KOVRR AI INTERACTION DATA FABRIC IN ACTION

Unauthorized ChatGPT Use Detection · Sensitive Financial Data Upload Prevention

Shadow AI in Finance

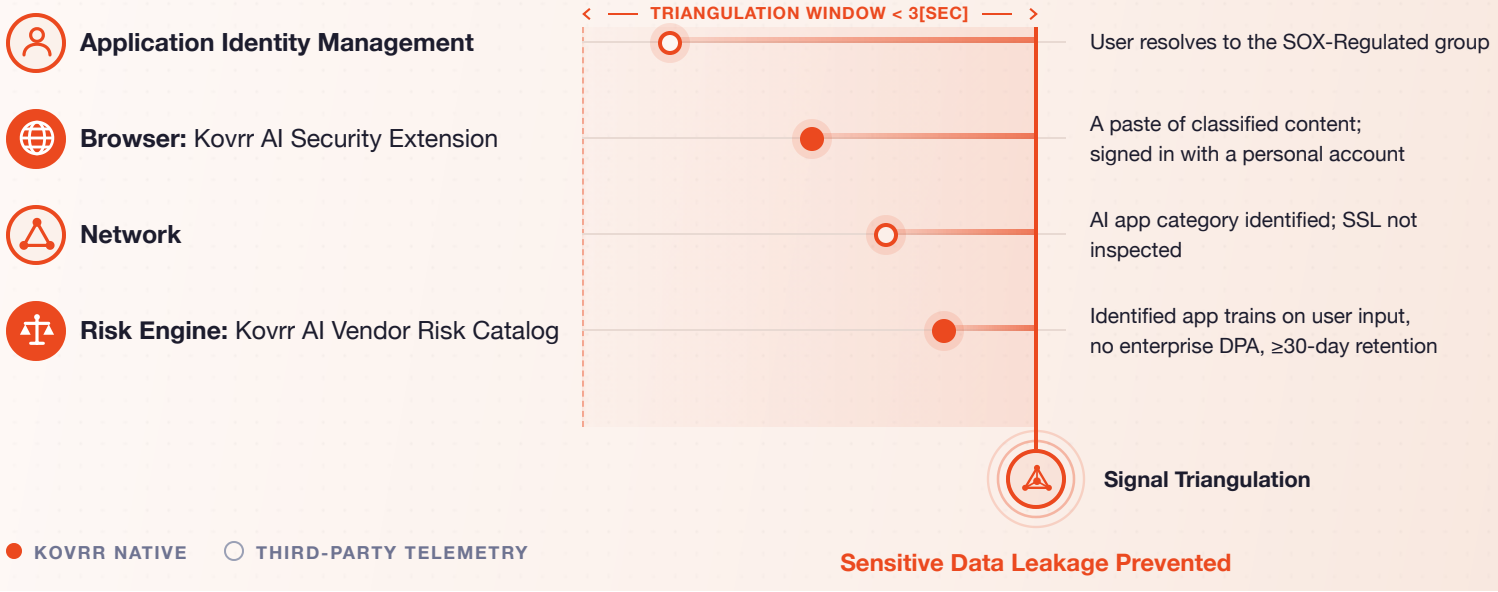

RISK SCENARIO EXPLAINED

A finance-team member reaches for an unsanctioned consumer AI assistant to speed up month-end close. Finance identities working inside sanctioned applications such as the ERP and the approved copilot fall within normal behavior. A finance-role identity pasting classified financial data into ChatGPT falls outside it. The vendor profile for that tool records training on user input and retention of at least 30 days, with no enterprise data agreement in place. The activity qualifies as regulated-data exposure under GLBA and SOX. Against any single tool in the stack, it registers as ordinary web traffic.

EACH TOOL CAPTURES A SLIVER. TRIANGULATION SURFACES THE EXPOSURE.

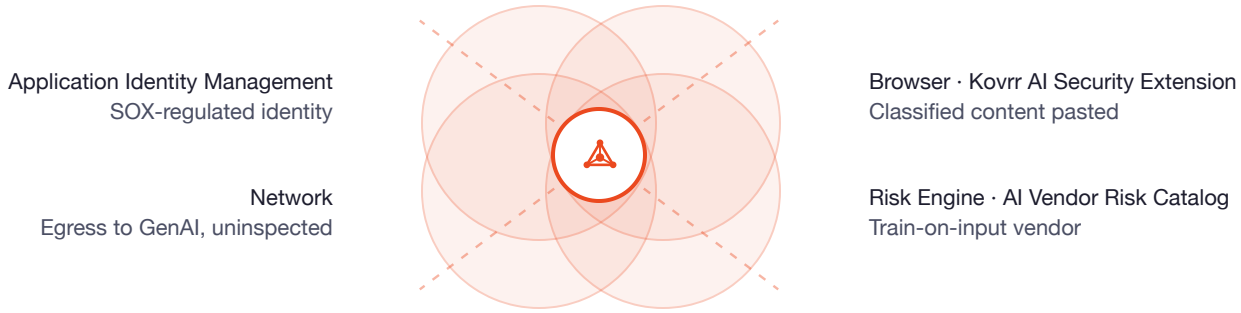

Each signal reads as benign on its own. The overlap is where the event becomes a reportable regulated-data exposure.

ALSO INTEGRATED IN THE FABRIC, NOT TRIGGERED THIS EVENT: ENDPOINT · DLP · AGENT TELEMETRY · CLOUD

WHY SIGNAL TRIANGULATION WINS · EACH SOURCE'S BLIND SPOT IS ANOTHER'S STRONG SUIT

- The **Network** source recorded a 41 KB upload to a GenAI domain, though cert pinning kept the contents opaque.
- The **Application Identity Management** source identified a SOX-regulated Finance identity on a managed device, with no view into what that identity did next.
- **AI Security Extension** read the paste as it happened in-page, with visibility limited to the browser session.
- **AI Vendor Risk Catalog** held the train-on-input classification for the vendor, with no visibility into who reached it or what they sent.

Only triangulation across the sources establishes that one identity pasted regulated content into a train-on-input consumer tool over an upload the proxy could not read. Signals that read as benign in isolation become one reportable exposure.