PALO ALTO NETWORKS

perspectives

JULY 2025

FEATURE

Why Culture Is the First Line of Defense in the Age of Agentic AI

by Wendi Whitmore

THE AI AGENDA

FROM DISRUPTION TO DOMINANCE

Securing AI Agents: Building the Landing Gear While Flying the Plane

A Blueprint for AI Research That Actually Works

A Leader's Guide to Supply Chain Chaos

Contributors

RICH CAMPAGNA is the SVP of Product Management at Palo Alto Networks.

AARON ISAKSEN is the VP of AI Research and Engineering at Palo Alto Networks.

NICOLE NICHOLS is a Distinguished Engineer at Palo Alto Networks.

HAIDER PASHA is the CSO of EMEA and LATAM at Palo Alto Networks.

HELMUT REISINGER is the CEO of EMEA and LATAM at Palo Alto Networks.

SAM RUBIN is the SVP of Consulting and Threat Intelligence for Unit 42 at Palo Alto Networks.

CHRIS SCOTT is the VP and Managing Partner for Unit 42 at Palo Alto Networks.

MICHAEL SIKORSKI is the CTO and Head of Threat Intelligence for Unit 42 at Palo Alto Networks.

KARIM TEMSAMANI is the President of Next-Generation Security at Palo Alto Networks.

WENDI WHITMORE is the Chief Security Intelligence Officer at Palo Alto Networks.

The Disruption Arrives: A World in Flux

Navigating the Challenge and Pivoting to Dominance

The Human Response: Culture, Talent, and Leadership

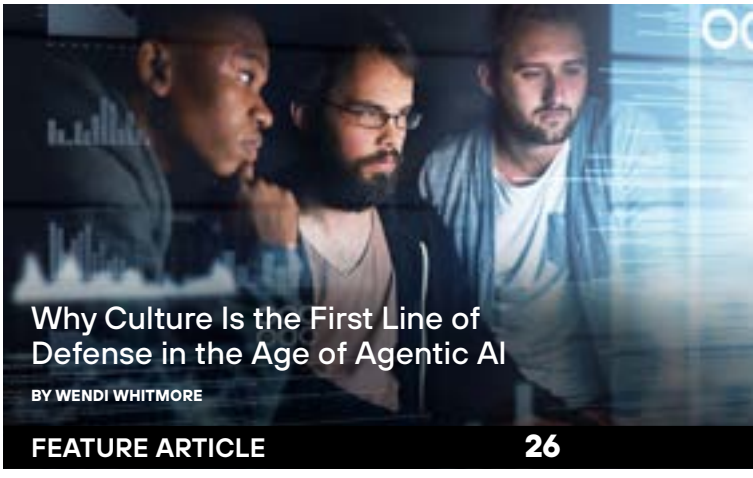

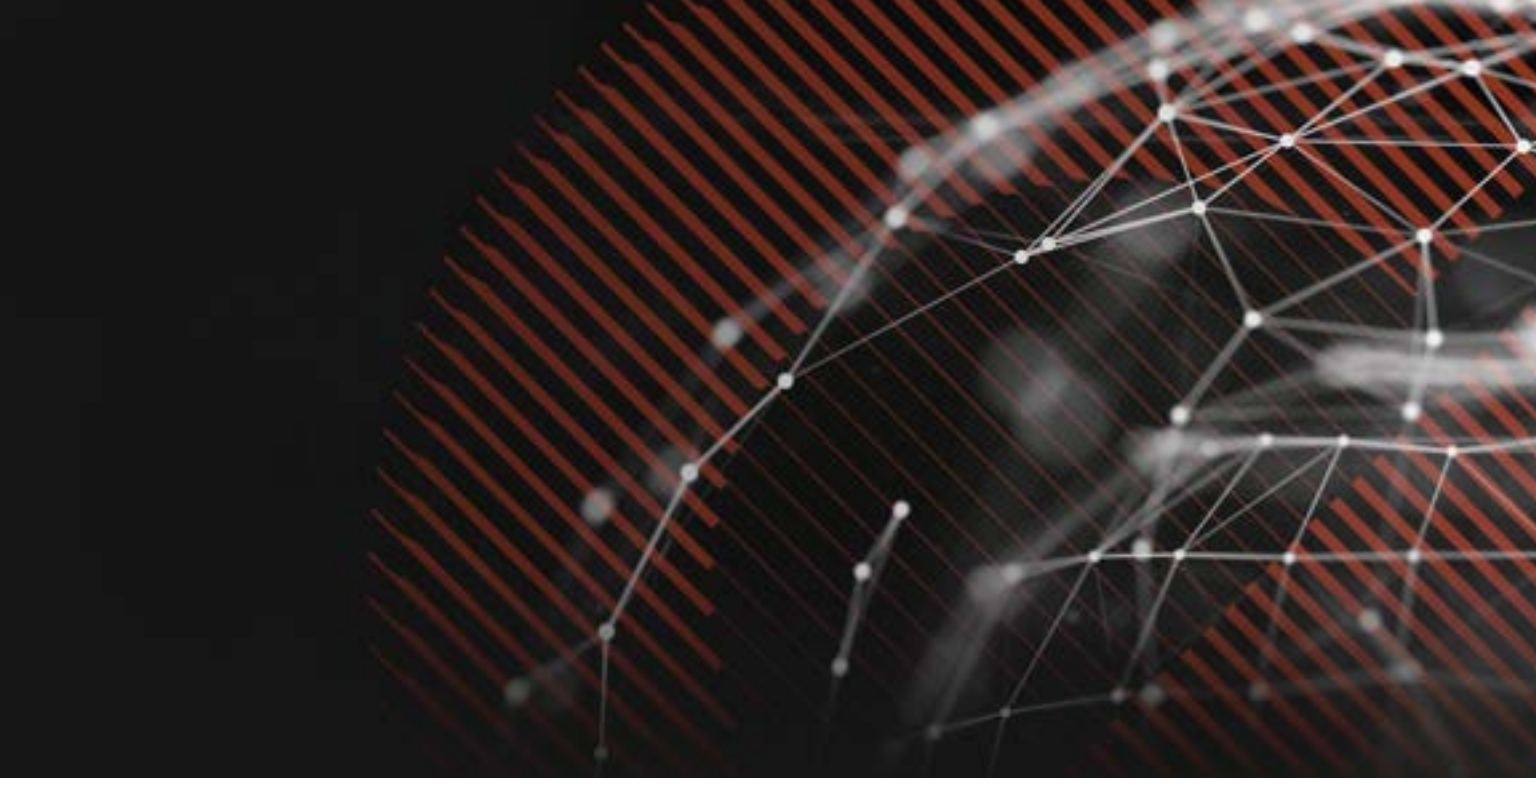

The Rising Stakes of Cyber Resilience

What the 2025 Global Incident Response Report Means for Business Leaders

BY SAM RUBIN

Cyberthreats were once isolated to breaches, technical risks, and financial extortion. In recent years, however, they have become a direct threat to business continuity. Ransomware has morphed into a cataclysm of multilayered extortion schemes; cloud vulnerabilities have become an enterprise-wide risk; and the speed of cyberattacks has outpaced traditional defenses.

So what can companies do? The findings in the 2025 Global Incident Response Report make one thing clear: security is more than just about preventing breaches. It's about ensuring that organizations can withstand, recover from, and outmaneuver cyber disruptions that are increasingly engineered for maximum operational impact. In short, business leaders must stop viewing cybersecurity as a function of IT alone — treating resilience not as a defensive measure, but as a core driver of growth, continuity, and competitive advantage.

The Three Defining Cyber Trends of 2025

This year's "Global Incident Response Report" highlights three defining trends that demand an immediate recalibration of how businesses approach security.

1. Ransomware Has Become a Business Disruption Weapon

Cybercriminals have evolved beyond locking up files and demanding payment. They are exfiltrating data before encryption, threatening to leak sensitive information, and intentionally disrupting business operations. The numbers are stark:

- 92% of ransomware incidents in 2024 still involved encryption.
- 60% also included data theft, amplifying reputational and regulatory risks.
- 13% escalated to harassment, with attackers targeting employees and customers to force payment. ››

Business leaders must stop thinking in terms of data loss alone. The real risk is operational paralysis, reputational destruction, and regulatory fallout. Organizations should assume their data will be stolen and, more importantly, plan accordingly. *Reactivity is not a solution.* Proactive security measures, like AI-driven threat detection, Zero Trust architectures, and rapid response playbooks, are now mandatory.

2. Cloud and Identity Are the New Attack Frontiers

With more businesses relying on cloud-first and hybrid environments, attackers have shifted their focus:

- 29% of all incidents in 2024 involved cloud infrastructure.
- 70% of the incidents happened on three or more fronts, underscoring the need to protect endpoints, networks, cloud environments and the human factor in tandem.
- In nearly half of cloud breaches, attackers exploited misconfigured identity and access controls.
- Threat actors are exfiltrating cloud data before destruction, ensuring they can still extort organizations even if they refuse to pay ransom demands.

The traditional security perimeter is gone, and with it, the idea that cloud security is simply "an IT problem." Because an identity breach isn't an IT failure — it's a business-wide failure. When a single compromised credential brings operations to a halt, security leaders must prioritize identity-first security strategies. They must also enforce least-privileged access, continuous monitoring, and AI-driven cloud security controls that operate at the speed of today's threats.

3. The Speed of Attacks Has Outpaced Traditional Defenses

The report also confirms a sobering truth that many have long believed: Cybercriminals have already weaponized automation and AI, launching attacks at speeds that human-led security teams simply cannot match:

- Median time from breach to data exfiltration: 2 days.
- 25% of cases saw exfiltration within 5 hours — three times faster than the exfiltration stats in 2021.
- Nearly 20% of incidents saw exfiltration in under an hour.

> Cyber resilience can no longer be siloed within security teams. It must be a C-suite priority, with measurable outcomes and clear accountability.

For organizations relying on manual detection and slow response times, this reality is a wake-up call. Cyber resilience is now about operating at machine speed — leveraging AI-driven threat intelligence, automated response systems, and continuous security posture assessment to outpace attackers rather than just reacting to them.

What Must Change: From Cybersecurity to Cyber Resilience

The traditional security playbook — detect, contain, remediate — is necessary, but not sufficient. Organizations must shift their mindset from defense to resilience by embedding cybersecurity into broader business risk management.

1. Make Security a Continuous, AI-Driven Function

Cyber resilience shouldn't be restricted to periodic audits or compliance checklists. The shift must be toward real-time, AI-powered security operations that detect, analyze, and neutralize threats before they escalate.

- AI-powered SOCs should function as autonomous detection engines, flagging anomalies and prioritizing risks based on real-time attack data.
- Zero Trust architectures must continuously validate access — ensuring credentials, identities, and permissions remain secure even as users and workloads move across hybrid environments.
- Continuous validation means security isn't a one-time exercise; it's a living system that evolves as fast as the threats do. ››

2. **Rethink Risk as a Business Decision, Not a Security Concern**

CISOs have long struggled to communicate cyber risk in terms that resonate with boards and executive teams. That must change.

