Building a Real-Time AI Inventory

How Continuous Discovery Replaces Manual AI Asset Tracking

THE REALITY

Every AI governance framework begins with the same fundamental requirement of knowing what AI exists across the organization. The EU AI Act, NIST AI RMF, and ISO 42001 all treat the AI system inventory as the foundation on which every subsequent control depends, from risk classification to conformity documentation. GRC teams have reacted by building AI inventories through the same processes that have served every other asset class, through surveys, stakeholder interviews, and manually maintained spreadsheets.

However, an organization's AI estate changes faster than any manual process can track. Research from IAITAM found that 98% of organizations report unsanctioned AI use within their environments. In the same vein, a shadow AI discovery report found that enterprises run 3.2 times more AI tools than their registries reflect. Additional research puts the operational picture even more bluntly, with 47.4% of organizations acknowledging that they have at most only a "partial" inventory of the AI tools in use across their workforce.

A hand-built inventory captures a snapshot of an environment that changes continuously, and the distance between the snapshot and reality grows from the moment the collection window closes.

THE BLIND SPOT

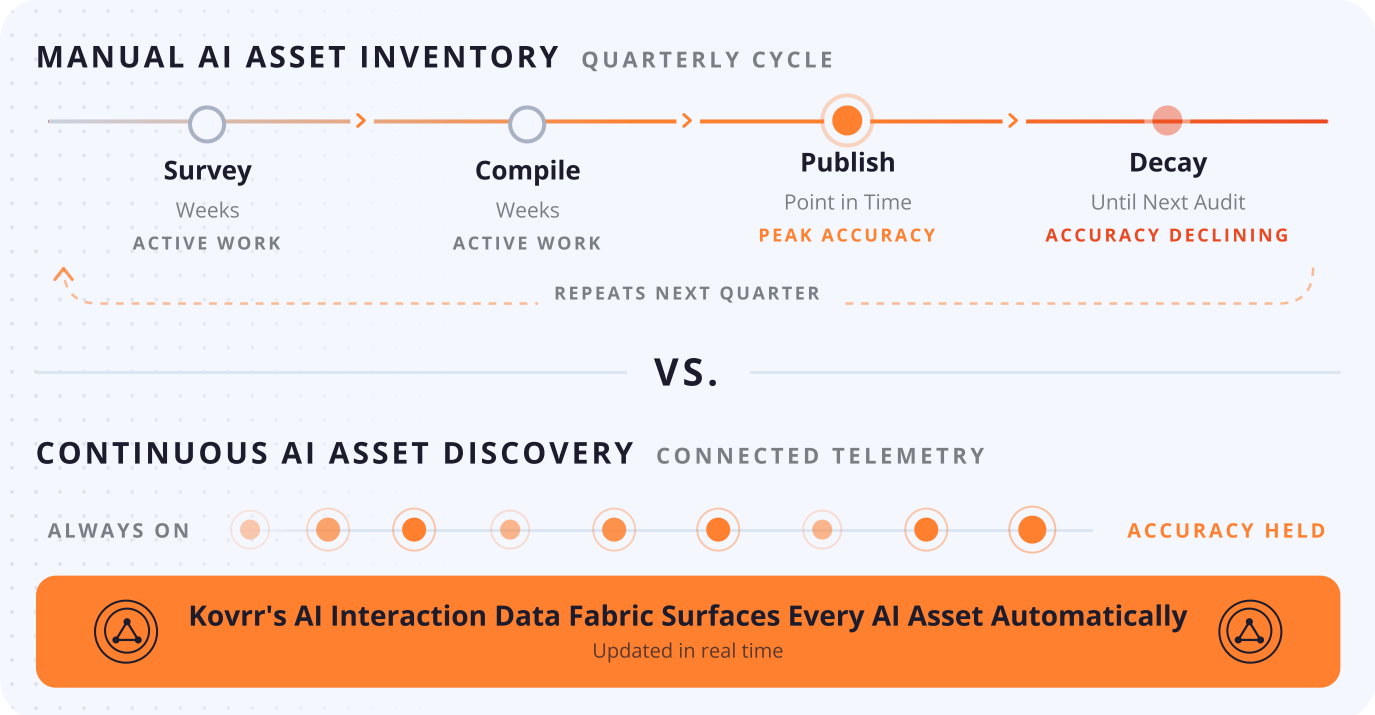

The first failure mode is velocity lag. A survey-based inventory captures a single point in time, and the AI environment evolves almost immediately afterward. Employees onboard new tools without procurement involvement. Departments license AI-powered software directly. SaaS vendors embed AI capabilities into existing products through silent updates, meaning an application that was AI-free at the last audit may be processing corporate data through a language model today. A quarterly review cadence produces a registry that spends most of its life describing an environment that no longer exists.

