ENDOR LABS

AI Model Governance

Govern open source AI models like the dependencies they are

Developers are pulling open source AI models from Hugging Face directly into production applications — without the review every other open source dependency gets. AI Model Governance treats AI models as software dependencies and applies a full lifecycle — discover, evaluate, enforce — inside the same platform AppSec teams already use for SCA.

50+
risk signals scored on every Hugging Face model

3 steps
to govern models: discover, evaluate & enforce

1 platform
for AI models and every other open source dependency — same policies, same SBOM

How it works

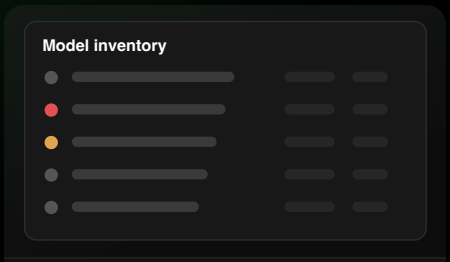

DISCOVER

Know which models are in your code

- Detect open source AI models automatically as part of every scan
- Report AI model usage in your SBOM alongside every other dependency
- Keep an audit-ready record of which models are in which applications

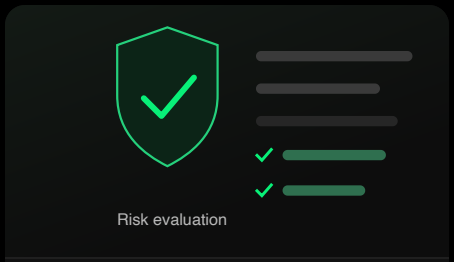

EVALUATE

Score risk before it ships

- Score models across 50+ signals spanning security, licensing, popularity, activity, and quality
- Catch unsafe file formats like pickle and unverified PyTorch that execute code on load
- Flag typosquatting, restrictive licenses, and abandoned models before developers commit

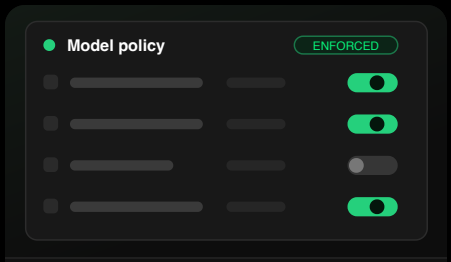

ENFORCE

Policy instead of ticket queues

- Warn developers in the IDE or block risky models from reaching CI
- Use pre-built policies or write custom rules with the same Open Policy Agent engine
- Let policy handle routine approvals instead of reviewing models one ticket at a time

WHY NOW

AI models are the new unreviewed dependency

Open source AI models are entering the codebase the same way open source libraries did a decade ago — pulled in by developers, with no security review. Traditional SCA tools cannot see them, so they never show up in your inventory or SBOM, and standalone AI security products don't connect to the rest of the application. AppSec teams are accountable for AI components either way. Governing models inside the platform you already run closes the gap without a separate tool, separate policies, or a separate workflow.