 tabled 23 Sep; announced 5 Oct; confirmed under Rule 72 on 7 Oct 2026. Council general approach 9 Jun 2026 (Council press release)	A negotiating mandate, not adoption
CIR (EU) 2025/2531	Confirmed: 16 Dec 2025; OJ 17 Dec 2025; **entered into force** 6 Jan 2026 (Art. 2), no separate application date	"in force since 6 January 2026"
CIR (EU) 2025/1569 Art. 8	Confirmed and expanded: status and version (8(3)(c)); authenticated inclusion and modification requests; governance including revocation mechanisms (8(3)(g)); sealed, readable and searchable catalogue (8(6)); technical specifications (8(7)); identifier (8(8)); recital 10 on preserving versions. Arts. 6–9 apply from 19 Aug 2026	Gap narrowed: Art. 8 does not itself specify catalogue-level withdrawal or deprecation transitions, or a maximum catalogue age for relying parties; the catalogue specifications and operations were not audited
CIR (EU) 2026/1735	Adopted 15 Jul 2026, OJ 22 Jul, in force 11 Aug 2026. It amends Arts. 3(2), 4 and 9 and the annexes; **Art. 8 unchanged**. The Art. 9 changes apply from 1 Jan 2027	Official Journal verification marker removed

Prior art added (S4):

- Hyperledger Fabric's execute-endorse architecture (arXiv:1801.10228);
- CoSi witness cosigning (arXiv:1503.08768);
- LegoSNARK commit-and-prove (IACR ePrint 2019/142);
- Pinocchio (DOI 10.1109/SP.2013.47);
- zk-creds (DOI 10.1109/SP46215.2023.00108);
- Camenisch–Lysyanskaya credentials with optional anonymity revocation (DOI 10.1007/3-540-44987-6_7);
- provable data possession (DOI 10.1145/1315245.1315318);
- proofs of retrievability (DOI 10.1145/1315245.1315317);
- the WCO Data Model.

S5.6 Fourth independent review (13 November 2026): the harness code

Scope and method. The reviewer received harness v0.4.0 at commit `b113a8f` as an archive (SHA-256 `f03c9ff7...`), with a focused pack. It reproduced every result (60 of 60 tests with swtpm, 65 model cases) and probed each component. It ran a 100,000-input differential corpus against the I-JSON check, and mutation tests on guards and test oracles. It reported 33 findings: 7 critical, 20 major and 6 minor. It also supplied an integrated patch with regression tests: 102 tests, 66 model cases.

Our verification before acceptance.

- We reproduced R4-01 (an admitted document expands with a blank node) and R4-05 (the effect gate executes another action, and executes twice) on the original code.
- We ran the reviewer's new regression tests against the original code. Every count matched its report:

Test file	Pass	Fail
CBOR-LD review	0	4
Replay review	1	4
Compliance lifecycle	0	3
Compliance graph	0	3
I-JSON review	0	3
LFSC pass 3	1	6
Replay effect	1	4

- Before running the patch, we read all of its source changes: no network access, no writes outside temporary directories.
- One change of ours: the reviewer's guard refused to launch the engine whenever `NODE_OPTIONS` was set, which breaks environments that set it harmlessly (ours sets a certificate option). We replaced it with an **allow-listed engine environment** (only `PATH` and the TPM connection), which also excludes `NODE_PATH` and `LD_PRELOAD`. A test includes a control showing that the preload does run under the inherited environment.
- Historical round-4 integration result: 103 of 103 tests and 66 model cases from a clean extract of commit `df61747`. The round-5 review reproduced that baseline and its WE-3, WE-3 R and WE-7 outcomes (20, 21 and 23 scenarios). After the round-5 fixes, the

reviewed candidate `2183b99` passed 121 tests with no failures or skips and retained 66 model cases (S5.7). The released harness v0.4.0 [7] carries the final source and clean-extract verification record identified in §5.