The second failure mode is lack of coverage. Manual discovery depends on what people know to report, and the most significant exposure sits in places no survey reaches. Teramind found that 89% of workplace AI use escapes enterprise governance, flowing through the approved platforms organizations deployed and trusted rather than through rogue applications. The same research found 67% of enterprise AI usage running through unmanaged personal accounts on corporate-licensed platforms.

Vendor-embedded models, browser extensions, and autonomous agents compound the problem. Indeed, the Cloud Security Alliance notes that a shadow AI agent actively initiates connections to external services and executes code, a fundamentally different exposure profile from a passive shadow SaaS application.

The third failure mode is escalating regulatory liability. An incomplete AI inventory, once an operational inconvenience, now carries enforceable legal and financial consequences. The EU AI Act ties binding obligations to a documented, risk-classified inventory of AI systems, with penalties for high-risk non-compliance reaching €35 million or 7% of global annual turnover. The financial exposure is already measurable, with industry analysis attributing an additional $670,000 in breach costs to organizations with high levels of ungoverned AI.

THE PLATFORM RESPONSE: AI INTERACTION DATA FABRIC

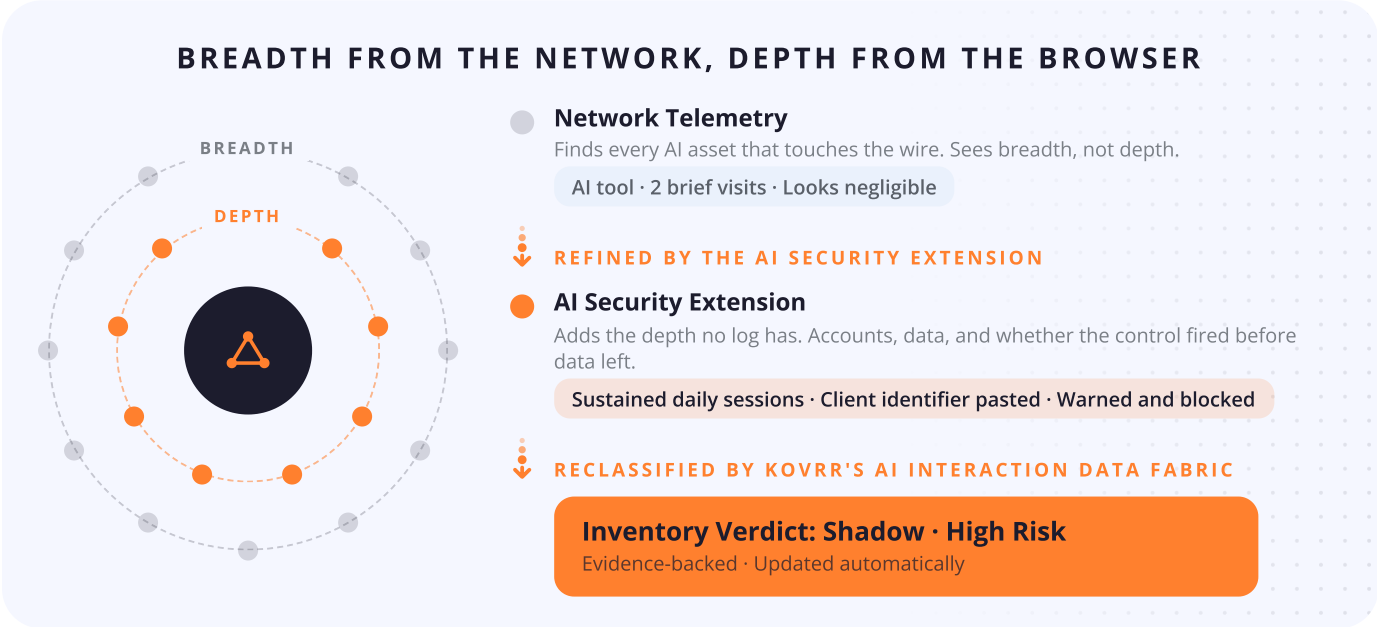

Kovrr's AI Security and Governance Platform treats the AI inventory as a living output of connected telemetry rather than a document maintained by hand. AI Asset Visibility, the foundation of the platform, collects signals from network traffic, identity systems, data governance events, endpoint logs, and browser and agent telemetry, alongside integrations with enterprise systems. Every AI asset operating across the entity surfaces automatically, whether sanctioned or shadow, internal or third-party, browser-based or embedded inside a vendor product.

