Mend AI

DATASHEET

Secure the layer where AI behavior becomes application risk

Continuous visibility, behavioral testing, and in-application guardrails for AI-powered applications — from prompts, agents, models, and MCPs to runtime interactions.

THE PROBLEM

Risk no longer lives in one layer

AI is expanding application risk beyond code — into prompts, agents, models, retrieval pipelines, tools, and runtime behavior.
The most dangerous risk emerges in the interaction between code, AI components, data, and agent behavior, where traditional AppSec and standalone AI tools have limited context.

More AI in more apps

Models, copilots, agents, MCPs, and prompts appear across application estates faster than manual review can track.

Risk between layers

The dangerous gaps sit between code, AI behavior, tools, and data flows — where point solutions lack context.

Less time to reduce risk

AI accelerates development and vulnerability volume, making automation and continuous enforcement essential.

WHAT MEND AI DOES

Discover, assess, test, and control AI apps

Mend AI extends Mend.io's application security platform to the AI layer — helping security teams manage AI-powered applications across the full lifecycle.

01

Discover

Inventory models, prompts, agents, frameworks, MCPs, and AI components in application code.

02

Assess

Surface weak prompts, risky components, license and malicious indicators with proprietary AIWE scoring.

03

Test

Run automated AI red teaming against real apps,
APIs, direct LLMs, and external channels.

04

Enforce

Apply in-app guardrails for user-to-agent and agent-to-tool interactions — no third-party proxy.

CORE CAPABILITIES

Everything you need to secure AI apps

AI-BOM & AI discovery
Inventory models, agents, and prompts.

System prompt hardening
Close weak, risky prompt paths.

Automated AI red teaming
Probe injection and data leakage.

Runtime guardrails
Enforce policy on agent calls.

AI-specific risk scoring
Rank weaknesses with AIWE scoring.

Unified governance
Tie AI risk to AppSec workflows.

WHY MEND AI IS DIFFERENT

One platform for the whole AI risk surface

- **Built for application context.** Tests prompts, agents, and workflows for behavioral weakness, with guardrails that screen risky prompts.
- **Covers the full AI lifecycle.** Discovery, hardening, testing, guardrails, and governance as one platform — not another point tool.
- **Meets teams where they work.** Fits existing AppSec and developer workflows without re-platforming your security model.
- **Protects the interaction layer.** Secures the code layer, the AI layer, and the risky space between them.

USE CASES

Where teams put Mend AI to work

Find Shadow AI
Surface unmanaged AI and prompts.

Harden apps before release
Test AI behavior pre-production.

Govern AI at scale
Policy, dashboards, and evidence.

Protect agentic workflows
Curb risky actions and exposure.

Prioritize AI risk
Focus on what increases risk most.

Extend AppSec to AI
One motion with code and OSS.

Mend.io is built for every risk, across AI and AppSec. By securing the code layer and the AI layer—and the interactions between them, where modern application risk now lives—**Mend.io** extends proven AppSec workflows to the models, prompts, and agents inside today's applications, delivering continuous protection across the entire AI application lifecycle.

Visit our website