All 33 findings accepted:

#	Sev.	Finding (short)	Fix
R4-01	critical	Excluded type spelled as an IRI passes admission; blank node	excluded terms rejected before the IRI fallback
R4-02	critical	CBOR-LD tables with duplicate codes: same manifest, different meaning	injective tables required
R4-03	critical	CBOR-LD issuer read from the compact field; aliases bypass it	semantic issuer after expansion, exactly one, equal to the signer
R4-04	critical	Co-signer commits a proposal changed after authentication	owned snapshot before any asynchronous check
R4-05	critical	Effect gate ignores action, single use, epoch from the receipt, expiry	receipt v2; gate checks all four
R4-06	critical	WE-7: genuine report placed under another report's identifier	reports bound to signed identity, product and batch
R4-07	critical	Soundness cross-check passes without evaluating documents	counts successful expansions
R4-08	major	Admission certificate with another processor profile accepted	profile compared
R4-09	major	Lone surrogates in string input repaired before checking	ill-formed strings rejected
R4-10	major	Negative cache survives policy and epoch changes	epoch in key, expiry, intrinsic failures only
R4-11	major	Malformed times fail open	strict timestamp profile; `validAt` obligation stated
R4-12	major	Unbounded did:key decoded before proof checks	length bound
R4-13	major	Proof type and purpose not enforced	`DataIntegrityProof`, `assertionMethod` required
R4-14	major	Profile B graph depends on unbound processor profile	graph-processing profile in manifest v2; binary integrity stated as trusted
R4-15	major	`NODE_OPTIONS` preload runs unmeasured in the engine	allow-listed environment (our variant)
R4-16	major	Engine rule ignores nested contexts	nested contexts give *unknown*; claim narrowed
R4-17	major	E2 passes when enrolment failed	E2 requires a certified key and the policy reason
R4-18	major	Missing or non-numeric times pass the co-signer	finite, ordered fields required
R4-19	major	Unbounded, never-expiring approvals	bounded, expiring, re-checked at signing
R4-20	major	Signature guards without independent tests	independent signature corruption tests

#	Sev.	Finding (short)	Fix
R4-21	major	Custodian replay supplied by the certifier; outcome unchecked	caller-provided replay removed; separately configured custodian replays and binds core and outcome; the example shares the evaluator and trust-context factory, so independent administration is not demonstrated
R4-22	major	Dispute ignores the private-context binding	context checked before replay
R4-23	major	Null or non-numeric results pass	finite non-negative values required
R4-24	major	Malformed coverage and request times	coverage and request times validated; round 5 additionally rejects authenticated test reports issued after the evaluation time
R4-25	major	N-Quads injection in the evidence graph	IRIs validated, literals escaped
R4-26	major	B10 never isolates the timestamp check	core re-signed, stale timestamp kept
R4-27	major	MC-6 never reaches the withdrawal guard	new schedule; 66 cases
R4-28	minor	Long valid numbers rejected	sticky full-text match
R4-29	minor	Missing vectors skip silently	vectors mandatory
R4-30	minor	Nondeterminism not fully declared	deterministic forged proof fixtures; round 5 completes the comparison contract for WE-6 timings, WE-5 timing-derived outcomes and environment/run metadata
R4-31	minor	Stale version banner in WE-1	installed versions
R4-32	minor	Release date and mutable citation	`CITATION.cff` records 2026-11-16 and the versioned `ssc-harness-v0.4.0` tag URL
R4-33	minor	W3C licence text not bundled	full W3C-2023 terms bundled; round 5 separates the citizenship repository notice and its referenced W3C-2015 terms, corrects vector-wrapper attribution, and leaves the citizenship context's per-file classification explicitly unverified

Enrolment and attribution (pass 9). The reviewer's proposal and ours are integrated in T11 (ii-a, ii-b), T11′, MON-8 and S1 (M24). It corrected two points in our pack:

- **RATS roles.** An enrolment service may perform the RATS Verifier and Relying Party roles and additionally issue application-level enrolment credentials. An Endorser supplies claims about an Attester (for example, a manufacturer's endorsement of a TPM). Issuing organisational authorisation is not automatically a RATS endorsement.
- **Legal qualification.**
  - A recognised enrolment authority and a qualified trust service provider are different designations.
  - Trusted-list status (eIDAS Art. 22) concerns the listed provider and service, not decision-class authority.
  - For a qualified electronic seal, Art. 35(2) provides presumptions of integrity and correctness of origin of the linked data, not substantive correctness or a mandate.
  - Sealing an enrolment credential does not make each receipt a qualified seal.

