acceligence

Breaking the inertia: rethinking cybersecurity operating models in the age of AI

Insight

March 22, 2026

Category: cybersecurity

Tags: communications, governance, inertia, outcomes, pace, strategy, tooling

Cybersecurity professionals have earned the respect they rarely receive. Over the past two decades, teams operating under relentless pressure have defended enterprises against nation-state adversaries, criminal syndicates, and disruptive insiders, often with limited resources, organizational skepticism, and incomplete information. They built frameworks that scaled across global infrastructures. They absorbed lessons from breaches that became industry case studies. They transformed cybersecurity from a technical afterthought into a recognized business function. That is a track record that few modern business disciplines can claim.

Artificial intelligence is not simply another technology wave; it represents a structural shift in the environment that cybersecurity teams must defend and nurture. Threat actors are already using AI to accelerate reconnaissance, personalize attacks, and compress the time between vulnerability discovery and exploitation. The scale and complexity of digital environments have expanded far beyond what traditional security operating models were designed to manage.

One of the most significant changes AI introduces is the compression of time. Tasks that once required human experimentation, including malware obfuscation, phishing customization, reconnaissance analysis, and vulnerability discovery, can now be iterated thousands of times automatically. Attackers can generate variations, test techniques, and refine their methods in hours rather than weeks. This acceleration does not merely increase the number of threats organizations face. It fundamentally alters the tempo of cyber conflict.

Organizations themselves are adopting new technologies at unprecedented speed. Cloud platforms enable infrastructure to be created and destroyed in minutes. Development teams deploy code continuously rather than through scheduled release cycles. Digital ecosystems extend beyond the boundaries of a single enterprise, connecting partners, suppliers, and service providers through complex data exchanges and APIs.

AI systems add yet another layer of complexity. They rely on data, model training, inference, and software supply chains that traditional security architectures were never designed to monitor, introducing new attack surfaces and new forms of risk that do not map neatly to existing frameworks.

Most practitioners recognize this shift. Security leaders understand that AI is changing the threat landscape. Technology leaders understand that digital transformation is accelerating the complexity of enterprise environments. Boards increasingly recognize cybersecurity as a strategic risk.

The challenge is not awareness. It is adaptation.

Security teams are working harder than ever: deploying new tools, integrating threat intelligence, expanding monitoring capabilities, and experimenting with AI-driven defensive technologies. Yet many organizations still struggle to translate that effort into meaningful improvements in resilience.

The obstacle: *inertia*.

Not the kind that appears as resistance to change. The kind that emerges passively from operating models that once worked well but are gradually becoming misaligned with a rapidly evolving environment.

That inertia tends to appear in four key places:

- **Operations:** how security operations are structured and measured
- **Communication:** how cyber risk is communicated to leadership
- **Adaptation:** how quickly organizations adapt their defensive architectures as technology evolves
- **Cognitive:** dependency on security tools that can displace the critical thinking and problem-solving skills that once defined the profession at its best

Understanding these forms of inertia and what it takes to overcome them may be one of the most important leadership challenges facing cybersecurity today.

The weight of what works

Inertia in cybersecurity stems from the accumulated weight of decisions that worked.

The modern security operating model took shape in an environment that was far more contained than what organizations deal with today. Enterprise infrastructure was centralized. Network boundaries were clear enough to define and defend. Applications moved through structured release cycles, and most critical data sat inside systems the organization directly owned.

Under those conditions, layered – and mostly centralized – defensive controls fit the problem well. Firewalls handled the edges of the network. Endpoint security focused on individual devices. Vulnerability management programs

surfaced weaknesses that could be addressed methodically. Security operations centers pulled everything together so analysts could investigate signals in a structured way.

Frameworks such as NIST CSF, ISO 27001/27002, and PCI DSS added consistency across organizations. They offered a shared way to measure maturity and demonstrate accountability. These approaches were shaped through real incidents and hard-earned lessons, not abstract theory. Over time, security leaders refined monitoring, incident response, and vulnerability management because those efforts consistently reduced risk.

As those practices matured, they settled into codified operating models. Vendors built platforms to support them. Training programs and certifications reinforced them. Governance processes began to assume they would remain the foundation of enterprise cybersecurity.

Tools continued to evolve, but the structure beneath them stayed largely intact. Detection capabilities improved. Endpoint tools became more sophisticated. Cloud security solutions were added as infrastructure expanded. Identity platforms grew in importance as authentication risks increased.

