The agentic AI security market

MARKET RESEARCH

Comparing 50 vendors against 12 capabilities and 40 skills

Cem Nisanoglu

Version 1.0 · November 11, 2026

Author

Cem Nisanoglu

Founder & CEO, Paneau

Cem Nisanoglu is an international revenue and go-to-market executive with 20+ years across software engineering, cybersecurity, consulting, enterprise sales, and entrepreneurship. He's currently the Founder & CEO of Paneau, an AI-native enterprise context and policy layer designed to help move cybersecurity from human-led operations toward agentic execution. He also advises early-stage startups and scaleups on go-to-market strategy.

License

© 2026 Cem Nisanoglu. The article text and my original figures and tables published with it are licensed under Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). Third-party material, including vendor logos, is excluded. For commercial permissions, contact me at: cem.nisanoglu@gmail.com

Executive summary

AI agents are changing both how enterprise software is used and how it needs to be secured. Even Salesforce, a pioneer of SaaS, is going headless. As agents gain access to enterprise data, tools, and workflows, organizations face new security challenges, and cybersecurity vendors see a significant opportunity.

While advising an agentic AI security startup on its go-to-market (GTM) strategy, I examined both sides of this market. This article covers the supply side: what vendors offer and how their offerings differ. The demand side, meaning who buys and why, is not covered. For transparency, the startup I advised is not included in this comparison.

In a nutshell:

- **Funding is concentrated for agentic AI security startups.** A total of $1.64 billion was raised by the 21 AI agent security specialists as of October 25, 2026. The top five in terms of funding account for 55% of total.
- **Visibility is common, response is rare.** Nearly every vendor documents or claims the ability to show what agents are doing. Of the 50 vendors, 12 document or claim the ability to investigate an incident, seven to interrupt a running agent, and two to reverse its changes.
- **APIs or connectors are the most common deployment method.** They appear for 28 vendors, ahead of gateways or proxies (20).
- **Vendor groups emphasize different capabilities.** Of the four groups defined in this article, *AI agent security specialists* lead on monitoring, and *identity and access specialists* lead on mapping access. *Established cybersecurity vendors* collectively cover all 12 capabilities. All four *cloud and AI infrastructure providers* document or claim all four enforcement skills.

Methodology: I built a custom framework and assessed these 50 vendors against its 40 product skills, grouped into 12 capabilities, for a total of 2,000 vendor-skill assessments. I used public product documentation and explicit vendor claims. The evidence snapshot is October 25, 2026. Targeted scoring reviews concluded on November 5, and the report covers market developments through November 11.

Buyers can use the framework to compare products against the controls they need. **Implementers** can see where those controls connect to the agent stack. **Investors** can use the evidence to test how a vendor's offering differs from its peers.

Contents

Part One. Background and method
Part Two. The framework and market findings
Capabilities at a glance · Market findings
Part Three. The 40 skills
Appendix A. Individual vendor profiles
Appendix B. Funding sources

Part One. Background and method

Starting with the essentials

To understand the market, I first wanted to clarify the terminology. I've kept the definitions simple and applicable across different models, architectures, and deployment environments.

What's an AI agent?

For this article, an AI agent is a software program that uses an AI model to choose actions in pursuit of a goal, guided by instructions and constraints.

At its core, the agent operates through a repeating loop. The model determines the next step based on the goal, the information available, and the results of previous actions. What makes agents distinctive is that the model can choose how to proceed dynamically, instead of simply following a fully predefined sequence of steps.

The agent uses tools to gather information and carry out actions, while the surrounding software manages execution. The results feed back into the loop, informing the model's next decision. This continues until the task is complete or a stopping condition is reached.

The anatomy of an AI agent

The diagram below shows how the model, agent logic, and tools interact within this core loop.

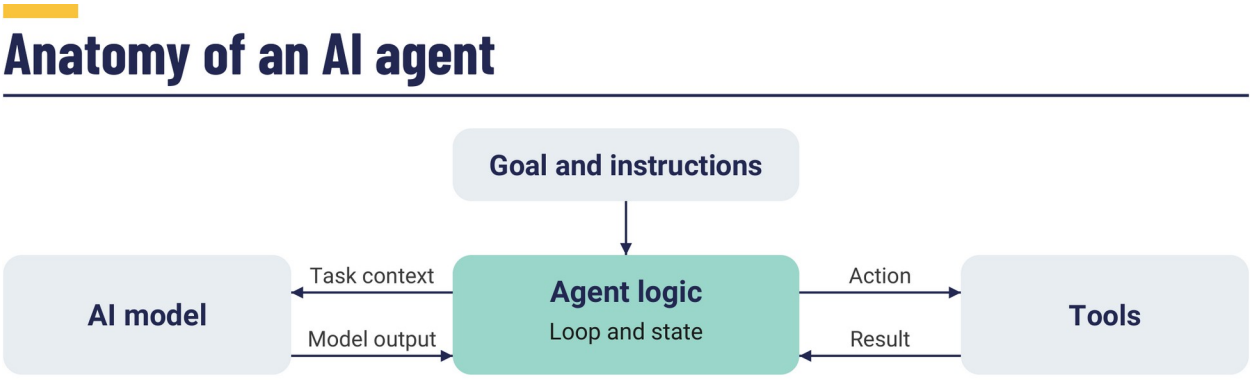


Figure 1. Example of a simple agent architecture, independent of model type or deployment.

Organizations can build and run their own agents, use agents embedded in products they buy, or combine both approaches. This affects what they can see and control. A company running its own agent may be able to change its software and execution environment, while a customer using an agent within a SaaS product may depend on the permissions, controls, and activity records the vendor provides.

In more complex deployments, the core loop sits within a broader system. The agent logic shown in the first diagram is often called the agent harness. It manages the loop and coordinates the model, tools, and task context. Memory retains information for use across steps or sessions, while runtime services support scheduling, checkpoints, retries, and resource management. These components can be packaged together or distributed across multiple services.

How do agents operate in practice?

Agents may be specialized for a particular task or domain, or general-purpose, handling a broader range of tasks. A specialized travel-booking agent might find flights and hotels that match your dates and budget, while a general-purpose assistant might help arrange that trip, compare broadband providers, or schedule a meeting in your calendar. What shapes their scope is how they are configured, including their model, instructions, the information they can access, the tools they can use, and their operating environment.

These differences affect how agents need to be secured. A travel agent that recommends flights and one that can book them using a company payment card may use the same model but carry different risks. The security requirements depend on what information and systems the agent can access, which actions it is authorized to take, and how independently it can act.

When several agents work together, organizations also need to understand what information they exchange and what authority they delegate. An agent passing a task to another agent introduces further questions: whose permissions apply, what can the receiving agent do, and can the organization trace their actions?

These practical differences help determine where visibility, access restrictions, approval requirements, and other security controls are needed.

What is agentic AI security?

For this article, agentic AI security means 1) protecting AI agents from compromise or manipulation, 2) controlling how their actions affect data, tools, and systems, and 3) maintaining oversight of their activity. The focus is on securing agents. Using agents to perform cybersecurity tasks is outside the scope of this article.

What can go wrong with AI agents?

AI agents can be compromised, manipulated, or given access they shouldn't have. They can also misuse legitimate tools or take unintended actions that expose information, change data, or disrupt systems. These failures can affect both the agent itself and the resources it interacts with.

The security challenge grows with autonomy and scale. Agents can carry out sequences of actions across enterprise systems, increasing the number of decisions and actions an organization needs to oversee. If an agent is manipulated, it may use legitimate access to carry out harmful actions, even while its software and credentials remain intact.

Consider an agent asked to summarize a confidential document. The document could contain malicious instructions that redirect the agent into sending its contents elsewhere, if its available tools and permissions allow it. The agent could therefore create a security breach while using access it was legitimately given.

OWASP identifies risk categories that help explain these failures, including goal hijacking, tool misuse, identity and privilege abuse, memory and context poisoning, and insecure communication between agents. These risks can arise across the architecture we've described, including the connections between its components. (OWASP Top 10 for Agentic Applications)

Where security risks can arise in an AI agent

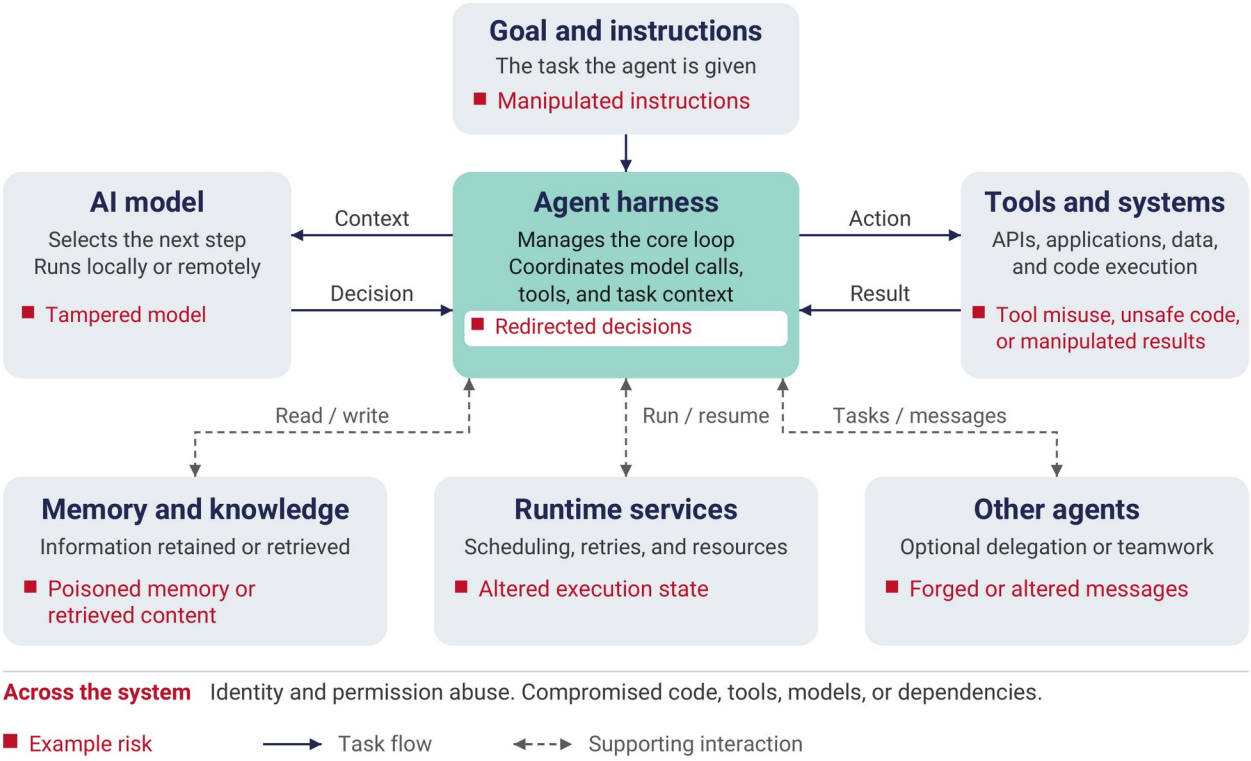


Figure 2. Illustrative security risks across an AI agent architecture. Based on the OWASP Top 10 for Agentic Applications (2026).

