Lawie Labs Technical Fact Sheet

Verifiable legal execution: what exists, what was measured, what does not exist yet

Technical fact sheet LL-FS-01
Version 1.0.0, 26 November 2026

Intelligence is not authority. Mission: Make high-stakes institutional work executable by supervised machines: completed, approved, evidenced, defensible.

In one paragraph. Lawie is building software in which supervised machines carry out legal work under recorded authority: decisive statements anchored to admitted official sources, consequential actions bound to approvals by named people on exact content, and evidence that a party who distrusts the operator can check. Its source repository implements verification, authorization and evidence components, and Lawie Labs has measured those verification components against pre-registered hypotheses (LLRB-1), negative results included. Lawie is not a lawyer and renders no legal judgment. Nothing described here is deployed in production, no product has a dated production acceptance, and nothing is independently audited.

The doctrine in five lines

1. **Intelligence is not authority.** A model's competence never confers the right to act or the warrant to be believed.
2. **Execution, not advice.** The product is a completed act with its authority and evidence, not a suggestion.
3. **No anchored source, no decisive claim.** A citation string is not an anchor; an admitted, applicable, supporting span is.
4. **We support professionals; we render no legal judgment.** Outputs are for the sign-off of a named professional.
5. **Refusal is a result.** Where authority is missing the system refuses and says why, and failures are published.

Architecture at a glance

Layer	What the source does at the basis commit
Sources	Official status only under an approved source contract with reviewed licence; observations record locator, retrieval time, hashes, licence and supersession. Production profile: twelve official sources, none admitted
Corpus and legal time	Provision text authoritative only as an exact UTF-8 byte span proven by hash; one span derivation (Czech statutes) admitted. Versions over half-open intervals from entry into force; dates of application not modelled
Retrieval	Dense and sparse fusion; jurisdiction, date and access filters; citations without a retrieved source dropped
Verification	Exact-span citation authority with typed refusals; unbacked citations withheld; citation gate; lexical grounding check; no calibrated probability of correctness
Agents and approvals	81 agent types; most restrictive capability mode; default-deny tenant grants; approvals bound to the exact call and checked for validity
Tenancy	Signed identity assertions bound to each request, single-use for state-changing requests; requested tenant compared with the authenticated tenant
Evidence	Hash-chained agent audit records; signed portable dossiers verified offline (Ed25519, ECDSA P-256, RSA-PSS) that state their own limits
Civic Guard	Deterministic screens that trigger, do not trigger or abstain; distinct investigator, reviewer and publisher; right of reply of at least seven days; quarantined

Status: First edition. Nothing described is deployed in production; nothing is independently audited. *Distribution:* Public. *Basis:* source repository `main` @ `ffa13d4bce` (2026-11-24); experiments of 26 November 2026; EUR-Lex texts retrieved 26 November 2026.

Headline measured results

Measurement	Result
LLRB-1 design	3,028 provisions of eight EU acts; 5,820 support items, 4,053 in the test split; labels fixed by construction; pre-registered; one confirmatory run
Lawie's grounding check	Accepted 1,959 of 1,959 decisively edited statements and 698 of 698 misattributed ones
Calibrated lexical verifiers	Accepted 27.1%, 54.2% and 51.8% of decisive edits (containment, token F1, Jaccard)
Typed research prototype	0 of 1,444 modelled edits accepted, 0 of 1,396 supported statements rejected; 456 of 515 (88.5%) held-out edits accepted; not a Lawie component
Reported confidence	0.973 (HIGH) for every statement; AUROC 0.5
Quotations and citations	Quote matcher passed 1,956 of 1,959 edited quotations given the text, 1.0% of authentic ones as deployed; 0 of 114 official-style citations resolved; 0 of 8 fabricated citations blocked
Legal time	In force read as applicable on 781 of 781 dates; schedules contradicted on 7–20% of provision-days, 2016–2028
AI Act amendment	82 of 83 article-level amendments found; precision 21.6% on rendered text, 86.8% canonicalised; postponement read as a new obligation
Authorization (BEN-02)	552 scenarios: 0 of 96 legitimate requests denied; every scenario of some attack classes refused and every scenario of others allowed (class-level outcomes restricted pending remediation)
Evidence integrity	Every tested in-place edit, deletion, reordering, insertion and deleted prefix detected (50 trials each); every tampered dossier failed

Components ran unchanged and offline at the basis commit, each run recording the hashes of the component files it loaded; no result is a production or customer measurement. This sheet reports research measurements at a stated code basis; it is not a security assurance, a compliance statement or legal advice, and no independent security, privacy or legal review of it has taken place.

Law and arithmetic the papers rest on

- As published in the Official Journal, Regulation (EU) 2026/1744 amended Article 113 of the AI Act; the amended text lists 2 December 2027 (Annex III) and 2 August 2028 (Annex I) for Chapter III, Sections 1–3, in place of 2 August 2026 (research labels from the official texts; not legal advice).
- A gate whose verifier has false-accept rate F and false-reject rate R releases outputs with error rate $eF/(eF + (1-e)(1-R))$ when a share e of its inputs is wrong; a verifier that accepts every decisive error cannot lower the error rate of what it releases.
- A screen that triggers on a share t of cases raises alerts of which at most p/t are true anomalies, for prevalence p; 40% of Czech contract awards had a single bidder in 2024, against an EU average of 28% (European Commission, Single Market Scoreboard, reporting period 2024).

State on 26 November 2026

Item	State
Production	Not deployed; awaits owner-controlled inputs and approvals; no product with a dated production acceptance
Official sources	None admitted for production
Civic Guard	Quarantined; no release stage admitted; five evaluation studies defined and blocked on external owners
Assurance	No independent audit, certification or expert review; benchmark labels not reviewed by a lawyer
Model evaluation	Evaluation laboratory has never run a model; generated-answer study defined, not run

Publications

LL-WP-01 Verifiable Legal Execution · LL-DOC-01 Make Law Executable: The Lawie Doctrine · LL-REL-01 The Reliability of Legal Artificial Intelligence · LL-EVP-01 Evidence, Provenance and Legal Reasoning · LL-GMA-01 Governed Multi-Agent Systems · LL-CG-01 Civic Guard: Evidence-Based Civic Intelligence · LL-BEN-01 Legal AI Grounding and Hallucination Benchmark · LL-BEN-02 Agentic Execution and Authorization Benchmark · LL-TMP-01 Temporal Reasoning in Legal Systems · LL-ADV-01 Adversarial Robustness in Legal AI · LL-RI-01 Machine-Assisted Regulatory Intelligence · LL-SEC-01 Security, Authority and Accountability · LL-DD-01 Technical and Investor Diligence (restricted) · LL-FS-01 this fact sheet.