- **Quantify cyber risk in financial terms.**

  If ransomware could cost $25M in downtime, lost revenue, and regulatory fines, leaders must calculate that as a business risk, not just a security risk.

- **Link resilience to competitive advantage.**

  Companies that recover faster from cyberattacks will outperform those that flounder. Cyber resilience isn't just about avoiding losses — it's about protecting market position.

3. **Align Cloud, Identity, and Security into a Unified Strategy**

Attackers don't distinguish between cloud and enterprise, so why should defenders?

- **Eliminate silos between cloud and SOC teams.**

  Identity security, runtime security, and endpoint protection must be operationalized as a single ecosystem.

- **Enforce intelligent identity controls.**

  With half of cloud breaches tied to misconfigured access, AI-driven security must continuously assess permissions and close exploitable gaps.

4. **Redefine Success: The Fastest Recovery Wins**

Security has long been measured in how many breaches were prevented. But in 2025, resilience is defined by recovery speed.

- **Shift KPIs toward resilience metrics.**

  How fast can you detect, isolate, and remediate an incident before it disrupts business operations?

- **Automate the recovery playbook.**

  Powerful incident response can shift from crisis management to a strategic advantage. The organizations that respond fastest don't just recover — they win.

Cyber Resilience Is a Boardroom Issue — Not Just a Security Concern

If nothing else, the report makes one thing clear: Cyber resilience can no longer be siloed within security teams. It must be a C-suite priority, with measurable outcomes and clear accountability. Here's how:

- **CIOs** must champion AI-driven security, embedding adaptive defenses that move at machine speed.
- **CISOs** need to shift from compliance to resilience, prioritizing AI-powered risk analysis and real-time incident response.
- **CFOs** must quantify cyber risk as a financial metric, aligning security investments with measurable business impact and ROI.
- **CEOs** must lead from the front, embedding security into organizational culture and making resilience a pillar of growth strategy.

The Future of Business Resilience Starts Today

Cyberthreats have become crucial, boardroom-level business concerns. The companies that survive and thrive in the next decade won't be the ones that simply react to attacks. They will be the ones that embed cyber resilience into the core of their business — ensuring security, continuity, and market leadership in an era where digital disruption is the new normal.

The C-suite used to ask: "How secure are we?" Today, they must ask: "How prepared are we to outmaneuver these inevitable attacks?" ◆

paloalto NETWORKS | UNIT 42

Threat insights that matter.

Trusted expertise from leaders in incident response.

EXPLORE FINDINGS >

Is There a Cyber Cold War? How Nation-States Are Reshaping the Threat Landscape

BY WENDI WHITMORE

We are already in a new kind of global conflict — a cyber cold war — and it's unlike anything we've seen before. Today's geopolitical tensions aren't playing out solely through sanctions or soldiers. They're unfolding invisibly, relentlessly, in the digital shadows. That's where ransomware, espionage, and AI-powered attacks are being deployed by nation-states to disrupt economies, sabotage infrastructure and destabilize societies. This is about stealing secrets and undermining operational continuity, sowing distrust and reshaping the global balance of power.

This backdrop of geopolitical uncertainty only increases the imperative of doubling down on a modern, cyber-defensive posture. Our adversaries certainly aren't sitting on their hands — and neither can we.

With cyberthreats representing potentially existential risks to commercial organizations' and militaries' ability to conduct their most fundamental operations, both CIOs and CISOs must be directly involved in their organization's cyberdefenses. That being said, CIOs must also keep in mind that this level of security defense and resilience isn't primarily an IT function. Rather, they need to focus on geopolitical intelligence and strategic planning, as well as using those tools to marshal support and direction from the rest of the C-suite and board of directors from a business and operational perspective.

The Rules Have Changed

In the original Cold War, the world's most powerful nations built up arsenals of nuclear weapons and played a careful game of deterrence. In today's environment, that deterrence has given way to digital aggression. Nation-states are gathering intelligence and working systematically to compromise infrastructure, steal intellectual property and trigger widespread disruption.

The usual players remain: China, Russia, Iran and North Korea. But, the tools of this war aren't tanks or missiles. They're malware strains, zero days, deepfakes, credential theft and artificial intelligence. At Palo Alto Networks Unit 42, we've investigated incidents where North Korean attackers posed as recruiters to deploy malware disguised as developer tools — and that is just one recent operation among many. ››

These operations are escalating. Cyber campaigns linked to nation-states are becoming more targeted, more coordinated and more emboldened. Our adversaries are moving beyond espionage toward sabotage.

Today's Target-Rich Environment

No organization is immune. Government agencies, power plants, financial firms, healthcare systems and tech companies are all in scope. The rise of distributed workforces, cloud migration and IoT has expanded the attack surface exponentially.

Nation-state actors are increasingly partnering with cybercriminal gangs to obscure attribution and share tools. This alliance of capability and deniability makes them harder to detect and disrupt. Even the most mundane endpoint — a smart thermostat, a printer, a contractor's laptop — could be the first domino to fall in the compromise of a whole network.

These threat actors are as creative as they are determined. The Unit 42 Threat Intelligence unit tracked activity from suspected North Korean cyberattackers posing as recruiters or prospective employers. Their trick? Asking potential "employees" to install malware that seems to be actual development software as part of the hiring process.

What Organizations Can Do in the Age of Geopolitical Risk

The cyber cold war is a real threat, with real implications. As such, it requires real-time and actionable solutions, as well as long-range planning. Complicating this dynamic threat landscape is the rise of a regulatory environment that requires businesses and organizations across all sectors to bolster their cyber resilience and better protect critical data.

Data protection and cybersecurity laws are proliferating throughout the world, led in large part by the European Union's landmark Global Data Protection Regulation. In addition, the Securities and Exchange Commission's new cyber disclosure rules require public companies to report breaches faster and more fully. This exerts more pressure on CIOs, CISOs and their teams to respond to rapidly changing regulations and the potential legal consequences of failing to comply with these emerging requirements.

Because this cyber cold war has been forming and transforming for a while, a blueprint of best practices is emerging for organizations' benefit. Some specific recommendations include:

- **Integrate geopolitical risk into business continuity planning.** This isn't optional. If your supply chain, customer data or cloud infrastructure spans borders, you're likely exposed to these transnational threats and the emerging regulatory efforts to counter those adversarial actors.
- **Shift from perimeter security to identity-first, AI-enabled defense.** In this new cold war, attackers move fast and hide well. Only AI-powered platforms can respond at machine speed — the way attackers already are.
- **Invest in cloud security with global supply chains in mind.** Nation-state attackers don't care where your workloads live. But they will exploit any misconfiguration, gap or delay in detection.
- **Operationalize threat intelligence.** Your teams need access to insights from groups like Unit 42, and not just the one-off threat report, but the continuous stream of intelligence to better inform your SOC, your infrastructure strategy and your updates to the board.
- **Rethink your role.** You are both the steward of systems and the strategist responsible for business resilience. That includes preparing for the geopolitical risks that now shape the global business landscape.

The Cold War May Be Digital — But the Consequences Are Real

The battlefield has changed, but the stakes are higher than ever. Full-scale disruption of your operations is no longer a hypothetical. The only question is whether you'll see it coming and whether you're prepared to respond.

CIOs who recognize the scale of this shift — and act decisively to modernize their defense posture — will emerge as critical strategic partners in the boardroom. Those who don't will face security failures and broader risks to your operational readiness and reputation, potentially exposing you to regulatory consequences.

The cyber cold war isn't looming. It's here. And now is the time to lead like it. ◆