What measures can you take to protect AI agents?

Protection needs to address two connected objectives: protecting the agent's operation from compromise or manipulation, and controlling how its actions affect data, tools, and other systems.

Some measures come from established security practices, such as securing software dependencies, limiting permissions, protecting credentials, and isolating workloads. Others address how agents receive instructions, retain information, select tools, and delegate authority. In the document example, protections could include inspecting content for manipulation and independently restricting where confidential information can be sent. Giving the model instructions alone does not enforce those restrictions.

Controls can operate at different layers. Software-based controls can inspect requests, enforce permissions, and restrict execution. Hardware-backed controls can add isolation from the agent and its host. NVIDIA's OpenShell provides software-based runtime controls. NVIDIA describes Sentry as a framework designed to monitor agents and enforce access policy from a hardware-isolated environment on BlueField DPUs. (NVIDIA technical overview)

AI agents and the regulatory landscape

Alongside the security risks, organizations deploying AI agents need to understand their regulatory obligations. These depend on where they operate, how the agents are used, what data they process, and the organization's role in developing or deploying the system.

Existing rules such as the GDPR can apply when AI systems process personal data. The EU AI Act sets requirements for systems classified as high risk, including risk management, technical documentation, record-keeping, human oversight, and cybersecurity. Which rules apply, and when, depends on the system, its intended use, and the organization's role. (European Data Protection Board, European Commission)

Enterprises also have internal policies and standards governing how AI can be used. These may impose additional requirements, such as restricting access to sensitive data or requiring approval for certain actions.

Together, these external and internal requirements help explain why vendor offerings extend beyond detecting and blocking attacks. Monitoring and logging can support both security and compliance, helping security teams identify and investigate problems while providing evidence for compliance reviews and audits.

How the comparison works

The comparison covers 50 vendors and 40 skills. Sources include public product documentation, integration guides, and announcements. Documented features (D) and explicit vendor claims (VC) receive equal credit. NE means I found no qualifying evidence for the skill in the reviewed material. For each vendor, capability coverage is the share of that capability's skills coded D or VC. Group coverage is the mean of those vendor percentages. The evidence snapshot is October 25, 2026. Targeted scoring reviews concluded on November 5, and the report covers market developments through November 11.

Assessments apply to named products or modules and their supported environments and integrations. The research records capture required editions, dependencies, and release status, including preview and announced features. Some skills depend on others or share a control mechanism, so capability percentages can reflect overlapping controls.

The evidence records retain the product boundaries, dependencies, and open questions behind each assessment.

Limitations of this research

This research is based on public information. The skills and capabilities are scored using product documentation published by vendors and their claims about what their products do. Both carry equal weight.

This provides a view of the features vendors are building and how they communicate them, which is what a go-to-market comparison needs. Because documentation and claims carry equal weight, coverage shows what a vendor offers or says it offers, not what has been verified. I have not tested the products in real-world deployments, and vendors have not reviewed or confirmed their assessments. The findings assess public feature presence, not technical effectiveness. I have no affiliation with any of the vendors assessed.

The 50 vendors form a purposive comparison of named offerings relevant to securing AI agents, supported by public documentation or explicit claims. This is not an exhaustive market census. Groups reflect my judgment about each vendor's commercial starting point. Vendor entries follow assessed offerings rather than a strict parent-company count, and may combine named modules with different prerequisites.

I built the custom comparison framework based on observed supply, so a lack of coverage does not establish unmet demand, and the framework isn't a complete account of security needs.

NE means not established in the reviewed material, not absence. This article and the accompanying framework do not provide an overall vendor ranking and should not be used to build one.

How to use the comparison

Buyers can turn agent risks into required skills and use the coverage tables to identify products to evaluate.

Implementers can use the deployment methods to identify where a control would connect, then verify the required integrations, permissions, and operating conditions.

Investors can compare vendors within a relevant group and use the skill definitions and source evidence to test differentiation claims.

Part Two. The framework and market findings

The 12 capabilities at a glance

Vendors approach agent security from different directions. Some focus on identity and access, while others focus on inspecting information, testing agents, or controlling their actions. To compare these offerings, we need a common language that describes what each product enables an organization to do.

To achieve this, I built a custom framework and organized it into twelve capabilities and forty skills, which define product features. Each addresses a practical security or governance need.

CAPABILITY	# OF SKILLS	QUESTION THIS CAPABILITY ANSWERS
1. See	6	Which agents exist, and what are they doing?
2. Identify	3	Who is behind each agent?
3. Map access	3	What is each agent allowed to access and do?
4. Test	5	Can the agent be exploited?
5. Inspect	3	Is information entering or leaving the agent malicious, unsafe, or sensitive?
6. Set policy	4	What rules should agents follow?
7. Monitor	3	Is the agent acting outside its intended purpose or policy?
8. Enforce	4	Can we apply the rules to agents?
9. Investigate	1	What happened, why, and what was affected?
10. Contain	3	Can we limit the damage?
11. Remediate	4	Can we correct the underlying security problem?
12. Recover	1	Can we reverse harmful changes made by an agent?

See the full list of 40 skills

Vendor list and groupings

I grouped the 50 vendors by their commercial starting point:

- **Group 1:** AI agent security specialists
- **Group 2:** Identity and access specialists securing AI agents
- **Group 3:** Established cybersecurity vendors
- **Group 4:** Cloud and AI infrastructure providers

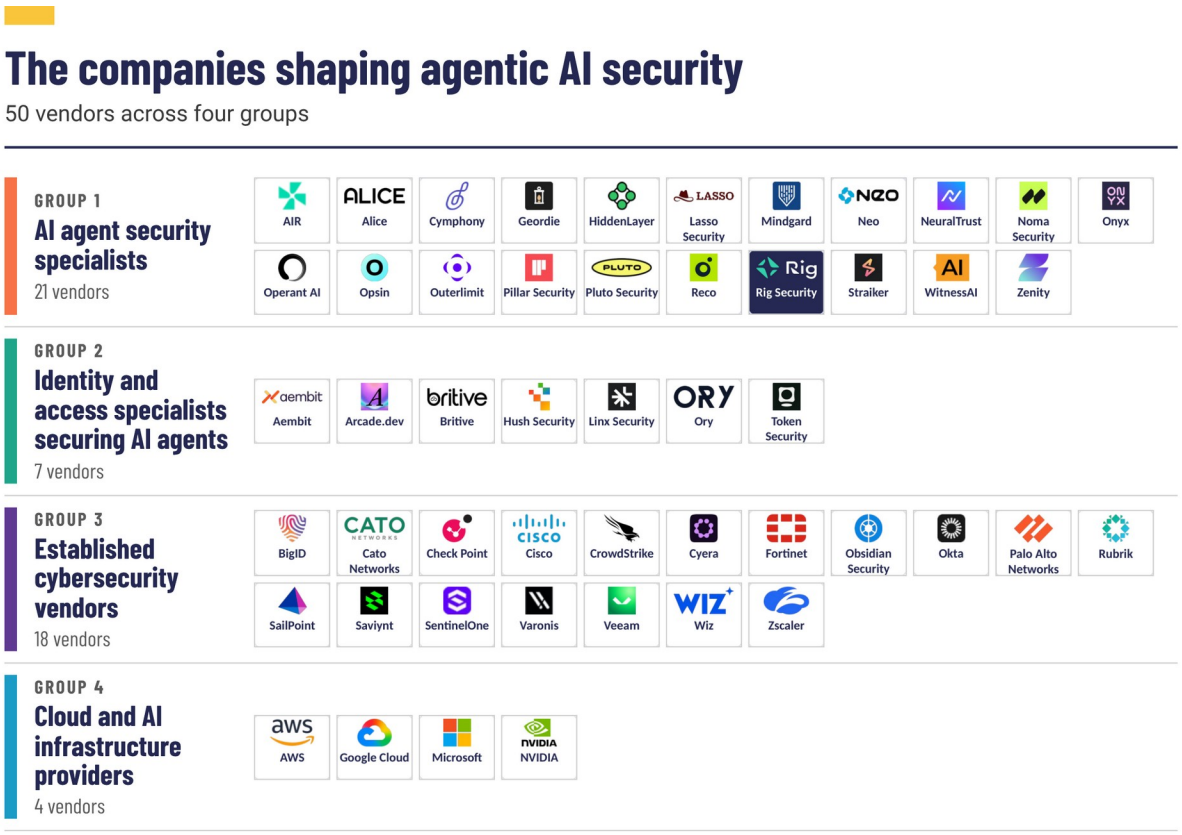


Figure 3. The 50 assessed vendors grouped into four categories.

What the vendor research shows

Coverage across the 40 skills

Visibility and policy are crowded. Six skills across three capabilities (See, Set policy, and Enforce) are documented or claimed by 40 or more of the 50 vendors. They cover discovering agents (1.1, 45 vendors),

observing and retaining activity records (1.2, 48 and 1.3, 46), authoring and scoping policies (6.1, 40 and 6.2, 41), and restricting access and actions (8.1, 46). I read these as table stakes for positioning.

Natural groupings of skills

Skills grouped by the share of the 50 vendors that document or claim them

TIER	SHARE OF THE 50 VENDORS	SKILLS	CAPABILITY MIX
Most common	80% or more (40 to 50 vendors)	6	See 3, Set policy 2, Enforce 1
Common	60 to 79% (30 to 39 vendors)	7	Test 2, Identify 1, Map access 1, Inspect 1, Monitor 1, Enforce 1
Occasional	20 to 59% (10 to 29 vendors)	21	See 3, Remediate 3, Identify 2, Map access 2, Test 2, Set policy 2, Monitor 2, Enforce 2, Inspect 1, Investigate 1, Contain 1
Rare	Under 20% (fewer than 10 vendors)	6	Contain 2, Test 1, Inspect 1, Remediate 1, Recover 1

Table 1. The four coverage tiers, with the number of skills and the mix of capabilities in each tier.

Most common skill in each capability

The skill documented or claimed by the most vendors within each capability

NO.	CAPABILITY	MOST COMMON SKILL		VENDORS	TIER
1	See	1.2	Observe agent activity	48/50	Most common
2	Identify	2.1	Attribute activity to an agent identity	34/50	Common
3	Map access	3.1	Map granted permissions	32/50	Common
4	Test	4.2 4.5	Find unsafe settings and permissions Classify and prioritize agent security risks	34/50	Common
5	Inspect	5.1	Detect attempts to manipulate the agent	32/50	Common
6	Set policy	6.2	Assign policies to a scope	41/50	Most common
7	Monitor	7.1	Detect unsafe agent actions	39/50	Common
8	Enforce	8.1	Restrict access and actions	46/50	Most common
9	Investigate	9.1	Investigate agent security incidents	12/50	Occasional
10	Contain	10.3	Quarantine an agent or component	17/50	Occasional
11	Remediate	11.1	Recommend a specific correction	21/50	Occasional
12	Recover	12.1	Reverse changes to connected systems	2/50	Rare

Table 2. The most widely documented or claimed skill in each capability. Tied skills are listed together.

