KEYCALIBER

AI MANAGEMENT

The AI inventory your other tools can't finish.

Every AI app, model and local runtime in use across your environment — including Shadow AI — discovered from the security tools you already run and **ranked by which critical systems it touches**, with the person behind it named wherever your identity data supports it. Exports as evidence an auditor will accept. No agents. No proxy. Nothing installed.

WHY MOST AI INVENTORIES ARE INCOMPLETE

Every discovery approach sees one slice of the organization. The AI you're missing isn't hiding — it's in the slice your tool was never built to watch.

WHERE AI ACTUALLY SHOWS UP	BROWSER EXTENSION	NETWORK GATEWAY	SAAS POSTURE TOOLS	KEYCALIBER
Web AI apps used in a managed browser	•	•	•	•
AI called from servers, scripts and APIs	—	•	—	•
AI features inside apps you already approved	—	—	partial	•
AI models running locally on a laptop	—	—	—	•
Which of that AI touches critical systems	—	—	—	•

SEVEN WAYS KEYCALIBER FINDS AI

01 **Network & firewall logs** — traffic to AI services, scored for confidence.

02 **DNS logs** — resolves destinations raw traffic data leaves unidentified.

03 **Web gateway & CASB** — full request detail, including the model in use.

04 **Single sign-on events** — names the person behind the login.

05 **Installed software** — local AI runtimes on laptops and servers, from your EDR and MDM. Almost nobody looks here.

06 **Embedded-AI catalog** — which approved vendors ship AI, and whether it's on by default.

07 **Manual entry** — add what your team already knows; discovery merges into it.

→ **One canonical list.** Found five ways, an app is still one row.

WHY THIS MATTERS NOW

43%

of security incidents now involve **shadow AI** — more than double last year's 20%.

IBM / PONEMON 2026

1 in 3

environments have **at least one endpoint with AI model files** sitting on disk.

TANIUM GUARDIAN 2026

40%

of enterprise apps will ship **task-specific AI agents** by end of 2026, from under 5% in 2025.

GARTNER

In force

EU AI Act **transparency duties and penalty regime** apply since Aug 2, 2026. Fines to €15M or 3% of turnover.

EUROPEAN COMMISSION

Sources: IBM Cost of a Data Breach 2026; Tanium Guardian endpoint research 2026; Gartner; European Commission. The AI Omnibus (in force July 27, 2026) deferred Annex III high-risk obligations to December 2, 2027 — the inventory work behind them was not. The 7% / €35M ceiling applies to Article 5 prohibited practices, not transparency breaches.

A LIST BECOMES A WORK QUEUE

Forty AI apps is a research project. The same forty, ordered by what they touch, is a short list your team can finish this quarter.

Local AI runtime on a finance workstation 1 person · asset identified as crticial · prohibited software	ACT NOW
Consumer AI app used from a payroll server 3 people · vendor trains on submitted data by default	ACT NOW
AI writing features switched on in an approved app whole tenant · enabled by default at the vendor's last release	REVIEW
Approved AI assistant, engineering group 61 people · inside the approved population	NO ACTION

ILLUSTRATIVE FINDINGS

FROM FINDING TO EVIDENCE

01 RANK

Every finding ordered by what it touches, not by a severity label.

02 GOVERN

- Unreviewed
- Evaluating
- Approved — everyone
- Approved — limited group
- Prohibited

One disposition per app. Authorize a named group; everyone else queues as unauthorized.

03 PROVE

Method, confidence, vendor training policy, decision and actor — the audit file.

Find the AI you don't know about.

Connect one endpoint tool and one network source. Most first runs surface local AI runtimes and approved-vendor AI features the customer didn't know were there.

keycaliber.com
hello@keycaliber.com