Even with those advancements, the way things operated on the ground changed very little.

Alerts continued to flow into centralized systems, where analysts work through queues. Vulnerabilities were still tracked by volume and remediation timelines. Risk reporting still leaned heavily on control coverage and compliance posture.

This model held up when change moved at a manageable pace. Then AI entered the picture and was democratized at the edges.

The environment the old model now protects moves very differently. Infrastructure spans multiple cloud providers and hybrid architectures. Applications are assembled from distributed services and updated continuously. Data flows across partners, SaaS platforms, and automated integrations, often without clear boundaries.

The security operating models built to defend relatively stable enterprise environments are now being asked to protect ecosystems that change constantly. The result is a growing tension between the stability security programs were designed to provide and the agility modern digital environments demand. That tension is where inertia begins to surface.

Where inertia actually shows up

Inertia rarely appears as stagnation. Most security organizations are extremely active and from the outside, the system appears dynamic or even innovative.

But activity is not the same as adaptation. Security organizations can be highly efficient at performing tasks while still struggling to evolve their operating models. The signs of inertia tend to appear not in the volume of work being done, but in the structure of how that work is organized and interpreted.

In practice, inertia tends to concentrate in four distinct patterns.

Operational inertia: output mistaken for outcomes

Many cybersecurity programs confuse output with outcomes. Organizations can point to an impressive array of defensive activities that generate a vast amount of output: vulnerability scanning, patch cycles, phishing simulations, endpoint monitoring, threat intelligence analysis, and incident response exercises. These activities are necessary and often well executed. But they rarely connect to explicit operational objectives.

Security teams track task completion rather than the achievement of defensive outcomes. Metrics emphasize activity levels: how many systems were scanned, how many alerts were processed, how quickly vulnerabilities were patched, rather than whether those activities materially reduced exposure to real threats. A vulnerability scanning program may identify thousands of issues each month, yet the organization may still struggle to prioritize which weaknesses actually expose critical systems to meaningful risk. A monitoring platform may generate millions of alerts each day, yet analysts may still struggle to identify the handful of events that represent genuine threats.

The result is a system optimized for throughput rather than outcome. Automation and AI promise to help manage this scale and significantly increase efficiency. But if security programs automate existing workflows without redefining their operational objectives, they risk accelerating activity without improving outcomes. Automation, in other words, can make inertia move faster.

Communication inertia: technical reporting that does not reach decision-makers

Cybersecurity carries a quieter form of inertia in the way it communicates risk.

Most security teams have become highly sophisticated in how they measure their environments. They track vulnerability counts, monitor detection coverage, analyze patch timelines, and process enormous volumes of alerts and threat intelligence signals across increasingly complex systems. Those measurements are necessary to run the function well. They just don’t translate cleanly into the decisions the business actually needs to make.

Executives and boards are looking for something different. They want to understand what could happen to the organization and how prepared it is to handle it. The mechanics of how security operations function are secondary to that question.

The gap shows up in how information is presented. Leadership may hear that thousands of vulnerabilities exist, but not which ones could realistically lead to a material disruption. They may see improvements in detection coverage without any clear sense of how that changes exposure to actual adversaries. The signal is there, but it does not land in a way that supports action.

This disconnect runs deeper than communication style. It reflects a difference in how risk is framed.

Security professionals are trained to think in systems, controls, and technical conditions. Business leaders tend to think in terms of operational impact, financial consequences, regulatory pressure, and reputation. Both perspectives are valid, but they operate on different planes.

When reporting stays anchored in technical language, the burden shifts to executives to interpret what it means for the business. That translation step rarely happens in a structured way. Even in organizations with strong security teams and engaged leadership, the gap remains because the underlying frames never fully align.

Adaptation inertia: governance that cannot keep pace with the environment

A third form of inertia shows up in how decisions get made.

Most security organizations still rely on governance models built for a slower, more predictable world. Decisions flow through defined checkpoints. Risk is evaluated at specific moments in time. Changes are reviewed, approved, and documented through processes designed to create consistency and control. That structure made sense when change itself was periodic. Today, it isn’t.

Infrastructure can be created, modified, or removed in minutes. Development teams ship updates continuously. New integrations, data flows, and AI capabilities appear as part of normal delivery, not as exceptional events. By the time a formal review takes place, the environment it was meant to assess has often already moved on.