Action and recovery are thin. Six skills are documented or claimed by eight vendors or fewer. They cover reversing changes to connected systems (12.1, 2 vendors), isolating the operating environment (10.1, 3), checking component integrity and origin (4.3, 7), interrupting agent execution (10.2, 7), assigning and tracking corrective work (11.2, 7), and detecting malicious links and code (5.2, 8). Across the 50 vendors, Recover, Contain, Investigate, and Remediate are the four least-covered capabilities, at 4%, 18%, 24%, and 26%. The median vendor has at least one documented or claimed skill in eight capabilities. The highest is eleven.

All 40 skills ranked by vendor count

Vendors (of 50) with each skill documented or claimed, colored by tier

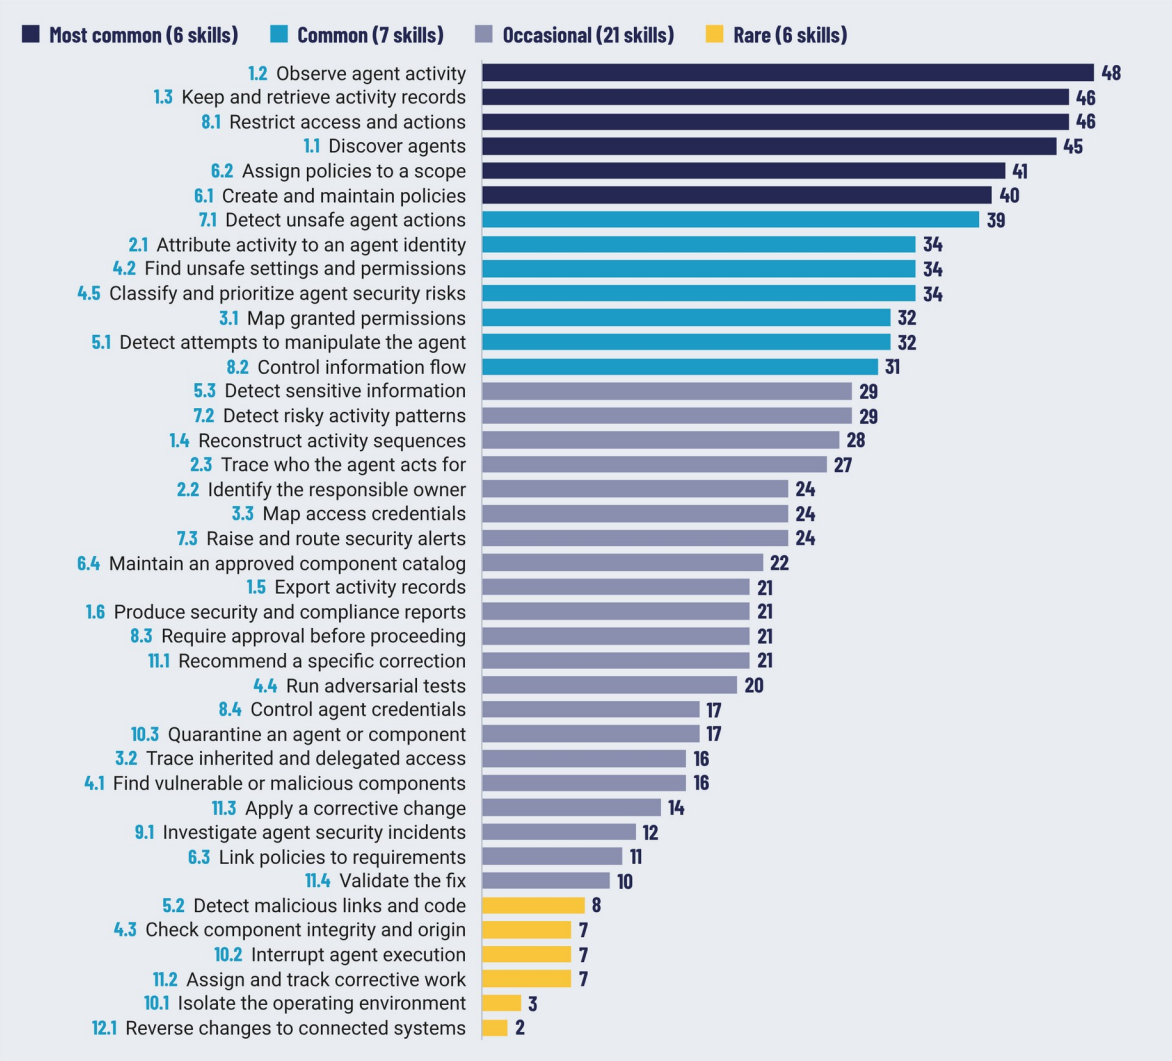


Figure 4. All 40 skills ranked by the number of vendors that document or claim them. Colors match the tiers in Table 1.

Thin supply leaves several possibilities. A skill may be underserved, a low priority for buyers, or already handled another way. Buyer research would distinguish those cases.

How the four groups differ

The group tables show the share of skills in a capability that a vendor documents or explicitly claims. Both types of evidence count equally. Darker cells indicate higher coverage. A zero means I found no qualifying evidence for any skill in that capability.

Tables 3 to 6 each show all the vendors in one group. The final row gives the group mean.

AI agent security specialists: The 21 specialists lead the four groups on Monitor (76%) and Test (51%). Their coverage is lowest in Recover (0%) and Contain (11%). The median specialist documents or claims 20 of the 40 skills, with a range of 7 to 27.

AI agent security specialists

Documented or claimed coverage (% of skills in each capability). n = skills per capability.

VENDOR	1 See n=6	2 Identify n=3	3 Map access n=3	4 Test n=5	5 Inspect n=3	6 Set policy n=4	7 Monitor n=3	8 Enforce n=4	9 Investigate n=1	10 Contain n=3	11 Remediate n=4	12 Recover n=1
AIR	17	0	0	40	33	0	67	25	0	0	0	0
Alice	83	0	0	40	67	50	0	25	0	0	50	0
Cymphony	33	33	33	40	0	0	0	25	0	0	50	0
Geordie	83	67	67	40	0	50	100	50	100	0	25	0
HiddenLayer	67	0	0	80	67	50	67	50	0	0	25	0
Lasso Security	100	33	33	80	67	50	100	50	100	0	75	0
Mindgard	33	0	0	60	0	0	33	0	0	0	50	0
Neo	50	33	0	20	33	50	33	75	0	0	0	0
NeuralTrust	100	67	67	40	100	100	100	100	0	0	25	0
Noma Security	83	100	33	80	33	75	100	75	100	0	0	0
Onyx	67	100	33	60	67	100	100	75	0	0	0	0
Operant AI	83	33	67	20	100	100	100	50	0	0	0	0
Opsin	50	67	100	40	0	25	100	25	0	33	100	0
Outerlimit	67	100	33	40	0	75	67	50	0	0	25	0
Pillar Security	100	33	67	100	67	75	100	50	0	0	50	0
Pluto Security	50	0	0	60	67	25	100	50	0	33	0	0
Reco	67	100	67	40	67	75	67	75	0	67	50	0
Rig Security	50	100	67	40	0	0	100	25	100	0	25	0
Straiker	83	0	0	80	67	50	100	50	0	67	50	0
WitnessAI	67	67	33	20	67	75	67	50	0	0	0	0
Zenity	67	67	67	40	67	50	100	50	100	33	25	0
Group mean	66.7	47.6	36.5	50.5	46.0	51.2	76.2	48.8	23.8	11.1	29.8	0.0

Table 3. Capability coverage for 21 AI agent security specialists. The final row is the group mean.

Identity and access specialists: The seven identity specialists lead on Map access (81%) and Identify (76%). The group averages 14% on Inspect and Monitor, 17% on Test, and 0% on Investigate and Recover. I read this as a profile built around agent access.

Identity and access specialists securing AI agents

Documented or claimed coverage (% of skills in each capability). n = skills per capability.

VENDOR	1 See n=6	2 Identify n=3	3 Map access n=3	4 Test n=5	5 Inspect n=3	6 Set policy n=4	7 Monitor n=3	8 Enforce n=4	9 Investigate n=1	10 Contain n=3	11 Remediate n=4	12 Recover n=1
Aembit	33	67	67	0	0	75	0	75	0	33	0	0
Arcade.dev	67	33	67	0	0	75	0	75	0	0	0	0
Britive	83	100	100	20	100	75	0	100	0	67	25	0
Hush Security	67	100	33	40	0	25	0	75	0	0	25	0
Linx Security	50	67	100	20	0	50	0	50	0	0	25	0
Ory	83	100	100	0	0	75	0	50	0	33	0	0
Token Security	83	67	100	40	0	75	100	50	0	0	75	0
Group mean	66.7	76.2	81.0	17.1	14.3	64.3	14.3	67.9	0.0	19.0	21.4	0.0

Table 4. Capability coverage for seven identity and access specialists. The final row is the group mean.

Established cybersecurity vendors: The 18 established vendors are the only group with documented or claimed skills in all 12 capabilities. Their mean coverage is highest on See (73%) and Monitor (61%), and they lead the groups on Investigate (33%) and Recover (11%). The two qualifying recovery paths (12.1) are Rubrik's **Agent** Rewind claim for Salesforce assets and Veeam's documented Undo AI function for supported Microsoft 365 files. (Rubrik, Veeam)

Established cybersecurity vendors

Documented or claimed coverage (% of skills in each capability). n = skills per capability.

VENDOR	1 See n=6	2 Identify n=3	3 Map access n=3	4 Test n=5	5 Inspect n=3	6 Set policy n=4	7 Monitor n=3	8 Enforce n=4	9 Investigate n=1	10 Contain n=3	11 Remediate n=4	12 Recover n=1
BigID	83	100	100	40	33	75	100	75	100	33	75	0
Cato Networks	33	0	0	60	33	0	33	25	0	0	25	0
Check Point	83	0	67	100	67	75	67	0	0	0	25	0
Cisco	67	33	33	100	100	100	67	50	0	0	0	0
CrowdStrike	83	100	0	20	100	75	33	50	100	33	0	0
Cyera	67	0	33	40	67	50	67	75	0	33	25	0
Fortinet	100	67	33	60	67	75	100	75	0	0	0	0
Obsidian Security	50	100	100	40	0	50	67	25	0	0	0	0
Okta	67	100	67	0	0	75	0	75	0	33	0	0
Palo Alto Networks	100	100	67	100	100	75	33	100	0	33	100	0
Rubrik	83	67	33	0	33	50	33	75	100	0	0	100
SailPoint	67	33	33	20	33	25	100	25	100	33	0	0
Saviynt	67	100	67	40	0	50	67	50	100	33	25	0
SentinelOne	67	67	33	20	67	75	33	75	0	33	0	0
Varonis	83	0	0	80	67	50	100	75	0	33	0	0
Veeam	83	0	0	20	33	75	67	50	0	0	50	100
Wiz	67	67	33	60	33	25	100	25	100	0	50	0
Zscaler	67	100	67	60	67	50	33	50	0	0	25	0
Group mean	73.1	57.4	42.6	47.8	50.0	58.3	61.1	54.2	33.3	16.7	22.2	11.1

Table 5. Capability coverage for 18 established cybersecurity vendors. The final row is the group mean.

