THE

AGENTIC SECURITY BUYER'S GUIDE

How to Evaluate AI-Powered Security Operations and Avoid the Most Common Mistakes

7AI

7ai.com

Why This Guide Exists

Buying agentic security isn't like buying traditional software.

With SaaS, you evaluate features, check integrations, negotiate price, and deploy. The product works the same way for everyone. Implementation is measured in weeks, maybe months. Success is binary: it either does what was promised or it doesn't.

Agentic security is different. AI agents need to understand your environment—your tools, your workflows, your exceptions, your definition of what matters. The same technology deployed two different ways can produce wildly different outcomes. And because these systems make autonomous decisions about security events, the stakes of getting it wrong are higher than a failed software rollout.

We built this guide based on what we've learned deploying agentic security in production environments—the questions that mattered most, the tradeoffs that weren't obvious upfront, and the patterns that separated successful deployments from stalled ones.

What Agentic Security Does

Agentic security uses AI agents to autonomously execute security operations—taking alerts from detection through investigation, and in some cases through response. The goal is to handle the repetitive, high-volume work that consumes analyst time, so your team can focus on strategic security challenges.

The Core Workflow

Detection: Ingest alerts from your security stack (for example: EDR, identity, cloud, email, SIEM) and understand what you're looking at and optimize detection.

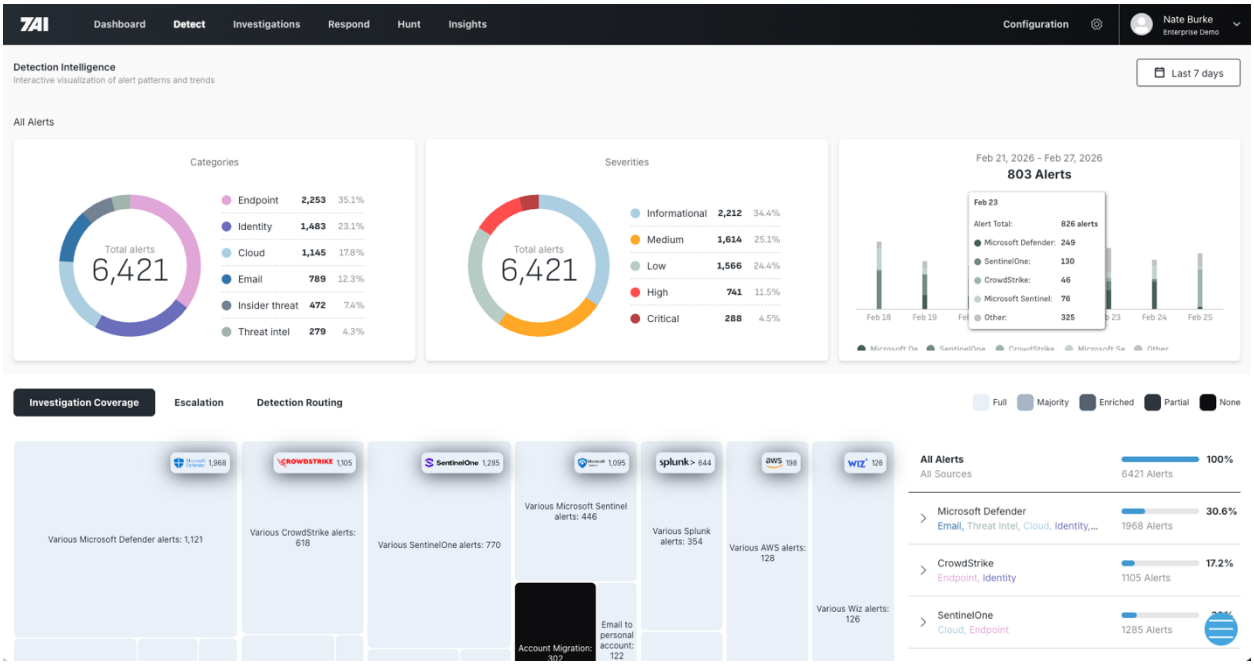

Investigation: Gather context from across your environment, correlate data, analyze the alert, and determine whether it posed risk to your organization. This is the work that traditionally takes analysts 30 minutes to 2+ hours per alert.

