ARBITIUM

ARBITIUM INSIGHTS

The Execution Gap: Why Cybersecurity's Biggest Problem Isn't Detection

John Remo, Co-Founder & CEO | January 22, 2026

I want to share something that took me fifteen years and every rung on the corporate ladder to fully understand.

When I was an operator running a SOC at 2 AM, I thought the problem was that we couldn't see enough. When I became an engineer, I thought the problem was that our tools weren't integrated well enough. When I managed a team, I thought the problem was that we didn't have enough people. When I became a director, I thought the problem was process. As a VP, I thought it was budget. As a CTO, I finally had the vantage point to see the real picture.

The problem was never detection. The problem was, and still is, execution.

We Solved the Wrong Problem Brilliantly

Let me be clear: the cybersecurity industry has done remarkable work over the past two decades. We've built detection capabilities that would have seemed like science fiction in 2005. Machine learning models that identify zero-day threats. Behavioral analytics that catch insider threats before data leaves the building. Threat intelligence platforms that correlate signals across billions of events.

And it's working. IBM's 2025 Cost of a Data Breach Report shows breach dwell time has dropped to 241 days, a nine-year low. We're finding threats faster than ever before.

But here's the number that should keep every security leader awake: 241 days is still 241 days. Even at a nine-year low, organizations are taking eight months from the moment a breach begins to the moment it's contained. Not detected. Contained. The gap between those two words is where the real damage happens.

The Anatomy of Inaction

To understand why this gap persists, you have to understand what happens between detection and resolution. I've watched this sequence play out at every organization I've worked with, regardless of size, maturity, or budget.

An alert fires. It joins a queue of somewhere between 960 and 3,000 other daily alerts, depending on the size of the organization. An analyst picks it up, if there's an analyst available. Nearly 90% of SOC teams report being overwhelmed by backlogs and false positives. The analyst investigates, which means correlating the alert with logs from five or six different systems, checking it against known indicators, and trying to determine if it's real.

If it's real, the analyst opens a ticket. The ticket enters a workflow. The workflow involves approvals. Approvals involve managers who are in meetings, or it's 3 AM and the escalation path leads to a

voicemail. Eventually, someone decides on a response. That response requires executing actions across multiple platforms: isolating a host in your EDR, revoking credentials in your identity provider, blocking an IP in your firewall, updating rules in your SIEM.

Each of those actions is manual. Each one requires access, expertise, and precision. Each one takes time. And while all of this is happening, the attacker is moving. CrowdStrike's 2025 data tells us the average breakout time (how fast an adversary moves laterally after initial access) is 48 minutes. The fastest they've observed? 51 seconds.

We're asking humans to outrun machines. That math has never worked, and it's getting worse every year.

The Headcount Trap

Here's where the conversation usually goes sideways. Leadership sees the problem. The response? Hire more analysts. But this creates a trap I've watched organizations fall into repeatedly.

Every new security tool you deploy generates more alerts. More alerts require more analysts. But good analysts are extraordinarily hard to find. There are 4.8 million unfilled cybersecurity positions globally. The ones you do find burn out fast. Industry data shows 70% of SOC analysts with five years or fewer of experience leave within three years. You're perpetually recruiting, training, and losing people to a cycle that was never going to scale in the first place.

I lived this cycle for years. And the uncomfortable truth I had to confront was that no amount of headcount was going to close the gap. You can't scale a manual process against an automated threat. It's like trying to bail out a sinking ship with a coffee mug. You can work harder, but the water's coming in faster.

What Actually Needs to Change

The industry doesn't need better detection. Detection is, for all practical purposes, a solved problem, or at least a mature one. What the industry needs is autonomous execution.

Think about what that means in practice. When a credential compromise is detected, the response shouldn't require a human to open a ticket, get approval, log into an identity provider, and revoke the session. The response should happen in the same timeframe as the detection, automatically, precisely, and verifiably.

This isn't a futuristic concept. The underlying technology exists today. AI systems can reason about complex, multi-step problems. They can plan sequences of actions across multiple systems. They can evaluate context, assess risk, and make decisions that account for business impact. The question isn't

whether autonomous execution is technically feasible. It's whether organizations are ready to trust it.

The Trust Equation

I understand the hesitation. I've been the operator who looked at an automated action and thought, "What if it gets it wrong?" That instinct is healthy. Security is a domain where a wrong action can be as damaging as no action at all.

But here's what I've come to believe after watching this space evolve: the risk of inaction now exceeds the risk of autonomous action. When your mean time to respond is measured in days or weeks and your adversary's breakout time is measured in minutes, the gap isn't just a performance issue. It's an existential vulnerability.

The organizations that will be most resilient in the next decade won't be the ones with the best detection. They'll be the ones that closed the execution gap, that built the capability to move from alert to action without a human in the critical path.

A New Measure of Security Maturity

For years, we've measured security programs by their detection capabilities. How many threats did you find? How fast did you find them? What's your detection coverage?

I think it's time we start measuring something different: execution velocity. How fast do you move from detection to containment? How many of your response actions are automated versus manual? What percentage of your incidents are resolved without human intervention?

These are the metrics that will separate secure organizations from vulnerable ones. Not because detection doesn't matter (it absolutely does) but because detection without execution is just expensive awareness.

And awareness without action has never stopped a breach.

John Remo is the Co-Founder and CEO of Arbitium, an autonomous AI execution platform for cybersecurity. He previously led engineering and security operations for global enterprises.

Sources referenced: IBM Cost of a Data Breach Report 2025; CrowdStrike 2025 Global Threat Report; Palo Alto Networks 2026 Cybersecurity Forecast; Gartner Top Cybersecurity Trends 2026.