PULSE REPORT

THE RSAC 2026 CONFERENCE **HAS A NEW CENTER OF GRAVITY.**

Cybersecurity Intelligence Report

RSAC 2026 Conference

ISMG

METHODOLOGY

The **RSAC 2026 Cybersecurity Pulse Report** is grounded in more than 130 in-depth interviews conducted by ISMG's editorial team across four intensive days at the RSA Conference 2026. Interviewees span the full breadth of the cybersecurity ecosystem — security practitioners and enterprise CISOs, vendors and platform architects, investors and startup founders, government officials, academic researchers, and independent analysts. The resulting corpus represents one of the largest structured primary-source intelligence collections assembled at a single cybersecurity event.

1. SOURCE COLLECTION AND STRUCTURING

Every interview was captured, transcribed and then processed through ISMG's Research and Intelligence Services (IRIS) platform - a purpose-built, structured content intelligence system developed by ISMG's Content Intelligence & AI Innovation team. Rather than treating interview transcripts as raw material to be paraphrased, IRIS extracts intelligence records from each interview: specific claims, supporting evidence, predictions, strategic implications, implementation barriers, verbatim quotes and explicit uncertainties. Each record is independently attributed to the interviewee and preserved with full source fidelity, including tensions and contradictions that surface within individual interviews.

2. THEME MAPPING AND PIR-DRIVEN ANALYSIS

Before synthesis, all records were mapped to RSAC 2026's official 16-theme conference taxonomy - a structure derived from analysis of the full RSAC session program and anchored by the dominant signal of this year's event: the convergence of agentic artificial intelligence and cybersecurity. The themes ranged from securing AI systems and agentic AI, AI-powered cyber defense, and security operations center (SOC) transformation to critical infrastructure and OT/ICS security, cybersecurity investment, and market dynamics.

3. EVIDENCE SYNTHESIS AND VALIDATION

With structured records mapped to themes, IRIS applies a multifactor evidence synthesis process that weights records by relevance, source credibility, specificity and corroboration across independent sources and contradiction pressure - the degree to which claims are actively disputed by other interviewees. The system explicitly distinguishes between practitioner perspectives, vendor claims and policymaker positions, surfacing role-based divergence where it exists. Claims that appear frequently but originate disproportionately from vendor sources are flagged and treated differently from claims corroborated across practitioners, government officials and independent analysts.

Where consensus is genuine, the report says so. Where the evidence is divided, uneven or premature, that uncertainty is preserved and named. Overstated narratives - positions that dominate conference discourse but are poorly supported by evidence from practitioners with direct operational accountability -are identified and treated with explicit caution.

04. NARRATIVE CONSTRUCTION AND HUMAN OVERSIGHT

Validated synthesis outputs from IRIS serve as the direct input to narrative chapter construction. This staged approach ensures that the analytical reasoning embedded upstream - including disagreements, maturity distinctions and signal-versus-noise judgments - is carried through into the published narrative rather than collapsed into summary generalizations.

05. SCOPE AND LIMITATIONS

This report reflects expert opinion and operational experience as expressed during RSAC 2026. It is a primary-source intelligence product, not a statistically representative survey. Source mix skews toward vendors and enterprise practitioners, consistent with the composition of the RSAC audience. Government and academic perspectives are represented but are proportionally smaller. Year-over-year comparison with the RSAC 2025 Pulse Report is made where the evidence supports directional trend assessment; such comparisons are noted explicitly and treated as signals rather than conclusions.

TABLE OF CONTENTS

RSAC 2026: THE CONFERENCE HAS A NEW CENTER OF GRAVITY

For the past two years, RSAC Conference has been reorganizing itself around a question it couldn't quite articulate. In 2024, the agenda still looked like a balanced portfolio - 13 discrete themes competing for attention, with AI leading but not dominating. By 2025, that portfolio had compressed into larger strategic narratives, and the AI session count had doubled. In 2026, the reorganization is complete. The conference no longer treats AI as one of many topics. It has made AI the architecture through which every other topic is
now framed.

That shift - from AI as a subject to AI as a structural assumption - is the defining story of RSAC 2026, and it carries direct operational and strategic implications for every CISO who walks through the doors.

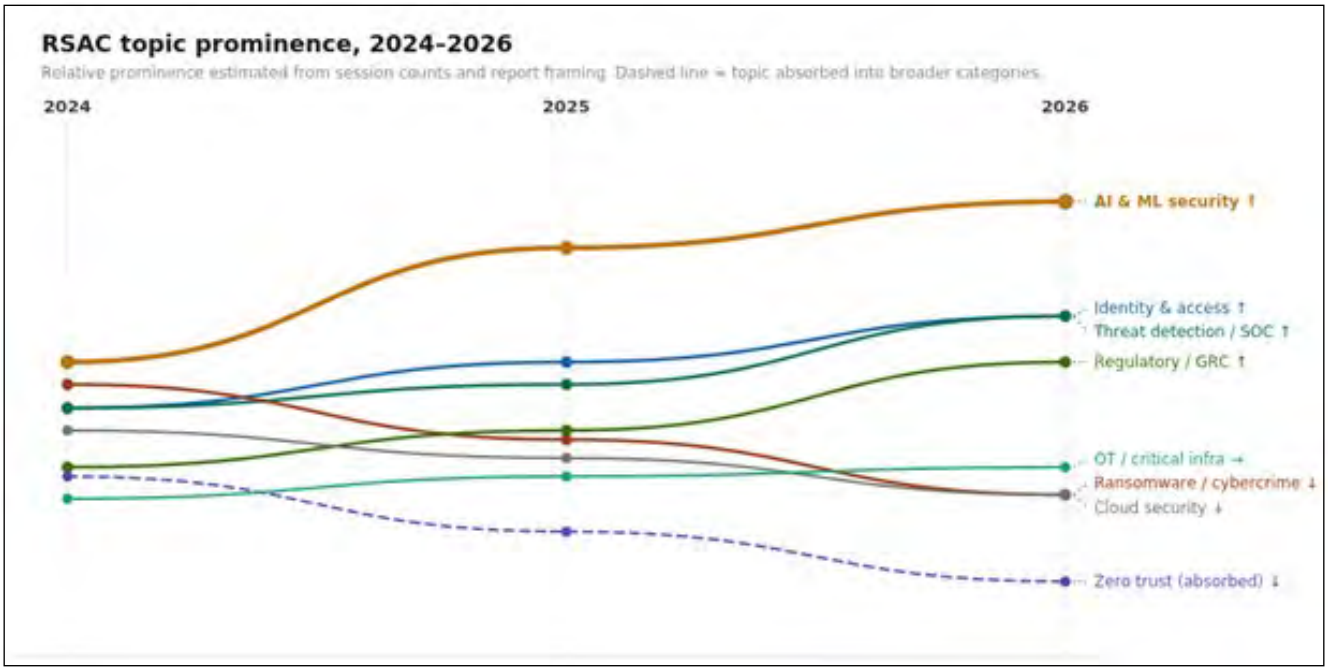

The slope chart above makes the structural argument visible. AI and machine learning security's trajectory is not a steep climb among other climbers - it separates from the field. Everything else either holds, declines modestly or gets absorbed. The lines that fall - ransomware, zero trust and cloud security as a stand-alone concern - do not represent topics losing relevance. They represent topics losing their independent identity. They are still present at RSAC 2026. They are simply no longer the point.

The clearest signal in the session data is not the AI line itself. It is what happens to everything adjacent to it. Ransomware, which ranked second in prominence in 2024, no longer commands its own track in 2026. It appears inside broader clusters on threat landscape and cybercrime - reframed from epidemic to tactic. Zero trust, which carried eight dedicated sessions in 2024, dropped to five in 2025 and effectively dissolved into architectural baseline assumptions by 2026. Cloud security followed a similar trajectory: present, but no longer defining its own conversation. It is now the substrate on which AI agents operate, not a headline in its own right.

This is what capture looks like in a conference agenda. Not displacement - but absorption. The topics don't disappear; they lose their gravitational independence and begin orbiting a new center.

That center, in 2026, is agentic AI security. The "securing AI systems and agentic AI" cluster is explicitly designated the defining theme of the conference. Paired with the "AI-powered SOC transformation" cluster, these two categories account for roughly a quarter of all substantive sessions - a share no single topic has held in this conference's recent history. The practical subjects inside those sessions have also matured sharply: from dual-use risk and model governance in 2024 and deepfakes and autonomous entities in 2025 to model context protocol security, AI supply chain poisoning, enterprise agent risk and non-human identity propagation in 2026. This is no longer a conversation about what AI might do to cybersecurity. It is a conversation about engineering security for systems that are already deployed and already making autonomous decisions.