Cloud and AI infrastructure providers: All four providers document or claim every skill under Enforce. The group also has the highest average coverage for See, Inspect, Set policy, Contain, and Remediate.

Three providers document a way to isolate execution (skill 10.1). AWS and NVIDIA provide an isolated place for the agent itself to run. Google Cloud's Code Execution provides a sandbox for the code the agent asks it to run. (AWS Bedrock AgentCore Runtime, Google Cloud Code Execution, and NVIDIA OpenShell)

Cloud and AI infrastructure providers

Documented or claimed coverage (% of skills in each capability). n = skills per capability.

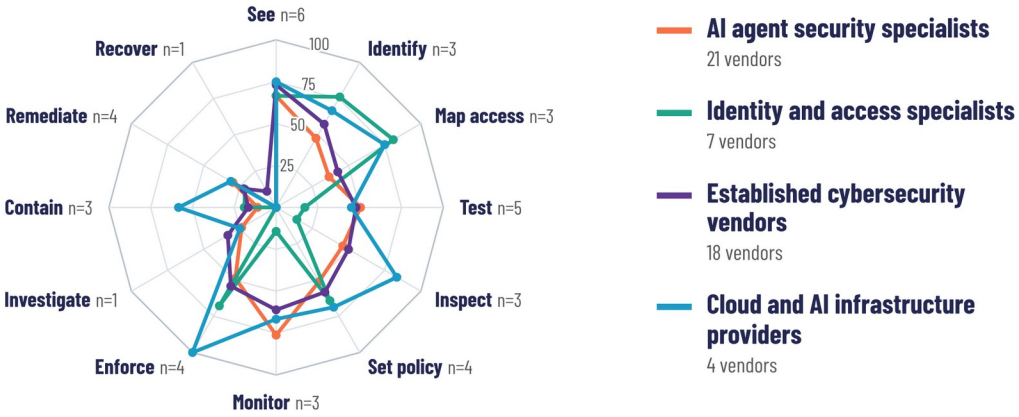
VENDOR	1 See n=6	2 Identify n=3	3 Map access n=3	4 Test n=5	5 Inspect n=3	6 Set policy n=4	7 Monitor n=3	8 Enforce n=4	9 Investigate n=1	10 Contain n=3	11 Remediate n=4	12 Recover n=1
AWS	83	67	67	20	67	100	0	100	0	67	25	0
Google Cloud	83	67	100	60	100	75	100	100	0	33	25	0
Microsoft	83	100	67	60	100	50	100	100	100	33	75	0
NVIDIA	50	33	67	40	67	50	67	100	0	100	0	0
Group mean	75.0	66.7	75.0	45.0	83.3	68.8	66.7	100.0	25.0	58.3	31.2	0.0

Table 6. Capability coverage for four cloud and AI infrastructure providers. The final row is the group mean.

Figure 5 compares the group means. Figure 6 shows how widely individual vendor profiles vary within each group.

Mean capability coverage by supplier group

Mean share of skills documented or claimed in each capability, from 0 to 100%

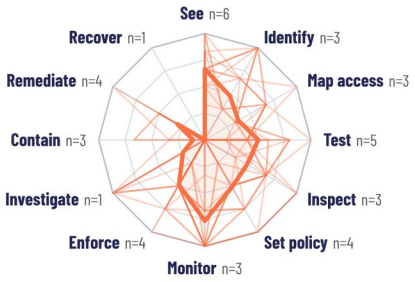


n = skills per capability.

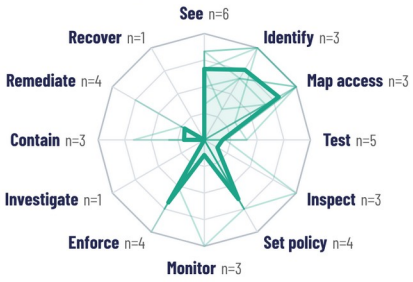
Figure 5. Mean capability coverage for each of the four vendor groups.

Capability profiles within each group

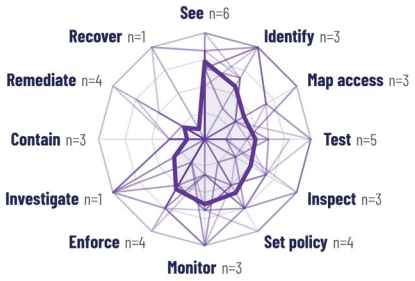
AI agent security specialists (21)



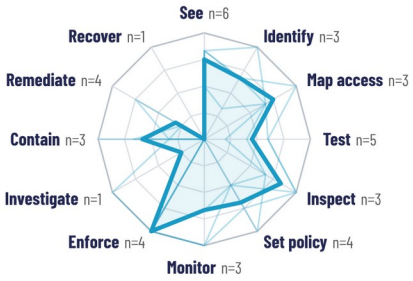
Identity and access specialists (7)



Established cybersecurity vendors (18)



Cloud and AI infrastructure providers (4)



Thin lines: suppliers. Bold line: group mean. n = skills per capability.

Figure 6. Individual vendor profiles and group means on the same scale.

How security products fit into an AI agent deployment

Capabilities describe what a product does. How it connects to the agent stack shapes what it can see and when it can intervene. In Figure 7, a gateway sits between the agent harness and the tools it calls. It can inspect a request before it reaches the tool and block it if configured to enforce a policy. A connector that pulls activity records typically sees the event after it happens.

I identified at least one deployment method for 46 of the 50 vendors. APIs or connectors were the most common (28 vendors), followed by gateways or proxies (20).

I grouped the methods vendors document or claim into seven overlapping categories. A vendor counts once in each category it qualifies for. Where a label uses "or," either form counts.

1. API or connector. Connects a security product to an existing system through an API or a packaged integration.

2. Gateway or proxy. Sits in the request path between the agent and a model or tool.

3. SDK or hook. Adds security through a library or runs a check at a specific point in the agent's execution.

4. Endpoint or runtime sensor. Observes activity on the host or in the environment where the agent runs.

5. Identity or credential broker. Mediates access by supplying or brokering the credentials an agent uses.

6. Agentless posture. Checks configuration or permissions remotely through existing interfaces without installing a local sensor.

7. Platform runtime or sandbox. Provides a managed environment that hosts or isolates the agent workload.

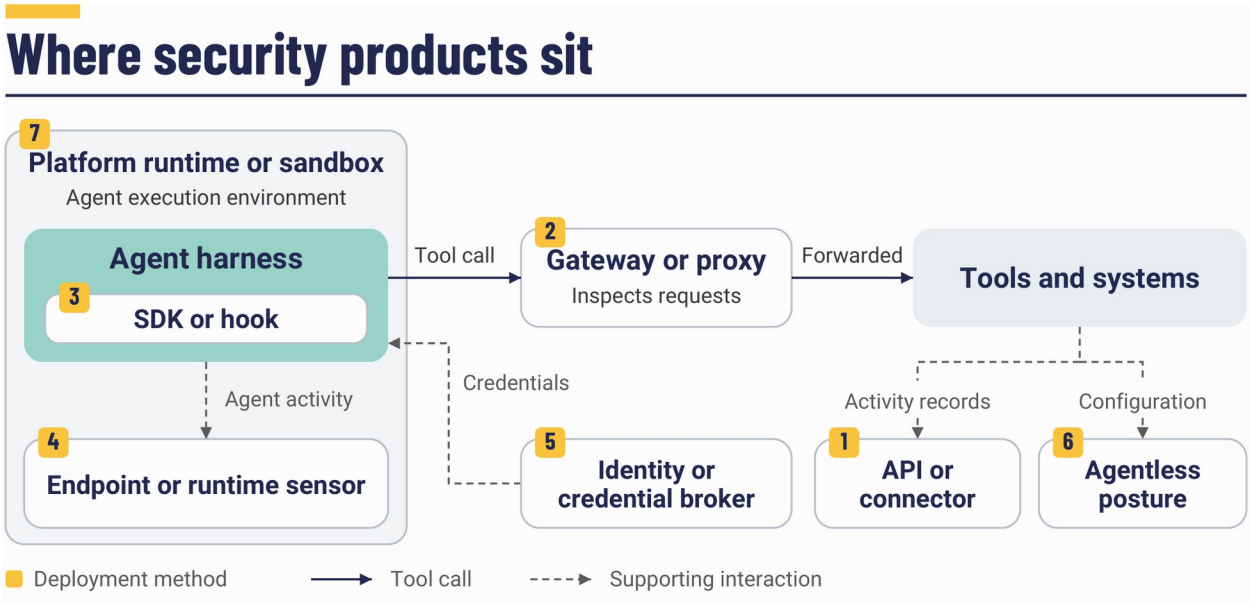


Figure 7. Example placements for the seven security deployment methods. A product can use more than one.

Figure 8 shows how often each method appeared across the 50 vendors.

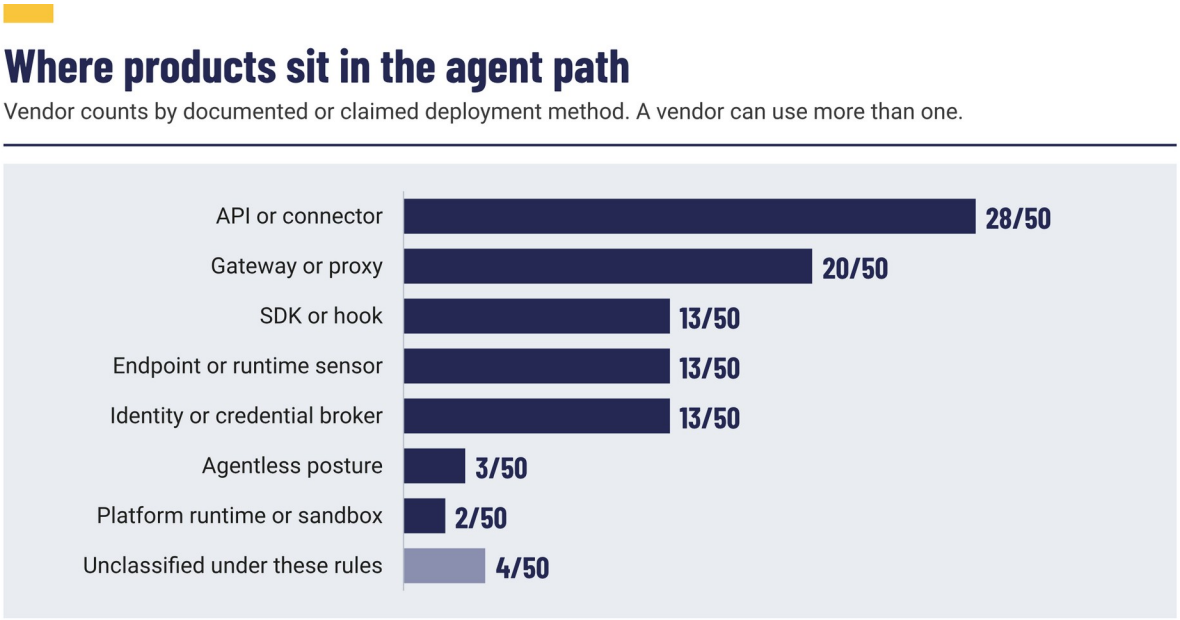


Figure 8. Deployment methods across 50 vendors. For four vendors, the evidence was insufficient to classify a method.

