ARBITIUM

ARBITIUM INSIGHTS

From Alerts to Action: How Multi-Layered AI Changes Incident Response

Arbitium Engineering | February 5, 2026

If you've ever tried to explain to a non-technical executive why your SOC can detect a threat in seconds but takes days to stop it, you know the frustration. The answer is surprisingly simple once you see it clearly, and it has everything to do with how we've architected AI in cybersecurity up to this point.

Let's walk through it together.

One AI Does Not Fit All

Here's something the marketing slides rarely mention: "AI" in cybersecurity has historically meant one thing: pattern recognition. A machine learning model trained on historical data, looking for anomalies. It's the digital equivalent of a security camera. Incredibly useful for seeing things. Not particularly useful for doing anything about them.

The reason is architectural. Traditional ML models in security are classifiers. They take an input (a log entry, a network flow, a user behavior pattern) and produce an output (normal or anomalous, with a confidence score). That's detection. It's a single-layer problem: observe, classify, alert.

Incident response is fundamentally different. To respond to a threat, a system needs to understand context across multiple data sources. It needs to reason about what's happening, not just that something is anomalous, butwhyit's anomalous andwhat it meansin the context of a specific environment. Then it needs to decide on a course of action, carry it out, and verify the outcome.

That's not a classification task. That's reasoning, planning, and execution. And the industry has been treating it like the first one for far too long.

Three Problems, Not One

The gap between detecting a threat and stopping a threat isn't a single failure. It's actually three distinct problems stacked on top of each other, and each one requires a different kind of intelligence.

Problem 1: Signal Overload

The detection layer works. It works almost too well. Enterprise security stacks generate anywhere from 960 to over 3,000 alerts per day. The ML models and deep learning networks powering modern EDR, NDR, and cloud security platforms are genuinely good at surfacing potential threats from massive telemetry streams.

The problem isn't that we can't see enough. It's that we seeso muchthat human teams can't process it all. Nearly 90% of SOC teams report being overwhelmed by backlogs and false positives. The detection layer

is generating signal faster than the human layer can consume it.

Problem 2: Context Fragmentation

When a human analyst investigates an alert, they don't look at it in isolation. They pull up the user's recent activity. They check whether the affected system is production or staging. They look at whether this pattern has appeared before. They consider time zones, holidays, change windows. They cross-reference with threat intelligence. They build a mental model of the situation.

This reasoning step is where most security workflows break down, not because analysts lack the ability, but because the context they need is scattered across a dozen different systems. Your identity data lives in one place, your asset inventory in another, your vulnerability data somewhere else, your business criticality mappings in a spreadsheet that someone updated six months ago. An analyst might spend 30 minutes just assembling the context they need to make a decision, and by then the attacker has already moved.

The industry needs AI that can reason across fragmented context in real time, systems that understand the relationships between users, assets, vulnerabilities, and business processes, and can synthesize that understanding as fast as the alerts arrive.

Problem 3: The Manual Last Mile

This is the big one. Even when detection works perfectly and context is assembled quickly, the final step, actually stopping the threat, is almost always manual.

Revoking a compromised credential means logging into your identity provider. Isolating a host means opening your EDR console. Blocking a malicious IP means updating your firewall rules. Creating a forensic snapshot means running commands on the affected system. Each action requires a different platform, different credentials, different expertise. And each one takes time.

CrowdStrike's 2025 data tells us the average attacker breakout time is 48 minutes. The fastest they've recorded is 51 seconds. If your response workflow requires a human to execute actions across four platforms sequentially, you're not in a race. You've already lost.

What the industry needs is the ability to execute precise, coordinated response actions across the full security stack autonomously, with the speed of the threat and the precision of your best analyst.

What “Good” Looks Like

Let's make this concrete. Imagine a credential compromise scenario.

An anomalous login is detected from an unusual geolocation. In isolation, it's a medium-severity alert, one of dozens that day. But when you layer in context, the picture changes: the user is a system administrator, they're supposed to be on PTO, and the source IP is associated with a VPN exit node in a country where the organization has no operations. This isn't medium severity. This is critical.

Now, whatshouldhappen next? The compromised session should be revoked immediately. The user's credentials should be rotated. Conditional access policies should be tightened for that account. Forensic evidence should be preserved for investigation. And all of this should happen in seconds, not the 30 to 45 minutes it takes a human to pull context, make the assessment, and execute actions across four different consoles.

Early adopters of AI-driven autonomous response are already seeing what's possible: up to 90% automation of Tier 1 analyst tasks and a 50% reduction in mean time to respond. These aren't theoretical projections. They're production metrics.

What This Means for Security Teams

Let me be direct about something: autonomous incident response doesn't eliminate the need for security professionals. What it eliminates is the need for security professionals to do repetitive, high-pressure, time-critical manual work that burns them out and drives them to quit.

The data here is stark. Roughly 70% of SOC analysts with fewer than five years of experience leave within three years. These aren't people who lack skill or commitment. They're people trapped in a workflow that was never designed for the scale of modern threats.

When the routine work is handled autonomously (the Tier 1 triage, the standard response execution, the repetitive remediation) analysts are freed to do work that actually requires human judgment. Threat hunting. Red team exercises. Architecture review. Strategic security planning. The work that makes organizations fundamentally more resilient, not just faster at playing whack-a-mole.

The Bottom Line, And What You Should Do About It

The path from alerts to action requires the industry to solve three distinct problems: signal overload, context fragmentation, and the manual last mile. Each one demands a different kind of AI capability, and solving any one of them in isolation doesn't close the loop.

The technology to address all three exists today. The question is whether your organization will adopt it before your adversaries outpace your ability to respond manually.

Here's what we'd recommend as a starting point: audit your current incident response workflow end to end. Time each phase, from initial alert to full containment, across your last ten incidents. Identify where the delays live. In our experience, you'll find that 80% or more of the elapsed time isn't in detection or even investigation. It's in the manual last mile: waiting for approvals, logging into consoles, executing actions one platform at a time.

Once you see it, you can't unsee it. And once you can't unsee it, the case for autonomous execution makes itself.

If you want to see what closing that loop looks like in practice,book a demo with our team. We'll walk through it with your stack, your alerts, your environment.

This article was authored by the Arbitium engineering team. We build autonomous AI systems for cybersecurity incident response.

Sources referenced: CrowdStrike 2025 Global Threat Report; IBM Cost of a Data Breach Report 2025; Gartner Top Cybersecurity Trends 2026; Google Cloud Agentic SOC Research 2025.