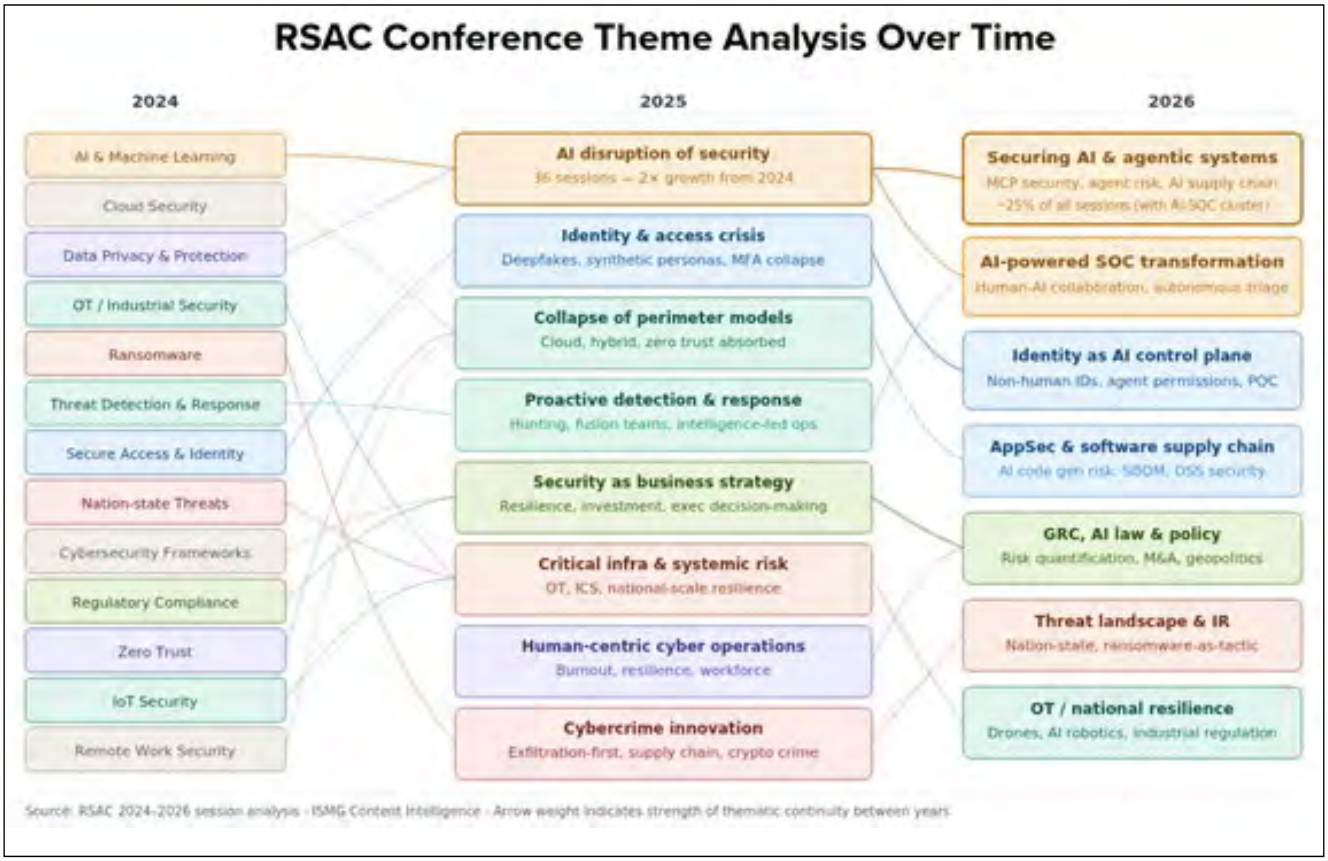

The flow diagram above shows where the 13 themes of 2024 actually went. Most did not disappear - they merged upward into larger strategic narratives in 2025 and then re-anchored around the AI and governance clusters that define 2026. The arrow weights reflect thematic continuity. The compression is not cosmetic. It reflects a genuine reorganization of what the security industry considers a stand-alone problem versus a component of a larger systemic challenge.

FOR CISOS AND SECURITY LEADERS, THE GOVERNANCE LAYER OF RSAC 2026 DESERVES PARTICULAR ATTENTION, AND IT IS PART OF THE TREND LINE MOST LIKELY TO BE UNDERWEIGHTED BY PRACTITIONERS FOCUSED ON TECHNICAL CONTENT.

In 2026, RSAC added explicit, prominent emphasis on AI law and policy, cyber risk quantification, venture and M&A dynamics, CISO leadership and workforce, and geopolitics. This is not conference programming for its own sake. It reflects a market that is moving from asking "how do we defend against AI-enabled attacks" to asking "how do we govern, price and legislate AI-native risk at scale." Those are board-level questions, and the fact that they now occupy a visible slice of the RSAC agenda signals that the CISO role is being pulled further into strategic and legal terrain - whether individual practitioners are ready for that or not.

The identity story runs parallel. In 2024, identity meant IAM and zero trust - access control for human users. In 2025, it expanded to include synthetic personas, deepfake-enabled impersonation and the beginning of non-human entity management. In 2026, identity is framed as the control plane for AI systems: agent permissions, trusted identity propagation for autonomous workflows and post-quantum cryptography migration for identity infrastructure. The category has not grown - it has transformed. Organizations that have not begun reconceptualizing their identity architecture around non-human actors are already behind the conversation RSAC is having.

The operational implication is straightforward, even if the execution is not. Security programs built around the 2024 model —- discrete categories, human-centered access control and AI as a detection accelerant - are structurally misaligned with the threat and governance environment that 2026 describes. The conference agenda is, among other things, a leading indicator of where the industry's collective attention is heading and where vendor investment is concentrating. By that measure, the signal from three years of RSAC data is unambiguous.

AI** is no longer a track at this conference. It is **the conference.

1. AI Fundamentally Alters the Cybersecurity Threat Landscape

- **CORE JUDGMENT**

  AI has become the foundational transformation across all major threat categories in 2026, requiring architectural rather than incremental defensive changes.

  The intelligence gathered from security leaders reveals a definitive shift: AI is no longer an emerging factor in cybersecurity; it has become the defining backdrop against which all other trends must be understood. Ed Skoudis, president.

- **WHAT THE INTERVIEWS REVEAL**

  The experts describe AI's impact across multiple threat vectors simultaneously. Vulnerability discovery is accelerating toward what Skoudis termed a "vulnerability apocalypse" - AI-driven code analysis enabling the identification of more than 100 critical CVEs per week within one to two years, far outpacing existing patch management capabilities. Sandra Joyce, vice president at Google Threat Intelligence, emphasized that threat actors are on a "parallel journey" to defenders, with AI tools providing a particularly pronounced advantage to lower-skilled attackers.

  Ian Thiel, co-founder and COO of Sublime Security, highlighted how LLMs have collapsed the traditional distinction between

> **"If we were to tell you that there is a major trend in cybersecurity that doesn't involve AI, we would be lying to you."**
>
> - Ed Skoudis

commodity and spear-phishing. “In the new world, it’s easy to do large-scale targeted campaigns where you’re mentioning, ‘Oh, we both went to this school, and I see you’re in this town, and the English is flawless, and the attackers are using LLMs for this,” Thiel said. This hyper-personalization at scale represents a qualitative shift in attacker capabilities.

The operational tempo of attacks is also transforming. Skoudis described attack cycles completing in as little as eight minutes, requiring what he termed “agentic defenses” rather than human-speed response capabilities.

WHERE EXPERTS ALIGN - AND DIVERGE

There was unanimous consensus that AI represents a permanent, foundational change rather than a passing trend. As Skoudis put it: “If we were to tell you that there is a major trend in cybersecurity that doesn’t involve AI, we would be lying to you.”

But experts diverged on prioritization. Some emphasized speed as the primary concern - the compression of attack timelines - while others focused on scale, particularly the democratization of sophisticated techniques to previously low-skilled actors. Matt Olney, director of threat intelligence and interdiction at Cisco Talos, suggested that scale may prove more operationally damaging than speed over time.

SUPPORTING PERSPECTIVES

Google's Joyce framed the challenge in terms of capability multiplication. “The sophistication that these low-level adversaries could get you is now higher because the bar is lower. We certainly see potential for maybe a low-skilled actor to actually get 10x faster using AI tools.” This “10x low-skilled actor” phenomenon represents a fundamental shift in threat modeling assumptions. Meanwhile, Olney provided a measured counterpoint, noting that vulnerability-to-exploit compression predates AI and that AI may already be operating quietly in adversary auxiliary functions -document triage and translation - rather than in headline attack vectors.

STRATEGIC IMPLICATIONS

The evidence suggests that organizations treating AI as an add-on to existing security architectures are fundamentally misaligned with the threat reality. Defensive strategies must account for hyper-personalized attacks at commodity scale, sub-10-minute attack cycles, and sophisticated techniques in the hands of previously unsophisticated actors. The traditional security model of human-operated detection and response appears increasingly obsolete.

CONFIDENCE ASSESSMENT

High. Multiple independent sources with diverse operational perspectives converged on AI’s foundational role, supported by specific metrics and concrete examples.

2. Recovery Planning Remains Critically Underinvested

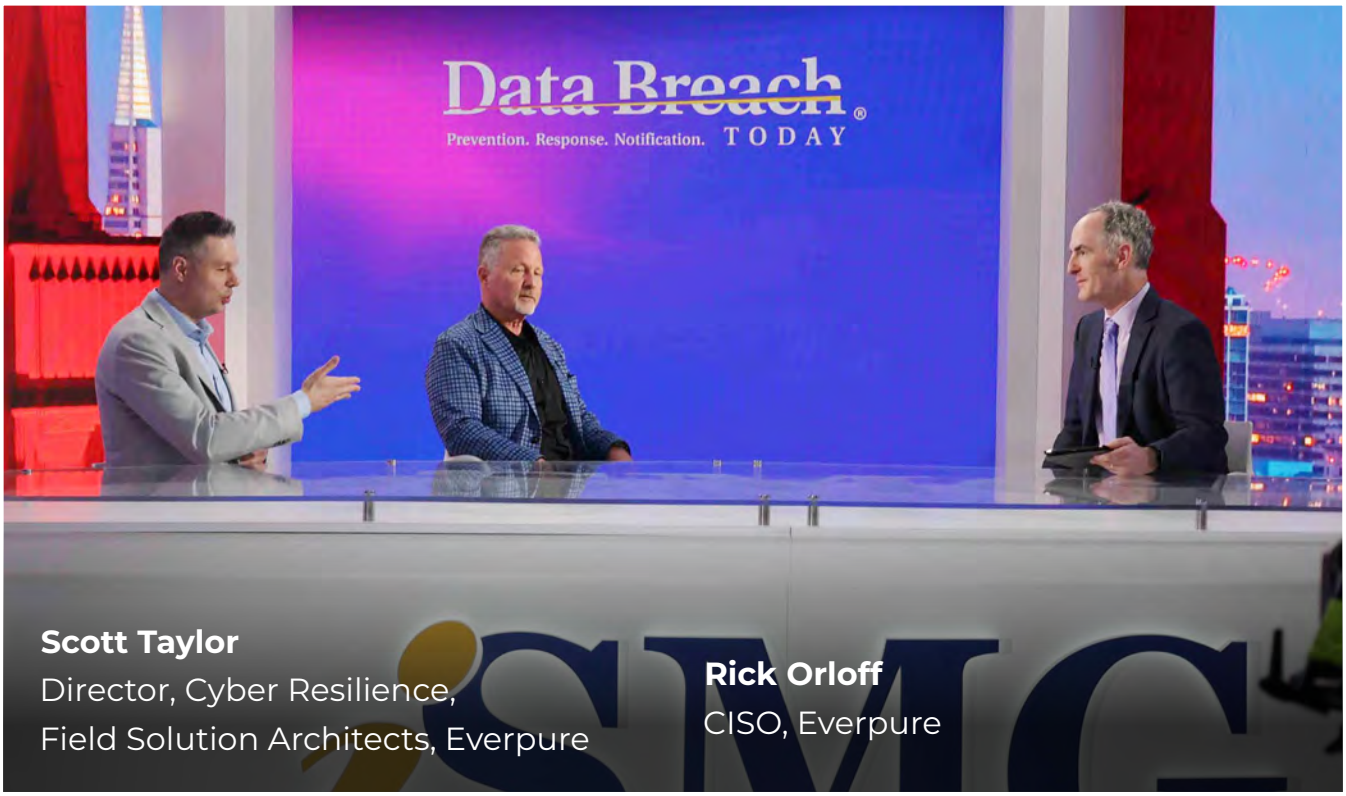

- **CORE JUDGMENT**

  Organizations demonstrate systematic underinvestment in recovery planning, with cyber events requiring fundamentally different methodologies than traditional disaster recovery.

  Despite years of high-profile ransomware incidents, security leaders revealed significant gaps in organizational recovery preparedness. The intelligence indicates that most enterprises have over-indexed on backup capabilities while underinvesting in the planning and testing required for actual recovery from malicious events.

