FUTURIOM
FUTURE OF CLOUD TECH

Cloud Market Trend Report

Enterprise Security for Agentic AI

October 2026



Sponsored by:

STELLAR CYBER®

FUTURIOM.COM

Highlights and Key Findings

- **Security is a large factor in the pace of enterprise AI adoption.** In the Futuriom 2026 Networking and SASE Survey, 94 of 144 respondents, or 65.28%, said infrastructure security concerns were slowing AI adoption.
- **The cybersecurity operating model is changing as the installed stack evolves.** We expect suppliers to gain ground by completing more security work across existing systems. An agent feature alone gives the enterprise little reason to reconsider a platform relationship.
- **Agentic security solutions can help integrate silo functions such as Extended Detection and Response (XDR), Security Information and Event Management (SIEM), and Security Orchestration Automation and Response (SOAR).**
- **Data economics increasingly affect security coverage.** Analysts and platforms need economic ways to easily access any type of evidence within a useful time and budget, otherwise autonomous workflows inherit the same blind spots as analysts. Querying completeness of data remains critical.
- **Agentic delegation needs an enforceable boundary.** Agent identities, limited credentials, and independent authorization could constrain actions even when a model follows a malicious instruction. Infrastructure assurance extends that requirement to the hardware and firmware supporting the boundary.
- **Risk reduction requires evidence after action.** A completed investigation or closed ticket does not establish that access was revoked, persistence was removed, or service was restored. Security leaders should require explicit verification and record residual uncertainty.
- **Specialists are entering through different enterprise constraints.** Key needs in the enterprise revolve around delivery of security operations, the cost of accessible evidence, and trust in the infrastructure that supports execution. These are distinct opportunities with different buyers and economics.

- **Companies mentioned in this report: 7AI, Cisco, CrowdStrike, Eclypsium, Exaforce, Formal, Google, Hinoki, Imply, Microsoft, NVIDIA, Palo Alto Networks, Stellar Cyber, Teleport, TENEX, Torq.**

STELLAR CYBER®

Agentic SecOps Platform

AI-Native. Human-Augmented. Full-Cycle. Integrated Security Operations Center (ISOC).

From vendor-agnostic telemetry to validated case verdicts, purpose-built for on-premises and cloud.

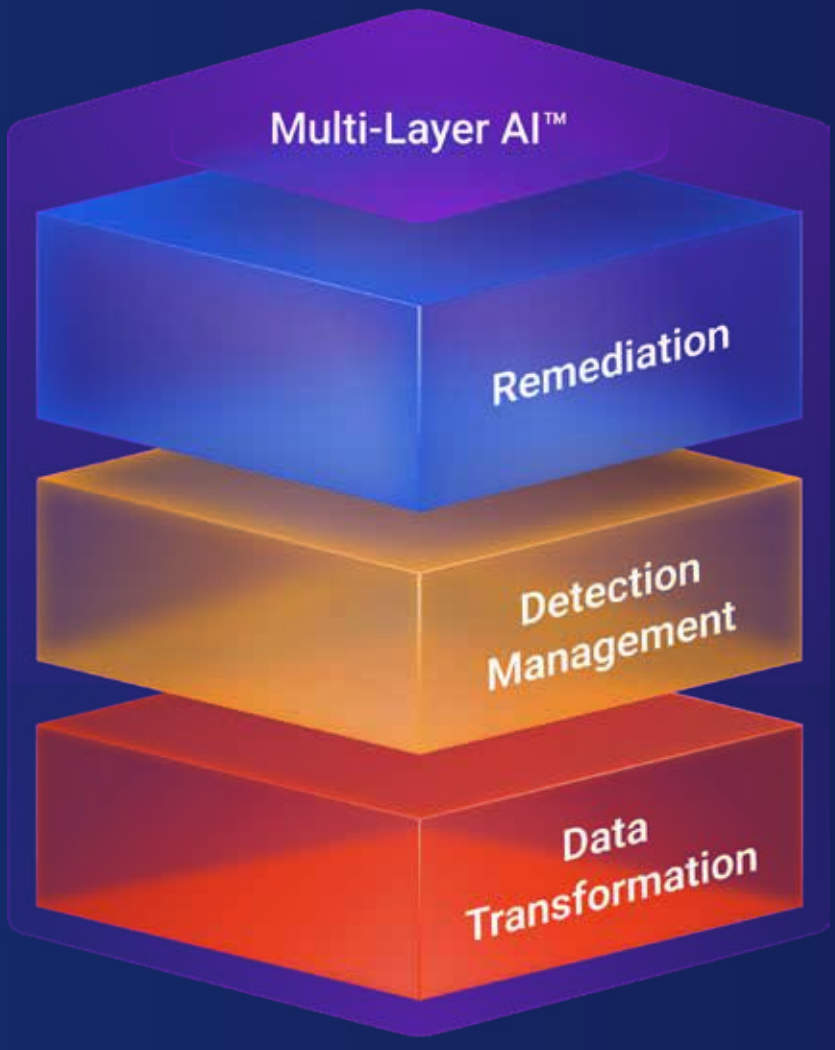


Table of Contents

1. Introduction 5

The Big Decisions 5
Our Definition of Agentic Autonomy 6
Market Development Around Existing Systems 7

2. Market Trends and Customer Challenges for Agentic Security 8

Trust and Identity Authority 8
Identity Governance: Making Delegation Traceable 9
The Autonomous SOC Needs Clear Operating Responsibilities 10
Tracking Where Data Is Stored 11
Supply-Chain Trust Extends Beyond Code 13
Investing for Successful Deployment 13

3. The SIEM Observability and Data Dilemma 15

Agents Change the Demands on Security Data 15
Building Rational Data Cost Models 15
The Splunk Question Spawns Solutions Such as Lumi 16
Preserving Operational Knowledge Already in the Platform 17

4. How Data and AI Will Change the Cybersecurity Market 18

An Opportunity to Measure Direct Results 18
Data Quality Shapes What an Agent Can Safely Do 18
Responding to Incidents 19
Explanations and Work Evidence Are Important 20
Infrastructure Assurance Extends the Market 20
Build vs. Buy Decisions Become More Complicated 21

5. Where Competition and Value Are Moving 22

Data Platforms Gain Leverage as Investigations Expand 22
Cisco and F5 Look to Connect the Enterprise Picture 23
XDR Incumbents Bring Installed Reach to the Agent Market 23
SOAR's Next Chapter Makes Agent Decisions Executable 24
Autonomous SOC Provides Managed Services Opportunities 24
Why the SOC Power Lives in the Operational Model 25

6. Conclusion and Enterprise Market Outlook 27

Investments that Shape the Transition to Agentic AI 27
Making Decisions That Matter 27

Appendix A: Additional Survey Research 29

Appendix B: Leading Vendor Profiles 30

Appendix C: Sources and Research Method 36

1. Introduction

Cybersecurity architectures are reaching a turning point in the evolution toward automation. With enterprises taxed by alert fatigue and a talent shortage, the only solution to building a more effective cyber response lies in automation, and of course AI.

Organizations would like to implement AI technologies to further automate their Security Operations Center (SOC). This can reduce response times, increase operational efficiency, and provide better protection of assets.

But it's not as simple as sprinkling AI pixie dust on the problem. Enterprise data and security architectures are complex and involve dozens of tools and data sources, if not more. Data management and costs are a significant consideration.

To understand the complexity, we can start with the acronyms. The big three—Extended Detection and Response (XDR), Security Information and Event Management (SIEM), and Security Orchestration Automation and Response (SOAR)—have evolved into silos that make the complexity of the solution part of the problem. Figuring out how to better integrate and automate these technologies is the core strategy of the automated, agentic SOC. In this report, we'll detail how agentic AI technology is influencing the modern SOC and what it means for these groups of legacy technologies. And we'll detail how agentic technologies can be used to reduce the complexity of IT environments and streamline cyber response.