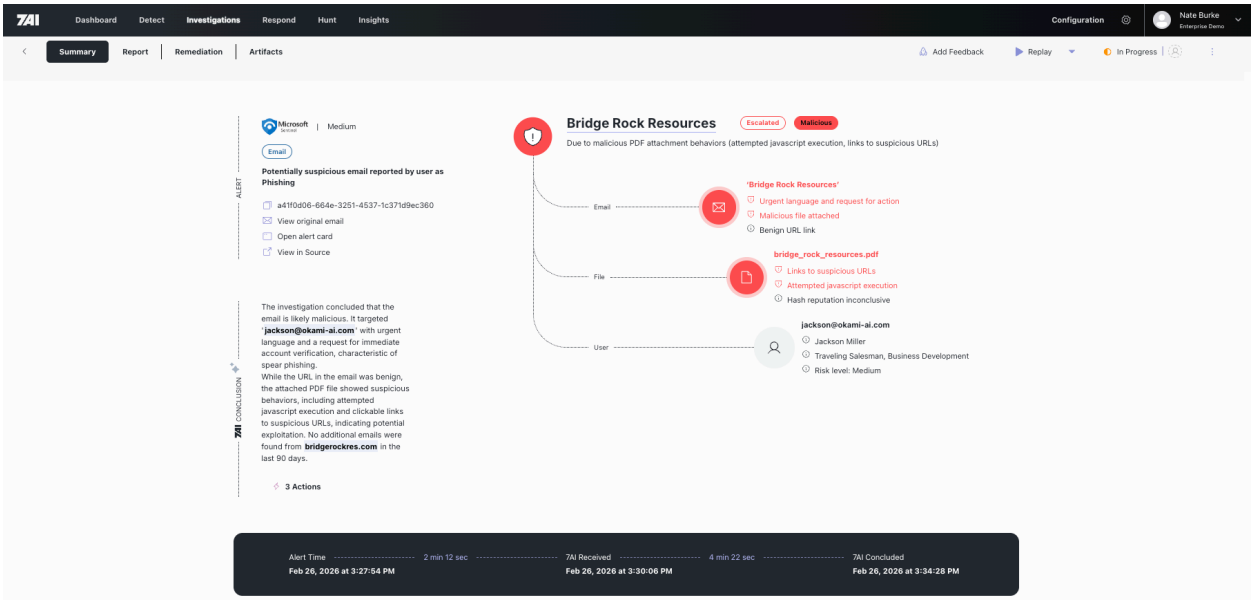

Conclusion: Reach an explainable determination—benign, suspicious, malicious—with the reasoning documented so humans can validate the decision.

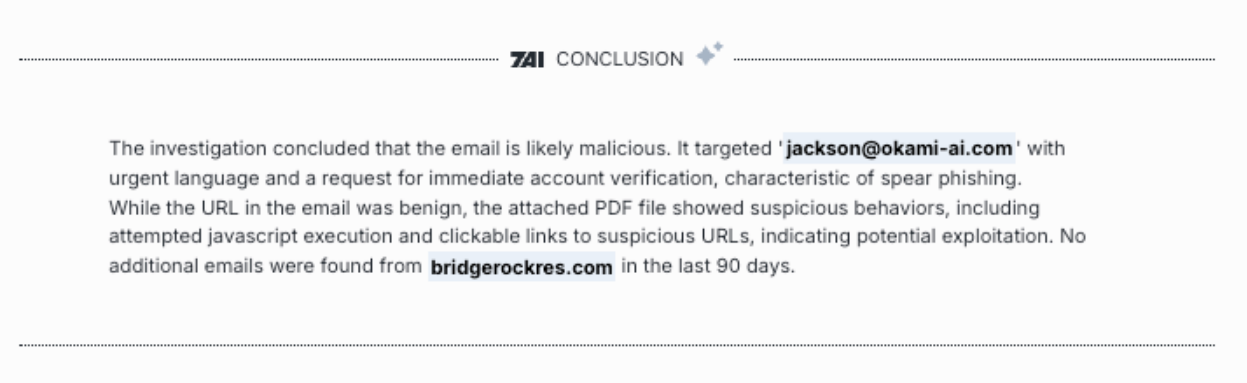

Response: Depending on your model, either hand off recommendations to your team or take autonomous action within defined guardrails.