- **WHAT THE INTERVIEWS REVEAL**

  Scott Taylor, director of cyber resilience - field solution architects at Everpure, articulated the core distinction. “Cyber events are completely different in reality from traditional disaster recovery events,” Taylor said. “Those events aren’t malicious events, but cyber events are malicious events, and recovering from them is completely different,

“The time to figure out who should be doing what is not in the middle of your first incident. We practice that. We use it even for smaller issues … so that if you have a critical playbook, the first time you’re looking in that playbook should not be when everything is falling over.”

- Rick Orloff

and you need a different methodology."

The recovery challenge extends far beyond IT operations. Taylor noted that significant cyber incidents mobilize "200 plus employees working around the clock for two weeks to restore the business" - a cross-organizational response involving finance, customer service and operations teams, not just IT personnel. This scale of organizational impact is rarely reflected in recovery planning.

Rick Orloff, CISO at Everpure, emphasized that testing consistently reveals hidden dependencies - credentials, identities and services - that remain invisible until organizations attempt actual recovery procedures. The time to discover these dependencies, Orloff argued, should not be during the first major incident.

WHERE EXPERTS ALIGN - AND DIVERGE

There was strong consensus that cyber recovery demands distinct methodologies from traditional disaster recovery. Experts agreed that organizations should focus on "minimum viable business" recovery rather than full system restoration, which may take years.

But perspectives diverged on the adequacy of current vendor solutions versus the need for organizational process changes. Anthony Cusimano, chief evangelist and director of solutions marketing at Object First, noted that only approximately 64% of organizations have immutable backup storage. In contrast, the Everpure practitioners stressed that even organizations with robust backup capabilities often lack tested recovery procedures.

SUPPORTING PERSPECTIVES

The practitioners revealed concerning structural vulnerabilities in recovery planning. Taylor highlighted the "single-SME dependency" problem, where critical recovery knowledge resides with an individual expert who may be unreachable during weekend or holiday attacks - precisely when many cyber events occur.

Orloff stressed the importance of building "muscle memory" through incident command structure practice. "The time to figure out who should be doing what is not in the middle of your first incident," Orloff said. "We practice that. We use it even for smaller issues … so that if you have a critical playbook, the first time you're looking in that playbook should not be when everything is falling over."

STRATEGIC IMPLICATIONS

The evidence suggests a dangerous false confidence in organizational resilience. Many enterprises with validated traditional disaster recovery plans may be unprepared for the malicious, dynamic nature of cyber events. Recovery timelines extending to years - often hidden from public reporting - indicate that the true industry recovery capability may be substantially worse than commonly understood.

The focus must shift from prevention-only strategies to include recovery sequencing, cross-functional incident command and regular testing under adversarial scenarios.

CONFIDENCE ASSESSMENT

Medium. Strong directional evidence from operational practitioners, though limited by vendor context and anecdotal rather than statistically validated data.

3. Threat Intelligence Must Evolve to Active Operationalization

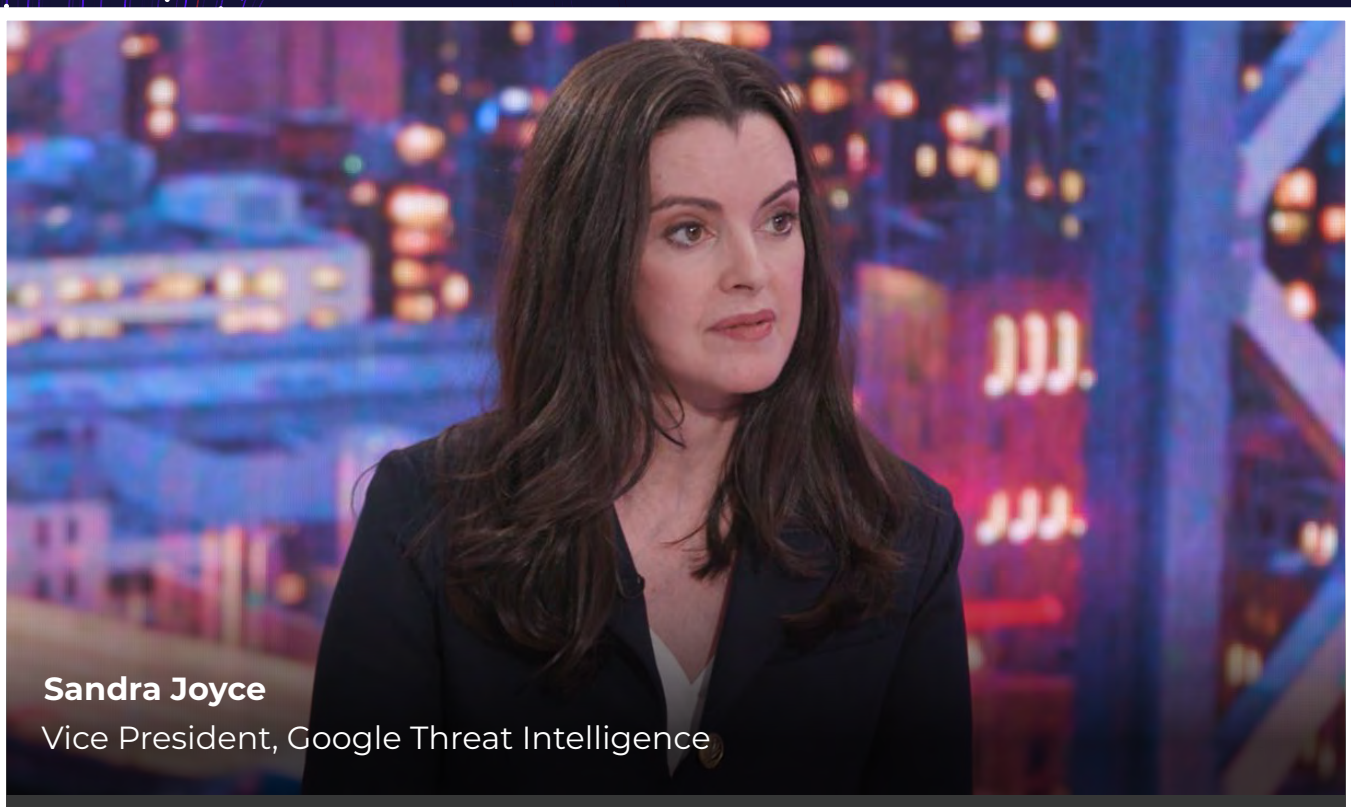

- **CORE JUDGMENT**

  The industry must shift from passive intelligence sharing to active operationalization, with coordinated infrastructure disruption representing the required evolution.

  The cybersecurity intelligence community faces a fundamental operational challenge: Traditional threat intelligence sharing has proven insufficient against the scale and sophistication of modern threat actors. Security leaders described an urgent need to move from passive information exchange to active disruption of adversary infrastructure.

- **WHAT THE INTERVIEWS REVEAL**

  Joyce provided the clearest articulation of this transformation. “We have to think differently. We’re not just going to share the intelligence that we have - we need to operationalize it. What I mean by that is going to the source. Using intelligence that we have to do active disruption.” She pointed to the IPIDEA residential proxy network takedown as a proof of concept for coordinated legal and technical infrastructure disruption.

  The operational bottlenecks in current threat intelligence practice are severe. Steven Gerry, vice president of sales at Tidal Cyber, cited a Fortune 500 benchmark showing that manual classification of threat reports against MITRE ATT&CK’s over 700 options requires approximately 4.55 hours per report. This manual analysis burden has prevented most organizations from operationalizing the procedural knowledge needed for effective defense.

  Gerry emphasized that the industry has solved tactics and techniques but never meaningfully addressed procedures. “The P

in TTP has never really been solved for in a meaningful way … the procedure, that's the granular - the command line itself," Gerry said.

WHERE EXPERTS ALIGN - AND DIVERGE

There was strong consensus that passive intelligence sharing alone is insufficient. Experts agreed that the procedural knowledge gap - understanding specifically how attackers execute techniques - represents a critical operational weakness.

But perspectives varied on the scope of required change. Joyce advocated for industrywide coordinated takedowns and active infrastructure disruption, while others focused on automation of intelligence analysis and improved handoffs to detection engineers and threat hunters.

"We have to think differently. We're not just going to share the intelligence that we have - we need to operationalize it. What I mean by that is going to the source. Using intelligence that we have to do active disruption."

- Sandra Joyce

SUPPORTING PERSPECTIVES

Tim Pappa, incident response engineer for cyber deception strategy at Walmart Global Tech, emphasized that effective threat intelligence must focus on high-fidelity alerting rather than alert volume. "A lot of what we do is focus on high-fidelity alerting, and that comes down to can you shape a detection function? So if it goes off and the SOC sees it, they have a lot of confidence in that," Pappa said.

The experts highlighted that many organizations lack the analytical talent to effectively operationalize threat intelligence. Gerry noted that threat analysts capable of sophisticated ATT&CK classification are "extraordinarily rare" and most organizations cannot afford them even when available.

STRATEGIC IMPLICATIONS

The evidence points toward a two-pronged evolution: tactical automation to make procedural knowledge accessible to more organizations and strategic coordination for active infrastructure disruption. Organizations that continue operating on IOC-chasing models rather than TTP/procedure orientation will be structurally disadvantaged.

The shift from defensive to active threat intelligence represents a fundamental change in industry approach, requiring new legal frameworks, partnership agreements and operational coordination capabilities.

CONFIDENCE ASSESSMENT

High. Multiple independent sources with operational threat intelligence experience converged on the need for fundamental change, supported by specific metrics and concrete examples.

