ARBITIUM

ARBITIUM INSIGHTS

# The 74% Problem: Why Security Teams Are Still Holding Back the One Thing That Could Save Them

John Remo, Co-Founder & CEO | March 10, 2026

The data doesn't lie, and right now it tells a story the security industry needs to sit with.

Seventy-four percent of organizations are deliberately limiting AI autonomy in their SOC. Not because the technology doesn't work. Not because they lack the budget. But because they don't yet trust it to act without a human sign-off.

Meanwhile, the average enterprise SOC is processing upward of 10,000 alerts per day. Mean time to respond, the metric that determines how much damage an attacker does before you stop them, is measured in days and weeks for most organizations. And the global cybersecurity workforce shortage just crossed 4.8 million unfilled roles.

We have built ourselves an elegant trap.

## The Asymmetry No One Talks About

There is a fundamental asymmetry at the heart of modern incident response that the industry has grown comfortable ignoring: adversaries operate at machine speed while defenders still operate at human speed.

An attacker who gets initial access doesn't pause to schedule a meeting. They don't wait for a ticket to be triaged, an analyst to come off PTO, or a CISO to authorize a containment action. They move. Within minutes, they're establishing persistence. Within hours, they're conducting reconnaissance. Within days, they've achieved their objective.

Meanwhile, your SOC is doing the best it can, which in most organizations means a 12 to 48 hour average response time on critical alerts.

IBM's X-Force Threat Intelligence Index for 2026 found that supply chain and third-party breaches quadrupled over the past five years. Not because attackers got dramatically more sophisticated. But because the window of exposure, the gap between intrusion and containment, stayed wide open. The tools got smarter. The response stayed manual.

## The Detection Trap

The industry spent the better part of a decade solving detection. We built better SIEMs. We layered in XDR. We trained ML models on billions of threat signals. We got genuinely good at finding the needle in the haystack.

But here's the uncomfortable truth: detection without execution is just a more efficient alarm system.

If your security platform surfaces a true positive, a real, confirmed threat, and the next step is a human reading an alert and deciding what to do, you haven't solved the problem. You've just moved the bottleneck.

The security community's obsession with detection precision made sense when false positives were destroying analyst trust and causing alert fatigue. But we've overcorrected. We optimized the front of the pipeline and left

the back of it, the part where breaches actually become costly, running on the same human-speed rails it always did.

Analysts confirmed this in 2025: 82% reported being concerned they were missing real threats not because detection failed, but because the volume of validated alerts was too high for humans to act on in time. That's not a detection problem. That's an execution problem.

## Why the 74% Hesitancy Is Rational and Dangerous

I want to be fair to the CISOs and security leaders pumping the brakes on autonomous execution. Their caution isn't irrational.

The concerns are real. What if an autonomous system quarantines the wrong endpoint? What if it blocks a legitimate business process during a critical revenue moment? What if it acts on a sophisticated adversarial prompt injection designed to make it do the attacker's bidding?

These are serious questions. The risk of AI agents becoming insider threats, systems that have deep API access, implicit trust, and the ability to act at scale, is one of the most underappreciated attack surfaces in enterprise security right now. Gartner estimates that 40% of enterprise applications will have embedded AI agents by end of 2026. Only 6% of organizations say their security posture has kept pace.

So yes, shadow agents, unvetted autonomy, and improperly scoped AI execution are legitimate threats.

But the alternative we've chosen, keeping humans in the loop for every action, is also a threat. It's just slower and quieter, and it shows up on the P&L as breach remediation costs, regulatory fines, and reputational damage rather than a specific incident in a post-mortem. The choice isn't autonomous AI or safety. It's which risks are you actually managing?

## What Responsible Autonomous Execution Looks Like

The organizations getting this right aren't removing humans from security operations. They're rethinking where humans belong in the workflow.

Humans are exceptional at ambiguity, judgment calls, stakeholder communication, and novel threat pattern recognition. Machines are exceptional at consistency, speed, fatigue-free execution, and operating across thousands of signals simultaneously.

