


Latio.

AI Security Market Report



2026

Table of Contents

Executive Summary	3
Survey Results	4
The Evolution of AI Security	7
Securing Agents	10
Emerging Endpoint Vendors	19
Innovative Capabilities	22
Buyer's Guide	27
Conclusion	32
Vendor Spotlights	33

Executive Summary

There are roughly 400 AI security start-ups on the market today, all competing for similar users, buyers, and budgets, and that number continues to grow. In the last 12 months, we've seen the AI security industry move from protecting browser-based tools such as ChatGPT and Microsoft Copilot to securing agents that can take action on behalf of users and organizations. Agents introduce a more complex security challenge because they can run across browsers, first-and third-party applications, endpoints, and cloud infrastructure, expanding both the security challenges and teams responsible for addressing them.

As a result, vendors that historically focused on different parts of the security stack are expanding into overlapping use cases, while teams approach AI security based on different priorities, existing investments, and AI strategies. The current priority is securing AI on end-user endpoints, but AI spans the entire enterprise architecture, making the market increasingly difficult for buyers to navigate. Understanding the differences between vendors and approaches is more important than ever.

The 2026 AI Security Market Report examines how this market is evolving, what organizations are prioritizing, and where vendors are investing, with our analysis focused on five key findings:

1. AI security is moving to the endpoint, meaning tools like Claude and Codex.

The primary focus for vendors is shifting from browser-plugin visibility and traditional DLP use cases towards governing AI agents running on workforce devices, such as Claude Cowork running on an end user's laptop.

2. Existing security platforms are expanding to secure first-party agents. First Party or "homegrown" agents can often be secured by existing categories that have expanded their capabilities to cover your cloud hosted applications.


3. Vendor selection depends on current and previous priorities. Selecting the right vendor to secure AI on endpoint devices depends on an organization's priorities, existing security stack, and overall security strategy.

4. Endpoint permissioning and proactive protection are key differentiators. Many of the larger AI Security vendors are investing in runtime and intent driven analysis, but the most meaningful developments are controlling the permissions and tools available to an agent.

5. Security isn't the only value driver. Emerging vendors aren't only focusing on making AI more secure to adopt across an organization, they're also building features ranging from cost optimization to employee AI marketplaces to make adoption simpler.

This report creates a clear framework for understanding where the AI security market is today, where it is heading, and how teams can make buying decisions with clarity and confidence.

Survey Results



Our surveys reflect the perspectives of practitioners and security leaders to accurately assess market demand, evolving priorities, critical use cases, and emerging trends - providing a real and comprehensive view of how the market is today and where it's heading.


Key Findings:

- Buyer demand for AI security tools is materializing, shifting from **8% to 37%** within one year.
- Security priorities have shifted to focus on endpoint agents such as Claude Code and Codex.

An Increase in AI Security Tool Demand

Has your team allocated budget for an AI security tool?

- **39%** No
- **37%** Yes
- **24%** Undecided




Security teams are moving away from treating AI Security as a secondary priority, instead making it an active budget item. Investments in this space aren't siloed to specific segments, it's expanding across organizations of all sizes, with ownership distributed across Product Security, Security Operations, and IT teams. Dedicated AI Security budgets are a strong signal of adoption and market maturity, and help explain the rush to introduce AI-focused security features.

Endpoints Are the Focus

What is your primary AI security concern?

- **45% Endpoint Agents**
(e.g. Claude, Codex)
- **23% First Party Agents**
(e.g. LangChain, Crew)
- **22% Browser Applications**
(e.g. ChatGPT, M365)
- **10% Hosted Agents**
(e.g. Agentcore, Foundry)




Our 2025 research highlighted two primary areas of concern: browser based data loss prevention (DLP) and securing AI infrastructure hosted in the cloud. Today, the shift is toward AI agents running directly on workforce endpoints, most prominently Claude and Codex. This shift changes the primary use cases for security teams. Endpoint agents have local permissions, can generate code, interact with developer tools, and operate with permissions that extend beyond the boundaries of a traditional browser session.

The Evolution of AI Security



The Evolution Timeline



The Two Eras of AI Security

Despite being a newer market, AI Security has already entered a distinct second era. The first era was defined by a wider range of use cases that consolidated into proxy and browser based controls for AI visibility.

Three outcomes supported by the first era:

- **Visibility and control over unapproved AI usage.** Identifying and preventing employees from using unauthorized AI applications, commonly referred to as "Shadow AI."
- **Sensitive data protection.** Detecting, redacting, or preventing sensitive information from being submitted to unapproved LLMs.
- **Runtime detection and prevention.** Identifying and blocking threats such as prompt injection, or manipulating an LLM to retrieve its system prompt.

Because the primary control points were the network and browser, many of the leading providers were acquired by network security vendors and integrated into existing firewall, secure web gateway, and browser capabilities. SentinelOne and CrowdStrike were notable exceptions, entering the category through acquisitions that brought AI runtime threat detection into their broader endpoint security platforms. These use cases and acquisitions laid the groundwork for the next phase of the market: securing AI anywhere it runs, rather than simply controlling how employees access it.

In 2026, AI Security entered a second era: one focused on securing Claude Code and Codex rather than ChatGPT. This shift has created a new market around developer endpoint protection, with startups building on the foundation created by Koi, acquired by Palo Alto Networks, while established AI Security providers race to deploy endpoint agents of their own. Last year endpoint agents were not priorities for vendors to build, but this year a platform is incomplete without one.

