PolicyLenz.

v2.2.1 · SEALED OFFLINE DISTRIBUTION

The compliance analysis engine that runs inside your perimeter.

PolicyLenz assesses security policy documents against **NIST CSF 2.0, ISO/IEC 27001:2022, SOC 2, and GDPR**. Every run produces a quantitative score per framework, clause-level gap findings, and audit-ready remediation, and the whole system ships as a **sealed Docker package** that keeps your documents on your own infrastructure.

AN ENGINE, NOT AN AI CHAT

A chat wrapper sends your document to a model and formats the reply. PolicyLenz is a deterministic pipeline in which the model is one bounded, replaceable component: the catalogs, the scoring, the merge rules, and the audit trail are all engine code, not prompts.

Every control, every time. The framework catalogs are enumerated structured data, not a retrieval index. All 296 controls are assessed on every run, so gaps can never be silently dropped.

Deterministic scoring. Fixed weights (full 1.0, partial 0.5, missing 0.0), a reproducible best-wins merge across document chunks, and no cross-framework averaging. The engine computes the numbers; the model never invents them.

An engineered LLM boundary. Structured output through forced tool use, computed token budgets, and an explicit retry ladder with model fallback. A failed framework degrades a run to partial instead of poisoning it.

The engine grades the AI. Policies drafted by the built-in generator are validated by the same analyzer that grades everything else, then revised in a loop until they clear the bar.

BUILT FOR PEOPLE WHO AUDIT THE AUDITORS

Immutable audit trail

Policy versions are immutable, runs pin the exact version they assessed and freeze their scope forever, and hard deletes are blocked while a run holds a reference. Dispositions and evidence overlay results without ever rewriting the raw verdicts.

Human-gated remediation

The AI proposes a change set, never a new version. Each change is accepted, edited, or rejected in a word-level diff review; apply is blocked until every change is decided, and the output DOCX carries a decision log naming each decider.

Sealed and local-first

Documents never leave the machine. Egress is default-deny with a single allow-listed host (`api.anthropic.com`) and a logged ledger of every attempt. No embedding or vector libraries ship in the image.

Signed content updates

Framework updates arrive as Ed25519-signed bundles verified against a pinned key, with per-file hashes, no downgrades, atomic install, and one-click rollback. Bundles carry content only, never app behavior.

Enterprise access model

Five application roles plus three project roles, invite-only onboarding, TOTP MFA with org-wide enforcement, last-admin protection, and a filterable audit viewer for admins and auditors.

Honest numbers

Effective and policy-only scores are reported side by side. Scope profiles make the score's denominator the engagement's intent, not the whole catalog, and cross-framework crosswalks are deliberately report-only.

4	296	214	29	13	266	5	2
FRAMEWORKS	CONTROLS ENUMERATED	CROSSWALK MAPPINGS	TABLE DATA MODEL	SCHEMA MIGRATIONS	AUTOMATED TESTS	SEALED RELEASES	ARCHITECTURES (AMD64 + ARM64)