The mature model emerging in leading security organizations in 2026 looks something like this. AI executes on well-defined, high-confidence response actions: isolating a compromised endpoint, blocking a malicious IP, revoking a suspicious OAuth token, or triggering a credential rotation. These are deterministic, reversible actions with well-understood blast radius. They should never require a human to approve in real time. AI escalates on judgment calls, surfacing ambiguous threat actors or potential business impact decisions to humans with full context, triage already complete, and a recommended action waiting for sign-off. Humans own

strategy, communication, and edge cases.

Organizations with this model in production are seeing MTTR improvements of 70 to 90 percent. What previously took hours takes minutes. What took days takes seconds. That's not an incremental improvement. That's a different category of security posture.

## Trust Has to Be Earned, Not Assumed

The most common mistake organizations make when evaluating autonomous security execution is treating it as a binary decision: either the AI acts or it doesn't. In practice, the path to autonomous execution runs through observation first.

The right starting point isn't full autonomy. It's a watch and analyze mode, where the system monitors your environment, processes your alerts, and shows you exactly what it would have done and why, without taking any action. Over days and weeks, your team builds a read on the system's reasoning. Trust isn't assumed. It's earned through a track record.

When confidence is established, execution can be turned on incrementally. Teams typically start with their lowest-blast-radius actions, IP blocks, credential rotations, session terminations, where the consequence of an error is minor and reversible. As the system performs, scope expands. More complex actions, deeper integrations, higher-stakes remediations become candidates for automation.

Throughout all of it, every action is accompanied by a full explanation: what the system observed, how it reasoned, what it did, and why. This isn't a log. It's a readable account designed for the analyst who needs to understand it without a manual in front of them. And every action carries a rollback. If a containment decision turns out to be wrong, reversing it is a single step, not a multi-system recovery exercise.

This is how trust is built between security teams and the autonomous systems they're asked to rely on. Not with a contract or a vendor pitch. With a transparent track record of getting it right, and the safety net to correct it when it doesn't.

## The Inflection Point Is Now

With RSA Conference 2026 kicking off in two weeks, autonomous security operations will be the dominant conversation on the floor. Every major vendor has a story about agentic AI. Every CISO has questions about governance.

The conversations will be good. But the industry has a habit of talking about transformation longer than it takes to actually transform.

The math is simple: 4.8 million unfilled security roles, 10,000 or more daily alerts per analyst, a threat landscape operating at machine speed, and a response posture still anchored to human availability. Something has to give. The organizations who figure out how to execute, not just detect, autonomously will not merely

perform better on security benchmarks. They will have a fundamentally different risk profile than their peers.

That's the real competitive moat of 2026. Not which platform has the best detection. Which organization has the best execution.

## What You Can Do This Week

Before RSA, take 30 minutes with your team and audit your current incident response workflow.

Map your MTTR by action type, not by alert category. Where are humans spending time on decisions that are actually deterministic? Those are your first automation candidates. Identify your highest-confidence, lowest-blast-radius response actions. Credential rotation, endpoint isolation, IP blocking are table stakes for autonomous execution. If they still require a human approval step, ask why. Review the scope of AI agent permissions in your environment. If you're deploying AI agents, are they operating with least-privilege? Do you know what they can touch? Agentic AI without proper identity scoping is its own attack surface.

The window between detection and response is where breaches live. Closing it is the most important security investment you can make in 2026.

If you want to think through what autonomous execution looks like for your environment, reach out to the Arbitium team at arbitium.com/contact.

---

John Remo is Co-Founder and CEO of Arbitium, an AI-driven cybersecurity company focused on autonomous incident response and security execution.

Sources referenced: IBM X-Force Threat Intelligence Index 2026; Help Net Security, 'Why SOCs are moving toward autonomous security operations in 2026,' February 2026; SentinelOne, 'Cybersecurity 2026: The Year Ahead in AI, Adversaries, and Global Change'; Palo Alto Networks, '6 Cybersecurity Predictions for the AI Economy in 2026'; Gartner 2026 AI in Enterprise forecast; Mitiga, 'Top Cybersecurity Trends for RSAC 2026.'