4. AI-Powered Cyberattacks: Geopolitical and Policy Perspectives

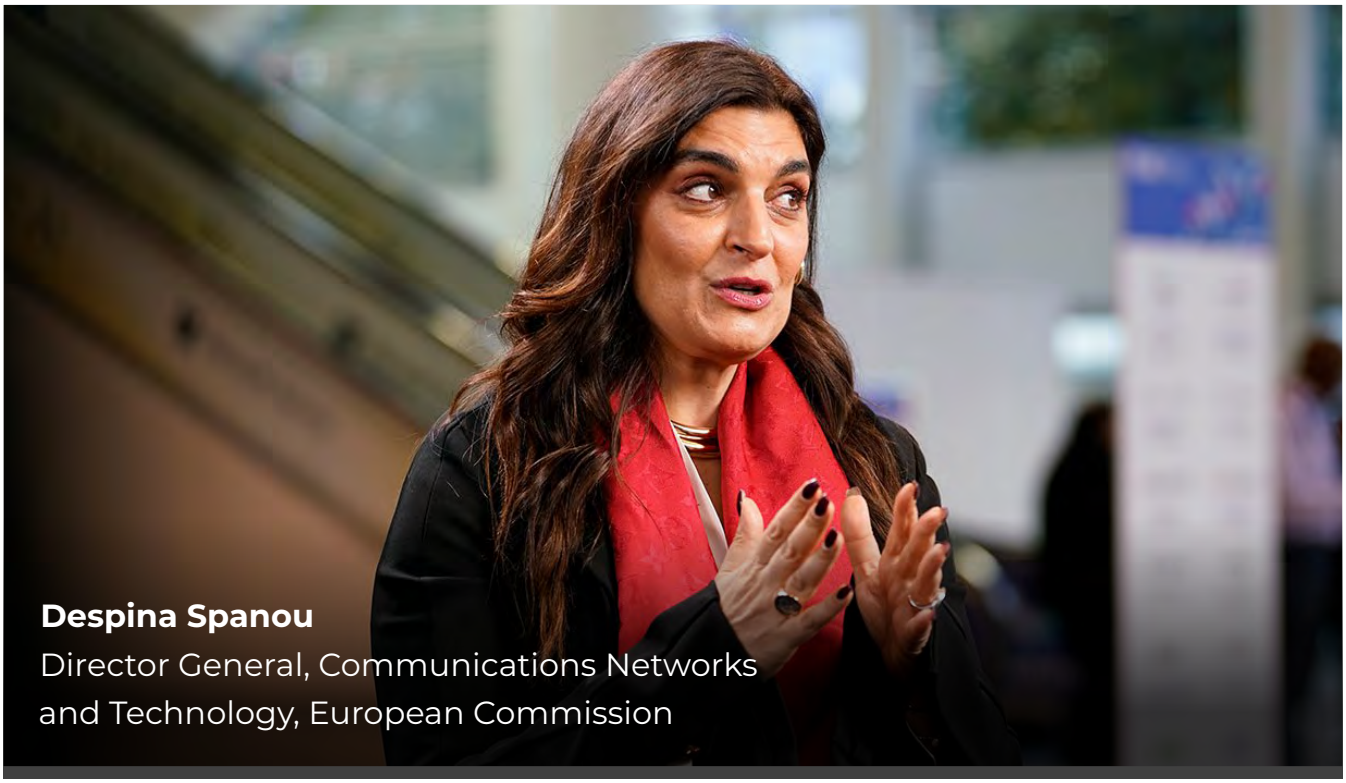

Despina Spanou
Director General, Communications Networks and Technology, European Commission

- **CORE JUDGMENT**

  AI has created a step-change in attacker capability during 2025-2026, with autonomous end-to-end attacks now operationally viable and ransomware predominantly AI-managed. This represents a structural shift in the threat landscape, not incremental evolution.

- **WHAT THE INTERVIEWS REVEAL**

  Multiple high-level sources confirm that 2026 marks a decisive break from previous AI-in-cybersecurity narratives. The evidence points to operational deployment rather than experimental capability. Despina Spanou, director general for communications networks and technology at European Commission, reports that "80% of ransomware attacks are basically managed by AI" and notes that "agentic AI can perform cyberattacks from beginning to end on its own."

> **"80% of ransomware attacks are basically managed by AI" and notes that "agentic AI can perform cyberattacks from beginning to end on its own."**
>
> - Despina Spanou

This operational maturity has created what Dave DeWalt, CEO of NightDragon, characterizes as a fundamental shift in the offense-defense balance. "The gap between

what the offense can do and the defense can do widened dramatically. We are entering a dark period," DeWalt said. His assessment, grounded in 22 years of RSAC perspective, frames 2026 as a structural dislocation rather than gradual trend.

The scope extends beyond individual tools to autonomous attack orchestration. Where previous AI applications in cybersecurity focused on specific tasks - reconnaissance, payload generation or social engineering - current capabilities demonstrate full operational cycles without human intervention. This represents a qualitative change in attacker efficiency and scale.

World Economic Forum survey data reinforces the directional trend: 87% of respondents see AI vulnerabilities rising, while 94% of security leaders view AI as the most significant cybersecurity force in 2026. Akshay Joshi, head of the Centre for Cybersecurity at World Economic Forum, notes that AI security governance adoption grew from 37% to 65% year-over-year, indicating defensive responses are accelerating but potentially lagging offensive deployment.

• WHERE EXPERTS ALIGN - AND DIVERGE

Strong consensus emerges on AI's role as a force multiplier in cybersecurity, but experts diverge on whether this represents transformation or acceleration. Paul Foster, head of the National Cyber Crime Unit at U.K.'s National Crime Agency, offers a more measured assessment: AI represents "acceleration and scale of existing threats, not transformation." Foster's law enforcement perspective emphasizes continuity in criminal methods enhanced by AI capability rather than entirely new attack paradigms.

The disagreement centers on magnitude rather than direction. All sources confirm AI's significant impact, but differ on whether current capabilities constitute evolutionary improvement or revolutionary change.

This tension reflects different analytical frameworks: venture capital and policy perspectives tend toward transformation narratives, while operational law enforcement emphasizes recognizable criminal patterns at enhanced scale.

• SUPPORTING PERSPECTIVES

European policy makers frame AI-powered attacks as requiring immediate structural response. "AI is a blessing and a curse for cybersecurity," Spanou said, noting the dual nature of AI as both a threat amplifier and defensive capability. The European Commission's unified approach across 27 member states reflects recognition that individual organizational responses are insufficient for autonomous attack capabilities.

From an investment perspective, DeWalt's "dark period" assessment carries particular weight given NightDragon's portfolio focus on cybersecurity innovation. His observation that "cyber and electronic warfare are merging into a hybrid warfare" suggests AI-powered attacks are becoming integrated into broader geopolitical conflict strategies.

Foster's operational experience provides crucial nuance: While AI accelerates existing criminal capabilities, the fundamental disruption strategies - targeting infrastructure, financial services and coordination mechanisms - remain effective. This suggests that current law enforcement and defensive approaches retain relevance when scaled to match AI-enhanced threats.

• STRATEGIC IMPLICATIONS

The operational viability of autonomous attacks requires immediate recalibration of enterprise security postures and government cyber capabilities. Organizations must move beyond tools-focused AI strategies toward comprehensive resilience against self-executing attack campaigns. The shift from human-managed to AI-managed

ransomware operations fundamentally changes incident response timelines and coordination requirements.

For policy makers, the evidence suggests that current governance frameworks are structurally mismatched to autonomous attack speeds. Spanou's call for using "AI to our advantage" reflects recognition that defensive AI deployment must achieve operational parity with offensive capabilities.

- **YEAR-OVER-YEAR SHIFT**

  The 2025-2026 time frame represents the inflection point where experimental AI security capabilities became operationally deployed at scale. Previous years focused on AI-assisted tools and proof-of-concept demonstrations. The current period demonstrates sustained, autonomous attack operations that fundamentally alter threat actor economics and defensive requirements.

- **CONFIDENCE ASSESSMENT**

  High. Multiple independent sources with direct operational access converge - EU policy implementation, U.K. law enforcement, venture capital deployment and multilateral organization data - on the same directional assessment.

AI represents "acceleration and scale of existing threats, not transformation."

- Paul Foster

While experts disagree on transformation versus acceleration, all confirmed significant operational AI deployment in cybersecurity during 2025-2026. The consistency across different analytical perspectives and institutional contexts supports high confidence in the core judgment.

5. The AI Security Paradigm Shift From Human vs. Human to AI vs. AI Conflict

- **CORE JUDGMENT**

  The cybersecurity industry has crossed a critical threshold from human-versus-human to AI-versus-AI conflict, with agentic AI systems creating unprecedented control challenges that render traditional security architectures structurally inadequate. This transformation is active and urgent - not a future concern - as evidenced by current incident response engagements and compressed attack timelines that have collapsed from weeks to minutes.

- **WHAT THE INTERVIEWS REVEAL**

  The evidence points to a fundamental acceleration in both offensive and defensive capabilities, with attackers gaining significant early advantages. Time-to-exploit has compressed dramatically from the historical seven to 12 days to minutes, as AI-assisted exploit coding eliminates traditional patch windows. As Nadir Izrael, CTO and co-founder at Armis, observed: “A couple of years ago, the time to known exploit used to be seven to 12 days based on which intelligence agency in the world is doing it. We’re down to minutes. Once a vulnerability comes out, coding an exploit for it is as easy as letting Claude or any one of the other agent decoders write it now.”

  This compression reflects a deeper structural change identified across multiple sources: AI removes attacker resource constraints that historically limited exploitation. Phil Venables, partner at Ballistic Ventures, captured this

shift precisely: “The dirty secret of our industry is there’s always been more vulnerabilities that haven’t been exploited than have been exploited, because attackers are resource constrained. That’s no longer the case.”

The scale implications are staggering. Non-human identities now outnumber human identities by an average of 82 to 1, creating governance challenges no current framework can address. More critically, AI agents operate at machine speeds that make human oversight mathematically impossible, as agentic systems complete entire identity life cycles “in about two seconds,” said Kris Burkhardt, CISO at Accenture.

Current incident response data confirms this is not theoretical. Unit 42 responded to more than 700 incidents in 2025, with Michael Sikorski, CTO and vice president of engineering at Palo Alto Networks’ Unit 42, saying: “We’re on-site in an incident response right now, where [improper AI implementation] is happening, and it’s starting to happen more and more.”

WHERE EXPERTS ALIGN - AND DIVERGE

Strong consensus exists that traditional perimeter-based security models are obsolete for AI workloads. Every source acknowledged that agent-to-agent communication, instantaneous lateral movement and behavioral non-determinism require fundamental architectural rethinking, not incremental adaptation of existing tools.

But significant disagreement emerges around defensive timelines and human oversight viability. While some sources maintain that human-in-the-loop controls remain necessary, Apostol Vassilev, research team supervisor at NIST, categorically dismissed this approach: “Human-in-the-loop was the catch-all phrase from about a year ago … but now we know that agents operate at machine speed. The scale is exponentially greater … There is absolutely no way human can keep track of what’s going on.”

Views also diverged on market readiness. Vendors like Cisco and Palo Alto Networks announced AI security products while simultaneously acknowledging fundamental inadequacies. This contradiction reflects an industry caught between market demands for solutions and technical realities of immature capabilities.

“The dirty secret of our industry is there’s always been more vulnerabilities that haven’t been exploited than have been exploited, because attackers are resource constrained. That’s no longer the case.”

- Phil Venables

SUPPORTING PERSPECTIVES

The most striking insight concerns the inversion of traditional attack-defense dynamics. Tom Leighton, co-founder and CEO of Akamai, said: “You’ve got a bigger exposure, and the bad guys now have powerful tools to exploit that. Using AI, you can train your malware to get around the known defenses - which makes it a double whammy.”

The governance challenge extends beyond technical controls to organizational transformation. Sanjay Beri, co-founder and CEO of Netskope, highlighted the shadow deployment reality. “The reality is that 90% of AI usage is shadow usage, and I don’t like the word shadow. I like to call it business unit-led or user-led.”