Three outcomes enabled by the second era:

- **Runtime monitoring and intent-driven analysis.** Identifying suspicious agent activity and evaluating actions based on an agent's intended behavior.
- **Governance capabilities and supply chain malware prevention.** Controlling and monitoring Skills, MCP servers, and other tools that extend what an agent can access and execute.
- **User permissions and guardrail enforcement.** Defining which tools, resources, and actions an agent is permitted to use, and enforcing those controls as the agent operates.

This shift towards endpoint agents continues to accelerate as frontier model providers build new ways to promote security and controls. For example, Anthropic recently introduced inference hooks and the compliance API, giving security teams more ways to implement runtime detection. Integration steps like these are moving providers closer to becoming an inline control layer for AI agents rather than simply monitoring how employees use them.

In short, the first era of AI Security was around controlling unapproved activity in the browser, the second era is about controlling unapproved activity on endpoints.

Securing Agents



Securing First Party Agents

Posture and Testing

Most application security platforms provide some level of posture and testing. The vendors below are dedicated AI security tools that offer in depth testing.



Application Layer Runtime Protection

Most AI security providers integrate their runtime detection engine with first party agents. The vendors below offer AI detection alongside broader application layer detection.



Cloud Posture for AI Services

Most cloud security providers support posture for AI services. The vendors below are dedicated AI security tools that offer cloud posture coverage.



AI Data Center Configuration



Most of the noise in the AI Security market stems from existing security capabilities extending into securing cloud hosted agentic infrastructure. Despite the newer terminology, much of “agentic architecture” looks familiar: containers with applications making API calls. That means established categories like Cloud Native Application Protection Platforms (CNAPP), Cloud Application Detection Response (CADR), and Application Security Posture Management (ASPM) all also offer “AI Security.”

Because architectures have not majorly shifted in light of AI, a few existing tool categories are especially applicable to AI Security:

- **CNAPP:** Cloud Native Application Protection Platforms provide cloud misconfiguration checks for AI systems, including but not limited to Bedrock, Jupyter Notebooks, and Sagemaker
- **CADR:** Cloud Application Detection Response, or separate ADR tools, extend application layer visibility into AI systems, giving teams deep control of the runtime behavior of their AI applications.
- **Cloud Identity and NHI:** Most “agentic identity” challenges are really agents using existing access or tokens for their own access, taking advantage of overpermissioned identities in the process.
- **ASPM:** Application security platforms, or emerging ones built for AI generated code protection, often extend their runtime testing and supply chain capabilities into AI applications, building “AI-BOMs” and conducting “AI red-teaming”

Within these capability sets, teams can opt to extend their AI security platform into their first party applications, or rely on existing toolsets to provide enough visibility and security for their use case.

Below is a table covering the areas where dedicated AI tools can be most competitive with their existing tool counterparts

Capability	Existing Tool	Specialist features
Cloud Security	CNAPP	Secure configurations for cloud hosted agents, especially coverage for emerging services like Agentcore and Foundry
Application Security	ASPM	Advanced AI red-teaming, support for local models, and data poisoning use cases
Identity Security	IGA	Most specialist tools are focused on permissions structures for MCPs

Securing Third Party Agents

Securing Workforce AI Usage

The rapid adoption of Claude Code and Cowork has led to AI's abrupt pivot from the browser and into the employee endpoints - where credentials live and risks start. When AI was confined to isolated chat experiences, the primary risks were often limited to a set of issues around unintended data leakage or the spread of misinformation. Once agents have access to compute and access keys, they're effectively given unrestricted access to internal systems and processes.

The adoption of endpoint agents has carried with it a massive increase in organizational risk. Agents can rapidly combine overpermissioned access with vulnerabilities in order to execute unintended commands, either injected by threat actors, or more often unintentionally executed by an employee themselves. The challenge for security teams isn't only about governing what employees can access, but confidently understanding and controlling what AI agents are enabled to do.

The Five Primary Use Cases


The use cases for AI as a broader technology are quickly changing, especially as companies expand their use of third-party agents. Below are the five primary security and governance priorities security teams have for securing their environments.

- **Preventing supply chain attacks, especially on developer endpoints.** MCP Servers and Skills are an evolution and extension of broader software supply chain concerns. Many skills come bundled with scripts and packages, each of which can run malicious code on machines; most MCP Servers are distributed as NPM packages. This makes securing these packages little different from broader open source malware concerns.
- **Preventing unintended company damage.** Many agents have access to unknown levels of permissions and data access by utilizing an end user's credentials. This can lead them to taking sensitive action without a human in the loop.
- **Preventing indirect prompt injection.** Frontier model providers have reached strong levels of direct prompt injection prevention; however, advanced threat actors may leverage downstream text to manipulate an agent's path. For example, hiding a task inside a SOC alert to divert an AI SOC tool.

- **Governing AI blast radius.** Currently teams lack necessary governance capabilities to understand what local agents have access to, such as files, SaaS applications, or authorized commands to run.
- **Governing effective workforce AI adoption.** Security teams need ways for employees to safely adopt approved AI tooling, while giving enhanced visibility and cost management to executive teams.

This list of use cases is a framework for understanding the vendors covered in this section. Each vendor approaches where and when they apply controls differently; for some the emphasis is runtime detection and prevention, for others it's preventative controls around files, permissions, and agent configuration.