The friction shows up in subtle ways. Risk assessments arrive after key architectural choices have already been made. Security reviews slow down delivery without meaningfully shaping it. Governance forums become checkpoints that validate decisions rather than influence them. Over time, the system starts to lag behind the reality it is supposed to guide.

AI compounds the problem by increasing both the speed and the surface area of change. New dependencies appear in the form of training data, third-party models, and inference services that operate at scale and evolve quickly. These are not discrete systems that can be evaluated once and approved. They are living components that require ongoing judgment, yet governance processes still tend to treat them as static assets. The deeper issue is not tooling or visibility, it’s decision velocity.

Security organizations are often structured to make careful, defensible decisions. That instinct remains important. But when the pace of change accelerates, the cost of slow decisions increases. Risk is no longer introduced in large, infrequent steps. It accumulates continuously, often in places governance never meaningfully touches.

Meanwhile, attackers operate under a completely different model. They are not waiting for approvals, aligning to review cycles, or navigating internal dependencies. When a technique works, they use it immediately. When it fails, they adjust just as quickly. AI strengthens that advantage by enabling rapid experimentation and iteration at a scale that traditional governance processes cannot match.

Defenders have access to many of the same technologies, but they are constrained by how decisions are made. Until governance evolves to operate at the speed of the environment, security programs will continue reacting to changes they had the opportunity to shape earlier.

Cognitive inertia: the atrophy of critical thinking in a tool-dependent profession

The fourth form of inertia is the hardest to see and, arguably, the most consequential. It lives in the very way cybersecurity practitioners think.

The explosive growth of security tooling over the past decade has been largely beneficial, but a dependency has developed alongside these gains, one that has quietly displaced some of the human capabilities that effective security work requires most.

When every alert comes pre-classified, when every vulnerability comes with an automated severity score, when threat intelligence arrives as a feed rather than as a question, analysts can spend entire careers responding to what tools produce without developing a deep sense of why adversaries behave as they do, how attack paths connect across an environment, or what the business actually needs to remain resilient. The tools do not suppress that thinking deliberately. They simply make it unnecessary in the short run, and anything that becomes unnecessary long enough eventually atrophies.

The effect shows up in several ways. Security teams that cannot articulate their threat model with specificity, only describe the categories of threats in vendor documentation. Incident responders who are proficient in the mechanics of containment but struggle to reason about attacker intent or reconstruct what an adversary was actually trying to achieve. Analysts who escalate more because they are uncertain how to apply judgment to ambiguous signals, not because the signals are genuinely ambiguous. Organizations that cannot prioritize effectively because prioritization requires a judgment about what matters most, and that judgment was never developed.

AI amplifies this vulnerability directly. When AI-driven tools generate recommendations, summarize threat intelligence, and surface anomalies autonomously, the analytical demand on human operators decreases further. Practitioners who have already ceded much of their analytical work to existing automation will cede more to AI, and the gap between what the tools produce and what humans can critically evaluate will widen.

This matters especially because AI also arms adversaries with the ability to craft more sophisticated deceptions. Phishing campaigns that are more contextually accurate. Lateral movement that mimics legitimate behavior more convincingly. Intrusions that generate fewer anomalous signals and require more interpretive judgment to detect. The defenders best positioned to counter these techniques are those who have maintained the habit of thinking through what they are seeing rather than processing what the tools report.

Breaking the inertia: five shifts that are harder than they sound

Inertia rarely yields to awareness alone. Most security leaders already recognize that AI is reshaping the threat landscape. Boards increasingly recognize cybersecurity as a strategic risk. Recognition, however, does not automatically produce change.

What follows are not best practices in the conventional sense; they are structural changes, each of which requires displacing something that currently works well enough and that is what makes them difficult. The organizations that make these shifts will not do so because they were compelled by a crisis. They will do so because their leadership decided that well enough is not the same as ready.

1. Replace output metrics with outcome-based operational objectives

This is a direct response to operational inertia and requires confronting the measurement system around which the entire organization has organized.

The shift begins with a question: *what are we actually trying to prevent?* Not in the abstract sense of “breaches” or “incidents,” but with specificity. What are the two or three scenarios that would cause the most material harm to this organization? What are the attack paths that lead there? Which of those paths are we confident we would detect, and which are we not?

Answering those questions requires threat modeling that goes beyond vendor-supplied frameworks. It requires security teams to reason about their specific environment, adversaries, and the consequences of failure. That reasoning is hard to automate, which is precisely why it has been deprioritized in organizations where tool-generated metrics have become the primary language of security performance.