AI agent security specialists: Of the 21 vendors, 13 documented or claimed an API or connector. Seven named a gateway or proxy.

Identity and access specialists: Five of the seven vendors documented or claimed an identity or credential broker. Four named a gateway or proxy, and one named an API or connector.

Established cybersecurity vendors: Of the 18 vendors, 11 documented or claimed an API or connector. Seven named a gateway or proxy.

Cloud and AI infrastructure providers: Three of the four vendors documented or claimed an identity or credential broker. Among the offerings I assessed, only AWS (Bedrock AgentCore Runtime) and NVIDIA (OpenShell) provided a runtime or sandbox for the agent itself. For Google and Microsoft, I assessed Google's code-execution sandbox and Microsoft's identity and prompt-protection modules, not their separate agent-hosting products.

Funding and consolidation

Across the 21 vendors in the specialist group, I identified about $1.64 billion in cumulative company financing as of October 25, 2026. These are company-level amounts and can include capital raised for work outside agent security. Some figures are minimums or estimates. The five highest recorded figures account for 55% of the subtotal. The median is $42 million, and the range runs from $9 million to $280 million. Appendix B gives the amounts and sources.



Figure 9. Distribution of cumulative disclosed funding among the 21 AI agent security specialists.

Established cybersecurity vendors are expanding their offerings to secure AI deployments through both product development and acquisitions, while AI-native entrants are building products specifically to address these emerging requirements. On May 29, 2026, Palo Alto Networks announced that it had acquired Portkey, adding an AI gateway to monitor AI traffic and apply security controls. On August 17, 2026, Fortinet announced that it had acquired Virtue AI, adding capabilities to test AI agents for vulnerabilities and protect them during operation.

Acquisitions are not limited to the established giants. In September 2026, Cyera completed its $1 billion acquisition of Oasis Security, combining data security with identity and access management to help control which data agents can reach.

And consolidation may not be over. In an interview published in October 2026, Palo Alto Networks CEO Nikesh Arora described continuing to monitor agentic security startups as potential acquisition targets.

After the October 25 evidence snapshot, on November 9, 2026, Vanta announced that it had signed a definitive agreement to acquire Pillar Security. Pillar remains in this research as assessed at the snapshot.

Conclusion

Personally, this has been the fastest-moving technology I've ever worked with. New model capabilities and optimizations seem to emerge almost weekly, making this research a snapshot of a rapidly evolving market.

As it currently stands, vendors widely document or claim agent discovery, activity records, and policy controls. Far fewer document or claim the ability to investigate an incident (12 vendors), interrupt a running agent (7), or reverse the damage (2).

The four groups defined in this article have different strengths. AI agent security specialists lead on monitoring coverage, and identity and access specialists lead on mapping access. Established cybersecurity vendors are the only group with documented or claimed skills in all 12 capabilities when taken together. Each of the four cloud and AI infrastructure providers documents or claims all four enforcement skills in the offerings I assessed.

The questions for buyers are how they handle agent risk today, who owns the problem inside the organization, and what they would pay to improve. Answering those questions requires research with buyers. **Investors** should note that the supply side is already crowded: 50 vendors across four groups, with acquisitions under way. The open question is how newcomers will differentiate from established vendors. **Implementers** should expect the technical details to keep changing, because the market is moving quickly.

> **I'd like to hear from you if you're:**
>
> - **working on positioning, go-to-market, or market entry in agentic AI security;**
> - **rolling out controls to secure AI agents as part of your cybersecurity program;**
> - **looking to share a correction or additional public evidence on behalf of a vendor covered in this research.**
>
> **Email me at:** cem.nisanoglu@gmail.com

Part Three. The 40 skills

Reach out via email or DM me on LinkedIn for the standalone framework.

The 12 capabilities

What a security product can do for AI agents

1	See	7	Monitor
2	Identify	8	Enforce
3	Map access	9	Investigate
4	Test	10	Contain
5	Inspect	11	Remediate
6	Set policy	12	Recover

Table 7. The 12 capabilities, from seeing agents to recovering.

The framework: 12 capabilities and 40 skills

NO.	CAPABILITY	SKILLS: PRODUCT FUNCTIONS ASSESSED			
1	See	1.1	Discover agents	1.4	Reconstruct activity sequences
		1.2	Observe agent activity	1.5	Export activity records
		1.3	Keep and retrieve activity records	1.6	Produce security and compliance reports
2	Identify	2.1	Attribute activity to an agent identity		
		2.2	Identify the responsible owner		
		2.3	Trace who the agent acts for		
3	Map access	3.1	Map granted permissions		
		3.2	Trace inherited and delegated access		
		3.3	Map access credentials		
4	Test	4.1	Find vulnerable or malicious components	4.4	Run adversarial tests
		4.2	Find unsafe settings and permissions	4.5	Classify and prioritize agent security risks
		4.3	Check component integrity and origin		
5	Inspect	5.1	Detect attempts to manipulate the agent		
		5.2	Detect malicious links and code		
		5.3	Detect sensitive information		
6	Set policy	6.1	Create and maintain policies	6.3	Link policies to requirements
		6.2	Assign policies to a scope	6.4	Maintain an approved component catalog
7	Monitor	7.1	Detect unsafe agent actions		
		7.2	Detect risky activity patterns		
		7.3	Raise and route security alerts		
8	Enforce	8.1	Restrict access and actions	8.3	Require approval before proceeding
		8.2	Control information flow	8.4	Control agent credentials
9	Investigate	9.1	Investigate agent security incidents		
10	Contain	10.1	Isolate the operating environment		
		10.2	Interrupt agent execution		
		10.3	Quarantine an agent or component		
11	Remediate	11.1	Recommend a specific correction	11.3	Apply a corrective change
		11.2	Assign and track corrective work	11.4	Validate the fix
12	Recover	12.1	Reverse changes to connected systems		

Table 8. Agentic AI security framework: 12 capabilities and 40 skills.

Appendix A. Individual vendor profiles

The charts below show, for each vendor, the share of a capability's defined skills that is documented or claimed. They use the same scale and the same equal credit as the heatmaps, and a value of zero means not established, not absence.

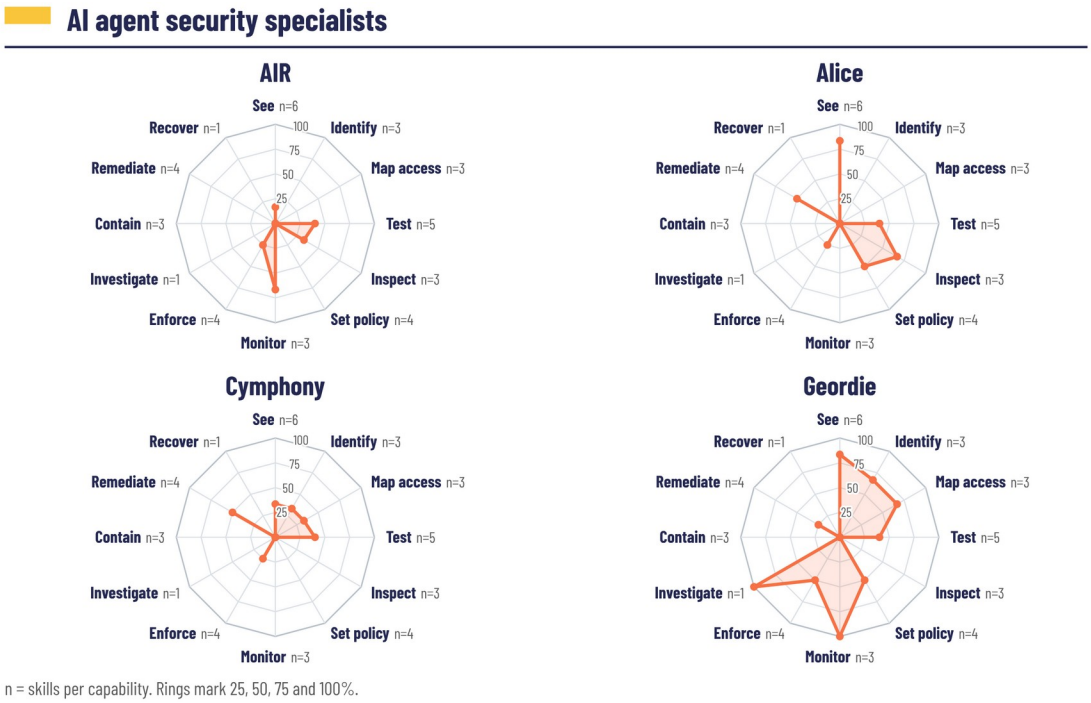


Figure A1. Capability profiles for AIR, Alice, Cymphony, and Geordie. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