Supply Chain Chaos in 2025

How Geopolitics Are Rewriting the Rules

BY HELMUT REISINGER

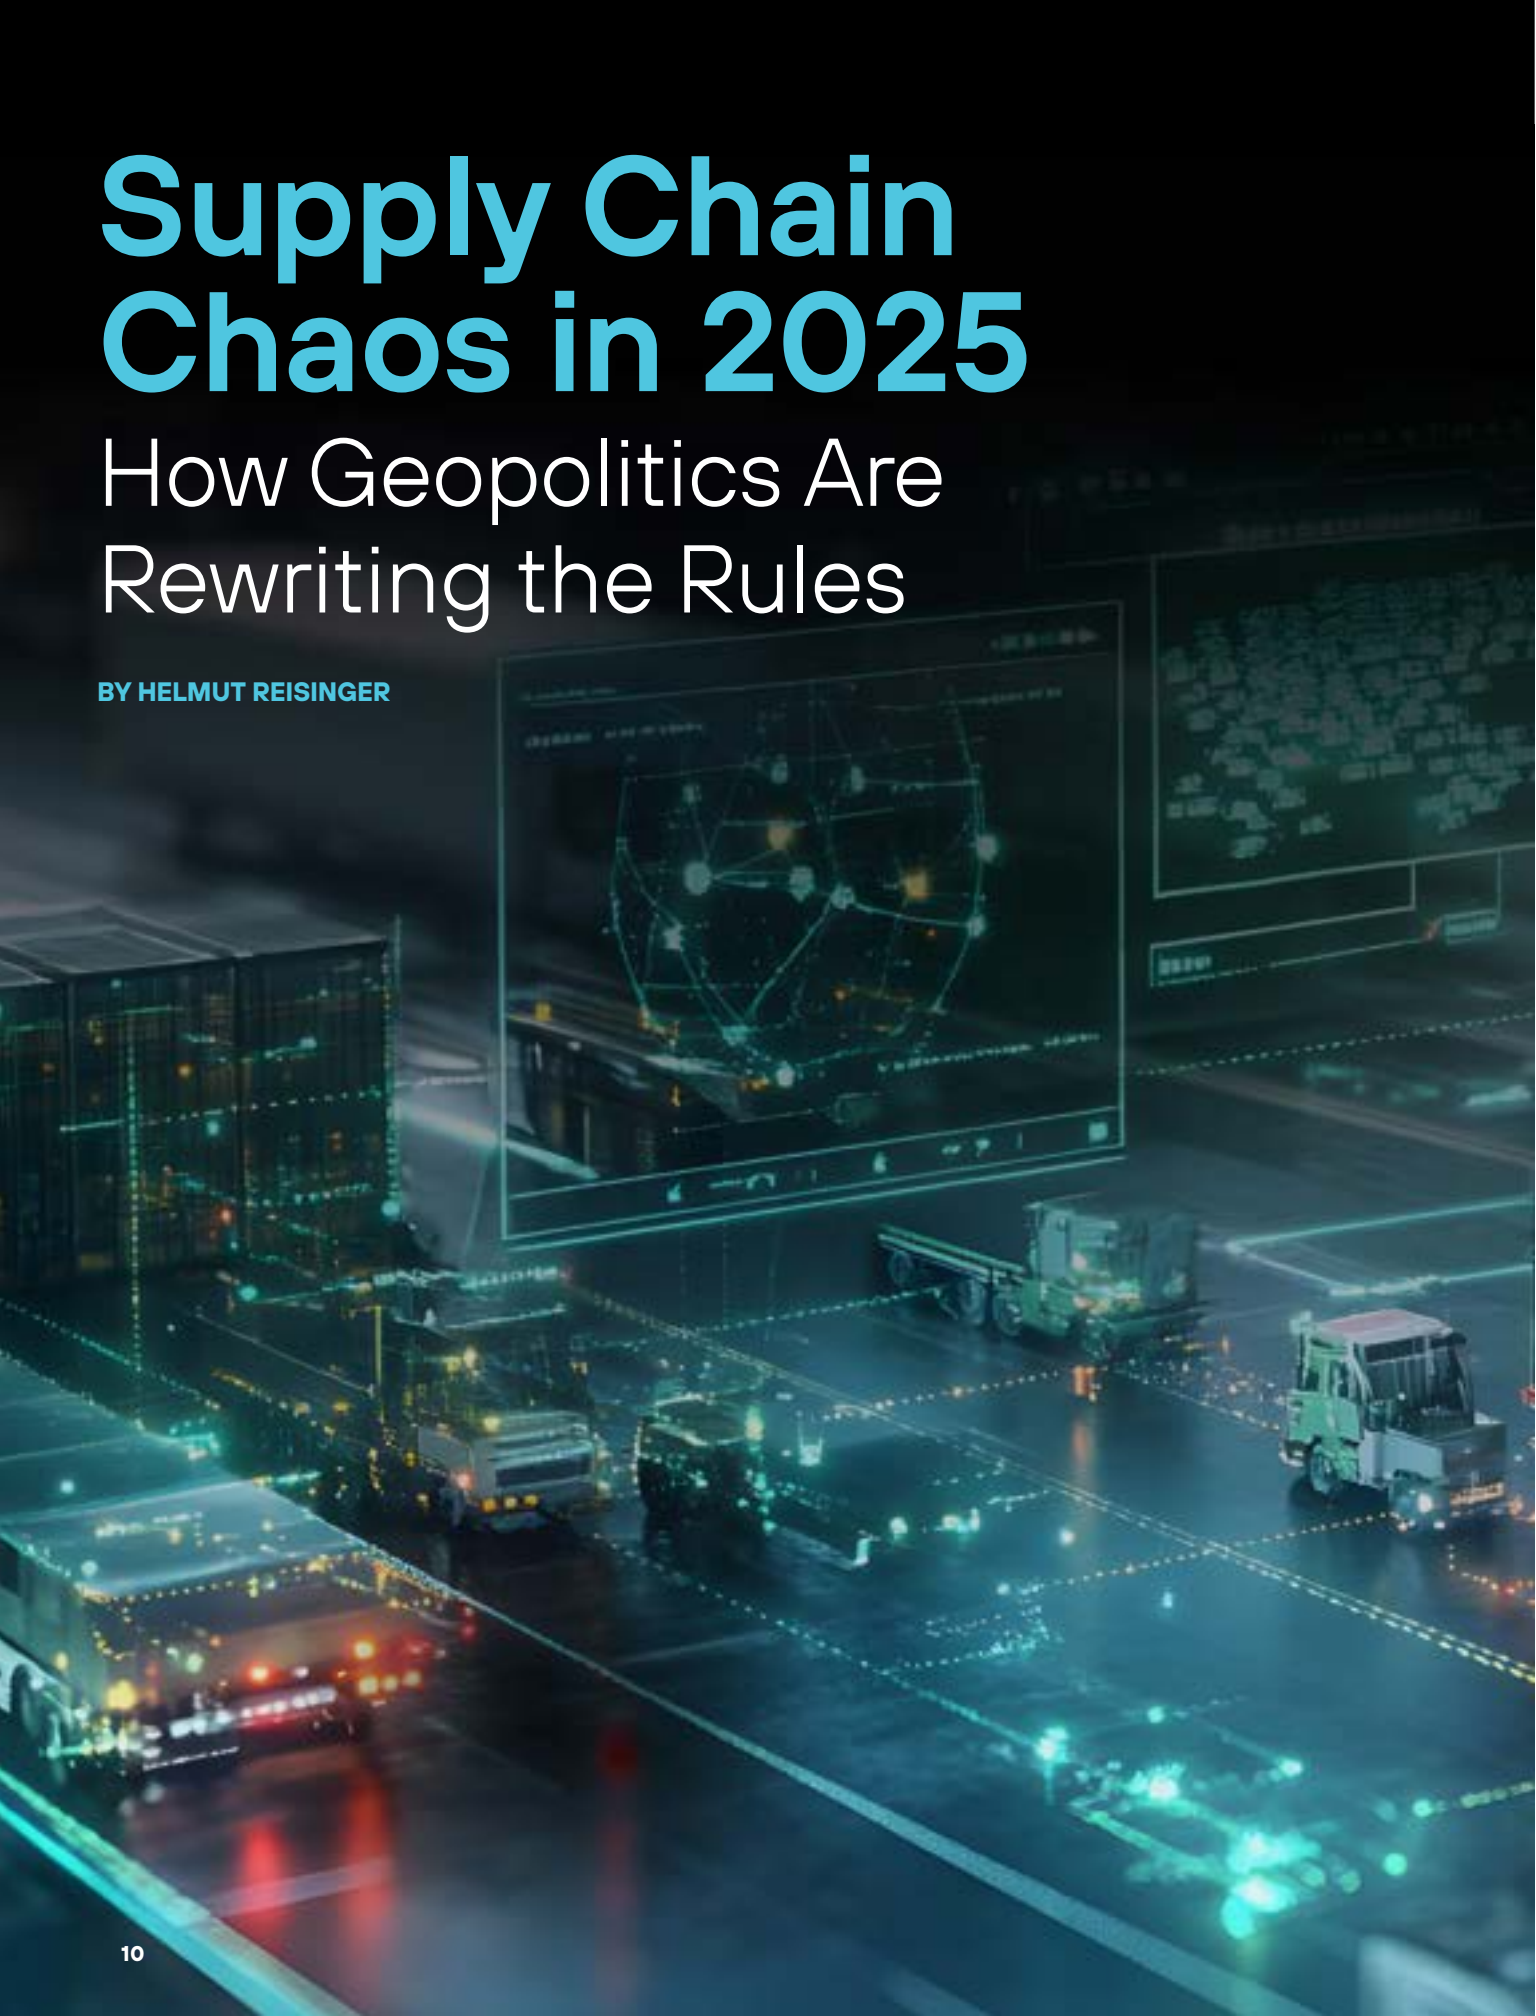

In 2025, global supply chains are no longer just operational marvels — they are geopolitical flashpoints. Once optimized for cost and efficiency, these complex webs of vendors, partners, and logistics networks have become prime targets in an era of escalating cyber aggression. As political tensions spill into cyberspace, state-aligned attackers are disrupting government systems and infiltrating the digital arteries of commerce itself. From ports to payment systems, supply chains are under siege. And the consequences aren't theoretical. They're operational. Financial. Existential.

Political unrest, sanctions, and digital sabotage have turned once-stable logistics networks into strategic liabilities. The old rules no longer apply. Organizations must confront a hard truth: Supply chain resilience can no longer be separated from cybersecurity — or geopolitics.

A Global Network Under Siege

Today's supply chains are vast, intricate ecosystems — sprawling across continents, supported by thousands of vendors, and stitched together by digital infrastructure that was never designed for geopolitical warfare. What once symbolized economic efficiency has become a strategic vulnerability.

The weakest link is no longer theoretical. As Palo Alto Networks reported, nearly one third of breaches in 2023 originated through third-party access. A single misconfigured device, a forgotten login, or a contractor with outdated credentials can offer adversaries a direct corridor into critical operations.

Nation-states and their proxies have taken notice. In an era of rising global instability — from armed conflict and economic sanctions to political fragmentation — supply chains have become a high-value target. These attacks are calculated, opportunistic efforts to destabilize markets, erode trust, and project influence far beyond the battlefield. In this new calculus, disruption itself has become the point.

From Cost Efficiency to Risk Efficiency

Global supply chains were once prized for their speed, scale, and cost efficiency. But in 2025, those same attributes have become liabilities. The world has changed, and the calculus has too. The real question for CISOs and chief risk officers is no longer: "How lean is our supply chain?" It's: "How fast can we isolate and recover when — not if — a trusted partner is compromised?"

This isn't a theoretical exercise. In regions like EMEA and LATAM, where commerce crosses borders, cloud adoption is accelerating, and geopolitical tensions are never far from the surface, supply chains are especially exposed. Risk now travels as fast as data, and too many organizations are still responding at human speed.

Security teams can no longer afford to chase yesterday's threats or rely on fragmented visibility. Resilience must be real time. Strategic. Executable. It demands investment in both technology and mindset — from the boardroom down.

How Regulation and Real-Time Security Are Forcing a New Playbook

Geopolitical instability and the regulatory response to it are driving urgency. Across the EU and beyond, data protection, resilience, and breach disclosure mandates are getting sharper, faster, and more unforgiving. Frameworks like DORA (Digital Operational Resilience Act) and NIS2 (EU's updated Network and Information Security Directive) now demand more than periodic assessments or written policies. They require continuous monitoring, real-time detection, and immediate reporting often within 24 hours of an incident.

Our platformized security approach gives organizations a strategic advantage. Our data security posture management (DSPM) capabilities help enterprises locate and secure sensitive data across sprawling cloud environments — a critical step for DORA compliance. Meanwhile, our XSIAM and XDR solutions enable AI-driven, real-time threat detection and automated response, supporting NIS2's aggressive disclosure timelines and ensuring incidents are detected and contained before they escalate.

This is the power of modular platformization: Organizations can start with the capabilities they need most — whether it's securing cloud data, protecting endpoints, or building SOC automation — and expand as new risks and requirements emerge. It's AI-first, real time by design, and architected for resilience. ››

The regulatory landscape is only going to get more demanding. Organizations that treat compliance as an enabler — not a box-checking exercise — will be best positioned to move with confidence in a high-stakes world.

What Playbook Do You Need Today? It's Not as Complicated as You May Think

You may ask yourself: What does a modern supply chain defense look like in practice? Well, it starts with a different playbook — one grounded in real-time visibility, AI-powered precision, and shared accountability. Instead of focusing on making their global supply chains more cost efficient, it is imperative that organizations place cyber resilience at the top of their modernized global supply chain strategy.

We've seen how today's most resilient organizations are rewriting the rules. The goal is no longer just defense. It's continuity under fire. Here's how forward-looking leaders are building security into the fabric of global supply chains:

- **Designing resilience from the start:** Zero trust can't stop at the enterprise boundary. The best organizations extend its principles across their vendor ecosystems, limiting access, enforcing segmentation, and continuously validating trust.
- **Using AI to match the speed of modern threats:** Adversaries are already exploiting AI to find and weaponize vulnerabilities. The countermeasure is precision — AI-powered platforms that automate detection, triage, and response before threats escalate.
- **Achieving visibility across complex ecosystems:** In a multicloud, multivendor world, fragmented security tooling creates blind spots. Platformized security enables unified intelligence and a single, actionable view of risk.
- **Making cybersecurity a core procurement function:** Security must be baked into global sourcing decisions. That means vetting vendor hygiene, enforcing measurable standards, and elevating cyber due diligence in M&A and expansion playbooks.
- **Collaborating across borders to stay ahead of global threats:** Security is no longer a regional responsibility. EMEA and LATAM leaders must engage in cross-border intelligence sharing, joint incident response, and regulatory coordination to outpace increasingly global adversaries.