– Qualified electronic attestations of attributes (Art. 45d, Annex V) and attestations issued by or on behalf of a public-sector body responsible for an authentic source (Art. 45f, Annex VII) have distinct requirements.

The reviewer's three questions, answered:

1. **CBOR-LD acceptance:** credentials only. Generic signed documents would need a separate profile.
2. **Admission soundness:** claimed only for grammar `ssc-admission/v1` under the bound processor profile (jsonld 9.0.0, JSON-LD 1.1, `rdfDirection` null). The evidence is the finite tests and regressions; an independent checker is future work, and the certifier stays trusted (T9.2).
3. **Implemented versus deferred** (stated in the README):
   - Implemented or partial: only the concrete clauses mapped to source and execution evidence in the conformance table. Scenario success is not proof of complete compliance with the conditional and optional requirements.
   - Deferred: MON-5 upgrades, MON-7 (with FROST), durable DEC-7 effects, a complete DEC-13 archive service, certified KEY-6 entropy, and MON-8 enrolment.

S5.7 Fifth independent review (14 November 2026): the release candidate and the paper

Scope and method. The reviewer received harness v0.4.0 at commit `df61747` (archive SHA-256 `752d3a8b...`), the round-5 pack and the complete claims note v0.9.5. It worked through eight passes:

1. reproduction;
2. survival of the round-4 fixes;
3. our deviations from its round-4 patch;
4. new defects in the changed code;
5. enrolment and attribution;
6. standards and legal statements against primary sources;
7. the paper against the code, clause by clause;
8. our own weak spots A–K.

It reproduced the baseline (103 of 103 tests with swtpm, 66 model cases; WE-3, WE-3 R and WE-7 at 20, 21 and 23 scenarios). It rebuilt the reviewed `b113a8f` sources from our diff and confirmed every round-4 fix present (table G of its report). It read the eIDAS articles in the official texts, the TCG specification, the RFCs and FDO 1.1; for IEEE 802.1AR only the public metadata and scope were available.

It reported 23 findings: 2 critical, 11 major, 10 minor. It also supplied:

- an integrated patch with regression tests: 121 tests, 66 model cases;
- 114 exact paper replacements;
- a clause-by-clause conformance map of M15, M17, M24 and M31.

Our verification before acceptance.

- We read every change in the patch before running it: no network access, and the tests write only to temporary directories.
- We ran the new regression tests against `df61747`. The counts matched the report:

Test file	Pass	Fail
LFSC round 5	0	3
Replay round 5	1	8
Compliance round 5	0	5

  The one passing replay test documents behaviour that does not change: historical receipt verification succeeds after the epoch advances, and the effect gate refuses.
- The launcher mutant (the engine started with the inherited environment) survived the shipped test and was killed by the new one.
- The round-5 patch was adopted **unchanged** in the reviewed candidate `2183b99`; the release includes these corrections, with its final commit and any subsequent release-gate changes recorded in [7].
  - Results were regenerated from that source and the README timing tables updated to that run.
  - `SHA256SUMS` lists 94 files.
  - Release check from a clean extract: checksums, `npm ci`, 121 of 121 tests with swtpm, 66 model cases; WE-3 20/20, WE-3 R 21/21, WE-7 23/23.
- The paper replacements were applied:
  - 111 were applied before the release;
  - the 3 held changes (P031, P043, P113), together with the remaining release-status updates, were applied after the tag, GitHub release and released-archive verification;
  - before release, two candidate sentences named `2183b99` instead of "the patch tree", and the not-yet-existing v0.4.0 release was not cited; the publication now cites the verified release [7].