Each discovered asset lands in a continuously updated inventory enriched with context that manual approaches never capture. The AI Vendor Risk Catalog scores every application across dimensions including model risk, business criticality, data and regulatory exposure, and company risk, drawing on a catalog of more than ten thousand AI applications. Discovery extends past the application layer to the interaction layer, revealing usage patterns down to the department level where available, what categories of data flow into each tool, and how usage changes over time. Entries become risk-scored, evidence-backed profiles that automatically update.

The inventory feeds every other capability on the platform. AI Compliance Readiness maps discovered assets against the EU AI Act, NIST AI RMF, and ISO 42001, turning the living inventory into audit-ready documentation. Downstream, discovered assets flow into risk scenario modeling and financial quantification, meaning the same telemetry that builds the inventory also powers the platform's broader governance and risk workflows.

THE RESULT

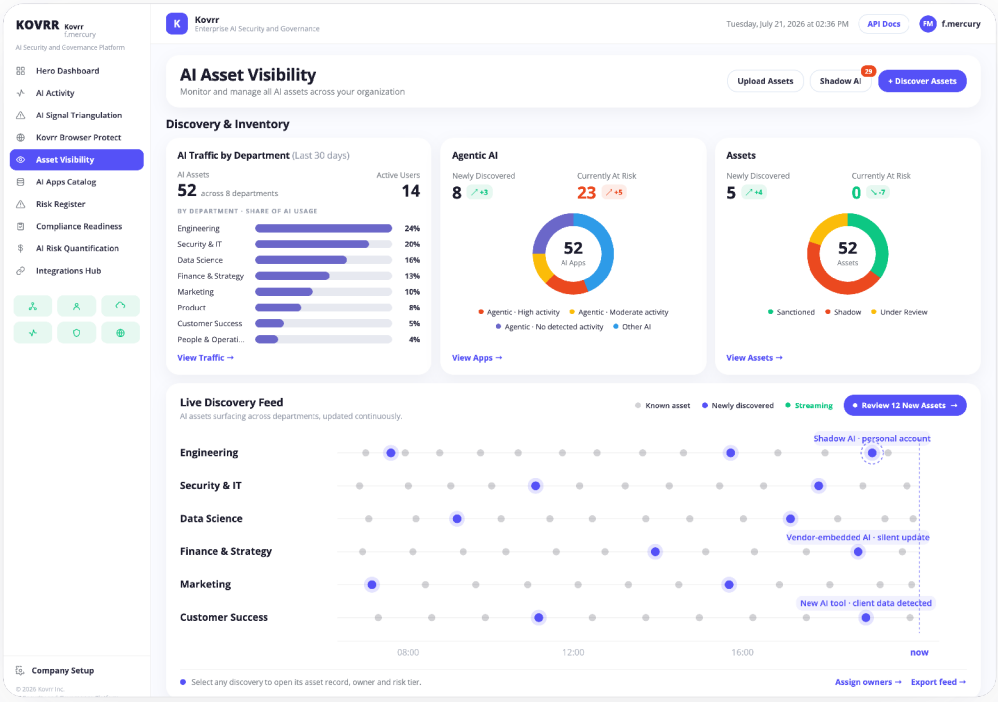

GRC and compliance teams can retire the quarterly audit and let the inventory maintain itself. A customer service AI tool might first surface as anomalous network traffic, gain context from endpoint and identity signals, and resolve into a full picture once browser-level telemetry reveals which teams use it, from which accounts, and with what data. The network establishes that an asset exists, and the additional signals explain how it is being operationalized across the organization.

Shadow AI surfaces as it appears, including the usage hiding inside approved platforms where most ungoverned activity concentrates. Regulatory readiness becomes a standing condition rather than a scramble before each audit, with every inventory entry carrying the risk classification and documentation that frameworks require. The same telemetry that builds the inventory powers risk quantification, meaning the assets carrying the greatest financial exposure are surfaced.

Replace the quarterly survey with an AI inventory that builds itself. Book a demo.

KOVRR

AI INTERACTION DATA FABRIC

Telemetry alone is data. Telemetry triangulated across sources is evidence.