

The evidence layer for regulated AI.

Model-agnostic. EU-sovereign. Every finding **cryptographically sealed** and independently verifiable — by an auditor, a counterparty, or a court, without trusting us.

SEALED EVIDENCE UNIT · FINDING_HASH

• ANCHORED

7e96b932b2d2132c2f2a17c7094ff5803cdea242388a8906f80093fd8731bf88

RFC 3161 TSA ✓ OPENTIMESTAMPS · BITCOIN ✓ MERKLE ROOT ✓

One harness, every frontier model.

Never tied to a single vendor. Route each playbook to the model that passes quality at the lowest cost — and escalate only where it matters.

- ◇ **AWS Bedrock primary**, Anthropic direct fallback, plus OpenAI, Gemini, Deepseek, Llama, Mistral, Amazon Nova, Cohere — **15+ models**, selectable per playbook or org.
- ◇ Production inference layer: **prompt caching, cross-region profiles, concurrency limiting, async + sync fallback.**
- ◇ High-stakes extractions run **≥2 models with cross-validation** — disagreement surfaces as "needs human review", never silently resolved.

MODEL CATALOGUE · LIVE

bedrock·opus-4.8**bedrock·sonnet**

bedrock·haiku

llama-4-maverick

mistral-large

nova-pro

command-r+

gpt-4o

gemini-1.5-pro

+ adversarial ✓

Two independent models cross-validate a DORA field. They agree on the notice period → auto-accepted. They differ on data-location → routed to a human. No silent averaging.

Your data stays in the EU. Your models stay your choice.

- ◇ Deploy on your **choice of sovereign EU hosting or AWS** — both supported, **data stays in the EU**.
- ◇ Fully containerized and portable: object storage via an **S3-compatible abstraction** (endpoint = env var), a **Postgres-based queue** — no SQS/Pub-Sub lock-in — secrets via env only.
- ◇ Sovereignty of inference too: run entirely on EU-region Bedrock, or bring your own endpoint.
- ◇ **Row-Level Security on every table**, per-tenant isolation, Supabase Auth + MFA.

RESIDENCY

EU-only data plane

Documents, embeddings, and evidence never leave the region.

PORTABILITY

No lock-in

Postgres queue · S3-compatible storage · env-var secrets.

ISOLATION

RLS everywhere

Every tenant table ships a policy + migration + test.

EXIT PATH

Containerized

Runs on your infrastructure — nothing proprietary underneath.

Encode an expert review once. Run it on every deal, forever.

Build one in a **drag-and-drop editor** — or let the AI builder draft it from a document. No markdown required.

- ◇ **Visual editor + AI builder.** Drag questions into groups by hand, generate a first draft from a source doc with the 5-stage agentic builder, or drop in a DOCX template and let it auto-detect.
- ◇ **Conditional decision trees.** Gate whole sections on a prior answer, a risk score, or the document type — the review adapts as it runs.
- ◇ **Forced output formats.** Every answer is typed — booleans, enums, **numbers** · **currency** · **%**, dates & specific **renewal dates**, **date ranges**, risk scores, LEIs — structured and machine-checkable, never free-text mush.

```
# SKILL.md – DORA ICT third-party review ## Group: Contractual
arrangements - Q: Is a notice period stated? [boolean] -
Q: Provider parent entity? [party_name] - Q: Data-processing
locations? [list] ## Group: Exit & substitutability
  IF_ANSWER critical-function = true: - Q: Documented exit
plan? [boolean] - Q: Substitutability rating? [risk_score]
# regulatory_config: DORA · pinned to source version
```

[versioned](#)[edit history](#)[re-run guardrails](#)[en · fr · de](#)

Your evidence, as one graph.

Compliance questions are about *relationships* — which vendor you depend on everywhere, which controls a finding also satisfies, what breaks when a law changes. The ontology answers them **without a graph database** — and now you explore it as one interactive graph.


- ◇ **Collect once, comply many.** Tag an Evidence Unit to one control; it surfaces against every mapped control across frameworks — DORA → NIS2 → ISO 27001, "via 2 mappings" and all.
- ◇ **Counterparty 360° + entity resolution.** One registry of every party; LEI & alias matches auto-link, the rest go to a review queue with **dual-model AI adjudication** — nothing merges silently.
- ◇ **Explore it live.** The whole ontology renders as an interactive graph you pan, zoom and trace — one click spotlights provider concentration or cross-framework reuse.