Mapping by Starting Technological Approach



Note: This graph maps where vendors started their integration approach. **Today, most AI security providers extend into each of these areas in order to provide broad coverage.**

There are three primary integration points where vendors began their AI security platform journey: network proxies, endpoint agents, and IDE extensions. Most AI security providers extend into each of these areas in order to provide broad coverage, but where they started speaks to their maturity. The initial integration points often define the capabilities a provider prioritizes, the telemetry they can collect, and the controls they’re capable of providing.

The market map above organizes providers by the benefits that come from the approach taken. It’s important to note, new integration points are continuing to emerge. Frontier providers are introducing capabilities such as inference hooks, for example, to equip Security Operations platforms with greater visibility into agent activity and creating new ways to provide runtime controls.

Because these agents operate on the endpoint, their activity can be logged and governed through different layers of a security stack. Each integration point offers a different combination of visibility, control, context, and coverage. Understanding these trade-offs is important when evaluating how to best invest in an AI security approach.

Understanding Technology Trade offs

Integration Point	Pros	Cons
Network Proxies	Flexible runtime security engines that can integrate with most types of agents	Lack visibility into permissions, thought processes, and local system access
Endpoint Agents	Deep runtime visibility and governance capabilities	Lack the flexibility to handle SaaS or first party AI applications
Hooks and Gateways	Visibility into sessions, and the ability to alter an agent’s context window	Not all agentic platforms support hooks, inability to alter device settings

There is no single best integration point for AI security. The right approach combines all three in order to deliver comprehensive security controls. When deciding their provider and deployment, teams should understand what each integration layer enables, identify where their highest-priority risk is, and make the decision that provides the coverage and control their organization actually needs.

Agent vs Agentless: Both Are Important


Numerous providers are revisiting Cloud Security's "agent vs. agentless" debate for AI endpoint protection. As with the cloud security market, the strongest approaches will support both deployment methods, enabling teams with the ability to mature into an agent based approach.

Agentless approaches commonly involve monitoring network traffic, AI gateway integrations, or use mobile device management (MDM; e.g. Jamf, Intune) or endpoint detection response (EDR; e.g. CrowdStrike, SentinelOne) tooling to deploy a script in order to capture relevant data. Though these capabilities are normally visibility focused, some providers extend into runtime protection by deploying agent hooks, or setting up a permanent proxy.

Agents based deployments excel at providing deep visibility into what the configuration and access of local agents actually look like, giving teams a clearer understanding of what an agent can access and how it operates. Especially as governance capabilities from frontier model providers are developing, endpoint sensors currently fill a much needed gap.

Leading providers will support both an agentless and agent based deployment method, where agent based monitoring is the long term goal, but agentless provides immediate visibility. While understanding deployment options is important, it does not paint the full picture of what outcomes providers are building to support.

Vendors Mapped by Outcomes



Note: This map captures AI security vendors based on the focus area and outcomes they support. Most vendors in this map cover more than the areas they are placed in.

Many AI security providers offer similar core features: some level of visibility and runtime protection for MCPs, Skills, and agentic activity. While these capabilities may overlap, focus areas speak to the provider's depth in a particular area of AI security. A vendor's focus area can lead to numerous points of differentiation - whether it's helping developers with secure code generation, deep runtime controls around agent behaviors, or monitoring both first and third party agents together.

For buyers, AI developments are happening so quickly that finding the right partner has more to do with their focus aligning to your own than the particular set of features they presently support. AI tooling is changing so quickly that it's important to choose a vendor that is aligned to your priorities and developing quickly to meet the market. For example, Anthropic released inference hooks one month ago, and they're quickly becoming a primary runtime enforcement point - most vendors in this report already support them.

AI Security Capabilities and Focus Areas


Capability	Primary Focus Areas
Endpoint Agent Specialists	Deep runtime monitoring for agents, especially mapping permissions and additional runtime context
Broader AI Security	Runtime monitoring across different types of agents, including endpoint, but with an emphasis on securing AI across the entire environment - such as SaaS based agents and first party ones.
Platform Providers	Tools with AI specialties that are consumed as extensions of larger vendor relationships. These approaches usually map to the “Broader AI Security Providers” categories, as they are often powered by acquisitions of these providers.
DLP and Insider Threat	Preventing unintentional data leakage to AI systems, or looking for intent driven issues such as insider threat.
Developer Governance	Protecting developer endpoints alongside providing tools for secure code generation.

The overall state of the market for securing third party agents is a rush from larger security platforms to build endpoint agents, while endpoint specialists broaden their coverage into the browser and SaaS. Last year, no one expected the endpoint to play such a pivotal role in AI security, so many vendors have needed time to build or extend their endpoint coverage outside of the browser.

Emerging Endpoint Vendors



The Emerging Endpoint Vendors




Over the last few months, many startups have launched focusing on “the agentic endpoint.” Most of these companies cover the same core endpoint AI security controls:

- Governing MCPs and Skills, with malicious scanning
- Runtime visibility into agentic activity, providing insight how agents are behaving and what they’re doing

The above spectrum illustrates how emerging vendors differ in their approach to securing the agentic endpoint. While many providers offer similar capabilities, the depth of governance and runtime protection tend to differ based on the vendor’s focus areas. The platforms on the left side of the spectrum have a feature set and focus that extends beyond AI security, with capabilities built to control endpoint behavior, manage approved applications, and enforce application approval workflows. On the right side are the platforms designed more specifically around agent behavior, with deeper monitoring of agent permissions, tool usage, and runtime activity.