But none of this transformation happens without imagination. As my colleague Haider Pasha recently wrote, "We are in greater jeopardy than ever at compromising our cyber resilience — our ability to rebound immediately and fully from a cyberattack with minimal operational impact — unless we stretch our imagination."[1] AI, analytics, and automation are essential tools, but they're not enough on their own.

Cyber resilience also demands leadership. Cybersecurity expert Ria Thomas underscores that resilience is not the responsibility of CISOs alone.[2] It must be driven by the full C-suite and board. That means the VP of operations or supply chain management can't go it alone.

Cybersecurity is a team sport. And safeguarding global supply chains requires 100% organizational alignment — from procurement to the boardroom.

Geopolitical conflicts may shift or fade. But the threat to global supply chains won't. The organizations that thrive in this era won't just adapt their networks; they'll rewire their priorities. Cyber resilience isn't a regulatory checkbox or an IT mandate. It's a strategic imperative.

Remember: Cyber Resilience Is Still a Board-Level Priority

This moment demands executive leadership. Supply chain risk can no longer sit solely within procurement, logistics, or even IT. It must be addressed at the C-level, with shared accountability across the organization. The goal is to both avoid disruption and build adaptive capacity in the face of it.

That's what resilience means: the ability to continue operating, serving, and growing — even amid geopolitical volatility. Because what once optimized commerce must now be what protects it. ◆

1. "When it Comes to Cyber Resilience and AI, Be Sure to Stretch the Limits of Your Imagination," Palo Alto Networks, March 2025.
2. *Beyond Compliance: The Human Element of Cyber Resilience*, Navigating the Digital Age, 2018.

Your Vendor's Cyber Failure Will Become Your Next Crisis

BY HAIDER PASHA

As organizations scale and modernize, their reliance on third-party vendors grows in parallel. From payroll processors and patent counsel to software providers and logistics partners, these external relationships have become essential to business operations. But each new vendor connection also opens a new door to cyber risk — and far too many of those doors are left unguarded.

Third-party cybersecurity risk isn't a theoretical concern; it's an escalating, enterprise-wide threat with the potential to trigger operational disruptions, reputational damage, and regulatory fallout. Yet despite its severity, many organizations continue to delegate vendor risk management to procurement, where the approach is often reduced to annual checklists and self-assessment questionnaires.

That's not just outdated — it's dangerous.

Research from PwC notes that only 31% of companies assess vendor cybersecurity risk through formal, organization-wide processes.[1] The rest are flying blind. In a world where even midsize firms may manage dozens (or hundreds) of vendor relationships — many with privileged access to sensitive systems and data — this status quo is untenable.

Third-party vendor cybersecurity risk represents an existential threat to nearly every organization. And it demands executive-level urgency. ››

Third-party risk isn't just a cybersecurity problem. It's an enterprise vulnerability — one that demands constant vigilance.

How Third-Party Vendor Risks Turn into Cyber Vulnerabilities

There have been numerous high-profile, headline-grabbing examples of third-party vendor risks turning into massive cybersecurity problems. The highly publicized SolarWinds attack from 2020 and the 2013 Target breach are vivid examples of what happens when third-party risk becomes an open backdoor to your business. These, too, are far from isolated cases. Across every industry, attackers are exploiting the weakest links in digital supply chains — often with devastating results.

Today's attackers aren't limiting themselves to traditional IT vendors. They're just as likely to target financial service providers, cloud and telecom partners, or even the power company. If a vendor connects to your systems — directly or indirectly — they're in scope. That includes software developers, OEMs, distributors, and increasingly, customers.

Once inside, attackers don't rush. They move laterally across networks, hunting for sensitive data: customer records, intellectual property, credentials. Many embed malware inside APIs, browser plug-ins, or update mechanisms — slipping past defenses by mimicking trusted processes.

The software supply chain is particularly exposed. In fact, research from Enterprise Strategy Group points out that 41% of organizations' software supply chains have been hit with zero-day attacks, exploiting new or previously unknown vulnerabilities in third-party code, while 40% of organizations report they've been hit with exploits of a misconfigured cloud service.[2]

These aren't edge cases. They're warning signs. And they confirm what many CISOs already know: Third-party risk isn't just a cybersecurity problem. It's an enterprise vulnerability — one that demands constant vigilance.

The High-Stakes Consequences of Vendor-Based Attacks

Third-party attacks often unfold quietly — and by the time they're discovered, the damage is already done.

A single compromised partner can expose sensitive data, disrupt operations, and force an organization into an extended cycle of forensic investigation, remediation, and recovery. In these moments, it's not just systems that go down. It's trust — with customers, regulators, partners, and the public.

The regulatory consequences alone can be staggering. From Europe's GDPR to Brazil's LGPD and the U.S. healthcare industry's HIPAA, data protection frameworks now hold organizations accountable for breaches, regardless of where the vulnerability originated. Financial penalties are one thing. Long-term reputational damage is another.

This is what makes third-party risk so dangerous: it scales with your growth. Every new vendor, every additional integration, every expansion into a new market increases the attack surface. Without real-time visibility into partner ecosystems, that surface becomes a blind spot — one that adversaries are increasingly skilled at exploiting. ››

What Organizations Should Do — Now

The rise in third-party risk demands more than a procedural response — it requires a mindset shift. Traditional approaches such as static audits, vendor questionnaires, and one-time compliance checkboxes no longer suffice in an era where the attack surface is continuously expanding through external relationships.

Here's what needs to change:

- **Classify vendors by business criticality.** Not all third parties carry equal risk — and your cybersecurity expectations should reflect that. Organizations should adopt a tiered model that assigns suppliers into risk categories based on their operational importance and level of system access. For critical vendors, enforce full adherence to your Zero Trust policies, including rigorous identity verification, segmentation, and continuous monitoring. For mid-tier or low-risk suppliers, ensure baseline controls are met, but scale the requirements proportionally.
- **Move from point-in-time to real-time risk assessment.** Third-party environments are dynamic — what's secure today may be vulnerable tomorrow. Risk evaluations must evolve accordingly. Regularly reassessing long-standing vendor relationships is just as critical as scrutinizing new ones.
- **Insist on Zero Trust principles — everywhere.** A Zero Trust model should apply across your extended enterprise, including vendors. If partners aren't segmenting access, validating identity at every point of entry, and monitoring anomalous behavior, then your defenses are only as strong as their weakest node.
- **Align vendors to your own policy architecture.** Too often, partners operate under looser protocols. Instead, organizations should require vendors to adhere to internal policy frameworks — from data handling to incident response — with no exceptions.
- **Use modern threat intelligence and automation.** Real-time visibility into third-party risk is possible today, but it requires investment in intelligent tooling. AI and machine learning can surface vulnerabilities before they're exploited, especially when continuously fed with live telemetry and threat intelligence.
- **Proactively share threat intelligence across critical suppliers.** Organizations must do more than secure their own perimeter — they must uplift the collective resilience of their ecosystems. At a minimum, critical suppliers should participate in bidirectional intelligence sharing, especially in industries like energy, healthcare, and retail where these practices lag behind those in financial services. A compromised vendor is still a breach on your books.
- **Extend accountability across the ecosystem.** Your third parties aren't just business relationships — they're extensions of your digital perimeter. And with that privilege comes responsibility. Make continuous security monitoring part of your procurement lifecycle, not an afterthought.

Ultimately, the question isn't whether your vendors are a risk. It's how quickly you can identify which ones are — and what you do about it. As regulatory scrutiny rises and cyberattacks grow more sophisticated, there is little room left for assumptions or trust-by-default.

Raising the bar on third-party risk isn't just about avoiding the next breach. It's about protecting the business you've built — and the reputation you can't afford to lose. ◆

1. "How SOC reporting can help assess cybersecurity risk management in third-party relationships—and beyond," PwC, 2022.
2. *The Growing Complexity of Securing the Software Supply Chain,* Enterprise Strategy Group, May 2024.

Securing AI Agents

Building the Landing Gear While Flying the Plane

BY NICOLE NICHOLS

When I began working on autonomous cyber agents in 2020, the timeline for real-world deployment was still measured in decades. At the time, these systems were seen as long-range bets — interesting but still mostly niche improvements for any near-term application.

Then, something changed. ››

While generative AI (GenAI) wasn't one, singular event, it unleashed an ongoing cascade of advances that are, to this day, causing development timelines to collapse at a continuously accelerating rate. This isn't just a case of moving the goal; the GenAI-driven wave is relentlessly bulldozing old benchmarks and redefining the frontier of what's possible, faster than we've ever experienced before. Capabilities once reserved for long-term research are now being integrated into live environments with astonishing speed.

Startlingly, but not surprisingly, agentic systems are being embedded in countless locations — company workflows, decision-making pipelines and even critical infrastructure — often before we've established how to govern or secure them. The year 2020 seems a lifetime ago considering we're no longer preparing for the arrival of agentic AI but responding to its continued and rapid evolution.

A Paper for a Moving Target

The workshop report I've co-authored, *Achieving a Secure AI Agent Ecosystem*,[1] is the product of a cross-institutional effort to make sense of this acceleration. Developed in partnership with RAND, Schmidt Sciences, and leading minds in agentic AI from across industry, academia, and government, the paper doesn't offer silver bullets but rather a different way to think about and approach agentic AI.

The crux of the paper outlines three foundational security pillars for AI agents and suggests where our current assumptions — and infrastructure — might falter as these systems evolve. Beyond simply acknowledging current realities, this argues for a profound mindset shift: We must recognize that the age of agentic systems is already upon us. Consequently, securing these systems is not a problem for tomorrow. It's an urgent challenge today that's intensified by the relentless pace of innovation, expanding scale, uneven risks for early adopters and the stark asymmetry between attack capabilities and defense goals.

One of the challenges in securing AI agents is that these systems don't look or behave like traditional software. They are dynamic, evolving and increasingly capable of executing decisions with minimal oversight. Some are purpose-built to automate tasks like scheduling or sorting email; others are inching toward fully autonomous action in high-stakes environments. In either case, the frameworks we use to secure traditional applications aren't enough. We are encountering problems that aren't merely variations on known vulnerabilities but are fundamentally new. The attack surface has shifted.

Three Pillars for AI Agent Security

This mindset shift is why the security landscape has been organized around three core concerns:

- **Protecting AI agents from third-party compromise:** How to safeguard the AI agents themselves from being taken over or manipulated by external attackers.
- **Protecting users and organizations from the agents themselves:** How to ensure that the AI agents, even when operating as intended or if they malfunction, do not harm their users or the organizations they serve.
- **Protecting critical systems from malicious agents:** How to defend essential infrastructure and systems against AI agents that are intentionally designed and deployed to cause harm.

