



What CISOs Must Ask AI Security Vendors

Get clarity before you commit. A vendor-neutral framework for evaluating AI security platforms across architecture, visibility, risk, innovation, and outcomes.

Why AI Security Needs Its Own Approach

AI is now embedded in everyday enterprise operations — from GenAI tools and API-based LLMs to open-source models and AI agents. The opportunity is significant, but so are the demands on security, compliance, and governance.


Boards expect AI-driven gains, putting CISOs under pressure to enable adoption while managing risks that fall outside traditional security tooling. This gap defines AI security as its own discipline, spanning four key domains:

AI User Protection

AI App Protection

AI Runtime Protection


AI Agent Governance



Mapping the AI Landscape

The four AI security domains span two broad categories of enterprise AI adoption: **AI You Use** and **AI You Build**. Understanding where your organization sits across this landscape is the starting point for any vendor evaluation.

Mapping the AI Landscape



AI You Use

- Internal use of third-party AI apps
- Integrating AI apps with internal datasets
- Connecting agents to enterprise systems

AI You Build

- Develop and deploy custom AI apps
- Move models into production
- Build and govern AI agents



How to Use This Guide

This guide provides vendor-neutral questions to help CISOs make informed comparisons when evaluating AI security platforms. It is designed for early-stage consideration, including internal preparation across GRC, AppSec, legal, and data teams, as well as exploratory vendor conversations and platform demos.

01

Internal Preparation

Align GRC, AppSec, legal, and data teams on AI maturity and security priorities.

03

Platform Demos

Validate vendor claims against your specific AI security requirements and business goals.

02

Vendor Research

Use these questions during exploratory conversations to assess platform fit.

04

Evaluation & Selection

Incorporate answers into your process to assess vendor suitability and alignment.

Section 1: Architecture & Design

ARCHITECTURE

Understanding how a vendor's platform is built is foundational to any evaluation. Push vendors to go beyond marketing language and describe their actual architectural approach.

Purpose-Built for AI?

Is the platform designed specifically for full-spectrum AI security, or is AI security bolted onto an existing product?

Converged Control Plane

Does the architecture enforce consistent policy across users, data, applications, and networks from a single plane?

Deployment Flexibility

What deployment models are supported — SaaS, API, on-prem, browser extension?

Data Privacy & Isolation

How is data privacy and tenant isolation ensured across environments and workloads?

Integration Interfaces

What APIs, connectors, and data pipelines are available to integrate with existing enterprise systems?

Section 1: Capabilities & Enforcement

ARCHITECTURE

Beyond architecture, CISOs need to understand the breadth of enforcement capabilities and how the platform handles real-world AI threats at scale.



Use Case Coverage

Does the platform address shadow AI discovery, DLP, prompt inspection, and runtime enforcement?



AI Types Protected

Are GenAI tools, LLMs, chatbots, agents, models, dev environments, and MLOps all in scope?



Enforcement Actions

What actions are available — block, redact, warn, allow, or context-driven prompts?



Global Policy & Scale

How are policies applied globally in real time, and how is latency managed at enterprise scale?



Also ask: How does the platform detect and mitigate jailbreaks, prompt injections, data exfiltration, and model manipulation?

Section 2: Visibility & Analytics

VISIBILITY

Effective AI governance requires deep visibility into how AI is being used across the organization. Vendors should demonstrate granular attribution, flexible reporting, and integration with existing security tooling.

AI Activity Attribution

Can activity be attributed to specific users, devices, or data sources? Can insights be segmented by department, geography, or identity group?

Dashboards & Reporting


What analytics are available to track adoption, usage, and risk posture? Are executive or board-level summaries available for AI governance reporting?

SIEM/SOAR/DLP Integration

How does the solution integrate with existing security platforms and export telemetry to observability tools?

Agentic vs. User Actions







Can analytics distinguish between user-initiated and autonomous agentic actions such as self-chained prompts or API calls?



Section 3: Risk Management & Compliance

COMPLIANCE

AI security vendors must demonstrate alignment with established frameworks and the ability to adapt as regulations evolve. Probe for specifics — not just checkbox compliance.

FRAMEWORK ALIGNMENT  Structured adherence to key standards.	RISK ASSESSMENT PROCESS  Identify, score, prioritize, and track risks.	DATA SOVEREIGNTY  Ensure jurisdictional compliance for all data usage.
ZERO-TRUST CONTROLS  Least-privilege applied to AI systems and data.	CONTINUOUS COMPLIANCE  Automated validation across multiple frameworks.	RESPONSIBLE AI  Transparency, fairness, and explainability in detections.

Also ask how the vendor adapts to evolving regulatory mandates and validates data provenance, model integrity, and version transparency.

Section 4: Threat Detection & Defense

INNOVATION

AI systems introduce novel attack surfaces. Vendors must demonstrate active, adaptive defenses — not just perimeter controls.

Adversarial Threat Detection

How does the platform detect and mitigate data poisoning, evasion attacks, or model inversion?

Malicious Model Blocking

What mechanisms identify and block compromised models, datasets, or AI components?

Agentic Runtime Protection

How is unauthorized tool use, prompt chaining, or data exfiltration by agents detected and stopped?