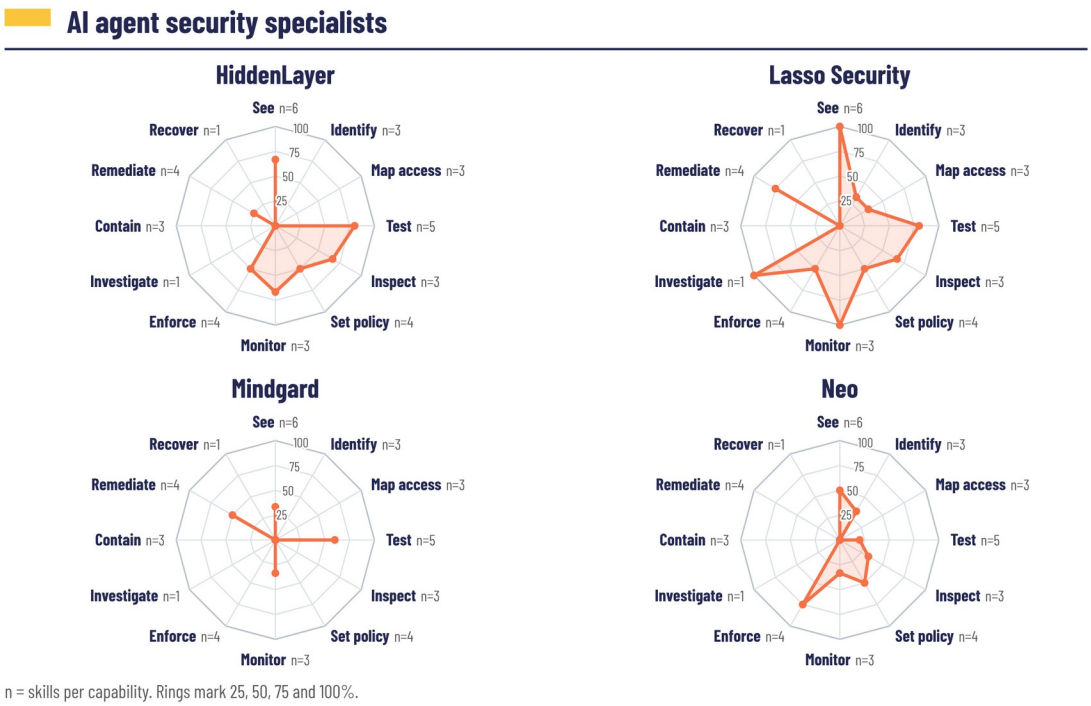


Figure A2. Capability profiles for HiddenLayer, Lasso Security, Mindgard, and Neo. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

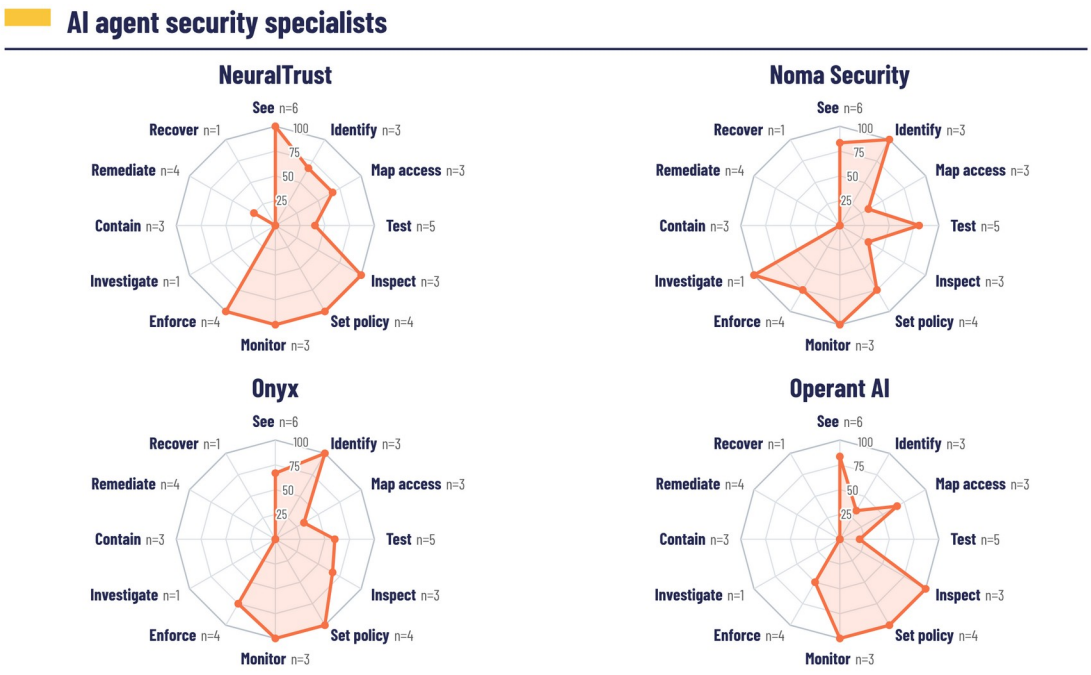


Figure A3. Capability profiles for NeuralTrust, Noma Security, Onyx, and Operant AI. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

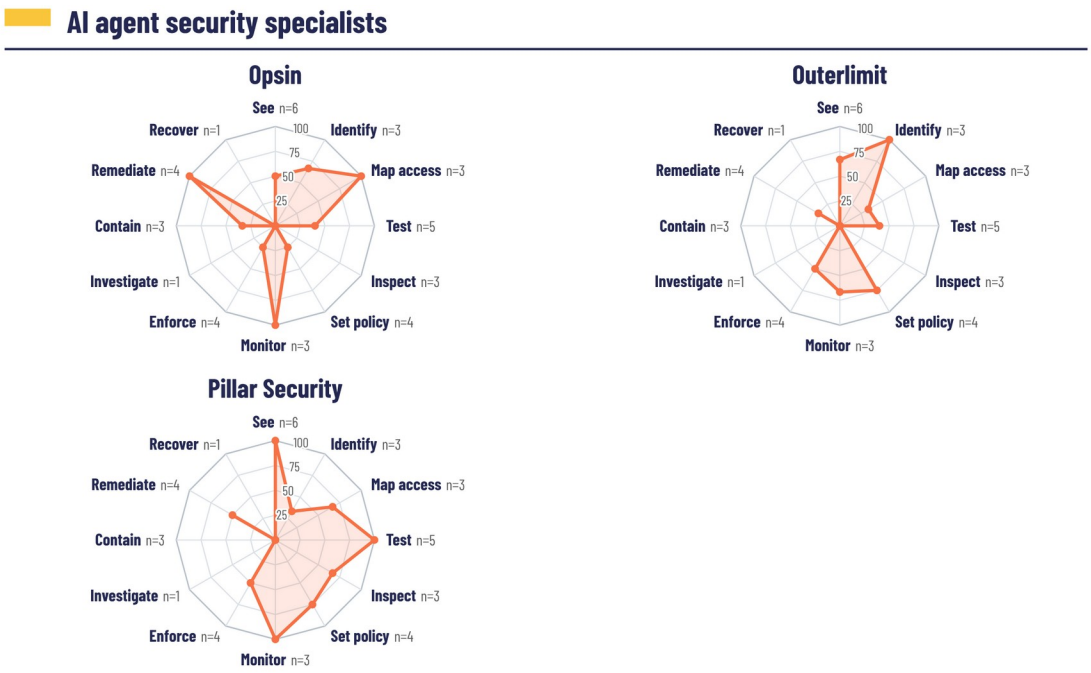


Figure A4. Capability profiles for Opsin, Outerlimit, and Pillar Security. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

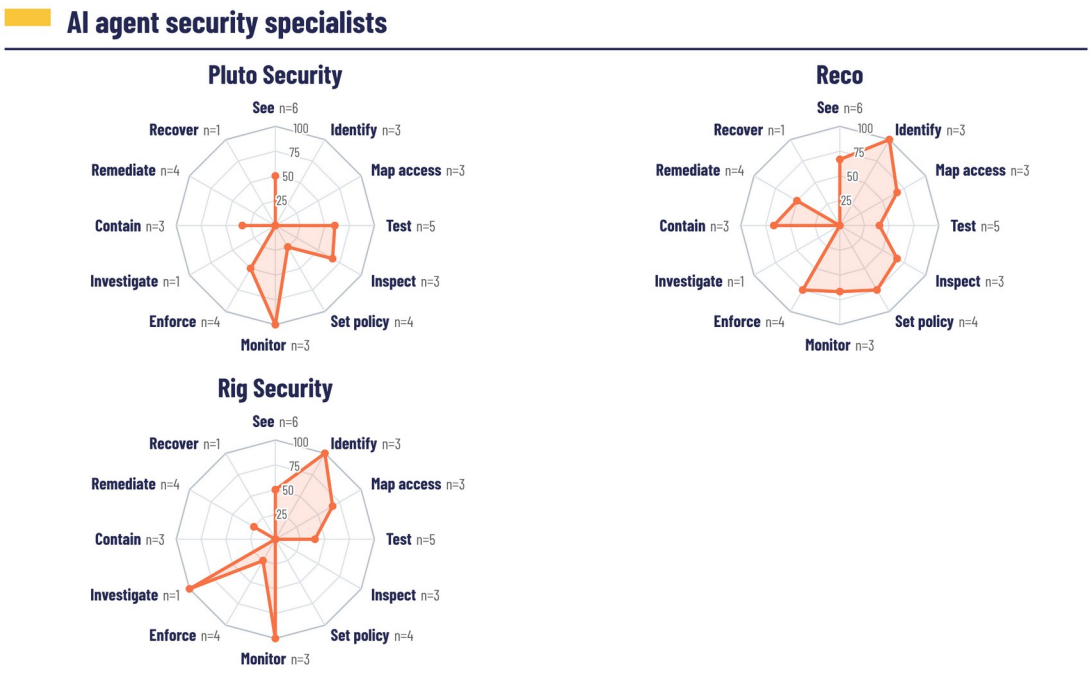


Figure A5. Capability profiles for Pluto Security, Reco, and Rig Security. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

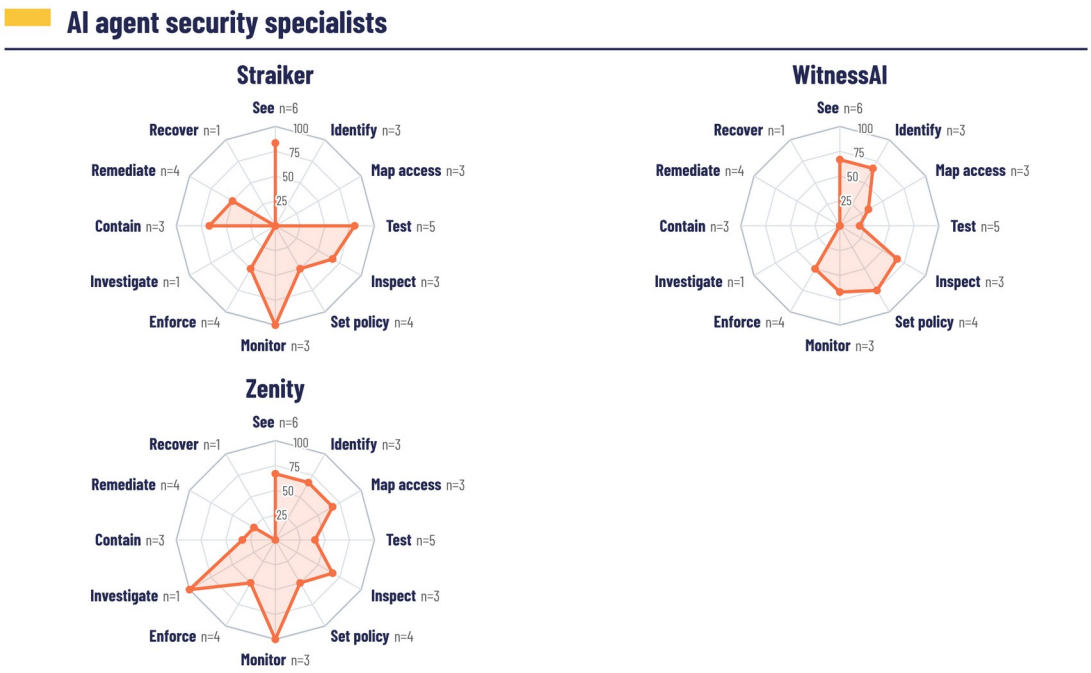


Figure A6. Capability profiles for Straiker, WitnessAI, and Zenity. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