Understanding the Endpoint Market



The endpoint security market has been stagnant for many years, with providers like Jamf and Microsoft Intune providing device management and configuration (MDM), and others like SentinelOne and CrowdStrike providing runtime attack prevention (EDR). However, this market stagnation led to several blind spots for security teams as technologies have developed: supply chain malware, application control, intent based DLP, and more. Below is a table of the major endpoint areas and the gaps emerging vendors are solving.

The Opportunity	The Existing Market Gap
OSS Supply Chain Security	Most tools do not provide adequate prevention and detection of supply chain malware such as NPM packages, IDE extensions, and browser extensions.
Proactive Endpoint Management	As part of broader supply chain security, traditional MDMs struggle to enforce application layer controls, or with creating approval processes for developer tools
Endpoint DLP and Insider Threat	Endpoint DLP has always been a struggle for organizations due to high numbers of false positives, AI tracking allows for better intent tracking across user behaviors
EDR	Newcomers expanding into the EDR space allow for making intent tracking and baselining a larger part of the detection mechanism

Most of the opportunities in this emerging market were uncovered by Koi Security's focuses across these unmonitored areas of endpoint security - browser plugins, IDE plugins, and supply chain malware. Many of the providers are seeking to reinvent what's possible in preventing and detecting misconfigurations and malware on endpoints.

Innovative Capabilities



Highlighting Market Innovations

Security platforms are building solutions to detect malicious activity and enable enterprise AI adoption. This section highlights and explains notable emerging capabilities, and is organized around areas where the most meaningful innovation is happening, from employee-built applications and agent infrastructure to MCP and skills marketplaces, agent sandboxes, granular permissions, on-premise deployments, and intent-based detection.

Employee App and Agent Builders



Employee applications and agent builders enable almost anyone in an organization to build, deploy, and operate applications and agents, without pulling engineers away from their intended workflows. From Replit to Lovable, AI has made it possible for a much broader set of employees to become builders by deploying internal agents and applications within their environments. This newly found independence has introduced additional security blindspots, as developer workflows now happen outside of the established SDLC.

Vendors across the industry are working to solve this challenge, from AI application pentesting for hosted applications (e.g. Wiz, Aikido), visibility into internal permission structures and data flows (e.g. Pluto, AIR), to runtime protection (e.g. Sweet, Pillar). As employee-built applications and agents become a core layer of how teams work, securing these environments will become an increasingly important capability for modern security platforms.

MCP and Skills Marketplaces



Organizations have moved beyond experimenting with local agents to finding ways to collaborate with them across teams. Most vendors in this report support some level of control over MCPs and skills at runtime, such as blocking the request if it contains a specific tool call; however, security teams need more ways to enable safe adoption across teams.

MCPs and Skills Marketplaces create a governed, self-service flow for discovering, approving, and sharing AI capabilities across the organization. This enables enterprise teams to easily share newly developed skills between their teams, while vetting them for security issues and creating a paved road for adoption.

Skills Detonation



Skills can be thought of as mini applications. In their most basic form, they're markdown files containing additional context for a prompt. However, more advanced skills can include custom scripts and open source packages, giving agents the ability to easily execute repeatable functions. As Skills become more complex, static analysis is less effective because a Skill can fetch information from external sources or manipulate expected behavior.

Skill detonation means executing a skill in a controlled environment at runtime to look for malicious behavior. Rather than static analysis, this introduces the ability to monitor behaviors such as network egress and process execution. While nearly every vendor in this report provides runtime security enforcement for skills, this added layer gives teams more insights.

Granular Permissions



Identity plays a large role in protecting agents of all types, as understanding the blast radius of an agent in many ways on understanding what credentials it has access to. For in-depth coverage of these capabilities, providers from non-human identity (NHI) to just-in-time (JIT) categories allow agents to access what they need, as they need it. From the AI security tool perspective, identities are mapped from local endpoints to their cloud counterparts, allowing teams to more effectively govern what agents have access to - from local file systems to their cloud IAM roles.

AI Data Center Configuration



As more organizations choose to self-host their AI Infrastructure or use neoclouds, securing the underlying infrastructure is becoming a greater challenge. Securing on-premise configurations covers use cases such as BMC access, GPU settings, and firmware management. This introduces a security solution that's not about the model itself, but securing underlying configurations and the controls protecting the hardware. This emerging market has vendors ensuring teams understand and protect their on-premise GPU footprint easily and at scale.

Intent Based Detections




Agentic security is challenging due to the indeterminate nature of agents - teams want them to accomplish their goal, but only within certain parameters, such as not hacking another company in order to obtain information. Intent-based detection monitors the behavior of an agent against its intended goal and organizational boundaries.

Most AI security tools are deploying their own versions of intent based runtime protection. These tools work by analyzing the system prompt, mapping it to various possible intents and guardrails, and then monitoring the session for conformity to that overall goal. In many ways, these tools are an evolution of early guardrails capabilities which focused on preventing specific questions or answers. From our hands-on time with tools, be sure to test a vendor's functionality here, as many runtime detection engines are not as powerful as promised.

Buyer's Guide

An AI security program can be designed around existing platforms or through a deliberately segmented approach. Our Buyer's Guide helps teams build a comprehensive program by focusing on coverage across first party and third party agents. Depending on priorities and goals, each section is mapped to the tools that an organization may already have in place.

A Guide to Securing First Party Agents



Note: Cloud, Product, and Application Security teams are typically responsible for securing first party agents.

Teams building programs for securing first party agents, such as homegrown applications instrumented with Langchain or Crew, should make their buying decisions based on their current coverage and risk exposure. Simple chat based applications, without access to real customer or internal data, pose very little risk that cannot be covered by more general providers.