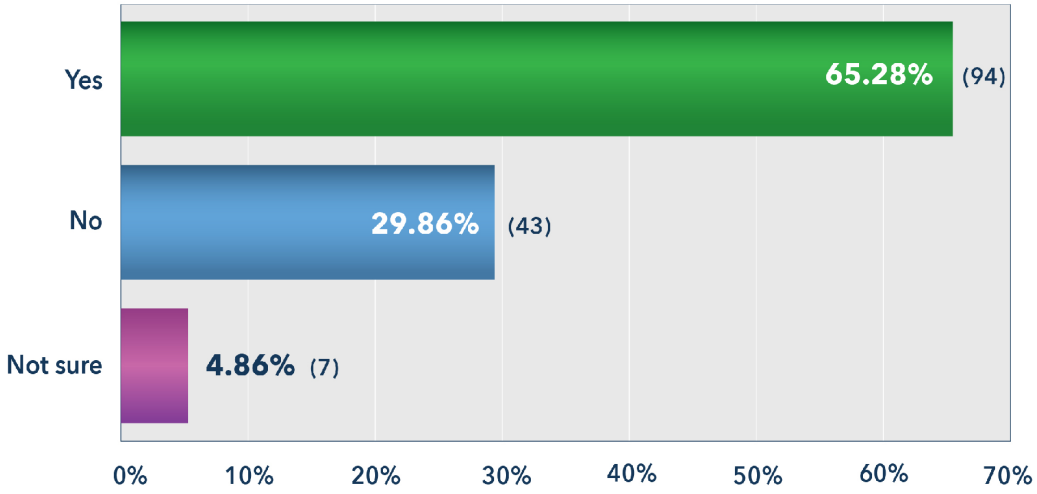
The Big Decisions

Enterprise security leaders face strategic decisions. They must determine how much security work to delegate to AI and how much authority to grant the AI systems being deployed elsewhere in the business. Those decisions increasingly impact production workflows. An agent investigating a suspicious account may access the same cloud services as an agent managing infrastructure. Both require reliable evidence and a controlled means of taking action.

Enterprise security investment will follow the efficacy of turning model capability into dependable operations. An enterprise can have capable models and still struggle to connect evidence, delegate authority, or establish the value of a response. Suppliers that resolve those constraints have a stronger proposition than those that simply add another agent to the console.

Some of our recent research in the SASE market demonstrates these constraints. Alongside the 65% reporting adoption friction to AI in our survey cited above, 60% selected consolidated data management and data sovereignty as a pressing AI security need. Agent and machine identity management followed closely at 59.03%. These multiple-response findings describe priorities within a networking and SASE audience, rather than purchase commitments. Our read from the field is that customers are concerned with the conditions for deployment as much as the capability of the agents themselves.

Are security concerns about enterprise infrastructure slowing down agentic AI adoption?



ANSWER CHOICES	RESPONSES	
Yes	**65.28%**	94
No	**29.86%**	43
Not sure	**4.86%**	7
TOTAL		**144**

FUTURIOM - 2026 Networking & SASE Survey N = **144**

For a large enterprise, replacing any stack is rarely the winning answer. Years of detection content and integrations represent real operating value. Realistically, we need evolution, not revolution. We expect customers to change the parts that constrain their internal AI team's efforts first while preserving the legacy parts that work. That gives specialists room to enter and forces incumbents to justify more of the architecture on its own merits.

Our Definition of Agentic Autonomy

We use agentic security to describe systems that can choose intermediate steps, retrieve evidence, and use tools within an assigned security task. A system that summarizes an alert is performing a narrower function than one that investigates related accounts and proposes containment. A system that executes containment has additional operational responsibilities. The word agent alone says little about those differences.

An autonomous SOC is best understood as an operating model with delegated execution. It may automate extensive parts of the investigation and response lifecycle while retaining human

authority over policy and consequential decisions. It does not imply an unstaffed SOC, unlimited production privileges, or a guarantee that every threat will be stopped. SIEM remains a distinct set of capabilities for collecting and analyzing security events.

We will discuss the consequences of delegation and how it creates customer challenges with security data architecture and changes in the supplier market. Infrastructure operations, data sovereignty, and supply chain trust are also part of the same enterprise problem. Agents cannot be governed only at the point where a prompt enters a model. Their access paths and execution environments matter throughout the task.

Market Development Around Existing Systems

Let's be real: Rip and replace is unrealistic. The majority of enterprises will adopt agentic security without replacing their existing stacks and uprooting their operations. AI SOC specialists want to apply agents to investigation and response while retaining human control over consequential decisions. Specialist data platforms can change storage and query economics around existing workflows. Others can address the integrity of the new AI infrastructure stack: hardware and firmware supporting AI workloads. These approaches enter through different weaknesses in the current environment, which is why the transition is unlikely to belong to one product category. [6][7][9]

The largest vendors are pursuing the opportunity. **Microsoft** supports security agents and agent identities, while **CrowdStrike** and **Palo Alto Networks** offer capabilities for building and coordinating agent workflows. Existing integrations and access to telemetry give them a useful starting position. We expect that position to translate into an advantage where it makes delegated work easier to operate across the customer's actual environment. Installed reach alone will not settle the competition. [12][13][14]

Companies with experienced engineering teams have another option: developing their own agents around tools they already operate. An agent can assemble account and endpoint evidence through APIs while an analyst reviews the case and asks follow-up questions. That changes where the team spends its time without immediately changing its underlying suppliers. This flexibility could make the interface a less secure source of competitive advantage and maintained operational expertise a more important one.

2. Market Trends and Customer Challenges for Agentic Security

Why do organizations need to adopt agentic AI? The simple answer is speed and efficacy. They need to implement agentic AI to speed up the response rate to threats, because their attackers are using the same technology to attack them.

On the organizational side, practitioners need to worry about more than the attackers. In addition to responding to threats using AI to accelerate reconnaissance and other parts of an investigation, they must also manage internal controls. For example, business and technology teams are introducing AI agents that can retrieve confidential records, execute code, and modify production systems. The security team has to respond to the first threat of attackers while helping the business move forward with the internal controls, often with the same people and tools already supporting day-to-day operations.

Based on existing data we have examined, there are several major challenges to enabling security team goals. Some of these key challenges include trust, identity governance, operational integration, and data integrity. Let's walk through some of the most important challenges.

Trust and Identity Authority

The key to enabling AI systems is to give them autonomy. To grant autonomy, you need trust. Organizationally, gaining trust is a major challenge to implementing agentic AI technology.

McKinsey recently conducted a study on AI trust and responsible AI (RAI) practices. In this study, two-thirds of respondents cited security and risk concerns as the top barrier to fully scaling agentic AI, well ahead of regulatory uncertainty or technical limitations. This suggests that organizations are less constrained by experimentation capabilities and more by a lack of confidence in their ability to safely deploy autonomous systems at scale.

Another conclusion of the McKinsey study was that trust is not a merely a tangential concern but a foundational requirement. In the same study, about one-third of organizations reported maturity of level three or higher in strategy, governance, and agentic AI governance. But the study concluded that some key areas remain open. These include:

- Who is accountable when an agent makes a bad call?
- Which actions need a human in the loop?
- How do these controls map to frameworks like the NIST AI Risk Management Framework?

Trust comes down to maintaining control over credentials and authority. When an attacker influences an authorized agent, a malicious instruction can be embedded in a document or tool to persuade the agent to misuse access it legitimately holds. For the practitioner, that complicates a familiar question: Is an authorized identity performing an authorized task? The credentials may be valid even when the action serves an attacker's purpose.

Anthropic's November 2025 account of a disrupted cyber espionage campaign shows why defenders should take this development seriously. The company reported activity against roughly 30 targets. It assessed that AI performed 80% to 90% of the campaign work. It also documented mistakes, including fabricated credentials and incorrect claims about stolen information. These findings describe one campaign rather than attacker productivity across the market, but they give security teams a concrete reason to test whether their detection and response procedures can keep pace with automated attack sequences. They also remind us that speed and reliability are different qualities. [2]

As mentioned, trust is not just about discovering unauthorized actions by attackers, but also about maintaining internal controls over agentic AI. Internal agents can also create incidents through ordinary mistakes, including choosing the wrong resource or repeating a change after a timeout. Open Worldwide Application Security Project's (OWASP) agentic application guidance captures why this market extends beyond malicious prompts. The commercial need includes governing legitimate access and controlling execution, areas where existing application and identity security capabilities remain relevant. Prompt filtering addresses only part of that need. [3]

Identity Governance: Making Delegation Traceable

If trust is a major obstacle to effective AI, account identity is key. For security teams, it can be difficult to manage service accounts whose ownership and purpose have become unclear. Agents add another layer to that problem. A single task can involve several tools and additional agents. When something goes wrong, responders need to follow the work back to the person (or machine) who authorized it and understand which credentials were used along the way.

Shared accounts and inherited user permissions make the history harder to reconstruct. This creates demand for identity systems that can preserve an agent's owner, assigned authority, and activity as work crosses tools. In our view, traceable delegation will become a condition for expanding production use, particularly where the organization must explain an incident to business leadership or regulators.