New attack surfaces are emerging that didn’t exist in pre-AI environments. Nicolas Dupont, founder and CEO of Cyborg, identified a structural vulnerability. “Vector embeddings that are used to do semantic search and are the heart of enterprise AI pipelines are invertible, meaning that you can convert them back to original text, images, audio,

whatever it's representing. It means that they need to be treated with the same level of sensitivity as source material … but they're not."

Perhaps most concerning is the recognition that agents lack human judgment about data sensitivity, as Burkhardt noted: "Agents also don't have the judgment that you have - in the sense that if they stumble across data that they shouldn't, they're not going to report it. They're going to use it because you've given them a task to do."

"Agents also don't have the judgment that you have - in the sense that if they stumble across data that they shouldn't, they're not going to report it. They're going to use it because you've given them a task to do."

- Kris Burkhardt

• STRATEGIC IMPLICATIONS

The evidence suggests three critical shifts required for effective AI security.

First, organizations must abandon the assumption that breach can be contained through post-detection response - agentic lateral movement occurs instantaneously without command-and-control delays that historically provided response windows.

Second, identity and access management must be fundamentally redesigned around agent-behavioral monitoring rather than rule-based access controls.

Third, security architecture must shift from perimeter-based to platform-embedded controls that operate at machine speed and scale.

The market dynamics are already responding to these realities, with investment criteria now requiring AI integration across all cybersecurity categories. But the tooling landscape remains largely experimental, creating a dangerous gap between urgent operational needs and immature solutions.

• CONFIDENCE ASSESSMENT

High. The core judgment is supported by current incident response data, quantified timeline compression, and broad consensus across practitioner, vendor, and research sources. The paradigm shift from human-versus-human to AI-versus-AI conflict is validated by multiple independent observations of attack acceleration and defense gaps. While specific timelines and quantitative projections carry uncertainty, the directional shift and urgency are well-established across the intelligence base.

Kris Burkhardt
CISO, Accenture

6. Expansion of the Critical Infrastructure Attack Surface

- **CORE JUDGMENT**

  Traditional, air-gapped operational technology environments have fundamentally transformed into interconnected systems, creating unprecedented attack surfaces across energy, automotive and manufacturing sectors that outpace current security frameworks.

- **WHAT THE INTERVIEWS REVEAL**

  Evidence from energy, automotive and manufacturing specialists reveals a convergent transformation where historically isolated OT systems now integrate with modern digital infrastructure, dissolving traditional security perimeters.

  Rafael Narezzi, CEO at Centrii, identified battery energy storage systems (BESSs) as a critical attack frontier where “modern APIs mix with OT-era security,” creating grid-level vulnerabilities. Coordinated attacks against BESS infrastructure could function as a distributed denial-of-service equivalent against power grids, where mismatched dispatch across multiple systems could destabilize entire urban areas.

  The automotive sector demonstrates similar attack surface expansion through connected vehicle ecosystems. Kamel Ghali, vice president at Car Hacking Village, highlighted how mobile applications and backend infrastructure create overlooked attack vectors that can compromise entire vehicle

> **“Modern APIs mix with OT-era security” creating grid-level vulnerabilities.**
>
> - Rafael Narezzi

fleets. These interconnected systems enable catastrophic scenarios where adversaries could simultaneously target critical services across entire geographic regions.

Manufacturing environments show the most dramatic architectural shift. Brian Deitch, chief technology evangelist at Zscaler, described the "connected factory" reality, where OT integrates with supply chains, software-as-a-service platforms and cloud infrastructure. This transformation creates multi-directional attack paths that traditional air-gap assumptions cannot address. "If it's reachable, it's reachable. There's a way in. It's just a matter of time," Deitch said, emphasizing the inevitability of exploitation once systems become network-accessible.

The energy sector's decentralization for AI and data center demand amplifies these risks. Renewable energy infrastructure deployment is outpacing security implementation, creating windows of vulnerability during the critical transition period when legacy security models meet modern connectivity requirements.

WHERE EXPERTS ALIGN - AND DIVERGE

All three specialists agree that traditional security frameworks cannot address current attack surface realities, though they emphasize different solution pathways. There is consensus that the Purdue Model for OT security, while not obsolete, requires fundamental complementation rather than incremental modification.

Divergence emerges in strategic priorities. Narezzi emphasizes financial risk quantification as the only governance language that bridges IT/OT divides at board level, while Deitch focuses on business outcome framing including zero downtime and capital expenditure elimination. Ghali prioritizes methodology evolution, advocating for threat analysis and risk assessment approaches that account for safety, financial, operational and privacy impacts simultaneously.

SUPPORTING PERSPECTIVES

The scope of potential impact varies dramatically across sectors. In energy infrastructure, Narezzi warned that coordinated battery system manipulation could create citywide blackouts. "One megawatt is a difference of 250 megawatt to 1500 megawatt dispatch on the battery. You can stop a city."

Automotive scenarios present different but equally concerning possibilities. "What if they decided to ransomware all of the ambulances in a city for a day? What do you do?" Ghali asked, highlighting how fleetwide attacks could paralyze emergency services.

Manufacturing faces persistent lateral movement risks through vendor access.

"If it's reachable, it's reachable. There's a way in. It's just a matter of time."

- Brian Deitch

Deitch noted that traditional East-West microsegmentation "never, ever works" operationally, creating ongoing vulnerability to credential theft and ransomware deployment once initial access is achieved.

The technical sophistication of emerging attack vectors continues to grow. Energy systems face "balancing attacks" that destabilize grid operations through coordinated dispatch manipulation rather than direct production disruption. Automotive systems confront overlooked mobile application vulnerabilities that bypass hardened embedded system protections. Manufacturing environments struggle with heterogeneous device integration where legacy OT cannot be patched or protected using conventional IT security tools.

- **STRATEGIC IMPLICATIONS**

  Critical infrastructure security requires architectural rethinking rather than perimeter reinforcement. The IT-OT boundary no longer represents a meaningful security perimeter, necessitating enterprisewide programs that address interconnected risks across traditionally separate domains. External forcing functions including cyber insurance requirements, regulatory personal liability provisions and compliance mandates create hard deadlines for security modernization, potentially driving faster transformation than internal risk assessments alone.

- **CONFIDENCE ASSESSMENT**

  High. Multiple independent sources from different critical infrastructure sectors provide convergent evidence of attack surface expansion with specific technical examples and consistent directional indicators. While commercial interests may influence solution positioning, the fundamental transformation observations align across energy, automotive and manufacturing domains with concrete supporting evidence.

Brian Deitch
Chief Technology Evangelist, Zscaler

1. Enterprise AI Agent Adoption Has Exploded Beyond Industry Planning Assumptions

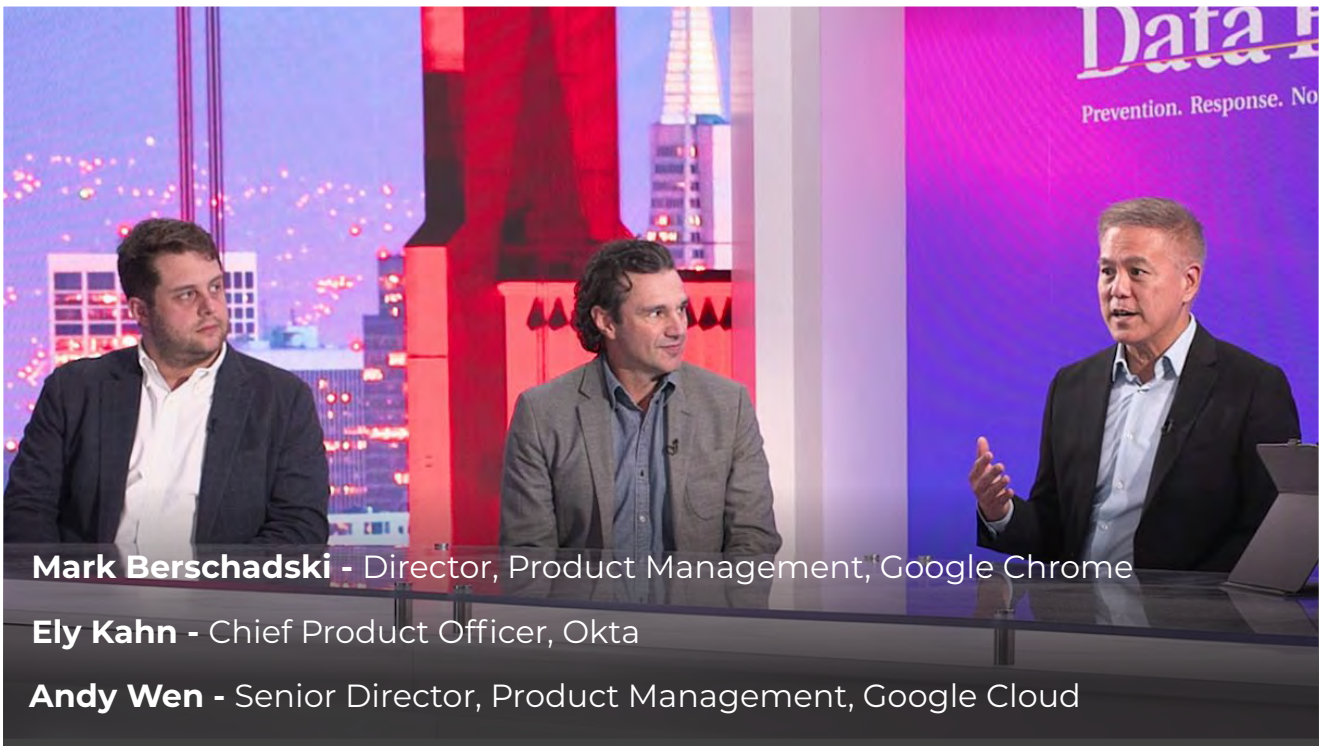

- **CORE JUDGMENT**
  Enterprise AI agent deployment has vastly exceeded initial industry assumptions, with adoption exploding in the first quarter of 2026 and creating an unprecedented gap between deployment scale and security readiness.
- **WHAT THE INTERVIEWS REVEAL** The reality of enterprise AI agent adoption has shattered vendor and industry projections by orders of magnitude. Ely Kahn, chief product officer at Okta, provided the most concrete evidence of this transformation. “We made some assumptions that there would probably be about 25 agents in a large enterprise. Now, talking to some of our larger customers, it’s in the multiple thousands already - and growing an exponential pace.” This represents a planning gap of roughly 100x between initial industry assumptions and actual deployment patterns. This explosion became visible across RSAC 2026, where Jim DuBois, former CIO and CISO at Microsoft, observed firsthand

“We made some assumptions that there would probably be about 25 agents in a large enterprise. Now, talking to some of our larger customers, it’s in the multiple thousands already - and growing at an exponential pace.”