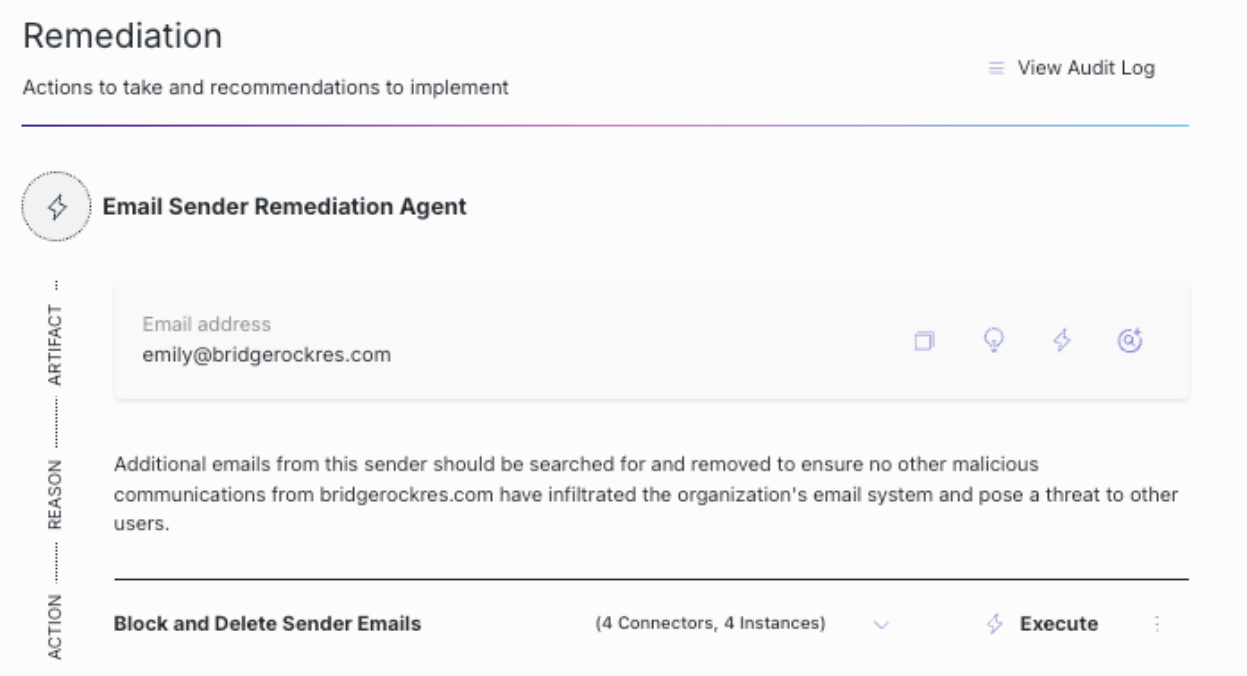

Threat Hunting: Based on new threats, proactively searching the environment for IoCs before an alert is triggered.

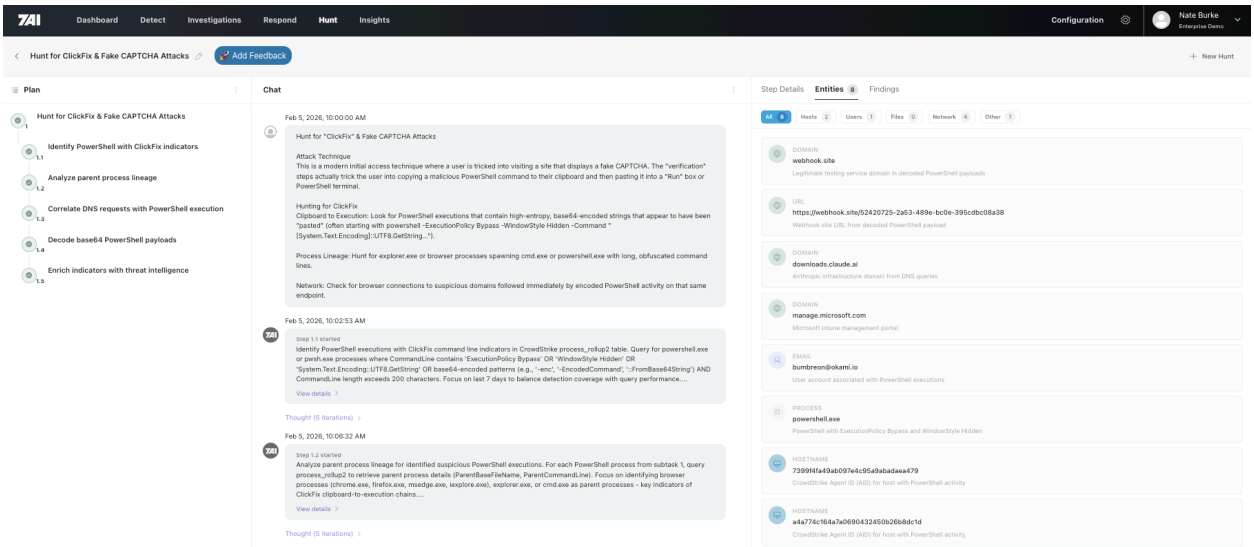

What This Replaces

Agentic security addresses the investigation bottleneck that has plagued security operations for years. Traditional approaches—whether in-house SOCs or MDR providers—rely on human analysts to manually investigate alerts. This creates predictable problems:

- **Volume overwhelms capacity.** Most organizations can only investigate 4-9% of their alerts.
- **Quality is inconsistent.** Investigation quality varies based on who's doing it, when, and how much context they have.
- **Response is slow.** SLAs measured in hours to days mean attackers have time to move laterally.
- **Talent is scarce.** Security analysts are expensive, hard to find, and prone to burnout.