Once operational objectives are defined with this level of specificity, activities can be evaluated against them. Vulnerability prioritization shifts from risk-score-driven queues to path-based reasoning: which weaknesses, if exploited in sequence, could enable an attacker to reach a critical system? Monitoring configurations shift from broad coverage goals to detection of behaviors that characterize the specific techniques adversaries use against this industry and this organization. Tabletop exercises shift from scenario rehearsal to genuine red-team thinking about

where current defenses are most likely to fail.

The institutional resistance to this shift is real. Throughput metrics are auditable, reportable, and politically comfortable. Outcome-based objectives require admitting uncertainty about what the organization can and cannot detect. Security leaders who make this change will, at least initially, produce reports that are more honest and less impressive-looking than the ones they replaced. That is a feature, not a defect.

2. Rebuild risk communication that infuses business context

This is the direct response to communication inertia, and it requires security leaders to step outside the professional vocabulary they have spent careers developing.

The fundamental problem is not that security teams communicate poorly. It is that they communicate accurately within a frame that most business leaders do not share. Vulnerability counts, control coverage percentages, and MTTD (mean time to detection) metrics are precise. They are also largely unintelligible to a CFO, COO, or CEO trying to make a resource-allocation decision, or to a board member trying to govern and oversee management strategy.

The shift requires translating observations into consequences. Not “we have 4,200 open vulnerabilities” but “three of those vulnerabilities exist on systems that, if compromised, would disrupt our ability to make product, ship product, or sell product for an estimated 72 hours.” Not “detection coverage is at 87%” but “there is a class of credential-based attacks that we would currently detect only after lateral movement had already occurred, and our clinical trial data environment sits in that gap.”

This translation is more difficult than it sounds for two reasons. First, it requires security leaders to reason about business consequences, not just technical terms, and that reasoning requires a level of business fluency that security training does not typically develop. Second, it requires accepting the discomfort of communicating uncertainty. Business-consequence framing often produces statements like “we believe” or “our best estimate is,” which feel less authoritative than precise technical counts. That discomfort is worth tolerating. A statement that is honest about uncertainty is more useful to a decision-maker than a precise number that measures the wrong thing.

Security leaders who make this shift will need to invest time learning the language and priorities of their CFO, general counsel, and business unit heads. They will need to understand which operational failures the organization fears most and frame the security posture in terms of how prepared the organization is to prevent or withstand them. This is fundamentally a strategic communication discipline, and it is not yet widely taught or valued in the profession.

3. Redesign governance for continuous change, not periodic review

This is the direct response to adaptation inertia, and it requires dismantling governance structures that were built for a different operating tempo.

Annual risk assessments and quarterly architecture reviews made sense when infrastructure changed slowly. They do not make sense when a single sprint can introduce new cloud services, third-party integrations, and AI-enabled capabilities that materially alter the organization's attack surface. The governance cadence must match the change cadence.

In practice, this means embedding security into the development and deployment process rather than reviewing it afterward. Security requirements are defined at the architecture stage, not appended at the review stage. Automated security testing is integrated into CI/CD pipelines, not run as separate assessments against finished systems. Threat models are updated as a standard part of product iteration, not as annual documentation exercises.

It also means restructuring how security organizations absorb new capabilities. Most security programs evaluate new tools through formal procurement processes that take months. By the time a technology has been assessed, budgeted, approved, and deployed, the threat environment it was selected to address may have shifted. Organizations that develop lightweight processes for evaluating, piloting, and deploying new capabilities, including AI-driven defensive tools, will be able to respond to adversary innovation more quickly than those that route every adoption decision through the same heavyweight governance cycle.

The harder change is cultural. Security organizations often prize stability, for understandable reasons: in high-stakes environments, undisciplined change introduces risk. But stability in governance processes is not the same as stability in security posture, because the former increases bureaucracy that is antithetical to a rapidly-changing business and technology environment. The organizations that maintain rigorous security while operating in dynamic environments are those that have learned to distinguish between what needs to be stable (principles, accountability structures, response protocols) and what needs to adapt continuously (detection configurations, architecture, tooling).

4. Deliberately rebuild the problem-solving and analytical capabilities that tool-dependency has eroded

This is the direct response to cognitive inertia, and it is the least comfortable recommendation in this article because it implies that some of what has been invested in tooling and automation has come at a cost that was not fully recognized when those investments were made. It's a return and a recommitment to the fundamentals.