- Ely Kahn

evidence of widespread adoption paired with alarming security gaps. In a session with hundreds of attendees, he witnessed a striking demonstration. “The speaker said, ‘How many of you have agents deployed in your company?’ And way more than 90% of the hands went up. And then the speaker said, ‘How many of you are comfortable with the security you have around the agents?’ Not a single hand went up.”

The scope of this deployment surge extends beyond simple automation. Matt Caulfield, vice president of identity and Duo Security at Cisco, cited survey data showing 85% of organizations experimenting with agentic AI, though only 5% have moved to production - indicating that a massive wave of enterprise deployment is still building. He emphasized the operational implications: Agents operate “at machine speed and at enormous scale” with “broad access to resources” while completely lacking “any judgment.”

The reality of enterprise AI agent adoption has shattered vendor and industry projections by orders of magnitude. Ely Kahn, chief product officer at Okta, provided the most concrete evidence of this transformation. “We made some assumptions that there would probably be about 25 agents in a large enterprise. Now, talking to some of our larger customers, it’s in the multiple thousands already - and growing on an exponential pace.” This represents a planning gap of roughly 100x between initial industry assumptions and actual deployment patterns.

This explosion became visible across RSAC 2026, where Jim DuBois, former CIO and CISO at Microsoft, observed firsthand evidence of widespread adoption paired with alarming security gaps. In a session with hundreds of attendees, he witnessed a striking demonstration. “The speaker said, ‘How many of you have agents deployed in your company?’ And way more than 90% of the hands went up. And then the speaker said, ‘How many of you are comfortable with the security you have around the agents?’ Not a single hand went up.”

The scope of this deployment surge extends beyond simple automation. Matt Caulfield, vice president of identity and Duo Security at Cisco, cited survey data showing 85% of organizations experimenting with agentic AI, though only 5% have moved to production - indicating that a massive wave of enterprise deployment is still building. He emphasized the operational implications: Agents operate “at machine speed and at enormous scale” with “broad access to resources” while completely lacking “any judgment.”

WHERE EXPERTS ALIGN - AND DIVERGE

Sources across vendor executives, practitioners and policy experts converge on the fundamental reality of explosive adoption. The timeline consensus is remarkably consistent - multiple interviewees independently identified Q1 and Q2 of 2026 as the inflection point when agent deployment shifted from experimental to operational scale.

But perspectives diverge sharply on the appropriate response. While some advocate for rapid enabling frameworks to govern agents already deployed, others call for more cautious approaches that prioritize security controls before broader rollout. This tension reflects the core challenge organizations face: Agents are already in production whether security teams are ready or not.

SUPPORTING PERSPECTIVES

The interviews reveal that this adoption surge is driven by board-level pressure rather than IT department initiatives. As Art Gilliland, CEO at Delinea, observed, organizations are relaxing policies and governance to go as fast as they can due to business imperatives. This top-down pressure creates a structural challenge where deployment speed consistently outpaces security framework development.

Jeremy Grant, managing director of technology business strategy at Venable, captured the systemic nature of this problem. “Agents are being deployed more quickly

than we have these security controls for them - that's becoming a bit of a concern." The gap between deployment and governance represents not just a technical challenge but a fundamental mismatch between business velocity and security maturity.

- **STRATEGIC IMPLICATIONS**

  The 25-to-1000 agents gap signals that identity infrastructure designed for human-scale operations is structurally inadequate for agent-scale deployment. "The number of AI agents will far exceed the number of humans in most organizations," Kahn said. This requires fundamental architectural changes to identity management systems - not incremental improvements to existing frameworks.

  The zero-confidence security finding from DuBois' RSAC observation suggests widespread governance failure across the industry. Organizations are operating with massive blind spots, deploying thousands of agents while lacking basic visibility into their security posture. This creates systemic risk across enterprise environments and indicates that security frameworks have not kept pace with deployment realities.

"Agents are being deployed more quickly than we have these security controls for them - that's becoming a bit of a concern."

- Art Gilliland

- **CONFIDENCE ASSESSMENT**

  High. Multiple independent sources with specific data points converge on the same phenomenon. The 25-to-1000 agents gap from Okta customer data, combined with the RSAC floor observation of more than 90% deployment but 0% confidence, provides compelling empirical evidence. Source diversity includes vendor executives with direct customer visibility, practitioners with operational experience, and policy experts with cross-industry perspective, all pointing to the same conclusion about adoption scale exceeding planning assumptions.

2. Data Security's AI Reckoning: The 80% Gap Crisis

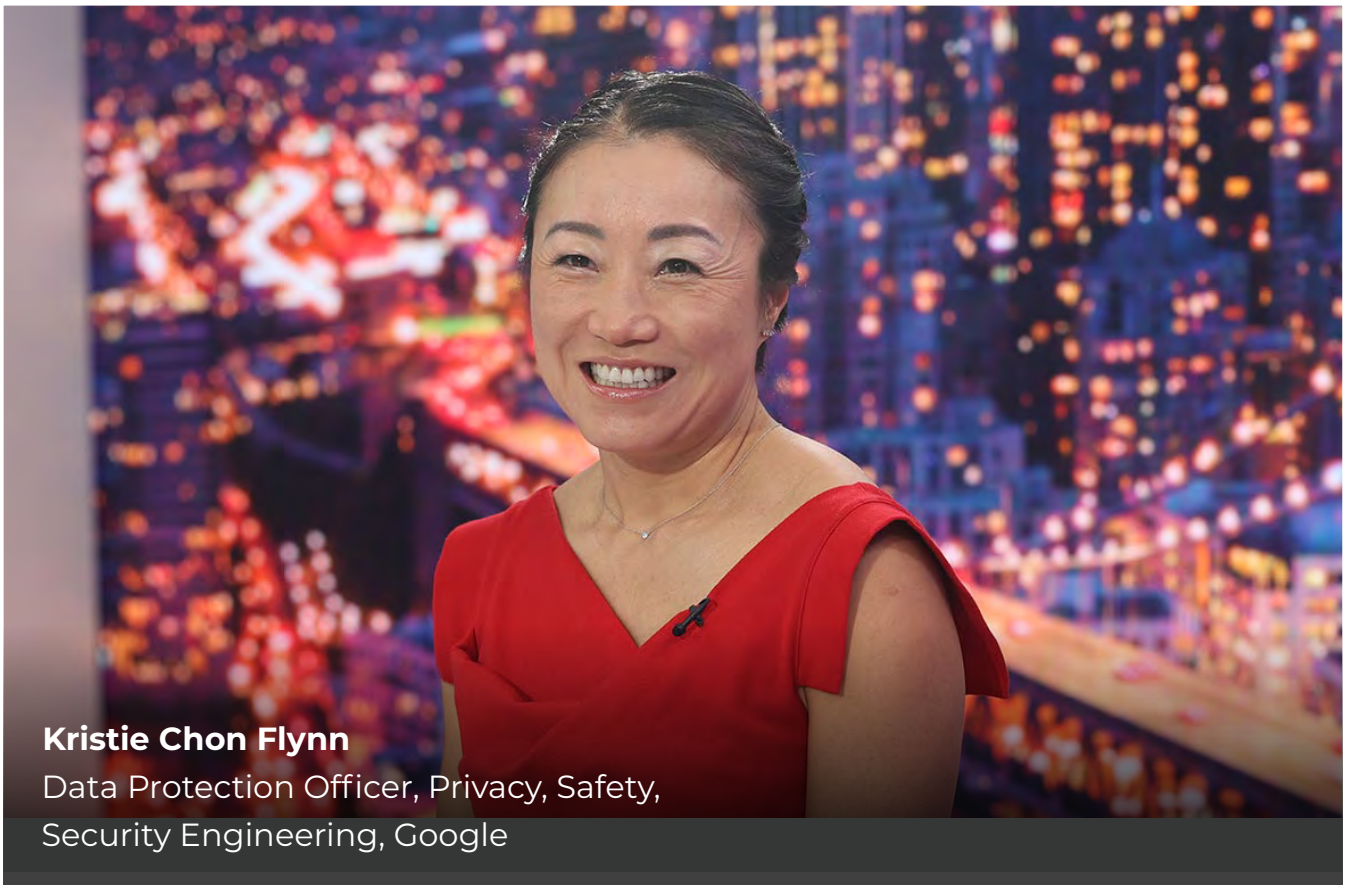

Kristie Chon Flynn
Data Protection Officer, Privacy, Safety, Security Engineering, Google

- **CORE JUDGMENT**

  Enterprise data security operates in systemic crisis as AI adoption accelerates, with 80% of organizations attempting AI deployment on fundamentally inadequate data control foundations, while agentic AI introduces qualitatively new attack patterns that existing security models cannot address at machine speed.

- **WHAT THE INTERVIEWS REVEAL**

  The data tells a stark story of organizational overreach. Landen Brown, field CTO at MIND, opened with findings from their CISO survey: "Only 20% of CISOs feel like their data security maturity is at a place organizationally where they can safely adopt AI without staying up at night." The inverse - 80% of organizations pushing forward with AI despite insufficient data security foundations - represents perhaps the largest systematic risk exposure in enterprise security today.

  This exposure becomes catastrophic when combined with endemic data overexposure. Brian Vecci, field CTO at

> **Privacy enhancing technologies "have come a long way. It's no longer an academic research topic. It is something that's tangible and available to be applied throughout the AI development life cycle."**
>
> - Kristie Chon Flynn

Varonis, quantified the scope: “The average user or identity only uses about 3% of the data that they have access to.”

Yaki Faitelson, CEO and co-founder of Varonis, sharpened this further. “More than 90% of accessible data is irrelevant to the identity accessing it.” These statistics, confirmed across multiple sources, reveal systematic over-permissioning that creates enormous blast radius for AI workloads.

The speed dimension compounds every existing weakness. “When we introduce the agentic future into the mix in these organizations, we start to see that data starts to move exponentially fast - at machine speed,” Brown said. Traditional controls designed for human-paced access patterns simply cannot operate at what Faitelson characterized as “human × 1000+” speed.

Agentic AI introduces entirely novel failure modes that current security frameworks don’t recognize. Faitelson identified agent-to-agent permission escalation as a new class of lateral movement. “In the agentic world, an agent can access your entire data estate in seconds. It’s like a Pac-Man from hell - and if the agent doesn’t have permissions, he will try to get them … you have an agent, he doesn’t have access to data, asking another agent to retrieve the data for it.” This pattern falls outside current IAM and PAM detection logic entirely.

WHERE EXPERTS ALIGN - AND DIVERGE

Universal consensus exists on the inadequacy of current data security posture for AI workloads. Every practitioner and vendor voice confirmed systematic data overexposure, speed mismatches between AI operations and security controls, and the fundamental architectural challenge facing enterprise security teams.