As first party applications increase in complexity, the need to run more sophisticated AI red-teaming, as well as more in-depth runtime defense becomes crucial. Additionally, most existing providers in well established categories are building towards supporting these use-cases.

On the posture side, the two key capabilities to look for are:

- Visibility into agentic architecture, such as systems prompts, tools, skills, local models, and libraries like pytorch that are frequently used in AI development. This is sometimes referred to as an “AI-BOM.”
- Runtime “AI red-teaming,” most often an extension of Dynamic Application Security Testing (DAST). This capability is used to test your application for unintended behavior before or after deployment.

Because these posture tools are making their way into most application security providers, be sure to check with yours to see if it’s on their roadmap and the product quality. If these tools are not being satisfactorily provided, consider prioritizing an AI Security Platform that already offers them in a mature form.

For protecting first party agents at runtime, guardrails are the only core capability. Guardrails are quickly being offered by most runtime providers: from major platforms such as Datadog and Wiz, to security startups across the spectrum. Our recommendation would be to take advantage of AI Security budgets to extend into Application Detection Response tools more broadly (listed in the image), as they cover more than AI Guardrails alone, and can often instrument without an SDK.


A Guide to Securing Third Party Agents

Are you prioritizing protecting agents on endpoints, or concerned about broader usage?

In-depth Governance of Endpoint Agents



BACKSLASH



or

Protection for Agents across SaaS, Endpoints, and First Party



Are you looking to fill gaps in your existing endpoint controls?

Proactive Endpoint Management
(MDM and Application Control)

Prevention at Install



Prevention at Execution



Endpoint DLP
and Insider Threat



EDR



OSS Supply Chain Security
(Browser, Packages, IDE)



Are you concerned about AI security specifically for developer workflows?



Note: IT and enterprise security teams are typically responsible for securing third party agents.

Buying decisions focused on third party agents map closely to an organization's priorities in securing agent usage. There are numerous approaches to securing third party agents, and every vendor is targeting different scopes.

- **The first major decision teams face is partnering with an endpoint specialist, or a broader AI security provider.** We'd suggest partnering with an endpoint specialist if AI security is being prioritized alongside needing greater control of employee endpoints. Alternatively, if you're treating AI Security as a broader initiative, being governed by a cross-disciplinary team, then the larger platforms offer more benefits that both teams can utilize.
- **Another consideration is if you're looking to fill visibility gaps across endpoints outside of AI security.** Especially for developer endpoints, traditional EDR and MDM tools lack support for the threats that most target these machines, from open source malware to IDE plugins. Getting AI visibility alongside these other benefits is a great combination.

For most teams that we're working with, AI Security for developers is being secured by extension of their application security platforms, but the rest of the workforce is being secured by Enterprise IT Security teams. Many of these teams end up with an endpoint overlap: using one endpoint tool for AI security for workforces, and then a specialist tool for developers. Depending on your comfort level with existing enterprise controls provided by frontier model providers directly, you may opt towards broader coverage, or more in depth endpoint support.

Opting for an emerging endpoint security provider who has strong AI capabilities has the added bonus of augmenting your general device controls, and creating better approval workflows for often unmonitored sections of endpoint usage. This should be weighed against the stronger runtime benefits of dedicated AI Security providers.

Conclusion

The AI Security market is rapidly changing, as frontier model providers define the capabilities that teams are looking to protect. The AI Security program of 2026 looks much different from 2025, with teams shifting from securing AI in the browser toward securing the endpoints where agents run. As these use cases continue to change, AI security programs will need to evolve to meet them rather than converge on a single static program.

Looking ahead to strategies that will be part of security programs in 2027, we expect to see the focus on endpoint control shifting to the cloud. This shift will come from hosted providers like Agentcore and Foundry taking up more of the market, and teams looking to centralize how employees deploy and interact with agents. The most advanced teams are already thinking about how they are going to centralize their agents into cloud platforms, allowing teams to build and deploy across their enterprise.

Our prediction: The “Hosted Endpoint” will be the 2027 AI security focus that the “Agentic Endpoint” is today.


Vendor Spotlights

All vendors were given the opportunity to spotlight their product by having the Latio team author a dedicated page explaining why they were awarded in this report.



AIR earned the Endpoint AI Security Leader badge by excelling at what buyers asked us about most: understanding what agents on the endpoint can actually access and do. The platform combines strong runtime protection and visibility for Skills, MCPs, and agent activity, alongside a deep understanding of what credentials an agent has access to. During an incident, AIR shows what data was targeted, when the attack was stopped, and exactly which permissions the agent had available, turning incident response from guesswork into actionable guidance.

End-to-end contextual runtime controls continuously analyze an agent's behavior for malicious skills, MCPs, and plugins. A key differentiator is AIR's filtering of every agent operation, such as loading a skill or running a webfetch, to protect before an incident occurs. Coupled with their deep understanding of permissions, this makes AIR able to spot the moment a session is hijacked,




AIR extends their endpoint visibility into cloud-hosted and SaaS agents, such as those running in Slack, giving teams one view of agent access from local file systems to cloud identities. AIR's capabilities combine agent blast radius with real time behavior to create an end-to-end security platform.

Air Security is Best For

teams looking to deeply govern endpoint agents, especially those concerned about their blast radius with local privileges.