All 23 findings accepted:

#	Sev.	Finding (short)	Fix
R5-01	critical	The approval returned to the caller was the co-signer's own record; editing it defeated the expiry, epoch and closure checks at signing	owned copy kept; four regressions
R5-02	critical	The engine-launch test re-stated the fork options and passed without the engine starting	test runs `launchEngine` and checks readiness and liveness; mutation test kills both mutants
R5-03	major	Receipt v2 parser: no exact array, type or canonical-byte check; signed key epoch ignored; malformed containers threw	one schema for encoding and parsing; exact length; canonical bytes; relying party's key epoch; containers rejected
R5-04	major	Malformed validity bounds and a missing or malformed `created` passed when no `validAt` was given	schema checked unconditionally; `validityChecked` reported

#	Sev.	Finding (short)	Fix
R5-05	major	WE-7: a report dated after the decision authorised it	future dates give *unknown*
R5-06	major	WE-7: the evidence graph lost statement keys and multiplicity, and evaluation depended on Map order	extraction/v2 commits role, key and statement; exactly one matching absence attestation; holder-only reports sorted
R5-07	major	T11's attribution conclusion lacked actual-use scope, relevant-time and compromise premises and a proof step	separate governance-attribution corollary with its own premises and argument
R5-08	major	MON-8 mixed verifier, issuer and governance duties, made optional attribution mandatory, and cited E3 as evidence	one verifier-owned acceptance predicate; optional module; owners in M24; evidence cell corrected; still 91 requirements
R5-09	major	The TPM → 802.1AR → BRSKI/FDO → enrolment chain is not one established sequence	profile-dependent assurance paths; mandate validated separately
R5-10	major	"Sealed record" extended the Art. 35(2) presumptions beyond qualified seals	qualified-seal condition stated; Arts. 45b–45f distinguished
R5-11	major	M15 and M17 descriptions differed from the code (envelope, archive key, digest names, signing hash, cache key, LFSC inner proof)	exact descriptions; v0.3 formats kept as history
R5-12	major	M24: implemented and proposed parts blurred; keys called certified; no replay audit log or operation-metered replay	implemented subset disclosed; encoding stated byte-exactly
R5-13	major	M31: custodian independence, full reconstruction and "checks all" acknowledgement fields exceeded the code; formula and cost wording	subset, formula and cost stated exactly
R5-14	minor	Receipt verification accepted a historically valid receipt after expiry or epoch change; non-effect reliance needed a boundary	historical verification labelled; current checks required at every use
R5-15	minor	The environment allow-list is not authentication of the runtime, tools or TPM	trusted configuration stated in code, README and M24
R5-16	minor	"Each fix has a failing regression" was false beyond R4-30 to R4-33	claim narrowed in README and paper
R5-17	minor	Historical timings mixed with current evidence; tense and counts	v0.3.0 measurements labelled; at-tag replacements prepared
R5-18	minor	"Every holding case is non-vacuous" omitted `MC6.offline.deferEffect`	exception stated at every occurrence
R5-19	minor	Declared nondeterminism omitted WE-6 timings, metadata and WE-5 timing-derived counts	comparison contract completed in the README
R5-20	minor	The citizenship context's licence category is not settled upstream	upstream notice and W3C-2015 terms bundled; uncertainty stated in NOTICE
R5-21	minor	Vector files called unmodified copies although the harness adds wrappers	attribution corrected
R5-22	minor	TCG edition and date, RFC statuses, and "credential issuer" as a RATS role	citations pinned; role wording corrected

R5-23	minor	"The operator answers" implied automatic liability; "exclusive use gives correctness" overstated	accountability as evidence under governance and law; execution provenance, not correctness

The reviewer's questions, answered:

1. **Is attribution mandatory?** In v0.9.6 it is an optional module (MON-8): technical key authentication stays mandatory under T11 or T11′, and organisational attribution is a separately claimed property. Whether a high-assurance profile should require it was then open; the author decided on 15 November 2026 that it does (§8, MON-8).
2. **Sources for mandate, scope, time, compromise and revocation.** Not fixed by the manuscript. A deployment profile must name them; MON-8 requires that it does.
3. **WE-3 R schema and threshold configuration.** The restricted text-scope schema stays the v2 receipt profile of v0.4.0. A FROST release (v0.4.1) has to define the authenticated group configuration and its share epochs (MON-7).
4. **Custodian guarantees.** DEC-13 and DEC-14 state the target. WE-7 demonstrates a subset, as S1 (M31) now says; a production custodian is not specified here.
5. **Citizenship context licence.** Open. Only its maintainers can settle it; NOTICE states the uncertainty.