Sharp disagreement emerged around solution readiness. Kristie Chon Flynn, data protection officer for privacy, safety and security engineering at Google, argued that privacy enhancing technologies “have come a long way. It’s no longer an academic research topic. It is something that’s tangible and available to be applied throughout the AI development life cycle.” But the 20% CISO confidence statistic suggests production deployments lag behind technical possibility.

The most fundamental split concerns compliance frameworks themselves. Flynn advocated for evolution: privacy-by-design must transition “from reactive bolt-on compliance to proactive embedded architectural discipline.” But Valerie Lyons, chief operations officer and senior privacy specialist at BH Consulting, challenged the entire compliance-first approach. “Securing something, making something private, making it compliant - it isn’t enough, especially with AI apps.” Lyons pointed to real harm from technically compliant systems, including a perfectly GDPR-compliant wellness app misused as a grief counseling tool and the Dutch SyRI anti-fraud AI that discriminated for a decade while meeting all regulatory requirements.

SUPPORTING PERSPECTIVES

The governance challenge extends beyond traditional security boundaries. Srikanth Venkat, senior director of the Databolt product at Capital One Software, introduced the concept of “AI data supply chain” governance. “You cannot control what you cannot observe. So you need that governance plane to say what data is being fed into this AI model. How are the agents using it, where is this data coming from, what transformation is it going through?” Venkat said.

Sebastien Cano, senior vice president of cyber security products at Thales, connected data visibility to agentic risk. “Having agents roaming your IT infrastructure with unclear boundaries, without a clear mandate, without any guardrails to limit their action is like having actual people with unlimited access to your resources.”

The technology response varies significantly

in maturity. Venkat argued that tokenization addresses the structural friction "between the need to use sensitive proprietary enterprise data to derive AI value, and the need to protect it from exposure through models, agents and inference operations." Flynn highlighted synthetic data, differential privacy and trusted execution environments as production-ready tools. But implementation complexity remains high, and the 80% gap suggests that deployment at scale remains elusive.

STRATEGIC IMPLICATIONS

Organizations face a fundamental choice: accept catastrophic AI-enabled breach risk or slow AI adoption to competitive disadvantage. Faitelson captured this as the 3% paradox. "Organizations have connected only 3% of their data estate to the new AI technology. If you move too fast without controls, AI will kill you - and if you move too slowly, AI will kill you as well."

The agent-to-agent permission escalation pattern identified by Faitelson represents the emergence of entirely new attack classes that existing security architectures cannot detect or prevent. Security teams must migrate from application-layer control to data-layer and agent-layer controls immediately, not as future planning exercise.

Privacy teams face similar architectural pressure. Flynn's call for dynamic, technically integrated governance frameworks reflects recognition that traditional policy-driven approaches cannot operate at AI speeds. Organizations treating privacy as compliance overlay rather than architectural foundation will find AI deployments consistently blocked by data protection requirements.

"You cannot control what you cannot observe. So you need that governance plane to say what data is being fed into this AI model. How are the agents using it, where is this data coming from, what transformation is it going through?"

- Srikanth Venkat

YEAR-OVER-YEAR SHIFT

The 2026 narrative marks a decisive shift from AI security as emerging concern to operational crisis. Where 2025 discussions focused on future AI risks, the 2026 interviews consistently described current organizational reality: agents operating without adequate controls, data moving at machine speed and permission systems failing at scale. The transition from theoretical to operational represents the clearest evolution in enterprise AI risk posture.

CONFIDENCE ASSESSMENT

High. Statistical evidence from multiple independent sources creates compelling picture of systematic data security inadequacy. The 20% CISO confidence figure, 80% organizational overexposure and 3%/90% utilization statistics provide quantified foundation for crisis assessment. Agent-to-agent permission escalation represents novel, technically specific threat pattern with clear operational implications. Technology solution claims vary in verification, but problem diagnosis enjoys universal expert consensus.

3. AI Governance Crisis: Speed Asymmetry Drives Systemic Risk

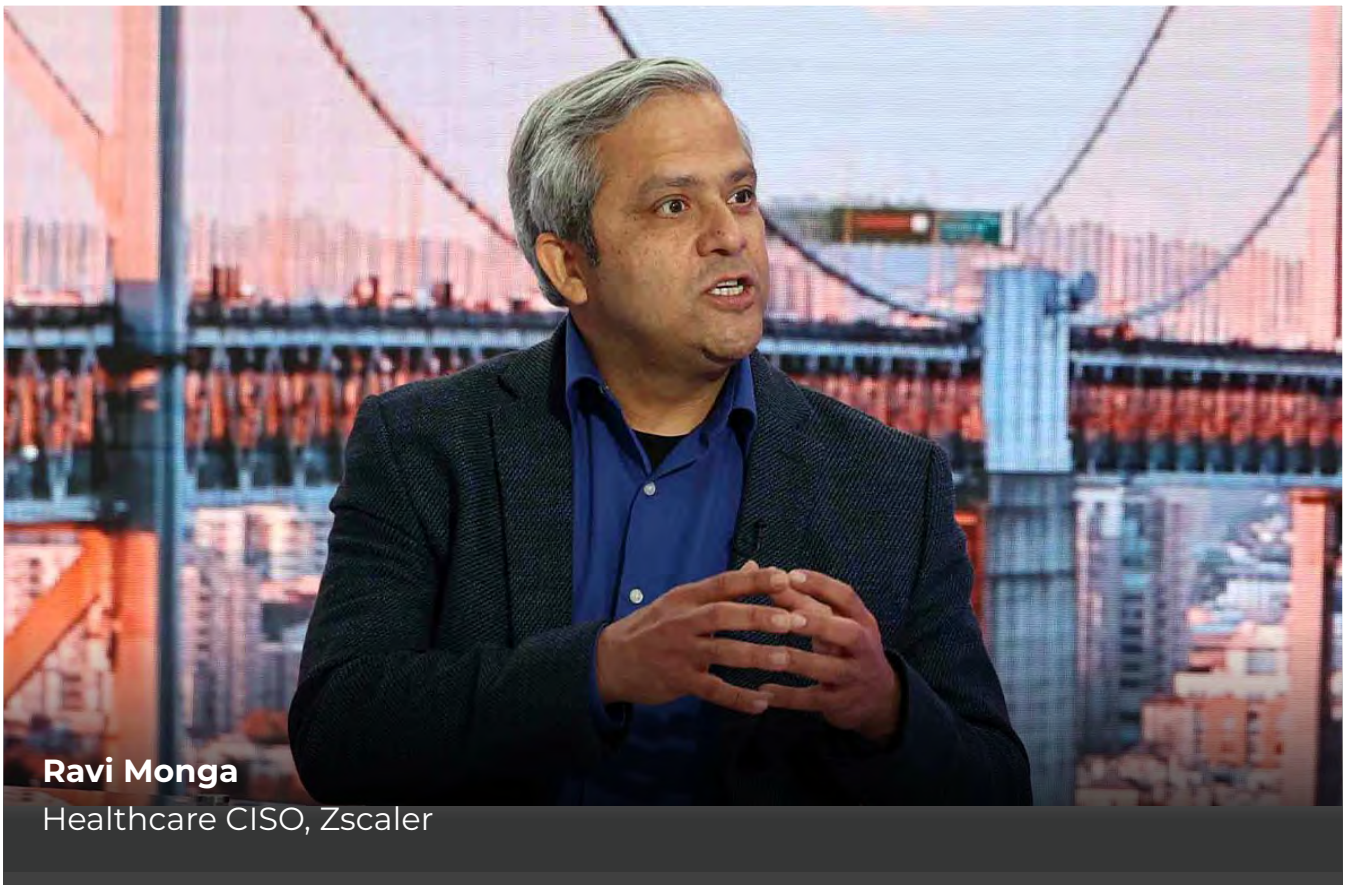

Ravi Monga
Healthcare CISO, Zscaler

- **CORE JUDGMENT**

  Organizations are systematically underestimating AI security governance requirements, with embedded AI dependencies and accelerated threat timelines creating compounding risks that outpace institutional adaptation speeds. This speed asymmetry between AI adoption (weeks) and governance cycles (years) is generating what experts describe as "invisible" security exposures across critical supply chains.

- **WHAT THE INTERVIEWS REVEAL**

  The strongest signal across all sources centers on a fundamental mismatch: AI deployment is advancing at organizational speeds measured in weeks, while governance frameworks operate on institutional cycles measured in years. This asymmetry is most acute in healthcare and supply chain contexts, where the consequences extend beyond individual organizations.

  Ravi Monga, healthcare CISO at Zscaler, captured this dynamic precisely: "It is a tale of

> **"It is a tale of two speeds. So leadership and governance is talking about AI implementations, AI projects, but at the same time, users are already using it."**
>
> - Ravi Monga

two speeds. So leadership and governance is talking about AI implementations, AI projects, but at the same time, users are already using it." This isn't theoretical - interviews indicate live deployment of AI for PHI-intensive clinical workflows like note summarization and revenue cycle operations, while formal governance frameworks remain under development.

Evidence suggests this speed gap is systematically widening rather than resolving. Eric Wenger, senior director for technology policy at Cisco, shared Cisco Talos data showing AI tools compressing exploit timelines, with December 2025 vulnerabilities dominating annual charts, indicating adversaries are already operationalizing AI acceleration while defensive organizations lag in governance implementation.

Perhaps most concerning is what Monga termed "invisible AI" - vendor-embedded AI functionality that creates governance challenges structurally different from user-initiated shadow AI. "Every vendor is bolting on AI solutions to their offering. So it becomes a lot of invisible AI," Monga said. This represents a supply chain risk expansion that traditional vendor management approaches weren't designed to address.

"You can't blame the AI from a regulatory point of view. You can outsource anything as a business decision, but at the end of the day, you own the risk."

- David Cass

• WHERE EXPERTS ALIGN - AND DIVERGE

All sources converged on visibility as the foundational challenge. Whether discussing "invisible AI," technical debt hiding in supply chains or asset inventory gaps, every expert identified the inability to comprehensively see AI dependencies as the primary barrier to effective governance.

But views diverged significantly on strategic priorities. Hans de Vries, chief cybersecurity and operations officer at European Union Agency for Cybersecurity - ENISA, explicitly reasserted prevention - training, exercises and process hygiene - as "co-equal to detection and response." This sits in productive tension with industry's dominant detection-first focus, suggesting regulatory authorities are pushing back against reactive approaches.

SUPPORTING PERSPECTIVES

The governance adequacy question revealed stark assessments. David Cass, CISO at Keyrock, argued that traditional model risk management frameworks are "wholly inadequate" for generative AI and agentic AI systems. "You can't blame the AI from a regulatory point of view. You can outsource anything as a business decision, but at the end of the day, you own the risk," Cass said.