EXPLORE IT LIVE • PAN • ZOOM • TRACE

DORA

ISO 27001

NIS2



- **Provider concentration** — who you depend on, transitively
- **Evidence reuse** — cross-framework pre-fill
- **Requirement impact** — a law changes → the blast radius


The DORA register — generated, cross-checked, filed.

Every EU financial entity must file a Register of Information on its ICT third-party arrangements. ComplianceOS builds it from your contracts — extracted, cross-validated, and checked against the ESA taxonomy.

- ◇ **Dual-model extraction.** Two independent models fill each register field — agreement auto-accepts, disagreement routes to a human. Never a silent guess in a regulatory filing.
- ◇ **Validation as data.** Required fields, date ranges, cross-field checks, and real **LEI validation** (ISO 17442 · ISO 7064 mod-97-10 checksum) — rules update without a redeploy.
- ◇ **Built to the ESA taxonomy.** Loads the official ESA templates (Reg (EU) 2024/2956) and exports the authority-ready packages — the per-template CSV set and the EBA **XBRL-CSV** report package — a sealed evidence trail behind every field.

REGISTER OF INFORMATION · EXTRACT		• 2 models agree
entity.name	Aetheris Bank SA	✓
provider.lei	9845 00XY...21	LEI ✓
arrangement.start_date	2024-03-01	✓
function.criticality	critical	✓
service.data_locations	IE / DE ?	review

Every finding is sealed, chained, and provable — without trusting us.



- ◇ Each finding gets a **SHA-256 hash**, chained into a **per-tenant tamper-evident chain** — append-only & immutable, enforced at the database.
- ◇ Chains are **Merkle-anchored and dual-timestamped**: RFC 3161 authority + OpenTimestamps on Bitcoin — verifiable with standard tooling.
- ◇ **GDPR-compatible by construction** — only hashes leave the system; deleting an analysis orphans a meaningless hash on-chain.
- ◇ Verdicts are **honestly tiered** — "independently anchored" vs "internally consistent" — RFC 6962 domain separation, no proof-forgery shortcuts.

The law is a live input, not a static PDF.

- ◇ Sync ingests official sources — **DORA RTS, ESA RoI, ReCyF, EUR-Lex, Légifrance, RIS, FMA** — with **content-hash version pinning**. Every EVU records *which version of which law* it was checked against.
- ◇ Playbooks inject **cached legal text as model context** at analysis time — built against the regulation, not model memory.
- ◇ Validation rules load **as data** — regulatory updates ship without a redeploy.

CONTINUOUS SURVEILLANCE · TRACK

Fingerprint-based monitoring auto-reruns or flags an analysis when any of three things change:

• documents

• playbook

• legislation

law v2024.1 → v2024.2

→ 4 analyses flagged non-compliant

→ linked to the exact amended clause

Meets your data where it lives.

Native connectors for your document stores, an open tool standard for everything else, and clean interoperability out.

DATA SOURCES · OAUTH

Where documents live

Google Drive · Notion · SharePoint / OneDrive · NetExplorer — with data-room sync & content-hash delta detection.

MODEL CONTEXT PROTOCOL

Open tool surface

Connect external MCP servers to give analyses live tools & data — OAuth / API-key managed, with tool discovery. Future-proof by design.

INGEST & CAPTURE

Get documents in

Inbound email → auto-analyzed · AWS S3 · Browserbase cloud screenshots · a Chrome side-panel for audit capture.

REST API

Python SDK

PDF

DOCX

Excel

CSV

JSON

Not a chatbot — a fleet of specialists.

The system decomposes a review into parallel specialized agents, cross-checks itself, and escalates only genuine uncertainty to a human — with a full audit trail of every step.

- ◇ **Agentic playbook builder**, multi-playbook × multi-persona **review panels** with weighted consensus, a **due-diligence orchestrator**, classifiers, and a **gap detector** ("4 founders named, only 3 CVs").
- ◇ Long-running work runs in a **Postgres-backed autonomous worker pool** — atomic claiming, retries, self-reclaim of crashed jobs.
- ◇ **MCP tool use** lets agents pull external data mid-analysis.

AGENTIC PLAYBOOK BUILDER · PIPELINE

- PARSE — extract all sheets / pages, no truncation
- ANALYZE — structure, hierarchy, patterns
- GENERATE — questions per group, full context
- REVIEW — an independent model checks quality
- MERGE — dedupe, assign IDs, version

Price it yourself — you pay for evidence, not tokens.