These categories are not static — they are points along a spectrum of capability and threat maturity. Today, most organizations that deploy agents are dealing with the first two concerns. But the third — malicious, autonomous adversaries — looms large. Nation-states were among the first to invest in autonomous cyber agents.[2] They may not be alone for long.

Navigating this new era of potent, widespread autonomous threats, therefore, demands far more than incremental refinements to existing defenses. It requires a foundational shift in how our expert communities must collaborate and innovate on security.

Historically, AI researchers and cybersecurity professionals often operated on parallel tracks, holding different assumptions about risk and architecture. Yet, the complex frontier of agentic AI security demands their unified effort, as neither community can tackle these immense challenges in isolation — making deep, sustained collaboration paramount. And while universal protocols and comprehensive best practices for this entire field are still maturing, the notion that effective turnkey ››

products for securing agents are scarce is, frankly, becoming outdated. Sophisticated, deployable solutions are now offering vital, specialized protection for critical agentic systems, signaling tangible progress. This further underscores the urgent need for adaptive, multilayered security strategies — spanning model provenance, robust containment and resilient human-in-the-loop controls — all evolving as rapidly as the agents themselves.

Interventions Within Reach

While robust and evolving product solutions are increasingly crucial in mitigating the immediate operational risks posed by agentic AI, achieving comprehensive, long-term security also necessitates dedicated industry-wide investment in foundational capabilities and shared understanding. Several such key directions, complementing product innovation, are well within our collective reach and warrant-focused effort.

For instance, a kind of "agent bill of materials," modeled after the "software bill of materials," is envisioned to provide visibility into an agent's components like its model, training data, tools and memory. However, its functional viability currently faces hurdles, such as the lack of a common system for model identifiers, which is crucial for such transparency.

Additionally, standardized, predeployment test beds could allow for scalable, scenario-based evaluations before agents are released into production environments. And communication protocols like MCP (Model Context Protocol) and A2A (Agent-to-Agent) are emerging, but few have security baked in from the start. However, even when security measures are integrated from the outset, the prevalence of "unknown unknowns" in these novel agentic systems means these protocols will require rigorous and continuous assessment to maintain their integrity and safety.

One approach our paper attempts to navigate is the critical challenge that an agent's memory, while essential for it to learn, improve, and crucially avoid repeating past mistakes, is also a significant vulnerability that can be targeted for malicious tampering. The strategy involves using "clone-on-launch" or task-specific agent instances. In this model, agents designed for particular operational duties or limited-duration interactions treat their active working memory as ephemeral. Once their specific task or session is complete, these instances can be retired, with new operations handled by fresh instances that are initialized from a secure, trusted baseline.

This practice aims to significantly reduce the risk of persistent memory corruption or the lingering effects of tampering that might occur within a single session. It is paramount, however, that such a system is meticulously architected to ensure an agent's core foundational knowledge and long-term learned lessons are securely maintained, protected against tampering, and effectively and safely accessible to inform these more transient operational instances. While managing operational states in this manner is not a comprehensive solution to all memory-related threats, it represents the kind of creative, systems-level thinking required for advancing agent security and robust containment.

A Call for Shared Commitment

Ultimately, securing agentic AI will not come from any single breakthrough but from a sustained, multistakeholder effort. These include researchers, policymakers, practitioners and industry leaders working together across disciplines. The threats are both technological and foundational. We are trying to secure systems that we do not yet fully understand. But if there's one thing the last few years have made clear, it's this: Waiting to act until the picture is complete means acting too late.

The evolution of agentic AI means our industry is developing critical safeguards concurrently with its widespread adoption. This simultaneous development isn't inherently a crisis, but a clear call for collective responsibility. Our success in this endeavor hinges on a shared industry commitment to building these foundational elements with transparency, rigorous standards and a unified vision for a trustworthy AI ecosystem. ◆

1. *Achieving a Secure AI Agent Ecosystem*, Palo Alto Networks, RAND, and Schmidt Sciences, June 17, 2025.
2. *Autonomous Cyber Defence Phase II*, Centre for Emerging Technology and Security, May 3, 2024.

paloalto NETWORKS | CYBERSECURITY PARTNER OF CHOICE

REPORT

Your guide to secure generative AI adoption in 2025.

BE PROACTIVE

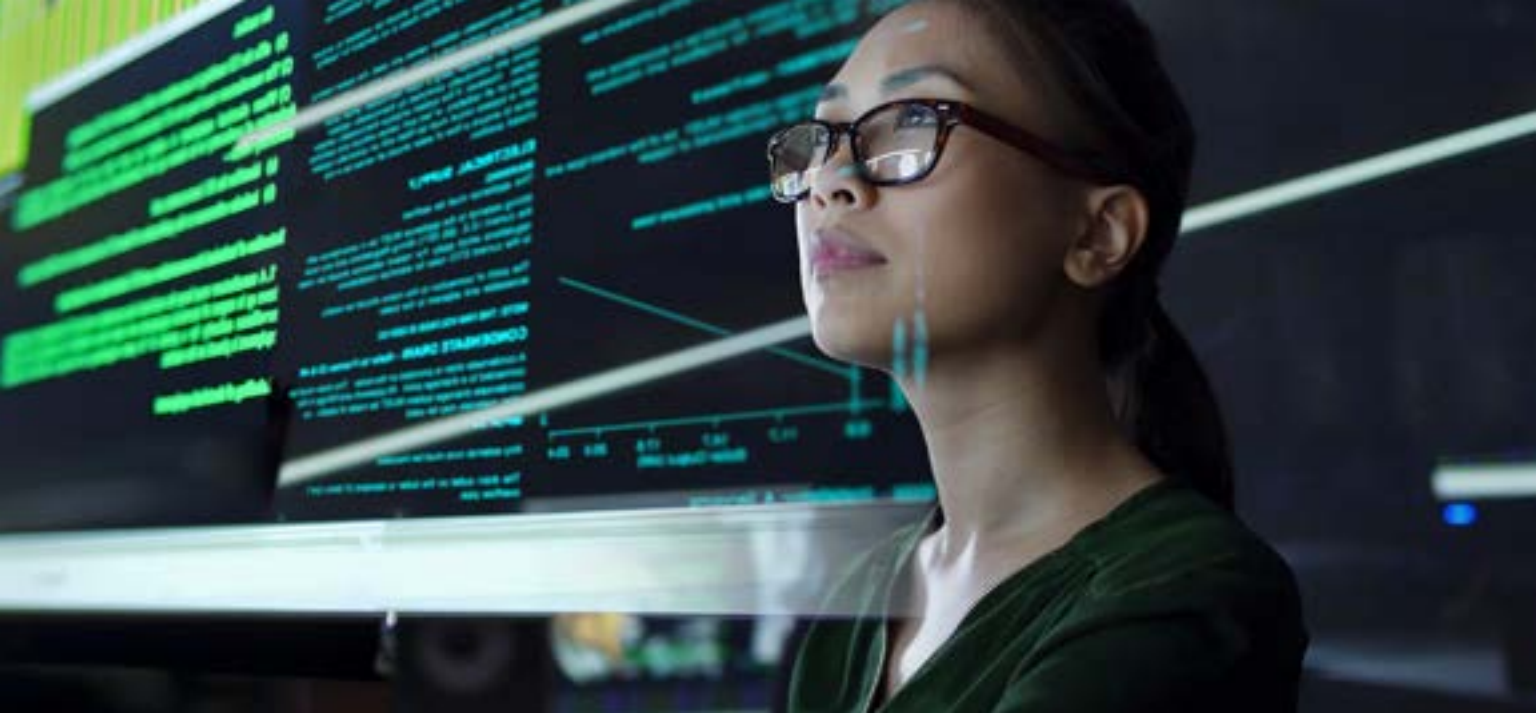

Hybrid Attacks in the Age of AI: How Cloud-SOC Convergence Is Our Best Defense

BY KARIM TEMSAMANI

Hybrid attacks — those that traverse enterprise and cloud environments with unsettling ease — have become a defining challenge of modern cybersecurity. These attacks are faster, more adaptive, and more complex than anything we've seen before. This isn't a theoretical concern; it's today's reality. To defend against these modern threats, not only products but security operations as a whole must evolve.

And within this challenge lies a pivotal opportunity. Artificial intelligence (AI), the very technology that attackers wield to automate and adapt, can become the foundation of a more resilient, responsive defense. The solution lies in unifying cloud and enterprise security operations through AI-driven automation and intelligence. When done by design, security teams gain the clarity and speed to outmaneuver adversaries. This evolution isn't just about defense — it's about empowering innovation and securing the foundations of tomorrow's digital growth.

New Threats, Blurred Boundaries

Hybrid attacks are growing more frequent and should be considered a major threat avenue for adversaries. Attackers move laterally across cloud and on-premises environments, exploiting fragmented defenses. Consider the modern cloud environment: workloads scale dynamically, applications deploy in real time, and data moves fluidly across geographic and organizational boundaries. Traditional security approaches, rooted in static perimeter defenses and isolated tools, are no match for this reality.

Not surprisingly, AI has amplified this complexity. Generative AI, for instance, helps attackers craft ultra-personalized phishing campaigns that evade detection, while machine learning algorithms identify and exploit cloud misconfigurations faster than teams can respond. Security leaders must accept that the line between enterprise and cloud infrastructure no ››

longer exists. We must stop thinking of cloud security as a separate domain and instead view it as an integral part of the broader security strategy.

Cloud-SOC Convergence: A Strategic Imperative

The convergence of cloud security with the security operations center (SOC) is more than a technical upgrade; it represents a fundamental rethinking of security architecture. Why? Because attackers don't distinguish between cloud and enterprise environments, and neither should defenders. When security teams work from a single, unified platform, they gain the context and capabilities needed to respond with speed and precision.

This unified approach offers several key benefits:

- **Real-time threat protection:** Cloud threats move fast — security must move faster. Cloud Runtime Security provides immediate, inline protection by detecting and blocking runtime attacks in real time, before they escalate. This agent-based approach prevents exploitation at the source, stopping attackers before they gain a foothold. AI-powered analytics then enrich security insights, identifying hidden patterns and correlating events across cloud and enterprise environments. This layered approach helps ensure organizations are proactively preventing attacks — not just detecting them after the damage is done.
- **Smarter prioritization with AI-driven context:** Not all vulnerabilities demand equal urgency. AI-powered prioritization leverages real-time runtime data, cloud posture insights, and active threat intelligence to distinguish between theoretical risks and real-world exploitation. By dynamically assessing which exposures are being actively targeted, security teams can focus on the vulnerabilities that matter most — reducing noise, eliminating guesswork, and accelerating response where it counts.
- **Automated response:** Speed is the currency of modern cyberdefense. In cloud environments, automated remediation must be immediate — isolating compromised containers, revoking credentials, and neutralizing misconfigurations before attackers can escalate. But cloud alone isn't the full picture. True resilience comes when automated response bridges cloud intelligence with the SOC, triggering enterprise-wide containment, forensic investigation, and adaptive policy enforcement.

