relyanceai

10 things to look for in AI Security

Not all AI security tools are built the same. Use this checklist to evaluate any AI-SPM or data security solution — and see how Relyance AI delivers on each one.

01

Always-current AI asset inventory across code, cloud, SaaS, and third parties

Continuously discovers every model, agent, MCP server, and SaaS AI integration — including shadow AI — across code, cloud, and third parties. A point-in-time map leaves you blind to the AI your organization is actually running.

Live AI inventory from code to cloud

Relyance AI maintains a continuously updated inventory across code repositories, cloud, SaaS, and third-party integrations — including Shadow AI.

02

Sensitive data flow tracing into, through, and out of every AI system

Traces exactly what PII, PHI, PCI, and proprietary data flows into training pipelines and inference, where it goes, and who or what touches it at every step. Metadata inspection is not enough.

Data Journeys™: sensitive data traced from source code to inference

Data Journeys™ is built on code-level data flow analysis — the only AI security capability that traces sensitive data from source code through model output, capturing risks that infrastructure-level scanning cannot reach.

03

Visibility into what every AI agent can access – and what sensitive data it can reach

Maps the complete delegation chain: which agent holds which API key, which service account, and what sensitive data it can ultimately reach. Overprivileged agents are among the highest-impact risks in production environments.

Identity-to-data intelligence — tracing what every AI agent can actually reach

Relyance AI maps permission chains from agent to API key to service account to sensitive data store — making the invisible relationships between non-human identities and sensitive data visible and actionable.

04

MCP server discovery and continuous risk assessment

Discovers third-party MCP servers — a primary channel through which AI agents access enterprise systems — and continuously assesses the risk they introduce. Each is a trusted access point; a misconfigured one is a direct path to enterprise data.

Purpose-built MCP server risk coverage

Relyance AI provides automated discovery and continuous risk assessment of MCP servers across the enterprise — mapping the sensitive data each can reach and the identities using them, across an attack surface most tools don't recognize yet.

05

Compound risk detection across data, identity, and AI behavior

Surfaces dangerous combinations invisible from any single dimension. An AI agent with administrative access to a sensitive database is a critical vulnerability no point-in-time scan would flag. The risk only appears when data, identity, and behavior are assessed together.

Compound risk detection at the intersection of data, identity, and AI

Relyance AI's Data Exposure Graph correlates data sensitivity, identity privilege, and agent behavior in real time — surfacing the lethal combinations that only emerge at the intersection of all three, and that no siloed tool can see.

06

Code-level data flow analysis – not just cloud infrastructure scanning

Most AI risks originate before data reaches a cloud resource. Tracing must start at source code and CI/CD pipelines, following data from a repository into a training pipeline or credentials embedded in agent scaffolding, all the way through to inference.

Code-level analysis where other tools start too late

Relyance AI's data-journey-first architecture traces sensitive data from source code through every transformation to model output — the foundational differentiator that infrastructure-focused tools cannot replicate.

07

Third-party and vendor AI risk – continuous, not questionnaire-based

Automates discovery and continuous risk assessment of SaaS vendors with AI features enabled, including validation of what data flows into those systems. Static questionnaires are out of date before they're filed.

Continuous vendor AI risk — no more questionnaires

Relyance AI discovers all vendor AI use, enriches DPIAs with live data flow context, and validates vendor runtime behavior against contractual AI usage clauses — cutting vendor AI risk assessment time by up to 70%.

08

Actionable explainability – a connected story, not isolated alerts

Assembles the full risk chain automatically: what is wrong, why it matters, who is involved, what data is at risk, and what to do — so teams can act without stitching context together across multiple tools.

Every finding arrives with the full story already assembled

Every Relyance AI finding includes the complete context chain. Lyo™, Relyance's natural language interface, lets teams interrogate risk and act without navigating dashboards — eliminating the investigation tax entirely.

09

Continuous AI security posture monitoring connected to live data and identity

Always-on and tied to live data and identity context. When models are retrained, agents gain access, or permissions expand, exposure surfaces immediately — not at the next quarterly scheduled review.

Always-on posture monitoring tied to live data and identity

Relyance AI continuously monitors AI security posture with out-of-the-box policies connected to live data flows and identity context — detecting drift as it happens, not after a periodic scan.

10

Live compliance mapping and audit-ready evidence – always current

Compliance evidence should come from a live AI inventory, not documentation assembled retroactively from multiple systems. Requires automated, continuous mapping to EU AI Act, NIST AI RMF, ISO 42001, and other frameworks.

Live compliance evidence — always current, never assembled after the fact

Relyance AI continuously maps live AI inventory and data flows to EU AI Act, NIST AI RMF, ISO 42001, GDPR, and HIPAA — generating audit-ready evidence in seconds. Policy validation time cut by up to 60%.