Every sealed finding is one **Evidence Unit**. Set your scope and we size the EVUs to service it at our *most rigorous* settings — consensus, adversarial review, screenshots, and a sealed hash-chain. Two components: a **one-time initial analysis**, then a **monthly monitoring bundle**.

FRAMEWORK


DORA · Art. 30

Due diligence

VENDOR AGREEMENTS MONITORED · 25



NEW / UPDATED DOCUMENTS A YEAR · 10



16 questions in the DORA · Art. 30 playbook

Auto re-check on legislation change · ~3/yr assumed

≈ 6.3 EVU per sealed finding $1.2 \text{ base} \times 3 \text{ consensus} \times 1.6 \text{ adversarial} \times 1.1 \text{ screenshots}$

Monitoring minimum · 250 EVU/mo keeps automated surveillance armed in quiet months

Illustrative · EVU weight = base 0.6–1.2 × mode & evidence factors (evu_ledger §4.2) · bundles on a sliding scale · free trial to start.

2,534

EVUS · INITIAL ANALYSIS (ONE-TIME)

250

EVUS / MONTH · CONTINUOUS MONITORING

€1,225 one-time for the initial analysis of 25 agreements, then €120 / month — a minimum bundle of 250 EVUs that keeps automated legislative & document-change monitoring armed and covers re-checks and updated documents. Sliding scale from €300 per 1,000 EVUs.

Trust, security, and architecture a CTO can sign off.

- ISO/IEC 27001 • information security
- ISO/IEC 42001 • AI management system
- GDPR by design
- EU data residency

TRACEABILITY Every answer defensible • Citations traceable to the source clause • Timestamped screenshot evidence (+RFC 3161) • 0.0–1.0 confidence · human-in-the-loop on low confidence • Quote / semantic / adversarial verification layers	SECURITY Isolation by default • RLS on every tenant table • Roles: user · org-admin · system-admin • Workspace-scoped collaboration, server-authoritative • Staged JWT verification · append-only audit trails • Aligned to ISO/IEC 27001 & ISO/IEC 42001	ARCHITECTURE Deployable = verified • Next.js + FastAPI + Supabase (Postgres · pgvector · RLS) • Decoupled, containerized, Postgres job queue • CI merge-gate: lint · tests · migration dry-run · docker • Full en / fr / de, multi-jurisdiction
--	--	--

Built for the people who sign off — inside and out.

A finding doesn't stop at "flagged." The people who own the risk — your internal legal & compliance teams, and the external auditors who check them — get first-class workflows, not a CSV dump.

- ◇ **Auditor portal.** Grant an external auditor or regulator scoped, read-only access to the sealed evidence for one engagement. They verify the hash-chain and timestamps themselves — **no ComplianceOS account**, and nothing outside the engagement is exposed.
- ◇ **Remediation workflows.** Escalate a flagged finding into a tracked remediation — owners, deadlines and decisions recorded in an **append-only ledger**. Overrides never mutate the AI's original finding; attach side letters and file periodic DOCX reports with deadline reminders.

ENGAGEMENT · SCOPED AUDITOR ACCESS

• read-only

hash-chain verifiable

no account

finding **flagged** → remediation opened
→ owner assigned · deadline set
→ override recorded (AI finding preserved)
→ side letter attached · report filed
→ auditor verifies the sealed chain ✓

AI findings a regulator will actually accept.

The market has plenty of "AI that reads documents." The hard part — and the moat — is making the output **provable, sovereign, and reusable**. That's the whole product.

Provable

Cryptographically sealed evidence, verifiable without trusting the vendor.

Sovereign

EU-only data plane, model-agnostic, portable — no lock-in.

Reusable

Collect evidence once; satisfy every mapped control.

Governed

Versioned playbooks, human-in-the-loop, full audit trail.

Collect once. Prove forever.

Let's run one of your real reviews through it — and hand you a sealed, independently verifiable result at the end.

• [Book a working session](#)

practicl.ai

PRACTICL AI SAS

16 Rue Saint Antoine du T, Bureau 3 • 31000 Toulouse,
France

SIREN 999 890 056 • SIRET (siège) 999 890 056 00016 • RCS
Toulouse • N° TVA intracommunautaire FR82999890056