Security as a Growth Enabler

Yes, this convergence is about stopping attacks, but it's also about enabling organizations to innovate with confidence. Enterprises that view security as a business enabler — not just a cost center — position themselves to capitalize on cloud-driven growth. With the right security foundation, enterprises can adopt AI technologies, deploy applications globally, and manage complex supply chains without compromising safety.

Moreover, we've talked at length about how platformization delivers measurable returns. And the recent study from IBM and Palo Alto Networks highlights how those organizations that embrace security platformization see stronger security outcomes, faster incident response, and better return on investment.[1] By consolidating cloud and enterprise security operations, companies can reduce tool sprawl, cut costs, and improve efficiency.

Preparing for an AI-Focused Future

The AI arms race in cybersecurity is well underway. Attackers will continue to refine their methods, and security teams must stay ahead by embracing the very technology used against them. Cloud-SOC convergence is a pivotal step in that defense. Because by unifying data, automating responses, and leveraging AI at scale, businesses can turn the tide against hybrid attacks. More importantly, they can build a resilient, adaptable security posture that supports innovation rather than stifling it.

In cybersecurity, the best defense has oftentimes been a well-informed offense. In this new era, it also requires an integrated, intelligent, and real-time approach. The future of cybersecurity isn't just in the cloud — it's at the intersection of cloud, AI, and enterprise operations. ◆

1. Mohamad Ali and BJ Jenkins, *Capturing the Cybersecurity Dividend*, IBM Institute for Business Value and Palo Alto Networks, 2025.

Zero Trust Isn't a Cybersecurity Luxury — It's the Cost of Doing Business

BY RICH CAMPAGNA

Walk through any cybersecurity conference or listen in on a corporate strategy session, and you are likely to hear two words: Zero Trust. The term has become an inescapable part of the business lexicon — shorthand for a modern, robust security posture. The data confirms this ubiquity. Gartner® says 63% of organizations worldwide have deployed a Zero Trust strategy[1] — a number clearly on a sustained upswing. Yet, beneath this veneer of consensus, Zero Trust is often an overused, but misunderstood term.

At its core, Zero Trust is not a product or a piece of software. Rather, it's a fundamental shift in how organizations approach cybersecurity. It's a model, a mindset, a strategy built on the simple yet profound principle: "Trust nothing, inspect every transaction." The best Zero Trust outcomes are achieved when a model is deployed in a way to protect everyone and everything in an organization: users (both internal and external), applications, data, and infrastructure.

In a digital landscape rife with threats, this is no longer an optional extra; it's core to a successful cybersecurity strategy.

Why Zero Trust Is Critical

Simply put, the digital world is a much more dangerous place today than ever before. Gone are the days when a simple password could safeguard an organization's data. The digital world has grown exponentially more complex and dangerous. Device proliferation, cloud computing, and the Internet of Things (IoT) — all have expanded the potential pathways for malicious actors. Hackers, armed with increasingly sophisticated tools like generative AI and machine learning, are relentless in their pursuit of vulnerabilities.

This has propelled Zero Trust into the mainstream. No user or device is automatically trusted, regardless of network location. Every access request must be verified, often through multifactor authentication. ››

If your organization hasn't started Zero Trust, begin **immediately**.

Zero Trust helps organizations manage an increasingly wide array of threats, especially zero-day attacks. It means exactly what it says: Trust no one and nothing when accessing digital assets. All requests are evaluated and validated.

In a Zero Trust environment, the goal is a robust security posture without hindering user experience. Balancing security with usability is key. Done correctly, Zero Trust delivers better security, simplifies infrastructure, lowers costs, improves security operations, AND delivers a better overall user experience.

Zero Trust Is Working

The important news is that organizations have seen tangible benefits from their initial and follow-on Zero Trust deployments. Take Zero Trust Network Access (ZTNA), a vital part of Zero Trust strategies to help employees and other users safely connect to network infrastructure, apps and services. According to Enterprise Strategy Group, 87% of organizations say Zero Trust has met all or most of the outcomes they expected when beginning their initiative, and the vast majority of those organizations use ZTNA for all of their remote access needs.[2]

Zero Trust's success stems from being an ecosystem of tools and processes, not a single product. It includes identity verification, least privilege access, continuous monitoring and device management.

Not all Zero Trust Deployments Are Equal

Zero Trust has become an essential baseline for security, but narrow adoption is not enough. The true differentiator lies in the quality of implementation, which significantly impacts an organization's readiness and resilience against modern threats. While many claim to have embraced Zero Trust, the reality often reveals wide variations. Some implementations lack the necessary depth of functionality and verification, failing to include all vital components, or are confined to specific use cases rather than the organization as a whole. This oversight leaves vulnerabilities exposed.

Similarly, endpoint coverage is crucial, yet frequently inadequate. Extending protection beyond traditional desktops and laptops to encompass smartphones, tablets, IoT devices and other connected assets is non-negotiable. Neglecting this expansion leaves a significant attack surface ripe for exploitation. Furthermore, the exclusion of third parties, such as developers and cloud providers, creates blind spots. These external entities often have access to sensitive resources and must be included in the Zero Trust framework.

Finally, any effective Zero Trust model must acknowledge the inherent weakness of passwords. Relying solely on them is a recipe for failure. Biometrics, digital certificates and passkeys are no longer optional; they are essential to a robust security posture. A thorough, comprehensive approach is the only way to ensure Zero Trust delivers on its promise of security.

What to Do in Your Organization's Zero Trust Journey

If your organization hasn't started Zero Trust, begin immediately. If it has started, it must accelerate its efforts. Cyber adversaries are continuously trying to bypass your defenses.

Here are the steps for a successful journey:

- Strong identity access management.
- East-west controls to limit lateral movement.
- Automated security responses to threats.
- Performance indicators to measure Zero Trust effectiveness.
- Continuous monitoring, evaluation and updates that align with threat intelligence.
- Clear communication of Zero Trust goals to all personnel. ◆

1. "*Top 3 Recommendations From the 2024 State of Zero-Trust Adoption Survey*", Gartner, March 18, 2024.
2. "*Organizations Should Prioritize VPN Replacement With Zero-trust Network Access*," Enterprise Strategy Group, May 28, 2024.

Beyond the Ivory Tower: The Blueprint for AI Research That Works

BY AARON ISAKSEN

Even after a career spent at the forefront of AI research and development, I can say with confidence that we are in an unprecedented moment. While AI is undoubtedly the most disruptive technology of our generation, in the world of cybersecurity, hype doesn't stop threats. Turning the immense promise of generative AI, deep learning and machine learning into tangible security outcomes requires more than just access to new models; it demands a disciplined and purposeful research philosophy. This is an essential component to Precision AI®.

Our philosophy rejects the traditional, isolated "ivory tower" of academic research and corporate R&D. Instead, it is relentlessly focused on real-world outcomes that are deeply embedded within the teams building the products, secure by its design and openly collaborative. This blueprint guides our work and is built upon four core principles that I believe are essential for making AI work for security.

1. Research in the Trenches

This philosophy begins with a foundational decision about structure. Instead of isolating researchers, we embed them directly within our product organizations for a simple reason: In cybersecurity, proximity to the place where security problems are solved is everything.

This structure ensures our researchers are focused on solving real-world problems, not abstract ones. Researchers sit alongside the engineers who build our products and the product managers who live and breathe our customers' challenges. This proximity makes the transfer of technology seamless and organic. It fosters a constant dialogue that ensures our long-term, high-risk research projects remain grounded in what will ultimately make our customers safer. ››

As the leader in cybersecurity, we have a dual responsibility: to build AI that promotes security, and to ensure the AI we — and our customers — build is itself secure.

2. Better Security Outcomes: The Only Metric That Matters

Many research organizations measure success by the number of academic papers they publish. We don't. Our primary metric is the tangible improvement our research delivers to our products and, by extension, to our customers' security.

This focus has two critical implications. First, it means that we train and evaluate our models on real-world security system data, not on sanitized "toy problems." This ensures our AI is effective in the complex, messy reality of a live security environment. Second, it gives our teams the freedom to fail. We encourage a "fail-fast" mentality, enabling us to quickly discard ideas that don't show promise and double down on those that do, without the pressure of a publication quota. Our goal is to build a portfolio of proven, effective AI for security, not to build a library of papers.

3. Security for AI, Not Just AI for Security

As the leader in cybersecurity, we have a dual responsibility: to build AI that promotes security, and to ensure the AI we — and our customers — build is itself secure. Our customers entrust us with their most sensitive system data, and protecting it is our highest priority.

This principle extends to our entire research operation. Our models are developed in highly secure environments, protected by our own best-in-class security products and secure-by-design frameworks. We meticulously vet our AI security and protect against model theft, prompt injections and other emerging threats. This may seem obvious, but secure AI is impossible without first having a strong understanding of AI security. Our commitment to this principle extends beyond our own walls; it is the core of our promise to our customers. The same security platform that protects our research is the one we offer to the industry, ensuring everyone can benefit from the lessons we learn on the frontlines of AI innovation.

4. Create a Community, Not a Fortress

Finally, we believe the best ideas come from collaboration. We actively avoid a "not invented here" mentality. Our teams are empowered to leverage the best innovations from the broader research community. These innovations include using LLM coding agents to make our own researchers more productive and generating synthetic data to make our models more robust.

We are committed to being active participants in the global conversation. We encourage our researchers to attend conferences, organize workshops, and continuously learn from what others have done. Progress in this field is a collective effort, and our goal is to contribute to and benefit from the shared knowledge of the entire ecosystem.

A Safer, Secure Future

Ultimately, these principles create a research engine that is both innovative and accountable. It's an approach designed to turn cutting-edge science into real-world security, ensuring that the power of Precision AI delivers on its most important promise: a safer digital future for everyone.

To see these principles in action and explore our deep research into securing this next wave, read our full paper, Achieving a Secure AI Agent Ecosystem.[1] ◆

1. *Achieving a Secure AI Agent Ecosystem*, Palo Alto Networks, RAND, and Schmidt Sciences, June 17, 2025.

Why Culture Is the First Line of Defense in the Age of Agentic AI

BY WENDI WHITMORE

The arrival of agentic AI is fundamentally altering how we approach cybersecurity. We're witnessing the ways attackers can reach us multiply as new tools and workflows open up fresh vulnerabilities across the board. In this new environment, where AI can effectively "take actions on its own," the sheer speed and cleverness of attacks demand more than just technological fixes. They call for a profound shift in our thinking — towards a security-conscious culture where trust and empowerment form our very first line of defense.

I firmly believe that every part of a business must embrace security as its own critical responsibility. This means ensuring our employees are well-equipped and empowered to make sound, secure decisions. It means fostering an environment where people feel comfortable speaking up when they spot something that doesn't seem right. And, critically, it means ensuring every leader across the business knows how to communicate and collaborate effectively should the worst happen and a breach occur.

The New Battlefield: Agentic AI and Our Widening Vulnerabilities

In my years specializing in computer crime investigations, including my time as a Special Agent with the Air Force Office of Special Investigations, I've seen firsthand how the front lines of cyber conflict shift. Today, it's clear that networks worldwide are the primary arena for those who wish to do harm — whether it's nation-states aiming to steal vital secrets, disrupt our critical infrastructure, or cybercriminals looking to cripple business operations for their financial gain.