Threat Corpus Updates

How is the threat corpus maintained for emerging vulnerabilities like prompt injection and embedding misuse?

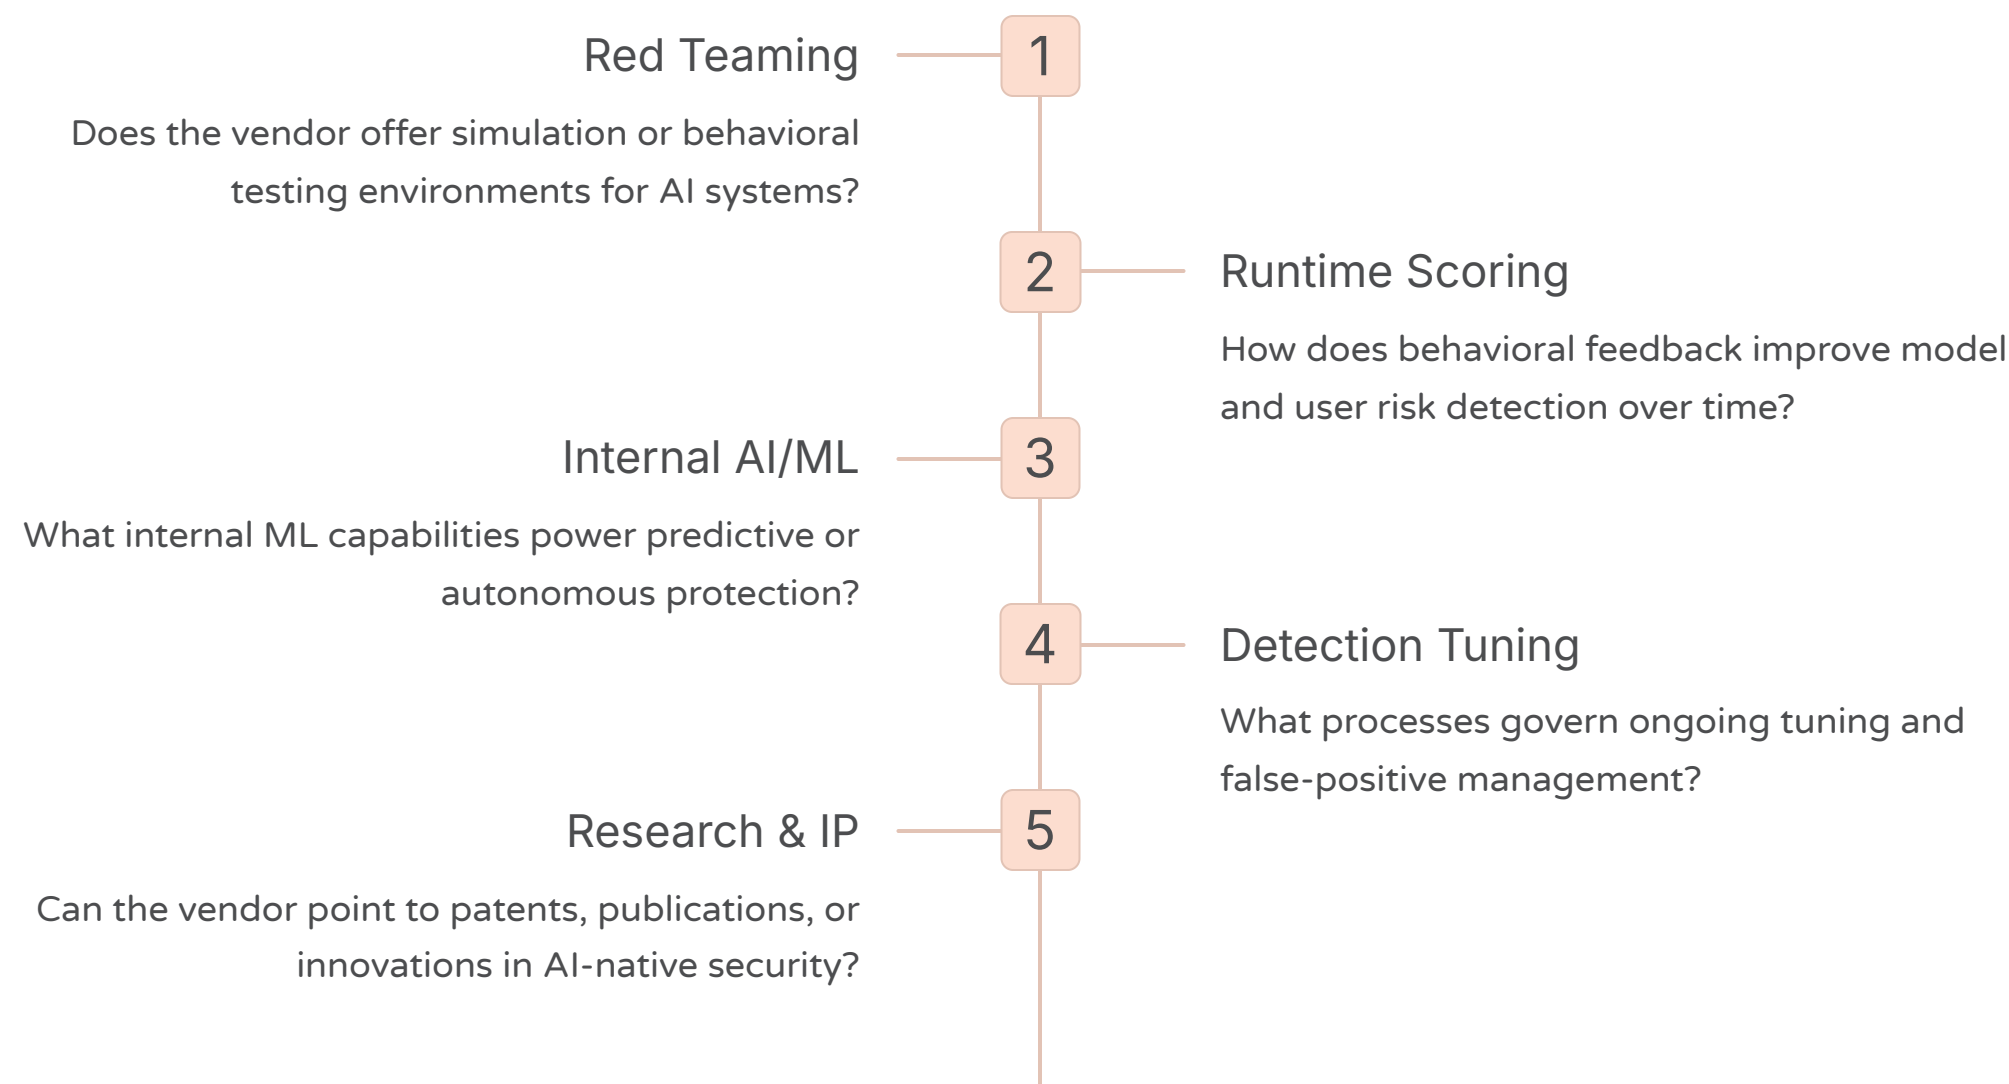
Self-Monitoring

Does the platform use AI to secure itself — including policy drift detection and AI-driven hardening?

Section 4: Testing & Continuous Improvement

INNOVATION

A vendor's commitment to continuous improvement reveals the long-term durability of their platform. Look for evidence of internal research, red teaming capabilities, and systematic tuning processes.



Section 5: Technical & Operational Requirements

OPERATIONS

Operational readiness is as important as capability. Vendors must demonstrate clear onboarding processes, enterprise-grade integrations, and robust SLAs.

Deployment & Integrations

- Onboarding and deployment methodology
- Identity system integrations (Okta, Azure AD, Google Workspace)
- API integrations and developer extensibility
- Secure sandboxing for prompt-injection or red-team testing

Scalability & Operations

- Scalability metrics: concurrency, throughput, enforcement latency
- SLA metrics for uptime, incident response, and policy propagation
- Update management, version control, and global config propagation
- Platform security hardening and validation processes

Section 6: Customer Success & Validation

VALIDATION

Vendor claims must be backed by real-world evidence. Ask for referenceable customers, independent analyst coverage, and clear metrics that demonstrate sustained value delivery.



Customer References

Request referenceable enterprise customers or anonymized deployments with measurable AI-risk reduction outcomes.



Analyst Validation

Ask for independent coverage from Gartner, Forrester, or similar, along with renewal and satisfaction metrics.



Roadmap & Adaptability