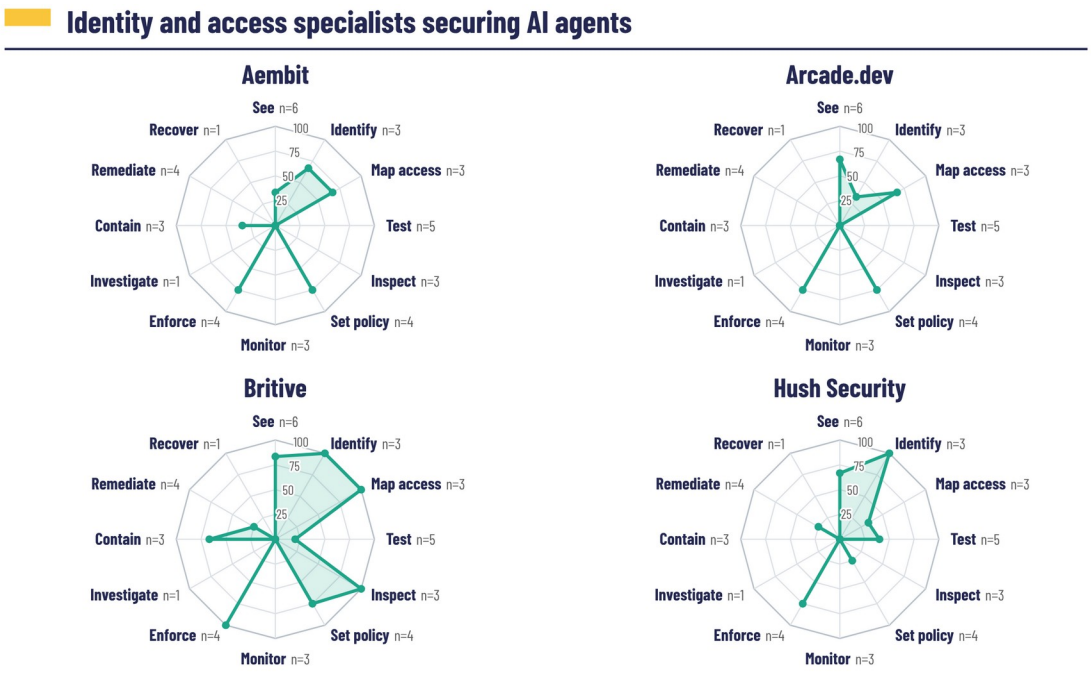


Figure A7. Capability profiles for Aembit, Arcade.dev, Britive, and Hush Security. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

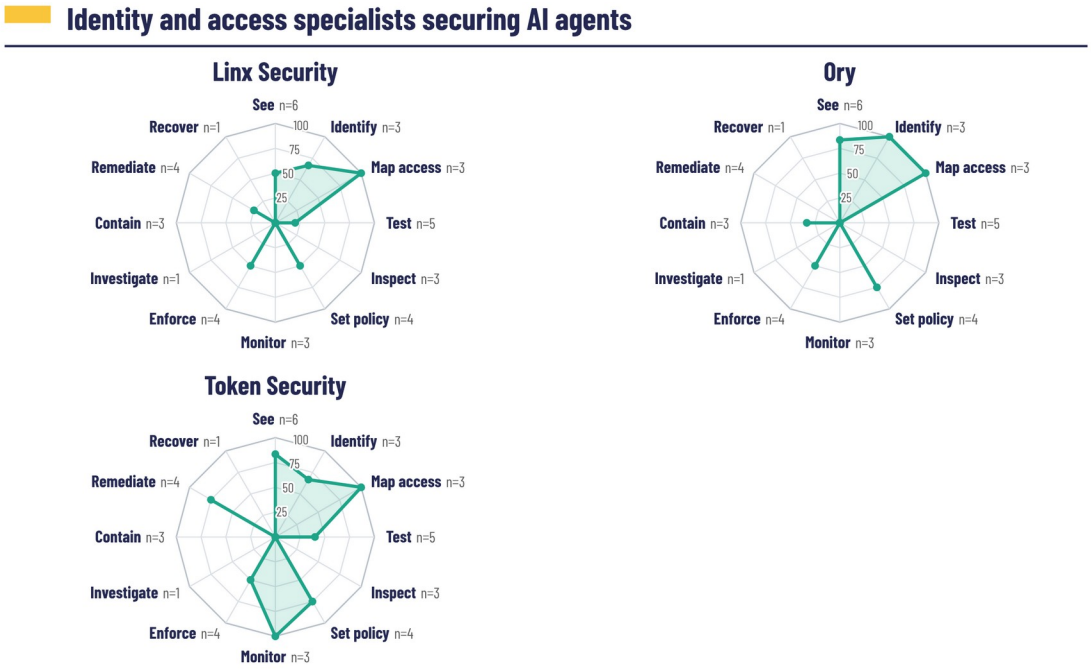


Figure A8. Capability profiles for Linx Security, Ory, and Token Security. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

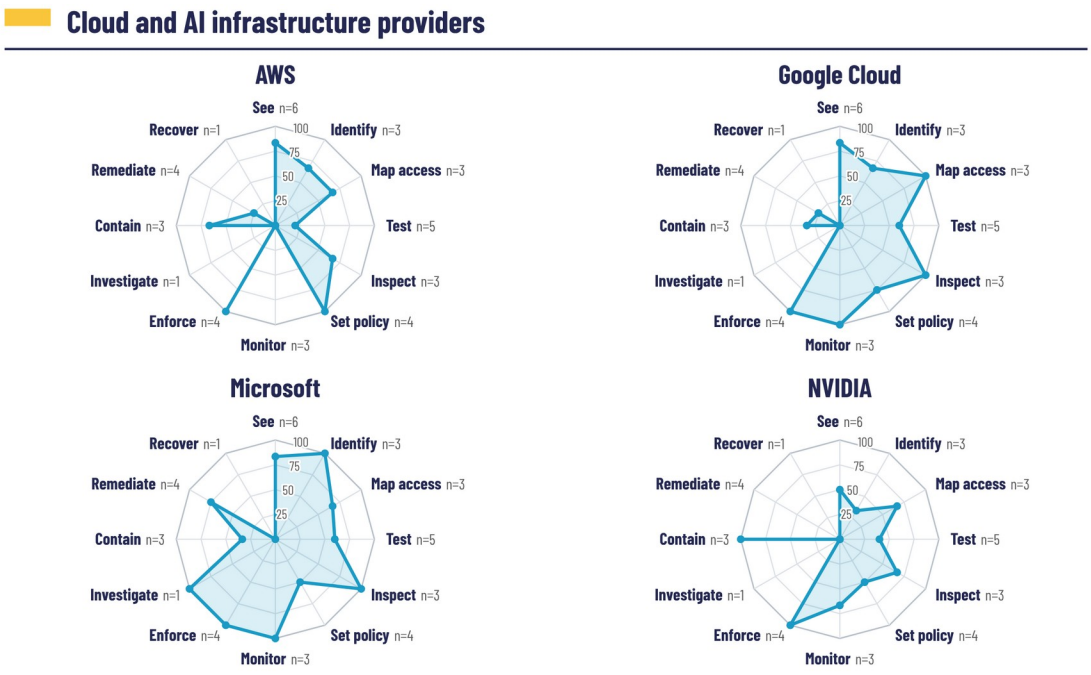


Figure A9. Capability profiles for AWS, Google Cloud, Microsoft, and NVIDIA. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

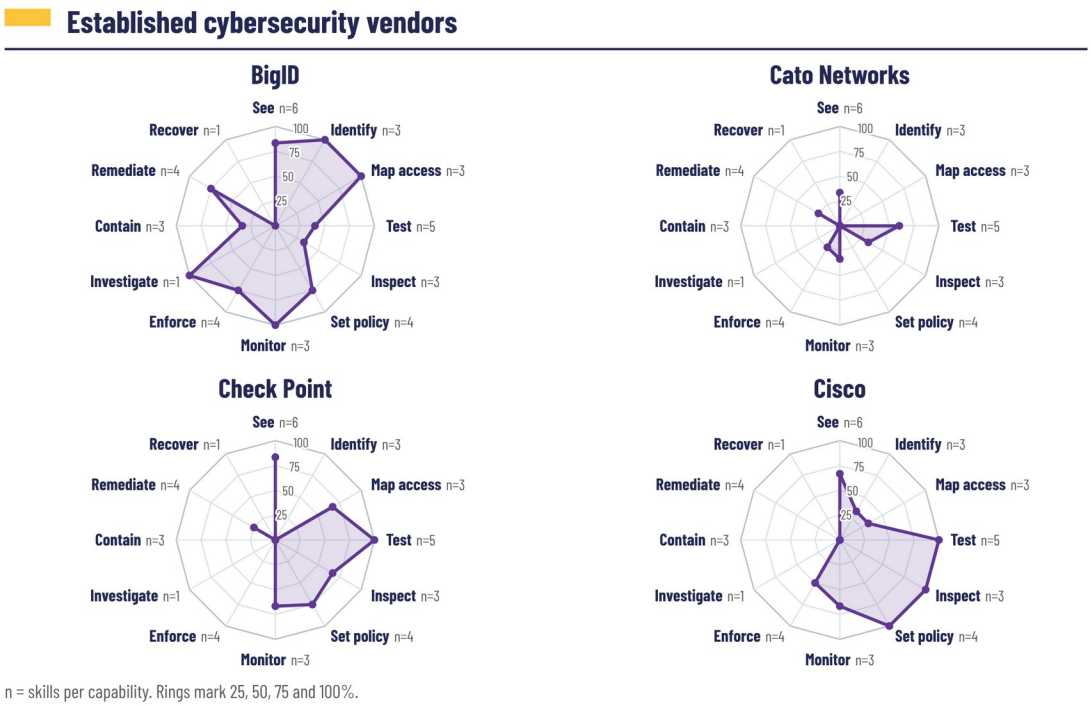


Figure A10. Capability profiles for BigID, Cato Networks, Check Point, and Cisco. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