Agentic AI magnifies this challenge considerably. When we talk about agentic AI, we're essentially describing AI that has been given its own "arms and legs" to take independent action — a powerful way to visualize it, as our CEO, Nikesh Arora, often describes. This reality propels us into what I can only describe as an "arms race." The question we must constantly ask ourselves is: will our defenses be nimble and smart enough to keep pace with those on the offensive, or will attackers gain the upper hand? At the heart of this race is speed—the speed with which attackers can use agentic AI to devise entirely new capabilities and coordinate their efforts with astonishing efficiency, and the speed with which we, as defenders, must detect these actions and respond effectively.

We can no longer think of our defenses like a fortress with a simple, hard outer wall. The "attack surface" — all the ways attackers can try to get in — is now much more fluid. It encompasses our mobile devices, our cloud computing environments, and what remains of our traditional networks. We need clear visibility and the ability to identify malicious actions at every conceivable point — from one computer to another, and between applications and the various layers of our digital infrastructure.

The Erosion of Trust: AI-Powered Deception

One of the things that truly concerns me about advanced AI is how cleverly it can be used for manipulation, adding another layer of complexity to our work. Attackers are already using AI in numerous ways, particularly in crafting social engineering schemes that are more convincing than ever. Language barriers, for instance, which once might have provided subtle clues of an attack, have been virtually eliminated.

This capability now extends alarmingly to voice and video. It's possible for attackers to take a mere 5 to 10-second snippet of someone's voice and then replicate it with frightening accuracy, making it incredibly difficult to detect fraudulent calls to a help desk or other deceptions that rely on voice. The rapid advancement into deepfake video capabilities further blurs the line between what's real and what's a manipulated imitation. Figuring out if you're really talking to a colleague or an AI-generated fake is going to get tougher and, I suspect, become a more common challenge.

This means we cannot solely rely on the ways we've traditionally verified identity. If an attacker's aim is to compromise someone's identity to access sensitive information, then it's paramount that all the subsequent steps in our processes are even more secure. Every transaction involving our important data — how it's accessed, changed, or moved — must have robust verification at every single stage. ››

Beyond Technology: The Enduring Power of Data, Process, and People

With the cost of data breaches now averaging nearly $5 million for organizations, being strong on cybersecurity is, without a doubt, a real business advantage. In my experience, success in this demanding environment hinges on having access to the right information at the precise moment it's needed to detect an attacker's activity. Then, almost instantaneously, we must determine: is this a legitimate action, or is it something malicious?

Organizations that do this well don't just have great people and effective technology; they also ensure that the visibility their technology provides is centralized. This allows their systems to automate much of the initial work of detection, freeing up their skilled employees to focus on investigating the most complex and nuanced situations. Conversely, a jumble of different security tools that don't talk to each other effectively creates inherent hurdles for our defenders—hurdles that attackers are all too quick to exploit.

One of the most pressing challenges I see organizations grappling with today is "Shadow AI." A frequent question I hear from CIOs and CISOs is, "How can I ensure we're using AI in our organization safely? How do I even get a handle on what AI applications are being used across different departments and what company data might be fed into them?" The larger and more distributed the organization, the more complex this becomes. This makes a clear, centralized AI strategy — complete with approved applications and strong measures to prevent data leakage — more critical than ever. We need the ability to specify which AI applications are approved for use and ensure that employees aren't inadvertently introducing new, unsanctioned applications into our environment.

However, even with these strategies, significant challenges remain. Stopping sensitive company data from accidentally being fed into public AI tools is something we're constantly working on. Ensuring our internal defenses can match the sophistication of AI-powered attacks is another ongoing effort. And, critically, we must address the challenge of how much we can trust the outputs of AI systems, which still often require human oversight and validation to guard against problems like "hallucinations" or simple inaccuracies.

Culture: The Ultimate Human Firewall

When I look at the kinds of cyber dangers we're dealing with now, they're faster, more intricate, and happening on a bigger scale than ever before. We're seeing nation-states borrow techniques from cybercriminal groups, and attackers exploit vulnerabilities across global supply chains within minutes of them becoming known. This situation highlights a simple truth I've come to learn through years on the front lines: technology by itself, no matter how advanced, isn't a magic bullet.

My ultimate advice, therefore, goes beyond just technology. It's not only about acquiring the latest tools or even about having brilliant people concentrated solely on the security team. It's fundamentally about cultivating a pervasive, deeply ingrained security culture within every organization.

What does this culture look like in practice?

- **Shared Responsibility:** From the legal department to operations, finance to HR, every single part of the business must recognize and internalize that security is their responsibility too.
- **Empowerment:** Our employees must be well-positioned and genuinely empowered to make secure decisions in their daily work. They need to feel it's not just safe, but encouraged, to raise their hand when they see something that doesn't look right.
- **Communication and Preparedness:** Our leaders across the business must clearly understand their roles and responsibilities. Crucially, they must know how to communicate effectively with one another and with security teams if a breach occurs. The more we practice and test our responses to various scenarios, the better prepared and more secure our organizations will inevitably be.

In this era where agentic AI is relentlessly speeding up the pace of cyber challenges, I believe a deeply ingrained security culture — one built on a bedrock of trust, shared responsibility, and continuous vigilance — is our most resilient and adaptable line of defense. It's about fostering an environment where every individual understands their vital role in protecting the organization. By doing so, we transform our entire workforce into an active, engaged, and ultimately formidable part of our collective security solution. ◆

The Weakest Link in Your Cybersecurity Isn't What You Think

BY MICHAEL SIKORSKI

CISOs spend countless hours thinking about defenses: fortifying networks, hardening endpoints, securing applications and safeguarding the cloud from an unrelenting wave of attacks. As an industry, we've collectively invested billions to keep pace with everything from traditional malware to the sophisticated onslaughts supercharged by generative AI (GenAI).

But there's a blind spot in even the best-prepared strategies — one that's increasingly dangerous because it's not fully within our control — our third-party ecosystems and their sprawling webs of suppliers, distributors, resellers, service providers and even customers. Collectively, they form the circulatory system of global commerce. And every node represents a potential point of entry for threat actors who understand a fundamental, uncomfortable truth: No matter how advanced our internal defenses, we are only as strong as the weakest link in our supply chain.

If you're thinking, "We've got this covered," I urge you to think again. Yes, many organizations include third-party risk in their audit checklists. Yes, they use compliance reporting as a measure of vendor hygiene. But let's be clear: That's not security, but rather periodic paperwork.

Nearly a decade ago, we were warned about this kind of complacency in another context — the false sense of safety in virtualized environments. Back then, unless every component of the architecture was equally advanced, the entire system was inherently vulnerable. The same principle applies here: Static, outdated approaches to third-party risk are both inadequate and dangerous. Unless we treat supply chain security as an urgent discipline, the whole enterprise is at risk. ››

What's Needed: Real-Time Vigilance and Data

Cybersecurity in the supply chain cannot be treated as a periodic exercise. Monitoring, managing and maintaining security must be an ongoing, always-on discipline. Static audits and annual compliance reviews might satisfy regulators, but they do little to stop a zero-day exploit or a fast-moving supply chain breach. If every element of the system isn't next-generation, the result is insecurity — and ultimately, inoperability. The critical nature of this continuous approach is underscored by findings in the 2025 Unit 42 Incident Response Report, which revealed that, in 75% of incidents, critical evidence of the initial intrusion was present in the logs. Yet, due to complex, disjointed systems, that information wasn't readily accessible or effectively operationalized, allowing attackers to exploit the gaps undetected. This highlights a crucial disconnect: The clues are often there, but traditional, periodic approaches fail to bring them to light in time.

We've seen this play out before and will see it for many years to come. Yet despite the hard lessons of these attacks, many organizations continue to treat third-party risk as a procurement checklist item — an annual vendor questionnaire rather than a living, breathing threat surface.

This is a dangerous misconception. Supply chain risks don't wait for audit season. They evolve in real time — and so must our defenses.

What We Can (and Should) Do About Supply Chain Risks

CISOs must champion a shift from periodic vendor checks to continuous, live risk monitoring across every third-party relationship. Anything less risks operational disruption and creates some uncomfortable conversations in the boardroom, as well as scrutiny from regulators asking hard questions about why known vulnerabilities went unaddressed.

The truth is, we've been relying too long on static audits and compliance checklists. These might have satisfied yesterday's risks, but today's threat landscape moves at machine speed. As such, we need hyperaccurate, real-time insights into supply chain vulnerabilities, updated as conditions change.

That's a significant ask. It demands real investment in tools, talent and time. Fortunately, CISOs have a powerful equalizer at their disposal: AI and automation. GenAI, predictive models and advanced machine learning are tailor-made for this challenge. AI can scan an expansive universe of data — past incidents, public disclosures, certifications, behavioral signals — to build dynamic security profiles for every vendor in your ecosystem. It can track changes in posture, flag emerging risks and generate meaningful, quantifiable risk scores.

Automation amplifies this further. Given the persistent shortage of skilled cybersecurity professionals, automation acts as a force multiplier — continuously evaluating third-party risks and accelerating response times when anomalies emerge. Sophisticated, contextually aware analytics ensure that attacks are identified and neutralized before they can move laterally across your environment.

This is about both efficiency and necessity. Attacks unfold in minutes, not months. Automated alert triaging can mean the difference between containment and catastrophe. When every second counts, you don't want human operators parsing spreadsheets. You want AI-enhanced systems that detect, decide and deploy defenses in real time.

The Bottom Line: Act Before the Breach

CISOs can no longer afford to extend blind trust to their vendors. The future demands something sharper: unparalleled visibility, real-time evaluation, and staunch accountability across every vendor, every partner and every link in the supply chain.

Third-party risk must be part of a holistic, board-level cybersecurity strategy. It can't sit in a silo owned by procurement or delegated to compliance teams. Boards must understand how supply chain security contributes to overall enterprise resilience — and ensure that it's tightly integrated with broader risk planning, business continuity efforts and regulatory readiness. Build resilience now because, in this new threat environment, hesitation is the surest path to disruption. ◆

Curious about your supply chain risk? Check out our Supply Chain Risk Assessment.

How AI Will Forge the Next Generation of Cybersecurity Talent

BY AARON ISAKSEN

Esteemed voices in the AI community are speaking of a near-term horizon where AI significantly reshapes the job market, potentially leading to major reductions in white-collar employment.[1] It's a conversation that understandably stokes anxiety and inevitably begs the question: Are cybersecurity jobs also at risk of being automated into obsolescence?

As someone who has dedicated his career to AI, navigating AI's evolution from theoretical constructs to the powerful agents we see today, my answer is emphatically "no." The narrative for cybersecurity is different — and far more empowering. AI will not lead to obsolescence; it will lead to evolution. In fact, it already has. While it will automate routine tasks, it will simultaneously elevate the demand for professionals who understand uniquely human skills.