How does the vendor's roadmap anticipate new AI threats, frameworks, and compliance demands?



Ongoing Engagement

What is the customer success structure? How is continuous value delivered after go-live through outcome tracking and proactive tuning?

Section 7: Outcome-Based Value Demonstration

OUTCOMES

Vendors should be able to articulate and demonstrate measurable outcomes — not just features. Tie evaluation criteria to business-relevant KPIs and real-world scenarios.

1

Top 3 Measurable Outcomes

Ask vendors to identify their top three outcomes — e.g., reduced AI data leakage, audit readiness, faster policy enforcement.

2

Before/After KPIs

Request benchmarks tied to AI adoption and risk posture: MTTD, MTTR, alert fidelity, and analyst productivity.

3

Automated Improvement

How does the solution use AI-based policy tuning or drift detection to automate continuous improvement?

4

Scenario-Based Validation

Can the vendor demonstrate performance through sandbox demos, red-team scenarios, or simulation testing?



What CISOs Should Take Forward

AI will continue its ascent across the enterprise, and vendor claims around AI security will dominate marketing messaging. This framework helps CISOs narrow the field for AI security platform due diligence and selection.

The questions keep the focus on whether a vendor meets core architectural and capability requirements — a practical way to keep AI security conversations grounded as the market evolves.

Stay Grounded

Use these questions to cut through vendor noise and focus on architectural substance and real-world capability.

Align Internally First

Engage GRC, AppSec, legal, and data teams before vendor conversations to ensure aligned evaluation criteria.

Demand Evidence

Require referenceable outcomes, independent validation, and scenario-based demonstrations — not just feature lists.