How to Measure Success

These are the metrics we've seen matter most in production deployments—and the ones that separate real outcomes from marketing claims.

Coverage

What percentage of your alert volume is being investigated? Traditional MDRs investigate a small subset, often 4-9% of alerts. Effective agentic security should investigate 100% of the alert types you prioritize.

Speed

How fast does an alert go from detection to conclusion? MDR SLAs typically promise to begin investigating within hours—not conclude. Agentic security should conclude investigations in minutes.

> *Benchmarks: Alert to conclusion in 4-10 minutes average. Time saved per alert: 30 minutes to 2.5 hours vs. human investigation.*

Accuracy

How reliable are the conclusions? You need to measure both false positives (benign events flagged as threats) and false negatives (actual threats missed).

- False positive reduction rate (target: 95%+ reduction)
- False negative rate in production (target: 0 known false negatives)
- Escalation rate (what percentage requires human review?)

Transparency

Can you see how the AI reached its conclusions? Black-box AI that delivers verdicts without explanation creates liability and erodes trust. Every investigation should produce an explainable report.

Time to Value

How long until you're seeing production results? Traditional security platform deployments take 6-12 months. Agentic security should demonstrate value in days to weeks.

But don't confuse speed to first value with a complete evaluation window. Agentic security evolves continuously—the platform you see in month three will be materially better than what you see on day one. Evaluate the trajectory and the team driving improvement, not just a snapshot in time.

The Three Approaches to Agentic Security

Not all agentic security is created equal. The market has three distinct approaches, each with different tradeoffs.

Approach 1: One-Size-Fits-All AI SOC Tools

What it is: Software products that promise autonomous alert investigation out of the box. Buy the license, connect your data sources, and let the AI work.

The pitch: "Just deploy it and it works."

The reality: These tools treat every customer the same. They don't understand your specific environment: your approved software, your user roles, your workflow exceptions.

The result is one of two failure modes: too many false positives because the AI doesn't understand your context, or missed threats because it applies generic logic that doesn't account for what's abnormal in your environment.

> *What you get: Binary yes/no answers without environmental context. Ticket-based support when something isn't working. Your team adapts to how the tool works, not the other way around. Months of tuning before reliable results.*

Approach 2: SOAR and Hyperautomation Platforms

What it is: Orchestration platforms that automate predefined playbooks. You define the logic; the platform executes it.

The pitch: "Automate your existing processes at scale."

The reality: SOAR solves a different problem. It excels at orchestration and response—executing predefined actions after a decision has been made. But SOAR doesn't make decisions.

SOAR handles the "now execute these steps" problem. It doesn't handle the "what is this alert and what should I do about it?" problem which is where analysts spend their time.

> *What you get: Automation of known scenarios with predefined playbooks. Significant development and maintenance overhead. No ability to handle novel attacks without new playbook development.*

Approach 3: People-Led, AI-Driven (PLAID)

What it is: AI agents that autonomously investigate alerts, customized to your specific environment by expert security engineers who work alongside your team.

The pitch: "AI that understands your environment, backed by humans who make it work."

The reality: This approach recognizes that AI agents are only as good as their understanding of your context. The key difference is who does the customization—expert security engineers who learn your environment, configure AI agents to match your context, and continuously optimize based on results.

> *What you get: AI agents customized to your specific environment and workflows. Expert humans handling configuration, tuning, and optimization. Days to production, not months. Continuous improvement built into the engagement model.*

Buyer's Checklist

Use this framework to evaluate any agentic security solution.

