AI EXPOSURE ASSESSMENT

An assessment you can run in a week — and build a practice on.

KeyCaliber reads your client's existing security tools and returns a complete AI inventory, ranked by which critical systems it touches. Every finding is scoped follow-on work.

THE ENGAGEMENT SHAPE

DAYS 1–2

Connect

Read-only API credentials to tools they already run. No agent, no proxy, no rollout.

FIRST DAYS

Find

A complete AI inventory across seven discovery channels, ranked by business impact. Unknowns identified.

WEEKS 2–3

Report

The evidence export is your appendix. The ranked findings are your roadmap.

THE FINDING THAT EARNS THE NEXT CONVERSATION

Shadow AI + Shadow IT:

“This laptop is running a local AI model. It belongs to someone in Finance. Nothing you currently own is securing it.”

AI models running locally on laptops and servers

Browsers can't see a binary; gateways can't see a process. The KeyCaliber platform reads the software inventory their EDR and vulnerability scanners already collect but don’t correlate.

AI features switched on inside vendors they already approved

Their SaaS vendors passed review before they included AI. The KeyCaliber platform shows what SaaS now uses AI, which is on by default, and training on customer data.

AI touching business-critical systems

Instead of guessing at asset criticality or ignoring it, the KeyCaliber platform provides a ranking list based on a computed impact score so clients see AI on critical assets.

AI called from servers, scripts and APIs

No browser, no installed app — just an IP in a firewall log. The KeyCaliber platform translates that into a named AI service.

SCOPING

What each finding opens

The assessment is the wedge. Every finding maps to a workstream you already deliver.

Local AI runtimes on endpoints	Endpoint policy, software allowlisting, acceptable-use rewrite.
Embedded AI in approved vendors	Vendor re-review, data-processing terms, a procurement gate.
AI in use on business-critical systems	A remediation sprint scoped by impact, not alert count.
No authorization model for AI	Governance: dispositions per app, approval by group, review cadence.
Nothing an auditor would accept	EU AI Act Article 50 and ISO 42001 readiness.

Then ongoing: managed AI exposure, re-run on a cadence or delivered continuously.

DELIVERING

What you can provide

Utilize built-in reports or create your own based on the platform's findings.

Complete AI inventory	Every app, model and local runtime, with discovery method and confidence.
Ranked exposure list	Ordered by business impact — the remediation plan, pre-sequenced.
Vendor posture	Data-training policy per app, dated and sourced.
Authorization state	Status per app, who set it, and who is outside it.
Evidence export	The full inventory as CSV — the artifact an ISO 42001 or Article 50 file needs.

QUALIFYING

What you need from the client

- **Read-only API credentials** to tools they already run.
- **No single feed is mandatory.** Coverage scales with what they connect.
- **Endpoint inventory alone** finds local AI runtimes — no traffic feed needed.
- **Where their identity data supports it,** findings name the person, not just the IP.

Run the first one with us.

We'll co-deliver your first assessment — you keep the client and the follow-on work. We ship no sensor and sell no remediation.

keycaliber.com

partners@keycaliber.com