A place to start might be threat modeling, not as a compliance exercise but as a genuine analytical practice. Security teams that cannot reconstruct, from first principles, the most likely attack paths through their environment, without consulting a vendor's framework or a tool-generated risk map, have a gap that no additional tooling will close. Rebuilding that capability requires protected time for analytical work that is not tied to operational tasks, facilitation by practitioners who can model the reasoning out loud, and leadership patience for a process that produces uncertain outputs.

Another example could be adversarial simulation that is designed to challenge assumptions rather than validate them. Most red team exercises and tabletop scenarios are structured around known attack patterns and expected defensive behaviors. They confirm that existing playbooks work. The exercises that actually develop judgment are the ones where the scenario is unfamiliar, the adversary behaves unexpectedly, and the team must reason under

uncertainty rather than execute a rehearsed response. Those exercises are harder to run and more uncomfortable to debrief. They are also where cognitive capability develops.

A third practice to push is explicit investment in what might be called adversarial reasoning: the habit of thinking from the attacker's perspective, not to predict specific attacks but to develop an intuition about what is actually at risk and why. This is the kind of thinking that distinguishes security practitioners who are dangerous to adversaries from those who are merely proficient with tools. It cannot be automated, cannot be purchased, and cannot be developed by analysts who spend their working days processing queue-driven alerts.

Organizations that take this seriously will need to make structural changes, such as moving analysts through different functions before specialization calcifies. Dedicated time budgeted for analysis that is not attached to operational deliverables. Hiring criteria that explicitly value analytical reasoning alongside technical certification. These are not standard practices in the profession, and the organizations that adopt them will be building a capability that most of their peers are systematically eroding.

5. Integrate security proactively into strategic decisions before risk materializes

This shift addresses a structural problem that cuts across all four forms of inertia: security is most often brought into decisions after the consequential choices have already been made.

AI adoption, digital product development, data platform strategy, and cloud architecture all determine the organization's future risk posture. When those decisions are made without security input, the risk profile of the resulting systems is determined by default rather than by design. Security teams then inherit environments they did not shape and must defend against exposures that could have been designed out at lower cost and effort earlier in the process.

The pattern persists because of how organizations are set up. Security is frequently positioned downstream from strategy, brought in to assess or validate rather than to influence direction. By the time engagement happens, the room for meaningful change has narrowed.

Addressing this requires a different placement in the organization. Security leaders need to be present when strategic technology decisions are being formed, not introduced once they are already in motion. That presence depends on authority, sponsorship, and the ability to contribute in ways that resonate with both business and technology leadership.

Participation at that level also calls for a different mode of engagement. Conversations about AI, product architecture, or platform direction revolve around tradeoffs, investment, and long-term impact. Contributing effectively in those settings means evaluating options, articulating consequences, and connecting technical choices to business outcomes to support decision-making.

Organizations that make this shift do so by building capability and trust over time. Security leaders become part of the decision process because they add value to it. The influence gained in those moments shapes risk before it takes hold, which ultimately matters far more than any control added after the fact.

The moment to meet the momentum

The cybersecurity profession has achieved something remarkable over the past two decades. Under sustained pressure, security teams built practices, frameworks, and institutions that transformed cybersecurity into a recognized business discipline. They introduced structured governance, operational monitoring, and systematic risk management into environments that once operated with little security oversight.

That progress should be respected, yet it should not become an anchor.

Artificial intelligence is accelerating the pace at which threats evolve, technologies change, and organizations innovate. The environment cybersecurity must defend is moving at an unprecedented pace, is far more interconnected, and way more complex. The question facing security leaders is not whether this change will continue. It already has and will continue to do so at even faster velocity than we imagined possible today.

The real question is whether cybersecurity operating models can evolve as quickly, and whether the practitioners who oversee them can maintain the analytical capabilities required for adaptation.

Organizations that succeed will not simply deploy new tools or automate existing processes. They will redefine what effective security operations look like. They will translate security intelligence into the language their business leaders can act on. They will build governance capable of moving at the pace of the environments they defend. And they will invest deliberately in the human judgment that no tool can replicate, and no adversary can simply automate around.

Cybersecurity has matured into a professional discipline. The challenge now is ensuring it does not become a bureaucratic one.

Breaking the inertia that accompanies success will require the same qualities that built the profession in the first place: clear thinking, operational rigor, intellectual honesty about current constraints, and the willingness to adapt when the environment demands it.