What remains outside the release: host-root isolation, hardware trust, FROST, MON-8 enrolment, durable effects, a complete archive service, and certified entropy. The README lists them as deferred.

S5.8 Sixth and seventh reviews (14 and 15 November 2026): the release gate

Sixth review (14 November 2026). The reviewer of rounds two to five examined candidate `2183b99` and v0.9.6 and reported four findings, all accepted.

ID	Severity	Finding	Disposition
R6-01	blocking	WE-7 accepted statement Map keys that were not strings. Distinct object keys serialised to the same committed entry, so one evidence commitment could open to different outcomes. The round-5 repair of R5-06 was incomplete	keys must be nonempty, well-formed strings; regression test; candidate `6e565ce`
R6-02	minor	"Evaluated once" omitted the mandatory deposit replay	18 text corrections in §0, §2.5, §4.7, §7, S1.5 and T19
R6-03	editorial	A cross-reference to a section that does not exist	corrected
R6-04	editorial	Mixed spelling	38 line-scoped corrections

Our own check of that repair found that a dispute over an archive holding such a key threw an exception instead of reporting a finding (R6-N1). Candidate `fc2088c` reports an invalid opening and refuses the deposit; `049af0f` corrects one README paragraph.

Seventh review (15 November 2026). Two reviewers examined candidate `049af0f` and v0.9.11 independently and in parallel: the reviewer of rounds two to six, and a different reviewer

(Claude Fable 5.2, Anthropic), which worked from the written record of the earlier rounds. Both reproduced the candidate from a clean extract (123 tests, every worked example, 66 model cases) with no difference from the archived results outside the README's comparison contract. The first reported fifteen findings (two major, ten minor, three editorial) and none blocking; the second reported seventeen (one blocking, two major, twelve minor, two editorial). Several findings overlapped. We reproduced the blocking finding, the structured-refusal finding and both test gaps by execution, and checked the disputed dates and specification statements at the primary sources. All findings were accepted. Where the two reviewers proposed different text for one passage, we chose on the primary source: EUDI TS11 for §1.5, Bitstring Status List §2.1 for the status-message mapping, and VCDM 2.0 §4.3 for the context obligation.

Finding	Reviewer	Severity	Disposition
A statement Map entry holding no statement object (`null`, `0`, `false`, `""`) satisfied the coverage check but was skipped by the commitment: two evidence sets with the same canonical graph evaluated to UNKNOWN and DENY, and the DENY was certified and opened validly	second	blocking	every entry must hold a statement object; three regression tests, each failing on `049af0f`; candidate `16e85e6`
The T11 remark required exclusive key use by an isolated engine for replay-validated receipts	both	major	separated: T11's premises for execution-bound receipts, T11′'s for replay-validated receipts, MON-8 attribution for both
§1.5 denied governed identifiers and machine-readable semantics in SD-JWT VC and mdoc, and stated a negative search result too broadly	both	major, minor	rewritten against EUDI TS11 §§2.1 and 4.3; the negative claim is withdrawn and the gap narrowed to receipts that bind the definitions used and to governed cross-type checks
The review record stopped at the fifth review	second	major	this subsection; counts corrected in the front matter, §2.5 and S5
Disputes, deposits and certification still threw on an archive without a salt or on evidence the extraction cannot represent	second	minor	structured refusals; two regression tests, each failing on `049af0f`
The refused deposit of R6-N1 was not exercised by any test	both	minor	tests for new and existing certificate identifiers; store unchanged
Current passages still treated attribution as optional	both	minor	corrected in S2.11 and S1.6; historical passages annotated
The EU building blocks were phrased as a trajectory towards enrolment authorities	both	minor	stated as legal duties and potential building blocks; the European Business Wallet as a proposal
Abstract and §0: computation cost without its condition, "how to close it", "map which mechanisms transfer", the profile read as a W3C work item, enrolment not marked as unimplemented	second	minor	qualified
All of T10 was called RDF-specific	first	minor	only profile I-RDFC is; profile B can inform a format-specific binding