Revocation can conceal other weaknesses. For example, an agent may be disabled while sessions or delegated credentials remain active elsewhere. Suppliers that can show how authority is withdrawn across the workflow have a more credible governance proposition than those that control only the visible agents.

Microsoft's Security Copilot documentation describes dedicated identities for supported agents and an alternative using an existing user's credentials. These arrangements have different permission and audit consequences. They illustrate why "agent identity" is becoming a substantive architectural choice within established platforms. [12]

Teleport is a unique in its approach to targeting user and machine identity security, including agentic identity. Named to our Futuriom 50 list in consecutive years, Teleport provides identity-based security for managing infrastructure, cloud, and AI access using a unified platform and real-time chain of custody. Teleport's approach provides a unified identity layer that can issue and track credentials for any user, whether it's a human, machine, or agent.

The distinction between an agent proposing an action and the resource authorizing it is important. Enterprises need to use increasingly capable models while keeping authority within approved limits. We expect demand for enforcement that survives changes in the model and orchestration layer.

The Autonomous SOC Needs Clear Operating Responsibilities

For an analyst who spends much of the day gathering information from several systems, the appeal of agentic security is easy to understand. Retrieving account history, examining related endpoints, and assembling evidence can consume time that would be better spent assessing the threat. An agent that performs this work reliably can help an experienced team cover more ground and approach difficult cases with the relevant information already in hand.

The pace of delegation follows the consequence of a mistake. Evidence collection offers a more accessible starting point than changing identities across a production environment. This creates a gradual adoption path and helps explain why credible investigation products can gain traction before customers embrace broad autonomous response.

SOC vanity metrics may disguise problems. A shrinking alert queue can conceal incomplete work if missing evidence is treated as a reason to close a case. Providers that can distinguish among attempted investigations, completed investigations, and verified mitigation will give customers a clearer account of the service delivered. We expect this distinction to become more important as commercial claims move from productivity toward outcomes. AI SOC vendors are moving from talking about saved analyst hours and/or filtering positives to focusing on overall security outcomes such as faster containment and breaches prevented. Outcome claims only hold up if the provider can prove investigations were finished, not just closed.

As in many cases of production technology, effective implementation is as much about change management as it is about the technology itself. To adopt effective agentic AI, organizations will need to take a careful look at their operations teams and workflows.

Infrastructure teams are interested in agents for many of the same reasons as security teams: too much repetitive work and too little time for the problems that require experience. An agent may

help patch software, adjust a cluster, or update a firewall rule. Once it can make those changes, however, its work becomes part of the enterprise's security and availability responsibilities.

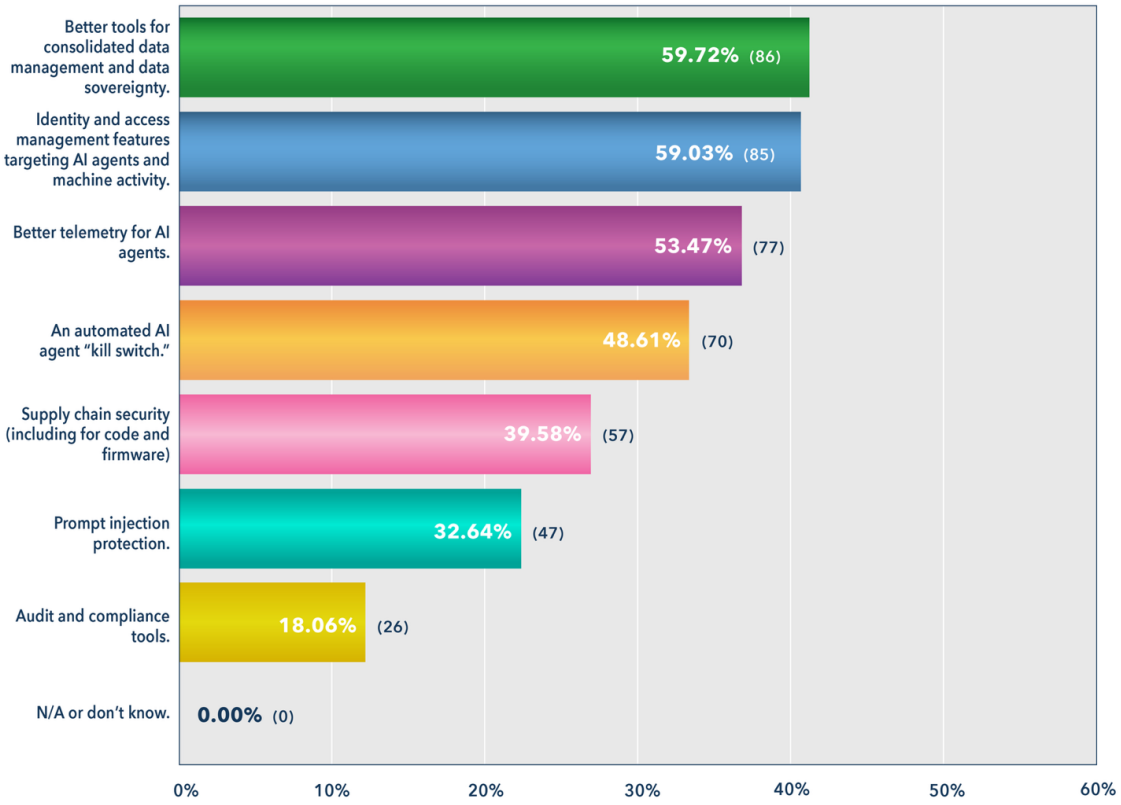
The complication comes when a technically valid action may be wrong for the business at that moment. A patch can interrupt a critical workload, while a firewall change can weaken an access restriction. Agentic operations make security and availability part of the same decision, creating demand for shared context and clearly allocated authority across teams.

Tracking Where Data Is Stored

Another operational issue is data sovereignty and how data is stored and accessed. Security leaders can carefully control where logs are stored and still lose track of where an agent sends information during an investigation. A query result may enter a model prompt, become part of a case summary, or pass to another service. Those records can contain sensitive identities and business activity even when the original events remain in an approved region.

Our recent survey of 144 security and IT professionals reflects this concern, with 60% of respondents saying there was a pressing need for better tools for consolidated data management and data sovereignty. Fifty-nine percent said their highest needs included identity and access management features targeting AI agents and machine activity, with 53% naming better telemetry for AI as a top need (respondents were asked to select all that apply).

What to you see as the most pressing needs in security functions for agentic AI? (check all that apply)



ANSWER CHOICES	RESPONSES	
Better tools for consolidated data management and data sovereignty.	**59.72%**	86
Identity and access management features targeting AI agents and machine activity.	**59.03%**	85
Better telemetry for AI agents.	**53.47%**	77
An automated AI agent "kill switch."	**48.61%**	70
Supply chain security (including for code and firmware)	**39.58%**	57
Prompt injection protection.	**32.64%**	47
Audit and compliance tools.	**18.06%**	26
N/A or don't know.	**0.00%**	0
TOTAL		**144**

FUTURIOM - 2026 Networking & SASE Survey N = **144**

This shows that data sovereignty can become a workflow issue as well as a storage decision. A regional log store cannot establish control over information subsequently sent to a model service elsewhere. We expect suppliers able to explain and constrain that movement to have an advantage in enterprises where data restrictions already influence architecture and procurement.

Supply-Chain Trust Extends Beyond Code

We also recognize that development teams are under even more pressure to deliver faster, and AI-assisted coding gives them another way to do so. That stress travels downstream: Security teams then must determine whether the resulting changes and dependencies are suitable for release. A package that appears to solve the immediate problem may still be malicious, unsupported, or inconsistent with the organization's requirements.

The NIST Secure Software Development Framework offers an established foundation for evaluating dependencies and controlling releases. Agentic development increases the value of a release process that can accommodate faster production without losing those checks. We expect suppliers to favor integrating controls into the development path, where they can influence what actually reaches production. [10]

Dependencies can extend beyond application code. Model artifacts and tool connectors become part of the workflow, and the underlying infrastructure brings its own firmware and component supply chain. These dependencies need different forms of scrutiny.