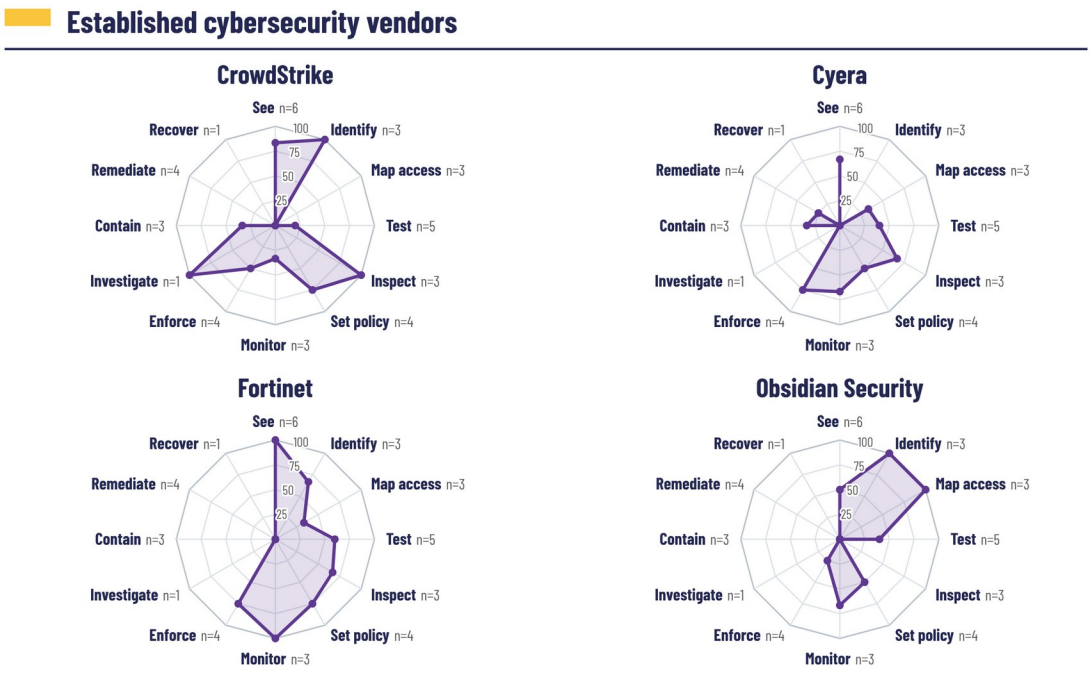


Figure A11. Capability profiles for CrowdStrike, Cyera, Fortinet, and Obsidian Security. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

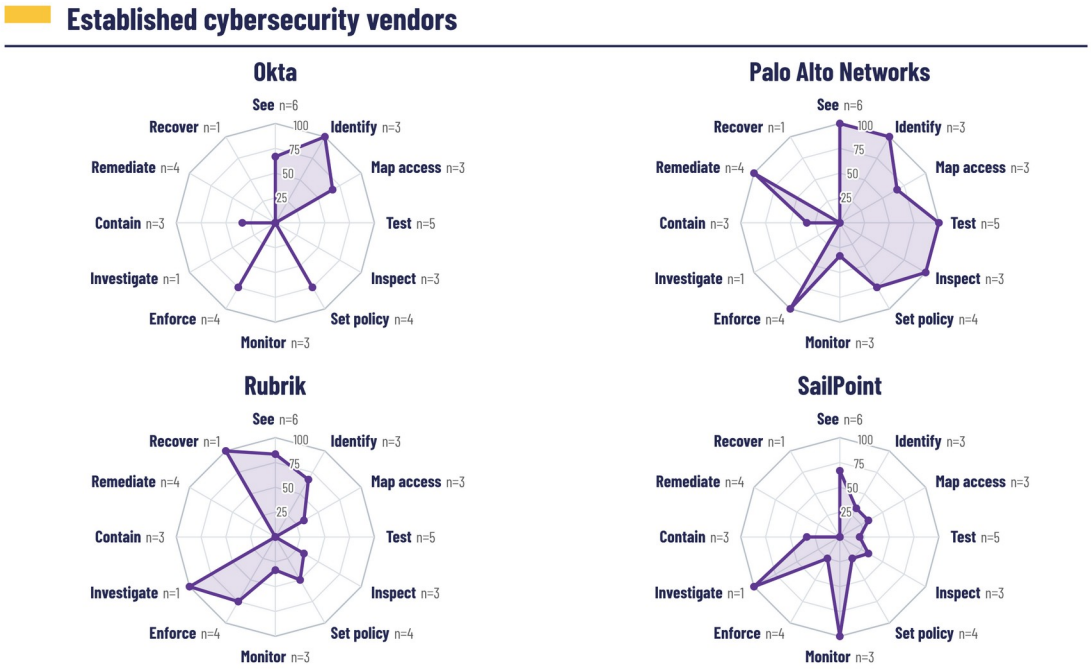


Figure A12. Capability profiles for Okta, Palo Alto Networks, Rubrik, and SailPoint. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

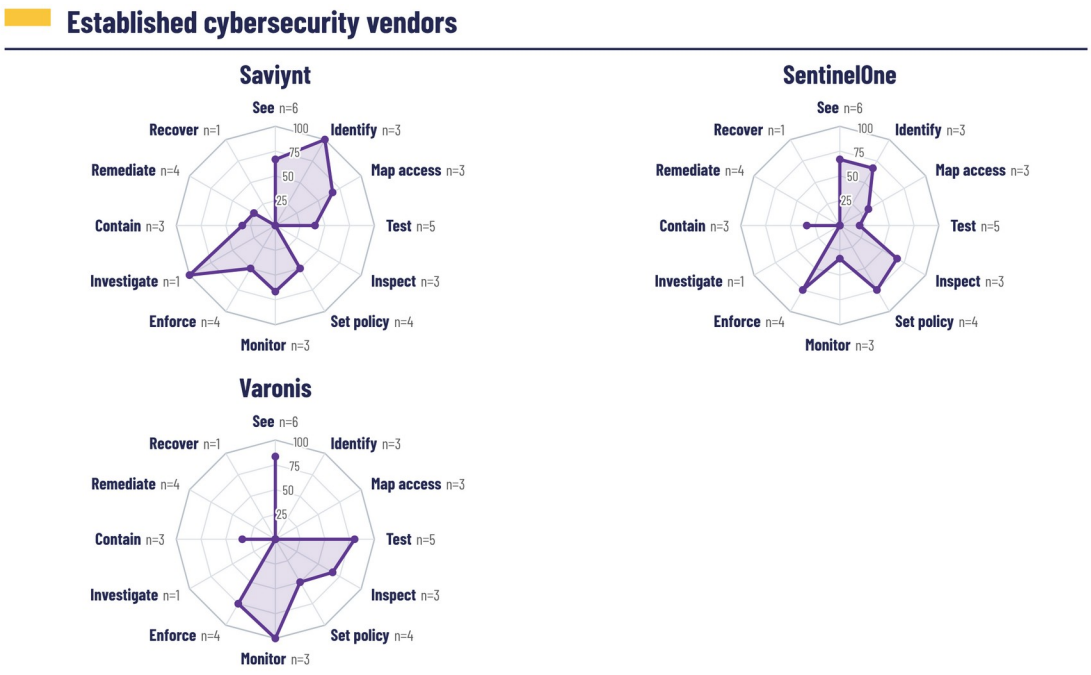
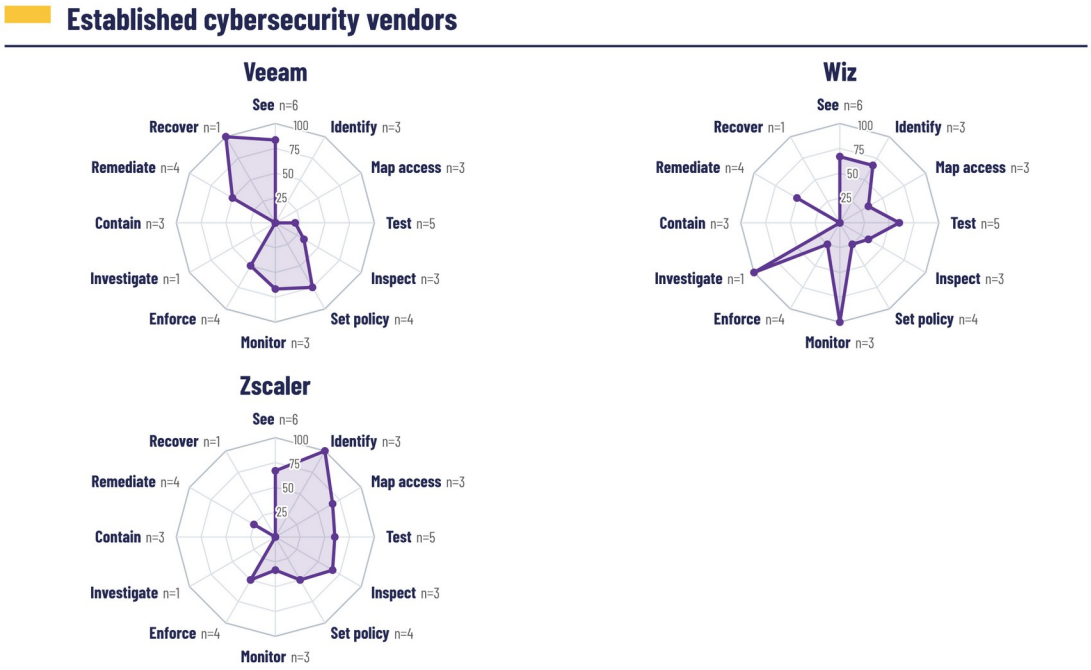


Figure A13. Capability profiles for Saviynt, SentinelOne, and Varonis. Each spoke shows the share of defined skills documented or claimed within the assessed scope.



n = skills per capability. Rings mark 25, 50, 75 and 100%.

Figure A14. Capability profiles for Veeam, Wiz, and Zscaler. Each spoke shows the share of defined skills documented or claimed within the assessed scope.

Appendix B. Funding sources

Company financing identified in public disclosures as of October 25, 2026. Figures are in US dollars and include financing for activities outside agent security. The subtotal uses the amounts shown, including estimates and disclosed minimums.

SPECIALIST	US$ MILLION	BASIS AND PUBLIC DISCLOSURE
AIR	50	Identified company financing. Source
Alice	280	Identified company financing. Source
Cymphony	30	Identified company financing. Source
Geordie	36.5	Identified company financing. Source 1, Source 2
HiddenLayer	156	Identified company financing. Source 1, Source 2, Source 3
Lasso Security	36	Disclosed minimum. Source 1, Source 2
Mindgard	42	Approximate total. Source 1, Source 2, Source 3, Source 4
Neo	100	Identified company financing. Source
NeuralTrust	20	Disclosed seed financing. Source
Noma Security	132	Identified company financing. Source 1, Source 2
Onyx	153	Identified company financing. Source
Operant AI	13.5	Identified company financing. Source
Opsin	21	Disclosed minimum. Source
Outerlimit	16	Identified company financing. Source
Pillar Security	9	Disclosed seed financing. Source
Pluto Security	42	Identified company financing. Source
Reco	140	Identified company financing. Source
Rig Security	12	Identified company financing. Source
Straiker	85	Identified company financing. Source
WitnessAI	85.5	Identified company financing. Source
Zenity	180	Disclosed minimum. Source 1, Source 2
Total identified	**1,639.5**	**21 assessed specialists**