[Backslash](#) was one of the first companies to focus on securing AI on the endpoint, building AI endpoint controls well before this year's rush of newcomers. That head start shows in the product's maturity: coverage spans the full agentic surface, from Claude Code and Cursor to the MCP servers and skills running on top of them, with robust capabilities for finding and governing malicious skills and MCPs at runtime.

Where Backslash stood out in our evaluation is in control depth. Beyond scanning for what's malicious, the platform provides in-depth control over endpoint agent settings and what's actually available to an agent in the first place. Backslash also provides real-time protection, alongside intent based runtime analysis.




Backslash also offers several features specifically for developer devices, including the ability to be deployed either with or without an agent. Backslash covers use cases that workforce-oriented tools tend to skip, like enforcing secure-by-default rules for AI code generation and sharing vetted skills across teams. For AppSec teams treating developer AI adoption as its own challenge rather than a subset of workforce governance, Backslash is one of the most complete options available.

Backslash Security is Best For

teams looking to govern AI adoption, especially for those most concerned with developer AI usage.

Bold Security has taken an endpoint-first approach to data security, offering teams granular context and intent-based controls for protecting the data on end-user devices. Their on-device AI blocks and coaches users on risky data movements as they happen, and it's built to spot the exfiltration and evasion attempts that slip past tools watching only the network or the browser. The real-time coaching aspect is especially helpful, correcting a user in the moment rather than silently killing their session and generating a ticket. Because intent analysis happens on device, Bold takes action quickly, and even offline.

Bold also shows what user credentials are using which agentic tools, turning opaque agent connections into actionable findings that help you understand how access connects. Users and policies are contextualized within the overall organization structure, enabling sensible policy enforcement depending on roles & responsibilities.




Bold surfaces AI alerts as an extension of the data protection platform, looking for suspicious actions whether a human or an agent takes them. Bold's combination of agent, human, and organization context enables flexibility in data loss detections, whether accidental or intentional. As agents and users blur into a single activity stream on the endpoint, tools built to watch both at once will age best.

Bold is Best For

teams looking to understand and govern sensitive data on their endpoints in real time, whether a human or an agent is moving it.

Ent is building an entirely new style of endpoint protection, one that granularly traces user activity across entire endpoint sessions. This enables detection of active attacks, data loss, and insider threats, creating a more flexible and usable layer above EDR. Where EDR focuses on kernel-level protection and lacks application-layer tracking of user activity, Ent fills that gap.

Tracing end-to-end user endpoint sessions is the innovation that powers the rest of the platform. The analysis runs on the endpoint rather than round-tripping to the cloud, enabling rapid detection while respecting data boundaries. Flexible collection and searching across session data helps deliver on the promises of activity baselining with cohort and peer analysis to spot deviant behavior. The ability to apply natural language intent rules across machine and user events means teams can create flexible detections and interventions molded to specific team's activities. Ent can warn a user, require justification, redact, or block an action, giving teams flexible response options.




For AI security specifically, Ent's value today is tracking suspicious or unapproved activity across applications and agents. This matters because agentic behavior can quickly come with unexpected ramifications, and session-level intent analysis catches what process-level telemetry misses. For teams that want AI coverage as part of a stronger endpoint protection layer, Ent is the newcomer to watch.

Ent is Best For

teams looking to extend their EDR by catching suspicious or unapproved activity, human or agent, through session-level intent analysis on the endpoint.

[Datadog's](#) expansion into AI security adds another layer to its runtime security capabilities, covering agents alongside clouds, general applications, and workloads. AI Guard, Datadog's agentic runtime security solution, protects against prompt injection, tool misuse, data exfiltration, and other threats. AI Guard is especially compelling due to its being integrated directly with Datadog's observability tools. This gives developers a significantly easier adoption path, and yields richer behavioral telemetry than many providers can collect, linking decisions directly to the surrounding LLM traces, APM data, and workload context.

That coverage is also expanding to coding agents with local device protection coming soon. That will give teams rich visibility and protection for MCPs, skills, plugins, and commands on developer machines, closing the gap between the hosted agents Datadog already sees and the coding agents driving this year's security priorities.




Datadog now combines AI, cloud, application, and workload protection in a single place, which matters most when it's time to operationalize these alerts as part of a runtime detection and response program. Analysts work an AI signal with the same workflow they use for everything else, instead of learning niche tooling for one alert type. As AI security is blending into runtime security more broadly, Datadog is already treating it that way.

Datadog is Best For

teams looking to secure their AI usage as part of their existing platform, without causing developer or implementation friction.

[Harmonic's](#) coverage of browser and SaaS tools stood out for the range of environments it covers, including major frontier model providers alongside embedded AI in tools like Canva and Grammarly that most vendors miss entirely. Their coverage is backed by some of the most mature data classifiers and controls in the market, so visibility comes with in depth control of what corporate data leaves your environment across endpoints, SaaS, and MCPs.

Harmonic also helps teams understand their employee's AI usage patterns and priorities. The platform tracks use cases and broader workforce activity to show what your organization's primary AI use cases actually are, helping teams accurately assess the risks and projects underway across their organizations. For most security teams, "what are our employees actually doing with AI" remains an unanswered question, and Harmonic answers it with in depth visibility.




Harmonic's endpoint capabilities are also developing quickly; extending their runtime detection engine and data controls into the endpoint more broadly, bringing intent-based classification to tools like Claude and Codex. For teams whose AI security priorities start with data protection, Harmonic is a great fit.

Harmonic Security is Best For

teams looking to protect data across workforce AI usage across endpoint, browser, and SaaS tools.