As discussed above, workflow changes mean that even more security controls are needed in systems controlled by agents. **NVIDIA**'s September 28, 2026, Open Agent Safety Platform announcement illustrates interest in enforcement outside the agent's execution domain. It combines OpenShell secure runtime software with the Sentry reference system design on BlueField-4 DPUs. Available software and a reference design represent different stages of readiness, but the architectural direction demonstrates that sensitive execution is drawing attention to controls in layers beneath the application. [8]

Independent enforcement also raises a question that infrastructure practitioners will recognize: How do we establish trust in the system enforcing the policy? Supply-chain security company **Eclypsium** describes checks on the integrity and security of hardware and firmware supporting AI systems. Other companies such as **Axiado** are pursuing similar problems with tools that provide secure root-of-trust in the underlying hardware in an AI system. We see infrastructure assurance as a complement to runtime controls because it examines the foundation on which those controls depend. Buyers should assess both the boundary around the agent and the controls for the infrastructure maintaining that boundary. [9]

Investing for Successful Deployment

So, how will investment in agentic security platforms and automated SOCs play out? The first deployments of agentic security operations are likely to help solve familiar bottlenecks: repetitive evidence collection, slow case preparation, and difficulty connecting the activity and communication across tools. This gives us a roadmap for adopting agents and judging their contribution to automating the SOC.

Identity ownership and access to evidence will also play large roles in how far many deployments progress. A demonstration can work on prepared data while the production workflow encounters missing permissions or incomplete records. This creates an opportunity for suppliers to take responsibility for those operating conditions alongside the agent itself.

3. The SIEM Observability and Data Dilemma

In AI, the data is key. Strong agentic AI solutions will depend on data quality. And the economics of data will be more important than ever before.

Security teams have spent years deciding which data they can afford to collect and keep searchable. Agents add a variable that ingestion-based budgets can miss: how much investigative work will be performed against that data. We expect this to reopen the economics of security data even where daily ingestion barely changes.

Agents Change the Demands on Security Data

For practitioners, data decisions represent an opportunity and a familiar concern. Agents could pursue questions that analysts rarely have time to examine today. But those investigations still depend on a data platform that can answer them without slowing ongoing detection or generating an unexpected bill.

This results in the SIEM, observability, and data dilemma. Customers want deeper investigation while containing the cost of evidence and retrieval. A cheaper archive can limit what an agent can examine in time; a fast query service can cover too little history. The market opportunity is to make more evidence usable at a cost that supports the work customers want to delegate.

A year of retained security data sounds reassuring until an investigation needs to use it. Some records may require restoration, a different tool or another team's intervention. The enterprise has paid to preserve the evidence, yet it cannot necessarily bring it into the decision when it matters.

Splunk's Enterprise documentation provides a useful example of how storage states affect access. Hot, warm, and cold buckets are searchable. Frozen data that has been archived must be restored to a thawed state before it can be searched. This is a specific storage lifecycle, so buyers should examine their own configuration rather than assume that all older Splunk data is inaccessible. The broader question applies across platforms: What needs to happen before retained evidence becomes usable? [4]

The commercial distinction comes when there is a negative finding and with an investigation that could not reach the right evidence. Platforms that point out that limitation give customers a more defensible result. Longer retention has limited operational value if its contents remain outside the normal investigation path.

Building Rational Data Cost Models

Let's talk more specifically about data costs. Security teams are accustomed to discussing data costs in terms of daily ingestion. That remains important, but it changes once agents begin conducting repeated searches and following investigative leads.

We expect pricing discussions to separate continuous detection from episodic investigation. The former may justify capacity that is always available; the latter can produce bursts of historical searches. Suppliers whose economics accommodate both patterns have a stronger case for supporting agentic operations.

Consider a planning scenario in which 1,000 daily cases each require an average of 20 query requests. That adds 20,000 requests a day before accounting for other searches. These are illustrative figures, but the calculation gives procurement a useful question: What happens to the proposed price when agents begin performing this amount of work?

Searches that return nothing still consume resources and can be necessary. If the commercial model makes those searches prohibitively expensive, customers may restrict the investigative coverage they hoped agents would expand. That is a meaningful limitation on the value proposition.

The Splunk Question Spawns Solutions Such as Lumi

It's no coincidence that **Splunk** (now owned by Cisco), one of the largest providers of SIEM, consistently returns to this conversation. Splunk became a leader in the data analytics space, but its cost model is now being questioned because it wasn't built with the agentic system in mind.

As we have pointed out, agentic systems can stress data cost models. For an enterprise with years invested in Splunk, reducing data costs might require a major migration or a reconsideration of the data architecture. Splunk is certainly aware of these changes and we expect adjustments to its commercial model going forward.

This provides opportunities to solve the problem with new data approaches. Some companies argue that a specialist data layer offers another route: changing the evidence economics while retaining the workflows.

Startup **Imply** describes its product Lumi as a data layer with adaptive schema handling, compressed event storage, and a query engine based on its distribution of Apache Druid. It integrates with existing tools, including Splunk. Lumi Loglake also has the capability to search unstructured logs directly in data lakes, without restoring them first. This is one of Lumi's biggest differentiators, giving agents access to more data without incurring the additional cost. Customers can change how the underlying data is stored and queried while retaining familiar investigation workflows. Detection and response still require the systems and procedures that perform those functions. [5]

Imply recently published a case study with a Brazilian Bank, BTG Pactual, which reports more than 60% lower Splunk-related observability costs, over 10 terabytes per day of additional queryable security data, and retention extended from 90 to 365 days. The bank retained its Splunk workflows. These are vendor-published results from a particular deployment, so they should be used as

evidence of an approach worth examining rather than an expected savings rate for every customer. [6]

For security leadership, the interesting aspect is what the bank gained alongside the reported savings: more evidence available for investigation and a longer searchable history without abandoning established workflows. That connects the cost discussion to coverage and gives the operational team a reason to participate in an evaluation that might otherwise be driven principally by finance.

The savings proposition depends on the customer's data and operational content. Compression, custom query compatibility, and permissions can all affect the result. A lower storage quote therefore needs to survive the cost of preserving the existing operation.

Preserving Operational Knowledge Already in the Platform

An established SIEM contains more than data and dashboards. Its detections, enrichment logic, and schedules often reflect years of work by people who understand the enterprise. A proposal to change the architecture should account for that investment.

Replacing the analyst's interface, moving the data layer, and replacing the SIEM are different decisions. They may eventually form part of the same program, but treating them as inseparable can turn a focused improvement into an unnecessarily large project. A new engine's ability to understand a query language does not establish that every saved search or detection dependency will behave as expected.

We expect selective changes to be more attractive than broad migrations where the existing operational content remains valuable. A different evidence layer can address retention and retrieval costs while effective detections stay in place. That leaves customers with more room to improve economics without accepting every disruption of a platform replacement.

The useful measure is time to a defensible answer across the relevant history, including retrieval failures. It connects the architecture to the security work. Suppliers that improve that result at a sustainable cost have a stronger claim on the budget than those offering query speed in isolation.

4. How Data and AI Will Change the Cybersecurity Market

Security leaders have long faced a frustrating question: After all the spending on tools and services, how much safer is the business? Vanity metrics are fun, but they may not offer the correct measure of business results. Teams can show alerts investigated, vulnerabilities identified, and cases closed. But connecting that activity to a reduction in risk is harder. An increase in detections might reflect better visibility or a worsening threat environment. A smaller queue might mean faster resolution, but it could also conceal incomplete investigations. Practitioners are left trying to explain the value of substantial investments through measures that tell only part of the story.

An Opportunity to Measure Direct Results

Agentic security creates an opportunity to reflect better on the results. An agent can gather evidence across products, investigate the affected systems, and carry out a response within approved limits. The important step comes afterward: checking whether that response achieved its purpose. Was the compromised account's access actually removed? Was the affected endpoint isolated? What exposure remains? We see this as one of the more meaningful promises of agentic operations. Security teams could gain a clearer connection between the work performed and the protection delivered, provided verification is built into the workflow.

This puts pressure on suppliers to provide better information about the handoffs among detection, investigation, and response. Customers are often responsible for the work to connect capabilities from different vendors. Agentic operations create an opportunity for the suppliers to take on more of that responsibility. When a provider can show both a completed action and evidence of its effect, its commercial value rises.

Based on Menlo Ventures' recent analysis of agents in security and Insight Partners' study, management and control planes offer useful ways to think about this transition. They describe a market in which software takes on more of the work that analysts have traditionally coordinated across tools. Enterprise value will depend on how well those pieces operate together and whether the team can verify the outcome. For practitioners who have spent years defending budgets with activity metrics, that would be a welcome change. It gives them a more useful question to bring to the next purchase: Can this investment help us demonstrate that a threat was addressed and explain the risk that remains? [11][15]

Data Quality Shapes What an Agent Can Safely Do

Security teams know how much judgment goes into interpreting an alert. A login that looks unusual may be legitimate. An apparently idle server may support a critical process that runs only at month end. Experienced analysts check those details before deciding what to do. As agents take on more

of that work, they need access to the evidence and business context that inform those judgments. Missing information can affect both the investigation and the action that follows.

