

// AI MODEL CARD · v1

# Vendric Labs

## Model Card

*Per-product AI provenance: which model is in the loop, what it sees, and where the human is.*

Document	Model Card · v1
Owner	Taylor Marr · founder
Last revised	2026-05-27
Scope	Reeve (live) + Vendric (pre-launch)

### 1 · REEVE

OSHA compliance · Live · production

Use case	Photo-to-hazard analysis (construction-site photos → hazard categorization with verbatim 29 CFR 1926 citations) and OSHA-rule retrieval grounding.
Underlying model	Anthropic Claude Sonnet (4.x family). Vendor-managed inference.
Fine-tuning	None. Off-the-shelf model with retrieval-grounded prompting.
Training opt-in	We do not opt into Anthropic's training-data program for this workload. Customer submissions to Reeve are never sent to any training pipeline (ours or vendor).
Future training corpus	Limited to de-identified public regulatory data (OSHA citations, 29 CFR text, court filings). Because those records are public, we cannot guarantee a specific customer's prior incidents are absent from the public source data.
What it sees	Construction-site photos submitted by the customer plus the relevant 29 CFR 1926 paragraphs surfaced by the citation index. Nothing else.
What it does NOT see	Cross-tenant data. Photos from other GCs. Worker PII outside the submitted photo. Anything from the audit chain unrelated to the active photo.
Output style	Verbatim citations only — never paraphrased. Every hazard claim is anchored to a paragraph and line in 29 CFR 1926.

---

<b>Human-in-the-loop</b>	Every penalty-bearing output is reviewed by the customer (the GC) before submission to OSHA. The system drafts; the GC signs. The system does not auto-file on the customer's behalf.
<b>Evaluation</b>	Methodology under documentation. Contact <a href="mailto:security@vendriclabs.net">security@vendriclabs.net</a> for current status.
<b>Versioning</b>	Model upgrades to a new Sonnet version are reviewed before customer-visible rollout. Pinning + regression testing posture under documentation.

---

## 2 · VENDRIC

Substrate layer · Live · provenance under Reeve in production

<b>Use case</b>	Substrate layer providing memory, cognition, prediction, and provenance primitives that downstream surfaces (Compli, Reeve) inherit.
<b>Underlying model</b>	Not a customer-facing inference surface. Substrate-layer building blocks only — no user-visible model runs at this layer.
<b>What it sees</b>	Internal data structures originated by Compli and Reeve operations. No direct user input.
<b>What it does NOT see</b>	End-user prompts or conversations. Vendric does not have a user surface.
<b>Training opt-in</b>	Not applicable — no inference endpoint.
<b>Persistence format</b>	Encrypted persistence format · spec frozen. Tamper-evident audit chain primitives live in production via Compli.
<b>Human-in-the-loop</b>	Not applicable at this layer. Human-review checkpoints live in the downstream product surfaces (Reeve today; future Compli verticals tomorrow).
<b>Status</b>	Live · substrate. Provenance layer beneath Reeve in production.

## 3 · METHODOLOGY

Every Model Card entry above is intended to be falsifiable. If a model changes, the entry changes. If a behavior shipped on Tuesday and was removed on Friday, the card reflects Friday. The card is part of the product — not a marketing surface decoupled from the engineering reality.

### 3.1 · What we will not claim

Per FTC guidance and the AI Trust + Security brief informing this rebuild, every AI claim on a Vendric Labs surface is backed by a specific output, customer type, or measurable behavior. We do not use the phrase “AI-powered.” We do not claim “revolutionary” capabilities. We do not benchmark on internal data and report the result as third-party validation.

### 3.2 · Human review checkpoints

Reeve: every form submission, every penalty estimate, every photo-to-hazard categorization is shown to the GC before any external action. The system does not auto-file with OSHA.

Vendric: substrate layer. No direct user actions originate here. All user-facing decisions happen at the downstream product surface (Reeve today; future Compli verticals tomorrow).

### 3.3 · Reporting concerns

If an output appears wrong, biased, or unsafe, email [security@vendriclabs.net](mailto:security@vendriclabs.net). We treat model behavior issues with the same urgency as code bugs.