Therefore, the future of our field will be defined by enduring archetypes:

- The **strategic risk translator**, who negotiates the critical trade-offs between risk and operational needs.
- The **AI-augmented hunter**, whose intuition and creativity outmaneuver intelligent threats.
- The **AI governance specialist**, who provides essential human judgment and accountability for these powerful new systems. ››

AI as the Great Enabler, Not the Great Replacement

Before we dig into AI's impact, let's first acknowledge a foundational truth: Cybersecurity is an incredibly demanding discipline. For decades, even with significant advancements in automation, we've grappled with a persistent shortage of experienced professionals capable of navigating its complex landscape. The sheer difficulty of doing cybersecurity right, and staying ahead of determined adversaries, means that human expertise remains an invaluable and often scarce resource. This isn't a gap that conventional automation alone has been able to close.

AI arrives into this challenging environment as the most potent amplifier for human ingenuity, not as its replacement. Its role is to augment our capabilities, allowing us to achieve more, faster, and with greater precision.

Imagine providing your security analysts with an intelligent partner that tirelessly sifts through petabytes of data, identifies the signals of a sophisticated attack, and prioritizes critical alerts with superhuman speed. Here, AI acts as a force multiplier, enabling teams to achieve what was previously impossible — coverage at both scale and depth. Instead of choosing between lower-resolution monitoring across the landscape or deep analysis on a few critical assets, your teams can now apply forensic-level scrutiny everywhere. This enables your teams to manage the massive queue of potential issues and protect the entire organization with unprecedented speed and precision.

AI is poised to revolutionize how we cultivate talent. For new team members, AI can act as an intelligent apprentice or a personalized mentor, guiding them through complex tasks and organizational specifics that once took months — if not years — to fully comprehend. This dramatically shortens the learning curve, enabling new hires to become effective contributors much faster and reducing the overall cost and time associated with talent development. For seasoned experts, AI offers a new dimension of insight. It can synthesize global threat intelligence in real time, identify emerging attack vectors previously obscured, and help veterans apply their deep expertise across an ever-expanding digital attack surface.

> The future I envision is one where human ingenuity, significantly augmented by intelligent AI, leads to a more resilient and secure digital society.

The Expanding Digital Universe: More to Secure, Not Less

While AI undoubtedly boosts efficiency, the digital frontier we must defend continues its inexorable growth. Every new technology adds layers of complexity, but none more so than AI itself. This creates a fascinating paradox: The tool we use to manage scale is simultaneously creating a vast new continent of risk that requires securing.

The discipline of securing AI systems is becoming deeply interwoven with traditional cybersecurity. It's difficult to fathom a future where we entrust the security of these powerful, autonomous agents exclusively to another AI. As these systems proliferate, the need for human oversight, strategic judgment and ethical governance will persist and intensify. AI, therefore, doesn't shrink the workload. It helps manage a rapidly expanding one while creating an entirely new domain that demands enduring human talent. ››

Judgment, Accountability and the Art of Security

For all its analytical power, AI is not a CISO. It cannot replicate the uniquely human attributes essential for true security leadership.

Yet, who is responsible when an AI makes a mistake or when the path forward is clouded by ambiguity? The accountability for that decision rests firmly with people. And this isn't a responsibility reserved for the C-suite alone; it is distributed throughout the entire security organization. It's present in the countless judgment calls made every day — from the frontline analyst validating an AI-flagged alert to the security engineer assessing the business risk of an automated response. In an AI-powered world, every security professional becomes a crucial steward of accountability and context, a role that technology enhances but cannot replace.

Moreover, security is rarely a binary decision. It often involves negotiating trade-offs with the business, balancing risk mitigation with operational needs and strategic goals. AI can inform these decisions with data, but it cannot navigate the nuanced discussions or make the value judgments inherent in these tradeoffs.

Human Interface: Where People Meet Peril — and Respond

We must also remember that cybersecurity is not solely a battle of machines. Many of an organization's most significant vulnerabilities lie at the interface between humans and technology. Social engineering, business email compromise, insider threats and ensuring data loss prevention and compliance are fundamentally human problems that require human understanding, training and effective communication as part of the solution.

And when an incident does occur, managing the response in a calm, effective manner requires working with people who are often under immense stress. This demands empathy, integrity, trust, relationships and leadership — qualities that are, for the foreseeable future, uniquely human.

Enduring Adversarial Dance

Finally, we operate in an inherently adversarial world, which is not a problem technology, even AI, will "solve" in perpetuity. As defenders get better and leverage AI, attackers will inevitably respond in kind, developing their own AI-driven tools and techniques to attack at scale or craft increasingly sophisticated exploits.

This new AI-driven landscape is inherently adversarial. As defenders, our strategic advantage comes from leveraging comprehensive data within a well-architected security platform to train superior defensive AI models. But this advantage doesn't mean the work is done but that it changes. The ongoing arms race shifts the burden from manual tasks to uniquely human ones: the ingenuity to anticipate novel attacks, the strategic insight to guide our AI counterparts and the adaptability to outmaneuver our adversaries.

Augmented Future

So, will AI diminish the role of cybersecurity professionals? My conviction is that it will do the opposite. AI will free us from the repetitive, data-intensive tasks that can lead to burnout, enabling us to focus on the strategic, creative and human-centric aspects of our work. It will amplify our expertise, accelerate our learning and empower us to tackle more complex challenges across a broader and more intricate digital world.

The future I envision is one where human ingenuity, significantly augmented by intelligent AI, leads to a more resilient and secure digital society. The demand for smart, adaptable and strategically minded cybersecurity talent will both persist and grow, as we collectively forge this new, more secure future. ◆

1. "Behind the Curtain: A white-collar bloodbath," Axios, May 28, 2025.

Inside the Mind of a Cybersecurity Crisis Leader: Lessons from the Frontlines

BY CHRIS SCOTT

Whether you're a CISO, CEO, or board member, you will eventually get the call — the one that jolts you awake in the middle of the night. Your heart will leap. Your mind will race. But amid the panic, it's important to remember: This is the moment to gather information, evaluate options, make decisions, and take action. In other words, it's time for leadership. I've been there. I've stood alongside teams navigating crises, and I've led teams through the chaos of recovery.

The challenge is never just technical — it's deeply human. Leaders must ask themselves: How do I demonstrate leadership in this moment? How do I make decisions with limited information and marshal every available resource to stabilize, confront, and resolve the problem?

The SONAR Method: A Model for Cybersecurity Crisis Leadership

In Cyber Crisis Response, my co-author and I distilled these hard-won lessons into a simple but powerful framework — the SONAR Method™. It's a method and model I return to every time there's a crisis, and is a framework built from real-world experience:

- **Stabilize:** Immediately contain the situation and regain control of critical systems.
- **Organize:** Assemble the right people with the right expertise — quickly.
- **Negotiate:** Balance the conflicting priorities of executives, legal, regulators, customers, and technical teams.
- **Articulate:** Communicate clearly and regularly with all audiences, internally and externally.
- **Remediate:** Execute a disciplined recovery plan that closes security gaps and restores business operations.

Having been in the trenches for more than two decades, I've learned that effective leadership during a cybersecurity crisis isn't about hoping for the best — it's about preparing for the worst. And like any discipline, it can be taught, practiced, and refined.

Stabilize First — Leadership Under Fire

Every crisis begins the same way: with uncertainty and disorder. That's why the first step is always to **Stabilize**. You won't have perfect information, but you must regain control. Whether isolating ››

compromised systems, containing adversaries, or protecting critical infrastructure, leaders must shift — temporarily — from collaboration to decisiveness. In a true crisis, there is no time for endless debate. Someone has to make the call, and as a leader, that someone is you.

Next, it's important to know that it's OK, and perhaps even necessary, to act like a dictator when a crisis hits. In that scenario, there isn't a lot of time for consensus, opinion, or discussion — you're under attack! Those thoughtful, collaborative traits usually are highly desirable for a leader, but remember that the first step is to right the ship — stabilize things immediately before the crisis spins out of control.

It's also vitally important that leaders remember that the number one priority will always be human life and safety. Fortunately, cybersecurity attacks don't often evolve to that level of peril, but when they do, that has to be your guiding principle. So if you have someone stuck in an elevator, your concern isn't the ransomware that's controlling the system that guides the elevator's behavior. It's about getting the person out of the elevator. Call 911, call the fire department, and get that process started. Then you can figure out, "OK, what are the next steps we need to do, from a technical standpoint?"

Organize and Negotiate — Teams Win Crises

So let's say you've stabilized the immediate situation. What next? It's imperative leadership quickly shifts to the next two actions: **Organize** and **Negotiate**. This is where teams matter most. It's easy to overlook in the heat of the moment, but no leader, no matter how seasoned, can recover alone. Success depends on assembling the right people, fast — technical experts, legal counsel, communications, and business leaders — and aligning them behind a common goal.

Negotiation here doesn't just mean external actors. It means balancing the competing priorities inside your own organization. The CEO wants business continuity. The legal team is focused on liability. Regulators expect timely disclosures. Every crisis involves conflicting agendas. Your job is to reconcile them without losing momentum.

Articulate Clearly — Communication Is Nonnegotiable

The fourth step, **Articulate**, is one of the most underestimated. Communication is not a soft skill in a crisis — it is an operational necessity. Too many leaders freeze, default to "no comment," or speak too soon without facts. In my experience, it is always better to admit what you don't yet know than to risk damaging trust.

Customers, employees, regulators, partners — they don't expect instant solutions, but they do expect accountability. Acknowledge the problem, commit to fixing it, and provide regular updates. Silence invites speculation. Worse, it can permanently damage the credibility of the entire leadership team.

Remediate — and Make Sure You Can

Crisis recovery is not just about technical remediation; it's about regaining confidence across the business. I've seen organizations unable to execute basic response plans because they stored the only copy of their documentation on the very systems now encrypted by ransomware. In my early days, we solved this the old-fashioned way: laminated wallet cards with key contacts and protocols. Today, it's about building resilient playbooks and practicing them under real-world conditions.

The fifth and final pillar, **Remediate**, is where leaders shift from response to recovery. This means fixing what's broken but also ensuring the business is better prepared for the next inevitable incident. It's the difference between surviving and emerging stronger.

Leadership After the Headlines

Surviving a cyberattack is one thing; leading through the aftermath is another. The best leaders take ownership, foster accountability, and conduct open, blame-free reviews to improve. They rebuild trust with boards, customers, and teams by demonstrating that the crisis wasn't just endured — it was learned from.

What sets great crisis leaders apart isn't just technical prowess. It's the ability to make decisions under pressure, communicate with integrity, and enable teams to adapt and act without hesitation.

Because when the next call comes — and it will — true leaders lead. ◆

When it comes to AI-driven cyberattacks, we've got every angle covered.

"The NHL brand must be protected at all costs. Having a partner like Palo Alto Networks allows us to enable the business to use the latest technologies, but use them safely and securely."

Dave Munroe
Vice President and Chief Information Security Officer
National Hockey League

paloalto® NETWORKS

READ THE STORY