Insight Partners' framework offers a useful way to organize those responsibilities. The data plane preserves events and makes evidence available under the appropriate access rules. The management plane coordinates the investigation and carries its context from one step to the next. The control plane enforces the resulting decisions through identity, endpoint, and other operational systems. For buyers, the framework helps reveal where a supplier's responsibility begins and ends, including the handoffs the customer will need to manage. [15]

The framework also clarifies the competition between broader platforms and specialists. A platform may simplify coordination, while specialists can offer more choice at an individual layer. The economic question is how much work remains with the customer. We expect integration quality to influence the value of either model as agents cross more systems.

Context also depreciates. Accounts change roles and applications acquire new dependencies. An agent using an outdated ownership record may act consistently with the data and still make the wrong business decision. Maintaining those relationships can become a durable service advantage, particularly as the ability to generate an attractive context graph becomes easier to reproduce.

Responding to Incidents

During an incident, a security team needs to know when it can safely move on. That is harder than it sounds, especially when a response spans several systems. An agent may report that it disabled a compromised account because the identity platform accepted the request. The next question is whether the attacker still has access through an existing session. Automating the action saves time, but the investigation remains open until the team has enough evidence to understand what that action accomplished.

Verification gives a response its operational meaning. An identity system can accept a request while an attacker retains an active session. A device can be isolated while persistence remains unresolved. Suppliers that establish the resulting state help the customer distinguish a completed action from an addressed threat.

Some incidents will leave residual exposure or require a business exception because immediate containment would interrupt an essential service. Those outcomes belong in the account of the operation. A provider that makes them visible gives leadership a more useful picture than one that treats every closed case as equivalent.

This is an area where human-assisted automated SOC can help. For example, **TENEX** connects investigation and response with enterprise context, auditable records, and customer policy. It distinguishes approved actions from consequential or uncertain decisions requiring human experts. That is relevant to an enterprise market where customers want more work completed without surrendering judgment over business consequences. [7]

Explanations and Work Evidence Are Important

One of the most persuasive things about an AI investigation is also one of the easiest to overvalue: a clear, confident explanation. It can turn a confusing collection of alerts into a story that appears to make sense. For a busy security team, that is immediately appealing. The difficulty comes when someone asks what supports the conclusion. Did the agent examine the relevant records? Did it encounter conflicting evidence? Was an important source unavailable? A polished summary can leave all those questions unanswered.

A stronger market proposition is making it easier to inspect an investigation. The analyst can follow the searches performed, examine the events behind the finding, and understand which procedure the agent followed. The summary then becomes a useful entry point into the evidence. This matters beyond checking the agent's accuracy. It allows another analyst to pick up the case, helps an incident commander understand an unresolved issue, and gives the organization a record it can revisit when new information emerges.

That distinction becomes more valuable in the less tidy cases that occupy so much of a security team's time. An unusual login may have a legitimate explanation, while a seemingly routine event may acquire significance only when connected to activity elsewhere. The useful agent is one that can pursue those connections and recognize when the available evidence leaves a question open. We expect suppliers to face greater scrutiny over this behavior as customers move beyond demonstrations and begin relying on agents during ordinary shifts. Fluent explanations may attract attention, but handling uncertainty credibly will matter more to sustained adoption.

This also changes the role of human oversight. Security leaders are looking for relief from repetitive work, so an operating model that sends every step back for approval limits progress. Yet a stream of approval requests with little supporting evidence can become another queue to clear. The more promising model gives agents room to complete well-defined work and brings people into the cases where judgment can change the outcome. When an exception arrives with the relevant evidence and a clear account of what remains unresolved, the analyst can spend time deciding what to do next. That is where automation has a better chance of making scarce expertise go further.

Infrastructure Assurance Extends the Market

An agent's permissions may be tightly controlled, but those controls still depend on infrastructure the enterprise can trust. Firmware and management controllers can affect a host or its telemetry. Moving enforcement to a separate component brings that component's software and supply chain into the security decision as well.

Eclypsium's September 2026 publication describes infrastructure assurance alongside NVIDIA's agent safety architecture. Its documented scope includes hardware and firmware inventory, vulnerability identification, and integrity checks. The enterprise implication is that evidence about the enforcement foundation should inform deployment and incident response. An infrastructure

alert should reach an accountable team with an agreed response, rather than become an isolated compliance record. [9]

We expect infrastructure assurance to matter most where enterprises operate substantial AI infrastructure or control the hardware enforcing sensitive workloads. Managed services shift some of that responsibility to the provider. As agent authority expands, evidence about the underlying foundation becomes more relevant to the enterprise's confidence in the whole operation.

Build vs. Buy Decisions Become More Complicated

For a technology team with strong engineering capabilities, there's now an exciting question: How much of the agentic system could we build ourselves? With capable models and accessible APIs, a team can create an agent that gathers evidence across its existing tools, answers follow-up questions, and coordinates an approved response. That opens a more manageable set of choices than the traditional decision to build or buy an entire security platform. The enterprise might develop a workflow tailored to its business while continuing to rely on suppliers for threat research, detection, and enforcement. There is real promise in that flexibility, particularly where established products leave analysts doing repetitive work between systems. The prototype is the beginning, though. Keeping it dependable as APIs, permissions, and investigation procedures change becomes part of the security program.

Specialists can retain value through capabilities that are difficult to sustain internally, including maintained investigation procedures, high-volume data services, and detailed infrastructure knowledge. Easier prototyping makes that distinction more important. The supplier is competing with the cost of operating the capability over time, rather than the cost of producing its first version.

We expect this flexibility to favor suppliers whose evidence and actions remain accessible as customers change their orchestration systems. A convenient interface can become restrictive when it is the only route to the underlying capability. Openness has commercial value where permissions and accountability remain intact across the connection.

5. Where Competition and Value Are Moving

The agentic security market is taking shape around a question that product categories alone cannot answer: Who will take responsibility for getting the work done? Data platforms, established security vendors, and newer operations providers approach that question from different starting points. Each has capabilities that could become more valuable as enterprises delegate investigation and response to software.

Our analysis focuses on where suppliers can take responsibility that customers currently carry themselves. Established platforms can extend their reach, while specialists can resolve constraints in data or infrastructure. We expect the strongest positions to emerge where a supplier reduces operating effort and can substantiate the result.

Data Platforms Gain Leverage as Investigations Expand

For years, security teams have had to decide how much evidence they could afford to keep readily available. Agentic investigation makes that compromise more consequential. An agent can pursue additional leads, search further back in time, and connect activity across more systems, provided the data remains accessible. The opportunity for data suppliers grows with the ambition of the investigation.

Observability platforms maintain an important role because application and infrastructure records can help explain what happened around a security event; serving both operations and security from shared infrastructure can also improve the economics. The challenge is accommodating their different requirements, including access to sensitive records, preservation of event detail, and dependable performance during an incident. Therefore, we're seeing a trend toward integration of observability platforms, whether you are looking at **Cisco**'s acquisition of **Splunk** or **Palo Alto Networks'** $3.4-billion purchase of **Chronosphere** in 2025, which has yielded observability platform Cortex XCOR.

The next generation of companies is looking at how observability can specifically enable agentic security, with lower costs and better performance. Startup **Imply** is taking the specialist approach to observability discussed in section 3 above: changing the economics and accessibility of the data layer while retaining established Splunk workflows. Other suppliers approach the problem through data lakes, search engines, or broader platforms. We expect competition to become more focused on the cost of usable investigative history. Storage pricing matters, but the larger question is what the enterprise can afford to examine when it needs an answer.

This could give customers more negotiating room. A company may value its existing SIEM detections while seeking a different way to store and query the evidence beneath them. Separating those purchases creates alternatives to a bundled renewal, although the integration and operating costs still belong in the business case.

Futuriom research helps explain why this discussion should extend beyond savings. In our recent SASE survey, improved security visibility and analytics was selected by 69% of respondents as a consolidation benefit in SASE platforms, while reduced operational complexity was selected by 58% and lower total cost of ownership by 41%. Within this audience, the appeal of consolidation includes making security easier to operate and understand. Our interpretation is that suppliers will need to connect their economic story to achieve better access to evidence and less work for the team. [1]

Cisco and F5 Look to Connect the Enterprise Picture

Because they have access to large amounts of enterprise infrastructure and data, large infrastructure incumbents are addressing ways to enable agentic security operations with integrations of their own platforms.