Finding	Reviewer	Severity	Disposition
§1.3 details: the VCDM §4.3 qualifier, RDFC-1.0's mandatory defence, the status-message location, VC-JOSE-COSE's dates, the JSON-LD API, the scope of the library finding	both	minor	corrected against the Recommendations and the cited paper
UC5 treated output provenance as direct evidence for AI Act Arts. 10 and 12; GraphRAG uncited	both	minor, editorial	provenance supports traceability; reference added
References and dates: ISO/IEC TS 18013-7:2025 and the 23220 parts, ARF v3.0.0 dated 21 July 2026, CIR (EU) 2025/1569 Arts. 7 and 8, two titles, the date of facts, two citation keys	both	minor, editorial	corrected; the last open verification marker in the references resolved
Tagging templates: stale line pointers; H04 and T05 wording	both	minor, editorial	refreshed and reworded
Four copies of one test fixture	second	editorial	one shared fixture

Candidate `16e85e6` passed 130 tests from a clean extract, every worked example and the model check, with no difference from its archived results outside the comparison contract.

Cross-check of the integration (15 November 2026). Both reviewers of the seventh round then checked `16e85e6` and v0.9.12 as integrated, each concentrating on the other's patches. Both reproduced the candidate from a clean extract and confirmed that it is `049af0f` with the three reviewed patches. The first reported no blocking finding and three minor corrections to the manuscript: the JSON-data premise of the extraction domain, the review history in §5, and the scope of a library finding in S4.2. The second found one further blocking defect, present since extraction/v2 was introduced in `2183b99` and not reported in the sixth or seventh round. The extraction turned the fields of a statement body into predicates in the same namespace as the committed `sourceEntry` tuple and took a statement's subject from its `id` field. Unsigned statements that named another entry's subject and carried the fields `sourceEntry`, `signer`, `signature` and `statement` therefore reproduced every triple of that entry, and evidence with and without the entry had one canonical graph. We reproduced it by execution on `16e85e6`: one commitment opened validly to GRANT and to DENY for two declarations signed by the same importer, and to GRANT and UNKNOWN with and without the absence attestation. The field name `sourceEntry` is now reserved in every statement body, the declaration and the coverage responses, with three regression tests that fail on `16e85e6`; canonical graphs, outcomes, the closure and the policy digests of valid inputs are unchanged. Evaluation reads only the declaration, the statement Map and the two coverage responses, each now committed in a tuple that no other entry can emit; for JSON data the canonical graph therefore determines the evaluation input (by reading of the code, and by a differential search over 128 evidence sets). An optional hardening from the same reviewer, also adopted, returns structured refusals in place of the remaining exceptions. Its other corrections concerned the review history, the scope of two statements and naming. All findings of the cross-check were accepted. Candidate `0dd6bc9` passes 136 tests from a clean extract, every worked example and the model check, with no difference from its archived results outside the comparison contract; its source archive is built with the commit time as file time, so its checksum can be reproduced from the commit.

S5.9 Harness release record

Harness v0.4.0 was released on 16 November 2026 from commit `2ee3f966aa99ce96e26c6764 3ff6ed52a6130bb5` (tag `ssc-harness-v0.4.0`) [7]. Its source and tests are those of the reviewed candidate `0dd6bc9` (S5.8).

