ARBITIUM INSIGHTS | THOUGHT LEADERSHIP

The Agentic Threshold: What RSAC 2026 Finally Got Right

John Remo, Co-Founder & CEO | April 1, 2026

The Agentic Threshold: What RSAC 2026 Finally Got Right

I have attended enough security conferences to know the difference between a theme and a reckoning. RSAC 2026, which wrapped in San Francisco on March 26, felt more like the latter. The conference had a throughline that was impossible to miss: **Security must evolve at agentic speed.** That phrase was not marketing language. It was an admission.

An admission that the detection-oriented model we have relied on for two decades is no longer keeping pace. An admission that the attack cycle, once measured in days, now compresses into minutes, sometimes seconds. And, crucially, an admission that human-paced response is the single largest attack surface in enterprise security today.

Detection Was Never the Problem

Here is a fact the industry rarely states plainly: we are exceptionally good at detection. The tooling is mature. The signal quality is high. Modern EDR, NDR, and SIEM platforms can identify a malicious process, a lateral movement pattern, or a credential anomaly within seconds of it occurring. The global SOC market reflects this investment, currently valued at approximately $41 billion and projected to reach $112 billion by 2035, according to industry analysis. We have spent a generation building extraordinarily precise early warning systems.

What we have not solved is what happens after the warning goes off.

The alert fires. A ticket opens. An analyst, juggling dozens of other tickets, eventually picks it up. They investigate, escalate, convene, document, and then, finally, act. By which point the adversary has had somewhere between 20 minutes and four hours to work freely inside your environment. Darktrace's 2026 Annual Threat Report documented this shift clearly: attackers have moved from exploit-driven breaches to AI-enabled credential abuse precisely because they understand the human lag in the response chain. They are not outsmarting your detections. They are outrunning your humans.

Agentic Is Not Assisted

This is where I want to be careful, because RSAC 2026 also gave us something less useful: a flood of vendors rebranding automation features as agentic AI. There is an important distinction that every security leader needs to internalize before their next budget conversation.

AI-assisted security means a human works faster. An analyst clicks a button and gets a pre-written investigation summary. They still make every decision. They still own every action. The latency is reduced, but the human is still in the critical path.

AI-agentic security means the system executes autonomously. It investigates, correlates, decides, and acts without waiting for a human to approve each step. The human moves from operator to commander, setting policy and reviewing outcomes rather than executing each response task.

Most of what was announced at RSAC 2026 is the former, marketed as the latter. One-click investigation summaries and AI-generated runbook drafts are valuable tools. They are not autonomous execution. The distinction matters because attackers are not moving at human-assisted speed. They are moving at machine speed.

The Numbers That Should Change Your Roadmap

Gartner estimates that by the end of 2026, 40% of enterprise applications will integrate task-specific AI agents, up from fewer than 5% just twelve months ago. Security operations is one of the highest-value areas for that transition. Analysts project that AI-native SOC architectures can reduce manual triage workloads by up to 70% and cut investigation time by 25% when agentic workflows, not just AI-assisted ones, are deployed. The SOC automation market itself is growing at a 10% CAGR.

More telling: by some estimates, 30% or more of SOC workflows at large enterprises will be executed by agents, not humans, before 2026 is out. That is not a prediction. That is a deployment sprint already underway at the most security-mature organizations.

The Leadership Imperative

I have been in this industry long enough to remember when threat intelligence was the transformative investment, then SIEM, then EDR, then XDR. Each wave improved the picture. None of them solved the execution problem. That is the wave we are riding now, and it is not incremental.

Security leaders who are still treating agentic AI as a future-state aspiration are already behind. The organizations that move first on autonomous execution will not just respond faster; they will reshape the economics of their security programs. Fewer escalations, shorter dwell times, lower breach costs. McKinsey estimates that productivity gains from agentic AI deployments could unlock up to $2.9 trillion in economic value by 2030, and security operations is one of the clearest early beneficiaries.

What RSAC 2026 got right is that it named the problem with unusual clarity. Attackers operate at machine speed. Defenses must match that tempo. The path to matching it is not more analysts or better dashboards. It is autonomous execution: systems that respond in the same timeframe that threats materialize.

Practical next step: Pull your last 10 incident response tickets and map each action to the person who took it. Count how many of those actions could have been executed by an automated system without human judgment. That ratio is your execution gap. It is also your most actionable security investment for

*the remainder of 2026. To learn more about building an autonomous response architecture, visit **arbitium.com/contact**.*

John Remo is the Co-Founder and CEO of Arbitium, an autonomous AI cybersecurity execution platform. He has spent his career at the intersection of security engineering and enterprise operations, from individual contributor to executive. He writes about the structural problems in security that tooling alone cannot fix.

Sources

RSAC 2026 Conference, March 23-26, San Francisco. Security Boulevard: "RSAC 2026 Day 1: Security Must Evolve at Agentic Speed."

Global Security Operations Center Market Report. OpenPR, 2025. Market size $41B (2025) projected $112B (2035).

Darktrace Annual Threat Report 2026. Industrial Cyber, 2026.

Gartner Predicts 2026: 40% of Enterprise Apps Will Feature Task-Specific AI Agents. Gartner Newsroom, August 2025.

McKinsey Global Institute. Agentic AI productivity value estimate. 2026.

SentinelOne Cybersecurity 2026 Blog: "The Year of the Defender." SentinelOne, January 2026.

SC Media: "RSAC 2026: AI Reshapes Cyber Defense and Threat Landscape." March 2026.