As we mentioned, there's significant discussion in the industry about Splunk. **Cisco's** acquisition of Splunk, completed on March 18, 2024, brought networking, security, and observability capabilities into the same portfolio. For customers trying to understand activity across a large environment, that combination has an obvious appeal. Network behavior can provide context that an endpoint alert lacks, while application and infrastructure records can help explain the business consequences of an incident. [16]

The value will emerge through the connections customers can use. If the combined architecture reduces the effort required to assemble evidence across domains, it could give analysts and agents a richer starting point for investigation. Cisco's reach across enterprise infrastructure and Splunk's established role in security operations make this an important development to watch.

Large customers will still operate heterogeneous environments. The opportunity is therefore broader than integrating products within the portfolio. It includes making those connections useful alongside other endpoint, cloud, and identity systems. The quality of that experience, together with the economics of retaining and querying more evidence, will shape how much of the potential customers realize.

In another example, **F5** has absorbed targeted startups to construct an automated discovery, testing, and remediation lifecycle within the SOC. This includes Fletch, integrated into F5's Application Delivery and Security Platform to leverage agentic natural language processing (NLP) to turn raw security logs into prioritized, real-time threat intelligence. F5 has also acquired SurePath AI for network visibility and shadow AI detection. F5 has also introduced an updated, agentic-ready AI Gateway embedded directly into the F5 AI Security Platform to establish control points.

XDR Incumbents Bring Installed Reach to the Agent Market

Established security vendors enter this transition with something newer companies must work to obtain: a place in the customer's operating environment. Extended detection and response, or XDR,

platforms already connect telemetry with controls across several security domains. That gives them a practical foundation for adding agents to investigation and response.

CrowdStrike's Charlotte AI AgentWorks ecosystem, **Palo Alto Networks'** Cortex AgentiX, and **Microsoft**'s Security Copilot agents show how actively large suppliers are pursuing the opportunity. We also closely follow **Stellar Cyber**, an XDR startup that pioneered the open SIEM approach (more on that later). Each approach differs, but each can build on existing integrations, product knowledge, and customer relationships. Cisco is also adding agentic capabilities with its Cloud Control Studio. [13][14][12]

We expect agentic advantages to matter most when they shorten the distance between identifying a problem and taking an authorized action. Customers may be more willing to delegate a task when the supplier already operates the relevant control and understands its behavior. At the same time, investigations often require evidence beyond any one vendor's coverage. The ability to work across that boundary will influence how broadly these agents can be used.

An agent builder and a supported operational workflow also represent different offers. One gives a customer the means to develop its own automation. The other takes on more responsibility for maintaining the procedure. That distinction will increasingly shape both purchasing decisions and the services surrounding these platforms.

SOAR's Next Chapter Makes Agent Decisions Executable

Agentic systems bring fresh attention to an old difficulty in security automation: Deciding what to do is only useful if the action can be carried out reliably. Security orchestration, automation, and response, or SOAR, has supplied much of the connective machinery for that work. Its integrations and procedures remain relevant as agents introduce more flexible ways to investigate.

The traditional SOAR market is being impacted as the landscape shifts toward agentic SOCs. Classic SOAR is no longer dynamic enough to adapt to the constantly shifting data and connectivity needs of agentic security, where a broken API can disconnect the whole system. Many SOAR technologies are being adapted or merged into SIEM and XDR platforms.

Established automation content is the asset. Organizations have already invested in procedures that reflect their tools, escalation paths, and operating constraints. Suppliers that can make those procedures available to agents, with controlled execution and clear recovery behavior, have a credible way to connect existing investments with new capabilities.

New companies are emerging with this type of approach. For example, **Torq** is a next-gen SOAR platform that allows customers to build their own agents.

Autonomous SOC Provides Managed Services Opportunities

"Autonomous SOC" has become a broad market label for offerings that improve work rates and workflows. Some companies supply software to an existing security team. Others combine

automation with analysts and managed operations. Some extend into detection engineering or response. Similar demonstrations can therefore conceal substantial differences in what the customer is buying.

In our broader observations, nobody expects the humans to go away. Agentic technology for the SOC is about making teams more effective and putting people in a better position to monitor and manage the operations. A customer seeking faster investigations may want software its own analysts can direct. A customer struggling to sustain coverage may place greater value on a provider that takes responsibility for operating the queue and handling exceptions. Both can benefit from agents, but their purchasing priorities differ.

We think this provides natural opportunities for companies that can augment managed security services, which represent a large piece of the enterprise security market. As an example, startup **TENEX** is relevant to the managed, human-led operations model. Its public positioning includes integration with **Google** SecOps and **Microsoft** Sentinel, placing it alongside existing security investments. Its proposition centers on coordinating work across the security lifecycle, with automation and human expertise operating within customer policy. [17]

Additional companies are positioning themselves as effective tools for the human-lead autonomous SOC. **Stellar Cyber,** which emerged out of the XDR side of the market, helps enable mid-market teams without huge budgets to build human-led autonomous SOCs. It has added agentic capabilities to a broad, flexible platform that includes choice of SIEM and data telemetry integration tools.

The area of autonomous SOC also provides opportunities to enable technology that uses agentic AI to speed up discovery and investigations. Additional companies to watch in this space include **7AI**, **Exaforce**, and **Torq**. **7AI** provides SIEM integration, investigative agents, and response—with the flexibility to bring your own SIEM and/or data lake. It has raised a total of $170 million. **Exaforce**, which this year raised a $125 million Series B ($200 million total), offers a security data platform combined with a suite of agents from detection through response. It can replace the SIEM or work alongside it. The agents use a realtime knowledge graph built by gathering a broad range of telemetry and providing a deep reasoning layer that includes deep context to the investigative process.

The autonomous SOC label is attracting attention, but it's a broad category with many variations based on the operating model. Details come down to the customers and how they expect results to be delivered, as well as the business model. This varies based on which work the provider owns, how uncertain cases are handled, and where the customer remains responsible. Data charges and storage models also vary. These points become key differentiators as the market matures.

Why the SOC Power Lies in the Operational Model

In autonomous SOC, the operational expertise can become an advantage. Investigation procedures need maintenance. Model updates need evaluation. Changes in permissions or available evidence

can alter the reliability of a previously successful workflow. A supplier that consistently manages those changes is providing value that is easy to underestimate during the initial demonstration.

We therefore see a stronger long-term proposition in the ability to sustain dependable operations than in access to a particular model alone. The same applies to specialist data and infrastructure providers. Their products become more valuable when they continue supplying reliable evidence to whichever orchestration systems the enterprise adopts next.

The best opportunities begin with solving the top constraints recognized by security leaders. Teams struggle to assemble evidence, retain enough searchable history, and establish confidence in the systems supporting sensitive operations. Agentic adoption gives those problems fresh urgency.

Operations specialists, independent data platforms, and infrastructure assurance providers address different constraints. Their opportunities will depend on a defined contribution to the wider security operation. Comparing them as interchangeable agent products doesn't accurately assess the full value.

6. Conclusion and Enterprise Market Outlook

Agentic AI will change what enterprises expect from security suppliers. Products have traditionally helped teams identify threats, assemble evidence, and coordinate a response. As software takes on more of that work, suppliers will face a harder question: How much of the operation can they reliably complete, and what responsibility will they accept for the result?

This market shift is consequential. Security leaders have spent years rationalizing substantial investment through activity measures that only partly justify the protection that is delivered. Agentic operations offer a chance to make that connection clearer. The commercial value will depend on whether suppliers can turn that promise into dependable work with an outcome the customer can verify.

Investments That Shape the Transition to Agentic AI

The focus on investigation may advance faster than autonomous authority over production systems. Gathering evidence and preparing a case can relieve an immediate burden without requiring the enterprise to delegate every response decision. Broader authority will develop unevenly because the consequences of an error differ sharply across organizations and workloads.

This means that most organizations will favor systems they already depend on. Established vendors bring integrations, telemetry, and access to operational controls. Newer suppliers can address gaps that those platforms have left unresolved. Neither position guarantees success. The advantage will come from reducing the effort required to complete security work while preserving the enterprise's ability to understand and govern it.

This SIEM is likely to remain in a central role even as analysts spend less time in its interface. It also means that specialist data layers may become more important as agents generate additional queries and demand longer investigative history.

In addition, Infrastructure assurance will gain relevance as software receives greater authority. These changes redistribute value across the architecture without requiring every enterprise to replace it.

With this massive architectural shift at hand, specialists can build durable positions when they solve a difficult problem and fit into the wider operation. At the same time, broader platforms have an opportunity to make those connections easier.