Integration & Data Access

Question	One-Size-Fits-All	SOAR Platforms	People-Led, AI-Driven
Native integrations to existing tools?	Varies—often limited	Requires custom development	Yes. Built for your stack
Requires SIEM centralization?	Sometimes	Yes	No—goes to where data lives
Handles custom detection rules?	Rarely	If you build the playbook	Yes—with customization
Time to production-ready?	Weeks to months	Months	Days

Customization & Context

Question	One-Size-Fits-All	SOAR Platforms	People-Led, AI-Driven
Understands your exceptions?	No—generic logic	Only if encoded in playbooks	Yes—enterprise insights
Accounts for business context?	Limited	If you build the logic	Yes—customized per environment
Who does customization?	You (or ticket support)	Your team	Expert security engineers
How are changes made?	Submit ticket and wait	Update playbooks yourself	Ongoing optimization meetings

Investigation Quality

Question	One-Size-Fits-All	SOAR Platforms	People-Led, AI-Driven
Investigates or just triages?	Varies—often shallow	Executes predefined steps	Full investigation to conclusion
Reasons through novel attacks?	Limited	No—requires playbooks	Yes—AI agents reason through unknowns
Conclusions explainable?	Varies	Shows playbook execution	Yes—full transparency
False positive reduction?	Varies widely	Depends on playbook logic	95%+ in production

Support & Partnership

Question	One-Size-Fits-All	SOAR Platforms	People-Led, AI-Driven
When something breaks?	Submit support ticket	Fix it yourself	Direct access to your team
Who optimizes over time?	You	You	Expert engineers in regular syncs
Dedicated team knows your env?	No	No	Yes
How is success measured?	You define and track	You define and track	Joint metrics with reviews

Choosing Your Consumption Model

Once you've decided on an approach, you need to choose how you want to consume it. This depends on your team structure and how much of security operations you want to own vs. outsource.

PLAID: People-Led, AI-Driven

- Best for: Organizations with existing SOC teams who want to augment their analysts with AI, not replace them
- Your analysts remain in control of security operations
- 7AI engineers customize the platform to your environment
- AI agents handle investigation at machine speed
- Conclusions delivered to your analysts for review and action

PLAID+Elite: 24x7 Managed Agentic Security Operations

- Best for: Organizations looking to replace their outsourced providers
- 7AI's elite security team operates 24x7 on your behalf
- AI agents handle investigation; human experts handle escalation
- Response coordination and strategic oversight included
- You handle strategic decisions; we handle everything else

Questions to Ask Every Vendor

Before you commit, get clear answers to these questions:

ON THE TECHNOLOGY

1. How does your AI reach conclusions? Can I see the reasoning for any investigation?
2. What happens when the AI encounters a situation it hasn't seen before?
3. How do you handle false negatives? How would I even know if you missed something?
4. What data sources can you integrate with natively vs. requiring custom development?

ON CUSTOMIZATION

5. Who configures the platform for my environment—me, or your team?
6. How do you learn what's normal vs. abnormal in my environment specifically?
7. What happens when my environment changes—new tools, new workflows, new processes?
8. Can you show me examples of customization for organizations similar to mine?

ON SUPPORT

9. When something isn't working, who do I call? How fast do they respond?
10. Do I have a dedicated team that knows my environment, or general support?
11. How often do we meet to review performance and optimize?
12. What does the ongoing relationship look like after deployment?

ON RESULTS

13. What results have you delivered for organizations like mine? Can I talk to them?
14. Can I speak directly with a production customer in a similar environment—not a hand-picked reference, but someone who's been running for 90+ days?
15. How fast can I get to production? What do the first 30 days look like?
16. How do you measure success, and how will I know if this is working?
17. What happens if it's not working—what's the path to resolution, or exit?

Beyond the Product: Choosing a Partner

Agentic security isn't a SaaS subscription you swap out next quarter. It becomes part of your operating model. That means the company behind the technology matters as much as the technology itself.

Visit their office. Meet the engineers who will work on your environment. Understand their roadmap and how it aligns with where your security program is headed. Ask yourself: do I trust this team to be in my environment making autonomous decisions about my security?

The investment here isn't just the contract value—it's the time, access, and trust you're extending to a partner who will be embedded in your security operations. If you pick wrong, the cost isn't just dollars. It's the time you lose and the ground you give up while competitors move forward.

Look for alignment on vision, not just features. The best partnerships happen when your security philosophy and theirs are pointed in the same direction.

Next Steps

Agentic security is no longer a question of "if"—it's a question of "how." The organizations pulling ahead are deploying AI that understands their environment, backed by people who make it work.

If you're evaluating agentic security solutions, we'd welcome the chance to show you how 7AI's PLAID model works. We'll deploy into your production environment on a specific use case—typically EDR or identity alerts—and show you results on your actual alerts within days. Not a sandbox. Not a simulation. Your environment, your data, our platform proving it works. That's not a marketing claim—it's how every 7AI deployment starts.

Ready to see agentic security in action?

Visit 7AI.com

Request a demo at 7ai.com/contact