As organizations move more AI workloads into production, the market is demanding runtime protection that can understand not just what an application is doing, but its intentions. [Lasso](#) is a standout solution for demonstrating some of the deepest intent-based runtime protection, with flexibility in their runtime engine categories and thought put into prioritization and classifiers rather than a one-size-fits-all guardrail. This culminates in a fast, flexible, and effective detection engine that can scale across platforms.

Agent discovery happens across first and third party applications alike, surfacing AI usage across enterprise environments. For first-party applications, Lasso pairs AI red teaming with intent-based runtime protection, and their visualizations of LangChain-style agentic workflows give teams an actual picture of what their homegrown agents are doing rather than a list of API calls.



The screenshot displays the Lasso Security interface for an agent named "Innovation Support Agent". The interface is divided into several sections:

- Agent Information:** Shows the API Key, Agent app, Session ID, and Trace ID.
- Agent Instructions:** A text box containing instructions for the agent, such as "A customer-support agent for an e-commerce retailer. May (a) answer product/policy questions via the RAG knowledge base, (b) look up order status by order ID belonging to the authenticated user, (c) create support tickets on behalf of the user. Must NOT access other customers data, modify accounts, issue refunds without explicit user request, recommend financial products, or discuss topics outside customer support."
- Intent:** A section showing the relationship between the agent and the model. It includes buttons for "User ↔ Model", "Developer ↔ Model", "Developer ↔ User", and "Developer ↔ External". The "Developer ↔ Model" button is highlighted with a red "Misaligned (!)" status.
- Trace content:** A section showing the JSON trace data, including fields like "tenantId", "sessionId", "eventIndex", "instanceId", "contentHash", "applicationTokenId", and "tool".
- Trace flow:** A section showing the flow of the agent's actions, including "Model Response product_handler", "Tool Input assistant", "Tool Output tool", and "Model Response product_handler".

Lasso has also expanded strongly onto the endpoint from the IT and enterprise security perspective, including the ability to red team local agents, something few vendors offer today. Lasso provides strong coverage of both first and third party agents from a single platform, with intent protection that goes deeper than most.


Lasso Security is Best For

teams looking for discovery, runtime protection and red teaming that spans both first and third party agents.

The rapid expansion of data centers and AI infrastructure are introducing new risks across every technical layer of the data center. Neoclouds, which commonly provide specialized GPU compute for training and deploying AI models, are considered especially risky, as industry experts warn that their scale is outpacing their risk mitigation, making them likely targets for agentic attacks. These types of non-traditional hyperscalers have gone mostly unmonitored by security teams, increasing the risk.

[Lava](#) is focused on securing neoclouds and datacenters (on-prem or hybrid) across network, compute, identity, virtualization, firmware, BMCs and storage layers.

Lava delivers the in-depth posture and misconfiguration capabilities that cloud security tools popularized back to data centers, which has never had an equivalent. That includes building toxic combinations and attack paths to data centers and neoclouds, giving teams real risk prioritization applied to GPU clusters and hardware that most scanners have never seen. By combining network and infrastructure information, teams are able to properly prioritize and patch their infrastructure.




As more companies rent neocloud capacity, build their own data centers or grow capacity, the amount of AI and data center infrastructure sitting outside any existing security tool's coverage keeps growing. Lava is a rare security product that's delivering meaningful innovation in an underexplored area of security.

Lava Security is Best For

teams looking to secure their neoclouds, neocloud capacity central, on-prem or hybrid datacenters, inference and the GPU infrastructure underneath their offerings.

[Noma](#) is one of the few platforms that covers first-party, third-party, and endpoint agents, combining strong runtime and posture protection to provide end-to-end enterprise governance for AI. By bringing posture and runtime capabilities together, teams can discover the agents, MCPs, and skills they didn't know they had and map permissions across providers to understand potential risks. They can then control what those agents actually do once they're running, whether on endpoints or in the cloud.

On the endpoint side, Noma's combination of posture and runtime protection helps teams identify toxic combinations and understand where they create risk. For example, teams can see when agents are allowed to execute Bash commands without proper sandboxing or perform destructive actions without confirmation. It can then enforce the protections around blocking the exploit of these issues, giving teams coverage while they understand their environment.




Noma enables flexible runtime enforcement across different agent types. Its observability capabilities cover the full behavioral chain of an agent session, from prompts and tool calls to data access and actions. Incident summaries then turn that activity into a clear narrative that analysts can understand without having to reconstruct the sequence themselves. For cloud-hosted agents, Noma's permissions tracking gives teams a clear picture of what an agent running in a cloud infrastructure can reach before anything goes wrong.

Noma is Best For

teams looking to cover every kind of agent they run with posture and runtime protection from a single platform.

Onyx has built an agentic control plane for securing agents across the enterprise, covering endpoints, browsers, SaaS applications, cloud platforms and code repos. Rather than stopping at raw telemetry, Onyx surfaces intent and common discussion themes from agent sessions, giving security teams visibility into how agents are actually being used across the organization, rather than guessing from inventories and access alone. Teams can also group users to correlate sessions, runtime alerts, devices, and users in a single place.

Onyx maintains a granular understanding of MCPs, tools, and skills, mapping an agent's capabilities and showing how it actually functions before something goes wrong. It pairs that visibility with in-depth MCP vulnerability detection, and gives teams the ability to block or upgrade to safer versions.