Item	Record
Release	harness v0.4.0, 16 November 2026
Commit	`2ee3f966aa99ce96e26c67643ff6ed52a6130bb5`
Tag	`ssc-harness-v0.4.0`
Source and tests	those of reviewed candidate `0dd6bc9`
Verification from a clean extract	136 tests; 66 model cases; WE-3 20/20, WE-3 R 21/21, WE-7 23/23
Release asset	`harness-v0.4.0-source.zip`, SHA-256 `0efbbca122a052246b4b59afb6102f10 5be47bdda3421815f67b512c7bee3f89`
Checksum manifest	`SHA256SUMS`, time-stamped by FreeTSA (RFC 3161) at 16 November 2026 12:25:22 GMT
Archive	Zenodo, DOI 10.5281/zenodo.24281131 (concept DOI 10.5281/zenodo.24168346)

Harness v0.3.0 was released on 12 November 2026; its manifest received an RFC 3161 timestamp at 12:12:09 GMT and web-archive captures.

S5.10 A discrepancy in the v0.3.0 release record

While the article was being prepared, its WE-4 figures were checked against the results archived in harness v0.3.0 (`results/we4.json`, run of 12 November 2026, 10:52 UTC). The v0.3.0 README table reports 1.2 s for the complete graph K_8 (1.7 kB) and 2.4 s for the 255-node symmetric tree (11 kB) under a verifier without a work bound. The archived run of the same release records 1,558 ms and 2,956 ms for these inputs. The other values agree: more than 20 s for the 511-node tree (21,668 bytes), and 0.02–0.11 ms for the lexical-first path on these families. The README evidently reports an earlier run that the archive does not contain. Harness v0.3.0 is published and archived (DOI 10.5281/zenodo.24168347), so its README is left as published; the article and Figure 4 use the archived values. The same kind of discrepancy, in the WE-5 paragraph of the v0.4.0 README, was corrected before that release.

S5.11 Pre-publication review (16 November 2026): the article, the supplements and the Executive Brief

Before publication, the article v1.0, the supplements S1–S5 and the Executive Brief were reviewed by two AI models working independently, OpenAI GPT-6 Astra and Claude Fable 5.2 (Anthropic), from the same package and prompt. The prompt asked for the brief's fidelity to the article, the accuracy of the new use case UC6 against the WE BUILD vocabulary, the fidelity of the restructure against the claims note v0.9.16, internal consistency, every legal and standards date at its primary source, readability and the typeset PDFs. Astra reported 19 findings (2 critical, 8 major, 6 minor, 1 editorial, 2 queries) and Fable 42 (6 critical, 11 major, 23 minor, 2 editorial). Both found the results, numbers and evidence levels intact and asked for wording corrections only.

Before accepting them, we checked the critical and factual findings at their sources: the two W3C drafts (no NFC, QR codes only in VC Barcodes, no carrier lifetime), the Commission's DPP Registry artefacts (five JSON-LD contexts, version 1.0.0 metadata and concept status, no expiry or artefact-level withdrawal), the Community Group's launch on 12 November 2026, the current revisions of two Internet-Drafts, the entry into force of Regulation (EU) 2026/2108 on 20 September 2026, the WTO notification of the draft battery-passport implementing regulation, the archived harness results (the 0.93 s list timing; the 1.1 s and 528 ms values of experiments X5 and X11) and the MC-4 enumeration in the model checker (12 reported cases, each over 24 combinations). The dispositions are recorded with the publication files; one proposal, renaming figure paths inside the review package, was not adopted because it concerned the package layout, not the article.

The accepted corrections: the abstract, the brief and the theorem summaries restore conditions that had been dropped (issuer binding; laboratory, JCS or envelope and unauthenticated-context conditions; premises of T1, T2.1, T6, T7, T8, T13, T16–T18, T19.5 and the governance corollary); statements attributed to third-party documents are corrected (carriers and lifetimes, DPP Registry metadata, the Community Group's status, CIR (EU) 2025/1569 Art. 8); UC6 separates semantic definitions from case facts and policy and states which mechanisms do not cover powers of attorney and derived master data; the closure-digest definitions in S1 follow S3.1.3, and M15's profile conditions follow T10; disclosure text lost in the restructure is restored (M2, M10, M17); harness versions, review-experiment figures and the scope of the enrolment claim are stated precisely; and readability, references and typesetting are improved.