Reach

Buyer's Guide: Evaluating AI for Security Architects

AI has become the most overused - and perhaps misunderstood - term in cybersecurity marketing. Vendors tout "AI-powered" everything, yet most implementations fail to deliver measurable improvements. As Gartner® notes, "The primary risk for AI cybersecurity projects is failing to deliver tangible value, as the emerging techniques may fall short of inflated expectations."

This guide helps security architects cut through the noise, focusing on the real questions to ask when evaluating AI for cybersecurity and how to identify solutions that operationalize, rather than simply advertise, AI.

1. Understanding the Spectrum of AI

Buyer Takeaway **Not all AI is created equal.**

AI in cybersecurity spans a spectrum, from simple rule-based automation and predictive analytics to generative models and emerging autonomous agents. Each requires different safeguards, oversight, and maturity to deploy safely.

> Gartner® notes, "AI is not a monolithic capability. It spans a broad spectrum of technologies, ranging from deterministic rule-based systems to probabilistic generative models and increasingly autonomous agentic architectures. Each class of AI introduces distinct behaviors, failure modes, and security implications. Treating them as interchangeable leads to misaligned expectations, overconfidence in tool efficacy, and inadequate and overly restrictive risk mitigation strategies.[1]"

What To Look For

- **Technical transparency:** Can the vendor clearly explain what kind of AI they use or how it works? What type of model is used? What is it trained or grounded on? Ask pointed questions. If you get vague or hand-wavey answers, buyer beware.
- **AI with Cybersecurity Expertise:** A generic LLM bolted onto a security product will have minimal impact. Look for evidence of cybersecurity domain-specific models (DS-LLMs, DSLMs, or non-language models) that understand security context (e.g., firewalls, SASE, endpoint security, email gateways) rather than just generic LLMs. Does the model understand security-specific data and context, such as firewall rules, directory configurations, or endpoint alerts?
- **Operational relevance:** Are AI-generated insights connected directly to actionable security controls and threat intelligence, or are they simply static dashboards?

Reach Perspective

Reach's MastermindAI™ is built on domain-specific language models trained with real security data and live control configurations. This enables a contextual understanding of risk across tools, people, and assets – bridging the gap between AI insights and operational security.

2. Evaluating AI Value: From Data to Decisions

Buyer Takeaway **Focus on measurable outcomes, not promises.**

Many cybersecurity AI projects fail not because the technology is poor, but because success was never clearly defined.

> As Gartner® states, "AI cybersecurity value is not guaranteed – without clear goals and measurable outcomes, initiatives often fail to deliver meaningful value.[1]"

What To Look For

- **Outcome alignment:** Evaluate how AI outputs map to your organization's specific objectives, such as reduced mean time to remediation (MTTR), improved control utilization, or decreased exposure surface.
- **Workflow integration:** Does the AI integrate with your existing systems – like ServiceNow, Jira, or configuration management tools – to make insights actionable?
- **Feedback loops:** Look for systems that learn from outcomes and improve accuracy over time.

Reach Perspective

Reach operationalizes AI by transforming insight into action. It prioritizes based on attack behavior and configuration context, and then creates or automates remediation workflows aligned to frameworks such as MITRE ATT&CK and NIST CSF – ensuring AI value is measurable and repeatable.

3. Governance and Trust: Building on What You Have

Buyer Takeaway **You don't need to wait for new AI governance frameworks to get started.**

AI adoption in cybersecurity often stalls over uncertainty around governance.

> Gartner® states : "Securing AI cannot wait for perfect governance structures to be in place. It starts by building on existing enterprise policies, processes, and structures.[1]"

What To Look For

- **Extension of existing controls:** The right AI platform strengthens your current identity, access, and data protection policies – without requiring new frameworks.
- **Operational transparency:** You should know where AI operates, what data it accesses, and what actions it takes.
- **Shadow AI detection:** Ensure visibility into unauthorized AI models or agents running within your environment.

Reach Perspective

Reach integrates directly with existing IT and security systems, including ServiceNow, Jira, EDR, and SASE platforms, to maintain governance continuity and ensure transparent AI operations.

4. Continuous Validation: Keeping AI and Security in Sync

Buyer Takeaway **Security posture, and AI performance, drift over time.**

AI's initial value can degrade quickly if it is not continuously validated against real-world changes.

> According to Gartner®, "AI risks continuously evolve as architectures change and new exploitation techniques emerge.[1]"

What To Look For

- **Posture validation:** The solution should continuously monitor configuration drift and confirm that defenses remain aligned to intent.
- **Adaptive learning:** AI should refine its models based on live outcomes and evolving threat patterns.
- **Operational resilience:** Continuous validation ensures that both human and machine responses remain calibrated to current risk.

Reach Perspective

Reach's ConfigIQ Drift continuously detects misconfigurations and validates that defensive controls remain aligned to evolving threats – helping teams sustain an optimized posture over time.

5. Conversational and Explainable AI

Buyer Takeaway **AI must be transparent and accessible to all stakeholders.**

Security teams increasingly rely on conversational interfaces for complex analysis, but trust is only possible when AI can explain why it reached a conclusion.

> As Gartner® notes, "Without baseline AI literacy, organizations are vulnerable to technical blind spots and social engineering tactics that exploit misconceptions about AI.[1]"

What To Look For

- **Natural language interfaces:** Look for platforms that enable security teams to ask questions in plain English and receive contextual, actionable responses.
- **Explainable logic:** Every recommendation should trace back to data, rules, or models that can be reviewed.
- **AI literacy tools:** Built-in education features can help teams understand AI's reasoning and boundaries, reducing misuse or overreliance.

Reach Perspective

Reacher™ , Reach's conversational AI assistant, delivers plain-language explanations of exposure, risk, and configuration drift. By making AI explainable, it empowers every team member – from analyst to architect – to engage confidently with security posture decisions. Reacher can also create custom rules across your integrated security stack. Simply tell Reacher what you want the rule to do, and it will create the rule in seconds.

The Path to Practical AI

By 2027, Gartner predicts that "90% of successful AI implementations in cybersecurity will be tactical – task automation and process augmentation – rather than role replacing." The message is clear: AI should empower people, not replace them. The most successful organizations will deploy AI that operationalizes security controls, drives measurable outcomes, and strengthens human decision-making – not tools that chase hype or promise autonomy they can't deliver.

1. Gartner®, Demystifying Common Misconceptions About AI and Cybersecurity, Craig Porter, Nader Henein, 23 September 2025 Gartner® is a registered trademark of Gartner®, Inc. and/or its affiliates and is used herein with permission. All rights reserved.

Buyer's Checklist

Before choosing an AI-driven cybersecurity solution, ask vendors to clearly demonstrate:

- [ ] Model type: What kind of AI is being used (LLM, DSLM, agentic, predictive, or hybrid)?
- [ ] Transparency: How was the model trained or grounded, and on what data sources?
- [ ] Integration: Can AI outputs directly connect to existing workflows and tools?
- [ ] Metrics: How is success measured – exposure reduction, MTTR improvement, drift detection?
- [ ] Scalability: Can you start small with a pilot, validate outcomes, and scale confidently?