That visibility feeds directly into a major focus of the platform: flexible runtime enforcement. Policies are written in natural language, and Onyx monitors each step of an agent's reasoning, blocking malicious requests and steering unintended actions. Onyx also correlates AI detections with EDR detections, creating more holistic endpoint protection than either signal provides alone. The endpoint and broader AI security markets are converging fast; Onyx is one of the vendors making that convergence practical.

Onyx is Best For

teams looking to understand how agents are actually being used, while governing their access with flexible policies.

Pluto has built a comprehensive AI security platform atop a strong foundation of endpoint security. Their enforcement and enablement capabilities across the endpoint extend from AI to broader software supply chain security, and their support for app builder tools like Lovable and Claude Code were among the strongest we saw, securing a growing blind-spot for security teams.

What makes Pluto a leading solution is the breadth of the platform alongside its deep capabilities on the endpoint. Pluto covers common endpoint blind spots: malicious IDE extensions, browser plugins, and supply chain security, alongside agentic runtime protection. For teams looking to enforce long-standing endpoint controls while enabling agent adoption, this means one deployment solves several challenges.




Pluto's runtime capabilities excel across skills and intent. Their runtime detonation of skills and attack paths create meaningful malicious skills analysis, while intent tracking can redirect or block agentic activities. Their community tools and research also raise the bar for security teams looking to better understand the space. Between benchmark performance, breadth of coverage, and research output, Pluto is a leader in the emerging agentic endpoint security category.

Pluto Security is Best For

teams looking to secure agentic activity across their environments, while increasing endpoint coverage for things like packages, MCPs, Skills, and browser plugins.

SentinelOne has built on its acquisition of Prompt Security to strengthen its position as a security platform leader across endpoints, cloud, and now AI. Prompt Security earned a Leader badge in last year's report for having one of the most thorough understandings of LLM behavior among the first wave of AI security providers, backed by some of the deepest research in the space. A year after the acquisition, the original team is still in place and doing great work building on that research, now with a major endpoint platform behind them.

What impressed us most is how well SentinelOne has kept pace with the market's pivot to the endpoint. They've extended their detection and governance capabilities onto the endpoint and are building out the in-depth permissions capabilities that define the endpoint specialists, including governing MCP tool calls and agentic interactions at a granular permissions level.



Where SentinelOne separates from newer entrants is coverage. Governance extends into the browser, covering over 15,000 AI sites like Copilot in addition to first-party applications, agents, and MCPs. SentinelOne provides a single platform spanning the browser, the endpoint, and homegrown AI applications, making it one of the strongest broader AI security options on the market, pairing first-era research depth with a strong second-era endpoint story.

SentinelOne is Best For

teams looking to cover browser, endpoint, and first-party AI from one platform, especially as an extension of an existing SentinelOne relationship.







 **zenity**

 **paloalto**
NETWORKS

ONYX

 **SentinelOne**

 **NOMA**

 **LASSO**



 **Runlayer**

AIR*

BACKSLASH

 **PLUTO**



WIZ*

 **harmonic**

 **Pillar**

 **Operant**



 **glow**

 **Bloom**

 **ent**

 **BOLD**

 **LAVA**

Definitions

AI Security Platform Leader

AI Security Platform Leaders demonstrated the broadest coverage of AI across the enterprise, including support for endpoints, SaaS agents, browser plugins, and homegrown applications. These leaders also demonstrated roadmaps closely aligned with leading use cases, and strong individual offerings in a particular area.

Endpoint AI Security Leader

Endpoint AI Security Leaders demonstrated the most in-depth controls for AI on the endpoint that we evaluated. This included the ability to map permissions, folder structures and approved commands, alongside runtime controls for MCPs, Skills, and Intent based tracking. These vendors represent the most in depth AI governance and runtime capabilities.

Technical Innovator

Technical Innovators offered larger visions for AI Security, alongside the highest scores within a particular assessment area of the report, such as SaaS or Cloud agents. These innovators have differentiated capabilities for particular use cases that make them a great fit for certain teams.

Market Disruptor

Market Disruptors were recognized for re-envisioning an existing lane, for this report mostly endpoint related controls. These disruptors are solving longstanding issues with incumbent providers, and demonstrating visionary leadership with a new technical approach.



Ever wonder: Am I using the right security tools for my business, or am I building the right product for the market?

Everyday companies are making decisions based on the information that is available to them, which is often incomplete and based on vibes rather than usage.

That's where Latio comes in.

Founded in 2023 by James Berthoty, Latio was built to solve a critical problem James was facing: there was no reliable, credible way to evaluate a vendor's capabilities until after an agreement was signed. Latio exists to make the buying and building processes better by getting accurate information to the most relevant teams.

We focus on the product, the practitioner, and the market rather than slides and hype cycles. We believe the greatest predictor of a great security tool and program is finding the right product fit for both vendors and buyers.

We are creating a future where every decision is based on tests, market insights, experience, and hard work, where it's easy to find the right product you're looking for.

Our mission is to help every team find the right security product. So we test every product, to make it easier for you to pick the right one.

A special thank you to everyone who has supported this mission, without you, none of this would be possible.


Learn more:

 latio.com

 [Schedule a security program sync](#)

 [Schedule a product briefing](#)

 [Follow us](#)

The background is a deep blue gradient representing outer space. It features several celestial bodies: a large planet with blue and white clouds in the upper right, a smaller planet in the upper left, and another planet in the lower center. Numerous small, light blue diamond-shaped stars are scattered across the background.

Latio.

The only analyst firm that tests products,
so you can find the right one.