Making Decisions That Matter

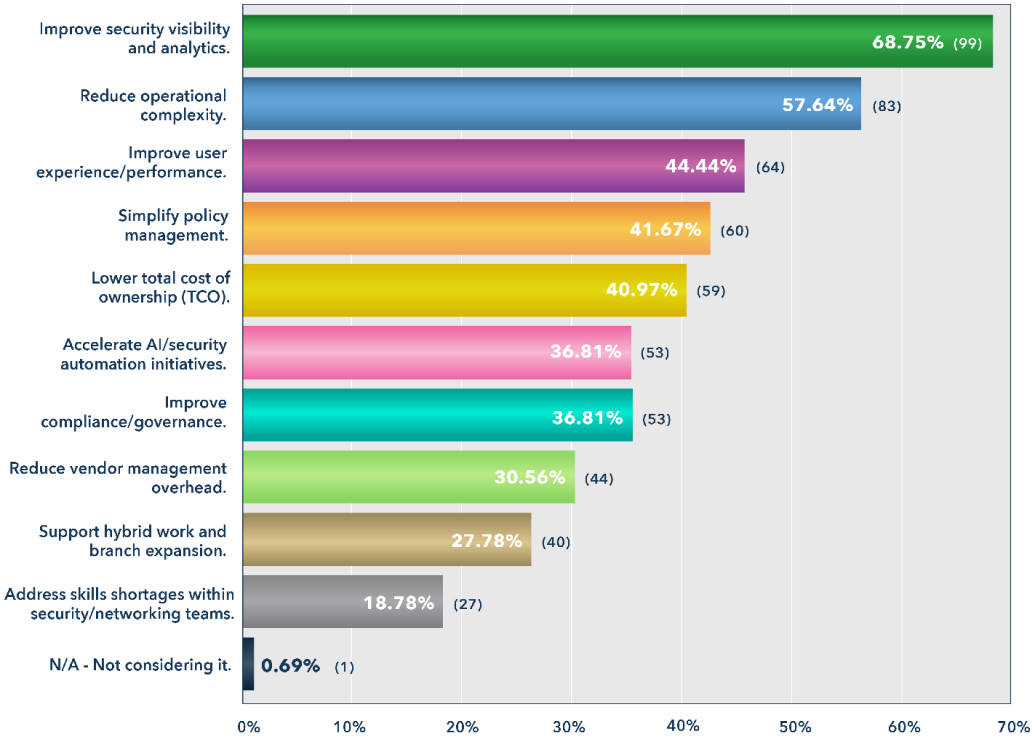
Security leaders have a stronger negotiating position when they distinguish access to an agent from responsibility for the operation. The former is becoming easier to offer. The latter depends on reliable evidence and execution in environments that rarely fit a prepared demonstration.

Our bet is that durable value will follow suppliers that take more of the operating burden off the customer with traceable results. Faster analysis is always welcome. A larger opportunity comes when the completed security work can be tied to direct business outcomes.

Appendix A: Additional Survey Research

Futuriom 2026 Networking and SASE Survey results. Each question has 144 responses. Multiple selections were permitted for security needs and consolidation benefits. The supplied graphics have not been independently audited. [1]

What do you perceive as the key benefits in consolidating security and networking tools onto fewer platforms? (select all that apply)



ANSWER CHOICES	RESPONSES	
Improve security visibility and analytics.	**68.75%**	99
Reduce operational complexity.	**57.64%**	83
Improve user experience/performance.	**44.44%**	64
Simplify policy management.	**41.67%**	60
Lower total cost of ownership (TCO).	**40.97%**	59
Accelerate AI/security automation initiatives.	**36.81%**	53
Improve compliance/governance.	**36.81%**	53
Reduce vendor management overhead.	**30.56%**	44
Support hybrid work and branch expansion.	**27.78%**	40
Address skills shortages within security/networking teams.	**18.78%**	27
N/A - Not considering it.	**0.69%**	1
TOTAL		**144**

FUTURIOM - 2026 Networking & SASE Survey

N = **144**

Appendix B: Companies to Watch

7AI

7AI delivers AI that Defends. Its Foundational AI Security platform puts AI at the core of security operations, with people directing the defense. Working across the tools and data an organization already has, 7AI agents detect, investigate, hunt, respond, improve detections, and report in a self-improving loop so protection compounds faster than threats evolve. Built for complex security environments, 7AI enables federal agencies and the systems integrators that support them to put autonomous AI agents to work with transparency, oversight, and human control. Proven in production, with more than 14 million security alerts processed and over 2 million analyst hours saved, 7AI frees security teams to focus on proactive defense, strategic decisions, and the work that requires human judgment. Founded in 2024 by Lior Div and Yonatan Striem-Amit, 7AI is headquartered in Boston.

Total funding: $166 million

Airrived

Airrived is an enterprise AI company pioneering autonomous operations through its Agentic AI Operating System (Agentic OS). The platform enables organizations to build, deploy, orchestrate, and govern AI agents and workflows across cybersecurity, IT, and business operations. With over 300 agentic applications, agents, and workflows deployed in production, Airrived demonstrates enterprise-scale adoption. Its model-agnostic architecture supports SaaS, private cloud, on-premises, and air-gapped deployments, ensuring security, governance, and data sovereignty. Recognized as an IDC Leader in AI SOC Platforms, a Gartner Cool Vendor, and a two-time Gartner Tech Innovator, Airrived is accelerating the transition toward autonomous enterprises.

Total funding: $6.5 million

Cisco

Cisco is building security for AI agents across identity, code and runtime. At RSAC in the spring of 2026, Cisco announced DefenseClaw, an “open source secure agent framework,” and also extended Zero Trust Access to AI agents through Cisco Duo IAM, which lets managers assign time-bound, task-specific permissions to agents. With Cisco Identity Intelligence and DefenseClaw, Cisco actively scans identity, actions, MCP servers, and assets. It also unveiled an Agent Runtime SDK that

embeds policy enforcement into agent workflows at build time, supporting AWS Bedrock AgentCore, Google Vertex Agent Builder, Azure AI Foundry, LangChain, and others.

Nasdaq: CSCO

CrowdStrike

CrowdStrike has spent the past six months positioning itself as the security layer for corporate AI agents. Most of the announcements came at its Fal.Con conference in September: Falcon Guardian, which monitors and blocks agent activity on endpoints; Agentic IdP, which gives each agent its own short-lived credentials; SafeMind, a pair of attack-and-defense models built on NVIDIA's Nemotron; and more automation in its security operations tools. A June identity product, based on its SGNL acquisition, laid the groundwork. In October it extended its NVIDIA partnership toward hardware-level enforcement. Many features are still only planned, and the performance claims are the company's own.

Nasdaq: CRWD

Eclypsium

Eclypsium, founded in 2017 and led by CEO Yuriy Bulygin and CTO Alex Bazhaniuk, is a supply chain security company. It builds trust in endpoints, servers, and network appliances across enterprise IT, cloud, data center, and network environments by identifying, verifying, and hardening the third-party software, firmware, and hardware inside each device. The company helps enterprises and government agencies reduce risk from complex technology supply chains, using an AI-powered platform to validate IT components, detect hidden threats, and respond with automated updates

Total funding: $100 million+

Exaforce

Exaforce is an agentic SOC platform that replaces or works alongside existing SIEM and SOAR tools. MDR is an optional service on top of the platform. High-growth companies use it to stand up an enterprise-grade SOC in days; larger enterprises cut total SOC costs by 60% or more. Agents cover all aspects of the SOC from Detection through Response. The agents use a real-time knowledge graph built from log, configuration, identity, device and resource data. Customers report 95% fewer false positives. Invisible Technologies relies on Exaforce for SIEM and 24/7 MDR, and Guardant Health uses it to widen coverage and speed up investigations. Founded in 2023 by serial entrepreneurs from Google, Palo Alto Networks and F5, Exaforce has raised $200 million in capital.

Total funding: $200 million

Formal

Formal is an AI security company that enables enterprises to give AI agents access to sensitive data and systems without losing control. Its network-level security platform enforces least-privilege access, allowing companies to control exactly what agents and employees can access and do across databases, applications and cloud infrastructure. Formal inspects every request, blocks unauthorized actions and provides visibility into agent activity, helping enterprises securely deploy AI agents across their operations. Headquartered in San Francisco, Formal secures billions of requests daily for leading technology companies including Cursor, Notion and Decagon.

Total funding: n/a

Google (Alphabet Inc.)

Google's Agentic SOC technology, augmented by its acquisition of Wiz, delivers autonomous "Agentic Defense" within Google Security Operations. Powered by Gemini AI models, the platform replaces playbooks with adaptive AI agents—including Triage, Threat Hunting, and Detection Engineering agents. These agents analyze threat profiles and reduce alert triage times from 30 minutes to 60 seconds, according to Google. By integrating Wiz's Security Graph and Attack Surface Management (ASM) with Google Threat Intelligence, the unified platform enriches telemetry in real time. The goal is to enable SOC teams to visualize complex cloud-native attack paths and prioritize exploitable vulnerabilities seamlessly across multi-cloud environments while retaining critical human control.

Nasdaq: GOOG

Hinoki

Hinoki is developing products in an emerging market for specialized AI intelligence. Its flagship model, Hinoki Offensive One (HO1), leverages more than 300,000 reinforcement learning environments to deliver advanced cybersecurity capabilities spanning threat hunting, malware analysis, and exploit development. Alongside specialized models like TypeSafe's Jev, Hinoki signals a broader industry shift toward purpose-built AI models optimized for specific domains. Founded by Johnny Wang, a former Google Offensive Security engineer and repeat cybersecurity founder, Hinoki equips enterprise security teams and cybersecurity vendors to validate vulnerabilities and conduct security testing, potentially challenging general-purpose frontier models and reshaping enterprise AI economics. Hinoki has a seed investment from Andreessen Horowitz.

Total funding: n/a (seed funding)

Imply

Imply is the data layer for observability, security, and AI. Imply Lumi is the data platform behind the agentic SIEM, giving security teams and AI agents access to more data and longer history. Because Lumi separates storage and compute, that history stays searchable at an infrastructure cost organizations can control. Teams use SPL and SQL to search indexed data and query unstructured logs directly in object storage. Lumi runs beneath existing security tools and emerging AI agents, so analysts can investigate threats and answer new questions without replacing their SIEM or changing how they work every day.

Total funding: $215 million

https://imply.io/

Microsoft

Microsoft's agentic SOC strategy is anchored by the integrated security operations center (ISOC) in Microsoft Defender. It delivers automated, predictive defense by combining SIEM and threat protection. Powered by Microsoft Security Copilot, the architecture utilizes a six-layer autonomous stack—spanning signals, context, models, and specialized agents to continuously execute protective actions. Autonomous agents handle complex correlation, automating roughly 75% of phishing and malware investigations. By using native telemetry and the Microsoft Security Store, specialized threat hunting and triage agents can turn natural language queries into automated hunting sessions to thwart live attacks.

Nasdaq: MSFT

NVIDIA

NVIDIA has introduced the Open Agent Safety Platform, a full-stack governance framework designed to mitigate risks associated with autonomous, action-taking AI. Shifting security boundaries to the system and silicon level, the two-layer architecture isolates and monitors agents in real time. NVIDIA OpenShell provides a secure, kernel-level software sandbox on Vera CPUs for zero-trust policy enforcement, while NVIDIA Sentry acts as an out-of-band hardware watchdog running on isolated BlueField-4 DPUs to quarantine non-compliant models. Broadly

adopted by enterprise cybersecurity giants like Palo Alto Networks and CrowdStrike, the platform establishes critical, automated physical boundaries for enterprise AI deployment.

Nasdaq: NVDA

Palo Alto Networks

Palo Alto Networks secures the autonomous enterprise through its dedicated agentic identity solution, Idira Agentic Identity Security, and a deep architectural collaboration with NVIDIA. Designed to manage the architectural risks of active execution, Palo Alto Networks implements strict identity governance, least-privileged access, and zero standing privileges to ensure AI agents do not abuse their delegated authority across connected enterprise systems. By natively integrating with the NVIDIA OpenShell reference architecture, the company delivers AI Runtime Security that runs out-of-band on BlueField DPUs, providing independent visibility and automated machine-speed threat containment directly within the AI factory infrastructure.

Nasdaq: PANW

Stellar Cyber

Stellar Cyber is a full-cycle AI-native SecOps platform purpose-built for MSSPs and lean enterprise security teams. With Stellar Cyber, organizations see every threat, know what matters most, and act with speed and confidence, always with the human in the loop. The strategy provides lean security teams with autonomy without surrendering judgment. Its 6.4.0 release added agentic auto-triage with humans in the loop, mid-year updates refined detections, and a Gartner NDR Magic Quadrant mention added to credibility.

Total funding: $95 million

https://stellarcyber.ai/

Teleport

Teleport helps secure critical infrastructure by extending identity-first access controls to autonomous AI agents. As enterprise workflows increasingly shift from human operators to AI agents, traditional static API keys and secrets represent an expanded attack surface. By implementing cryptographic, short-lived machine identities and real-time session recording, Teleport ensures that AI agents are subjected to the same rigorous zero-trust governance as human engineers. This transition from static credentials to dynamic, auditable authorization

within trusted runtimes is foundational for organizations seeking to scale AI automation securely without compromising their overall compliance or infrastructure integrity.

Total funding: $170 million

TENEX

TENEX provides a fully agentic, human-led SOC that enables specialized agents to investigate every alert while named analysts stay accountable for each decision. Through a partnership with WitnessAI, it extends that discipline to enterprises' own AI, policing prompt injection, data leakage and shadow AI. As a founding member of the Agentic SOC Alliance, TENEX is addressing agentic security with a human-led approached.

Total funding: $277 million

https://tenex.ai/

Torq

Torq provides agentic SOC capabilities with a flexible API orchestrator that sits on top of an existing security stack. It does not provide its own native SIEM or XDR telemetry. Instead, its Socrates engine acts as an execution controller that parses incoming JSON/YAML payloads from current alert tools. Torq's HyperAgents are task-specific, declarative execution bots. It focuses heavily on reducing Tier-1 alert fatigue by executing high-speed, multi-step programmatic workflows to filter out false positives before they ever reach a human analyst.

Total funding: $332 million

Appendix C: Sources and Research Method

Our analyst team consulted a wide variety of information for this report. This includes survey results pulled from recent Futuriom independent surveys, public surveys, investment analysis, and vendor studies and customer case studies. Vendor-published customer results are examples rather than comparative benchmarks. Market assessments and forecasts are editorial analysis; no customer interviews or field studies were conducted for this sample.

[1] Futuriom. 2026 Networking and SASE Survey

Supplied graphics in Sample Analyst Report.docx. N = 144. Licensing and full methodology pending.

[2] Anthropic. Disrupting an AI orchestrated cyber espionage campaign

November 2025. Publisher assessment of one campaign.

[3] OWASP GenAI Security Project. OWASP Top 10 for Agentic Applications for 2026

Published December 9, 2025. Public risk framework.

[4] Splunk. How the indexer stores indexes

Splunk Enterprise 9.1 documentation. Storage lifecycle and searchable bucket states.

[5] Imply. How Lumi works

Technical documentation. Architecture and integration descriptions.

[6] Imply. How BTG Pactual Scales Security Investigations Without Replacing Splunk

Current online case study reviewed September 30, 2026. Vendor-published customer results.

[7] TENEX. Eight Agentic Capabilities One Continuous Security Operation

August 5, 2026. Vendor description of connected operations and governed response.

[8] NVIDIA. NVIDIA Launches Open Agent Safety Platform to Secure Agents From Testing to Deployment

September 28, 2026. Announcement of software and reference system design.

[9] Eclypsium. Safe AI Agents Start With Trusted AI Infrastructure

Yuriy Bulygin, September 28, 2026. Vendor description of infrastructure assurance and NVIDIA architecture.

[10] NIST. Secure Software Development Framework Version 1 Point 1

SP 800-218, February 2022. Final framework.

[11] Menlo Ventures. Agents for Security Rethinking Security for the Age of Autonomy

Investor perspective. Used as a market hypothesis rather than product validation.

[12] Microsoft. Microsoft Security Copilot agents overview

Microsoft Learn, updated March 18, 2026. Identity modes and agent workflows.

[13] CrowdStrike. CrowdStrike Launches the Charlotte AI AgentWorks Ecosystem for Building Secure Agents

March 25, 2026. Vendor announcement.

[14] Palo Alto Networks. The Agentic AI Platform for the Agentic Workforce of the Future

October 28, 2025. Vendor description of AgentiX.

[15] Insight Partners. The Next Stack The future of AI in cybersecurity

Investor perspective. Data, management and control plane framework.

[16] Cisco. Cisco Completes Acquisition of Splunk

March 18, 2024. Acquisition completion and strategic rationale.

[17] TENEX. Fully Agentic Human Led Security Operations Platform

Vendor product description and integrations with Google SecOps and Microsoft Sentinel.