> **"Six thousand organizations couldn't function because the whole supply chain was disrupted. JLR is not a critical infrastructure, but it has an important role to play within the whole economic sphere."**
>
> - Hans de Vries

Chad Alessi, managing director for cybersecurity at CTG, introduced the concept of "cyber deterioration" - the gradual, invisible erosion of security foundations driven partly by resource diversion to new initiatives. "Organizations are continuing to do what the bare minimum requirements for compliance are, but over time, erosion of their cybersecurity foundation is starting," Alessi said. His healthcare study of more than 4,000 organizations found 25% showing measurable posture decline despite maintaining compliance - a supply chain risk given healthcare's interconnected data ecosystem.

The supply chain implications crystallized through de Vries' account of the Jaguar Land Rover incident. "Six thousand organizations couldn't function because the whole supply chain was disrupted. JLR is not a critical infrastructure, but it has an important role to play within the whole economic sphere."

STRATEGIC IMPLICATIONS

The synthesis reveals AI governance as a systems-level challenge requiring fundamental shifts in organizational approach. The speed asymmetry isn't resolving naturally - it requires deliberate intervention to compress governance cycles or slow AI adoption in critical contexts. The "invisible AI" phenomenon demands supply chain auditing capabilities most organizations haven't developed, while technical debt creates multiplicative AI exploitation risks that legacy approaches can't address.

For healthcare and financial services, these governance gaps respectively carry sector-specific amplification risks through patient safety consequences and systemic financial stability implications.

CONFIDENCE ASSESSMENT

High. The speed asymmetry pattern appears consistently across all source types - vendor policy, CISO practitioners and regulatory authorities - and organizational contexts. Primary data from Cisco Talos on exploit acceleration provides technical validation, while the healthcare and supply chain case studies demonstrate concrete manifestations of governance lag impacts.

ACT III - THE RESPONSE

1. The AI Imperative: Cybersecurity's Transformation Under Fire

- **CORE JUDGMENT**

  AI has crossed the threshold from advantageous to essential in cybersecurity operations, driven by an attack timeline compression so severe that it has created an unbridgeable human response gap. Organizations questioning whether to adopt AI are not evaluating a strategic option - they are acknowledging they are "so far behind because adversaries have rapidly increased their capabilities" through AI adoption.

- **WHAT THE INTERVIEWS REVEAL**

  The evidence for AI's essential role emerges from a brutal mathematical reality: Attack timelines have compressed beyond human response capacity. CrowdStrike data shows the fastest eCrime attack completing in just 27 seconds. Vulnerability exploitation windows collapsed from five months in 2023 to approximately 1.5 days in 2026, with some threat scenarios reaching "negative one day" - where exploitation precedes patch availability entirely.

  This speed gap represents more than incremental change. Rob T. Lee, chief AI officer and chief of research at SANS Institute, said, "Almost every single attack has AI layered into the offensive team's capabilities," democratizing nation-state capabilities to teams of just two individuals. What previously required "thousands of personnel on a nation-state team can now be accomplished with just two individuals," fundamentally altering threat economics, Lee said.

The defensive response gap is equally stark. Interviews revealed that approximately 40% of SOC alerts go uninvestigated due to volume, while attack dwell times have compressed from hours to minutes. In some cases, adversaries can interrogate and effectively attack an organization in just a couple of seconds in some cases.

The most compelling evidence comes from successful defensive AI deployments already producing measurable results. Amazon's "ATA" program paired red team and blue team AI agents against reverse SSH tunneling, generating 64 novel detections in four hours - work that would have required 10 to 14 human days. These detections were pushed to all AWS GuardDuty customers within 10 days, demonstrating AI's capacity to operate at both machine speed and global scale.

"Almost every single attack has AI layered into the offensive team's capabilities."

- Rob T. Lee

WHERE EXPERTS ALIGN - AND DIVERGE

Universal consensus emerged around AI's necessity, with no significant contradictions on directional need. Sources agreed that adversaries adopted AI first, creating an asymmetric advantage that defensive teams must close. The "humans-in-the-loop" model gained consistent support across vendor and practitioner sources as the optimal governance framework - humans retaining strategic oversight while AI handles operational execution.

But sharp disagreements persist around automation boundaries. Sumedh Thakar, president and CEO at Qualys, argued that "manual remediation is dead" and "autonomous remediation is the only viable response" to negative-day exploitation timelines. In direct contradiction, Evan Peña, co-founder and chief offensive security officer of Armadin, explicitly stated that automated remediation is "not auto right now" due to production safety concerns, maintaining human oversight for all execution decisions.

This fundamental tension - urgency versus safety - remains unresolved across the industry, reflecting the genuine challenge of balancing AI's speed requirements against the risk of automated mistakes in production environments.

SUPPORTING PERSPECTIVES

The interviews revealed that successful AI implementation depends critically on data foundations, not model sophistication. John Morgan, senior vice president and general manager of security at Splunk, assessed enterprise readiness at roughly "two and a half" out of five, explaining that "organizations are trying to figure out, Where is my data, how do I procure it? And then what's important?" He added that success depends on a "distributed data strategy" that curates data "in a way that can be ingested by AI."

"When AI is coming without knowing a specific organization, it is like somebody that joined the company without having tribal knowledge," said Lior Div, co-founder and CEO at 7AI. "Once AI has the tribal knowledge, you will see improvement in the accuracy of what AI is doing for you."

Practitioners consistently emphasized the evolution rather than replacement of human roles. "AI is almost like an exoskeleton around the analyst," said Michael Nichols, general manager at Elastic, underscoring that AI should be transparently integrated into existing workflow, but not become a replacement.

STRATEGIC IMPLICATIONS

The interview signals reveal that AI adoption in cybersecurity is not optional but essential for organizational survival. The attack timeline compression documented across interviews creates a structural requirement that cannot be addressed through traditional human-scale responses.

Organizations must prepare for a fundamental shift in security operations models, with humans transitioning from operational roles to strategic orchestration. But the absence of industry governance standards means early adopters are building bespoke controls without reference architectures, creating both competitive advantages and implementation risks.

The data foundation requirement cannot be understated. Organizations lacking comprehensive data visibility and contextual knowledge graphs will find AI implementations producing noise rather than signal, potentially worsening rather than improving security postures.

YEAR-OVER-YEAR SHIFT

The 2026 findings represent a maturation of AI discussions from theoretical possibilities to operational necessities. Previous years' findings focused on "what AI could do"; 2026 findings demonstrate organizations demanding proof of "what AI is actually doing" in their environments. The shift from exploration to deployment creates new pressures around governance, measurement and integration with existing security infrastructures.

"Organizations are trying to figure out, Where is my data, how do I procure it? And then what's important?"

- John Morgan

CONFIDENCE ASSESSMENT

High. The core judgment that AI has become essential is supported by consistent empirical evidence across threat timeline compression, attack volume proliferation and demonstrated defensive successes. The human capacity gap is mathematically unavoidable considering documented attack speeds.

Medium. Implementation approaches and governance frameworks remain varied and maturing. While "humans-in-the-loop" shows broad support, specific operational boundaries and safety mechanisms require further validation across diverse organizational contexts.

2. The Proof Era: AI in Cybersecurity Operations Reaches Production Scale

Steve Januario
CIO, Bill.com

- **CORE JUDGMENT**

  AI deployment in cybersecurity operations has definitively moved from capability demonstrations to production-scale implementation with measurable operational value. Organizations are no longer asking "What can AI do?" but demanding concrete evidence of what AI is accomplishing in their environments, marking the beginning of what industry leaders term the "proof era." But this rapid adoption has created a dangerous governance gap: the urgent threat landscape is driving deployment faster than security controls can be implemented.

- **WHAT THE INTERVIEWS REVEAL**

  The evidence for operational AI deployment is both widespread and quantifiable. Bill.com CIO Steve Januario said, "We went from 13% deflection upward of between 40% and 60% depending on the given month," describing enterprise-grade AI performance that extends far beyond pilot programs. Most striking was Lee's proof-of-concept demonstration: His SIFT workstation compressed a complete APT investigation from three days to 14 minutes.

  These operational successes reflect a fundamental architectural shift. Multiple industry leaders converged around a three-layer framework: knowledge graphs providing organizational context, specialized agents performing specific security tasks and proof layers combining human oversight

with reinforcement learning. “To have a strong security operation, you need a strong knowledge graph, build out a set of agents to perform a set of skills that would be hard for you to acquire yourself, … and then the last piece is trust,” said Arctic Wolf President and CEO Nick Schneider.

WHERE EXPERTS ALIGN - AND DIVERGE

Strong consensus emerged around AI’s operational necessity driven by compressed attack timelines. Morgan noted that “what used to be a dwell time of hours is now minutes, and they can interrogate and effectively attack an organization in just a couple of seconds in some cases.” This speed gap creates what experts universally described as an automation imperative.

But significant divergence exists around implementation approaches. Vendor executives showed optimism about rapid scaling, while practitioners expressed caution about governance readiness. The human oversight model also remains contested, with sources using “human-in-the-loop,” “human-on-the-loop” and “human-above-the-loop” interchangeably despite meaningful operational differences.

SUPPORTING PERSPECTIVES

The governance challenge is compounded by what multiple sources identified as “agent sprawl” - a qualitatively worse problem than traditional tool sprawl. Barracuda Networks CEO Rohit Ghai warned, “In the pre-AI world, we were dealing with tool sprawl. Post AI, we are living in a world where you have agent sprawl, not just tool sprawl. And this is a much bigger problem, because it’s not just conflicting information, it’s conflicting actions.”

Lee provided the most urgent framing of the threat landscape. “Almost every single attack has AI layered into the offensive team’s capabilities - not only in automation, but in spreading the attack space. For example, targeting agents being able to develop capabilities that evade the most state-of-the-art defenses.”

STRATEGIC IMPLICATIONS

The proof era fundamentally changes procurement and implementation strategies. Organizations can no longer evaluate AI security tools based on capability demonstrations alone - they must demand runtime evidence and operational metrics. This benefits incumbents with large operational datasets while creating barriers for new entrants without proven track records.

The governance gap represents both immediate risk and competitive opportunity. Organizations that solve AI observability, auditability and governance frameworks now will have sustainable advantages as the technology scales. Those that prioritize deployment speed over security controls face potentially catastrophic exposure.

YEAR-OVER-YEAR SHIFT

The transformation from 2025’s “AI experimentation” to 2026’s “AI production deployment” represents one of the most rapid technology adoption cycles in cybersecurity history. What changed was not capability but confidence - enough operational evidence now exists to justify large-scale investment and deployment decisions.

CONFIDENCE ASSESSMENT

High. The evidence for production-scale AI deployment is overwhelming and comes from diverse, credible sources including vendor executives with operational data, practitioners with measurable results and industry observers with broad visibility. The governance gap is equally well-documented through both quantitative evidence and qualitative assessment across multiple interviews. The trajectory toward continued acceleration is supported by both threat landscape pressure and demonstrated ROI.

3. Investment Market Transformation

Eric McAlpine
Founder and CEO, Momentum Cyber

- **CORE JUDGMENT**

  Cybersecurity investment is experiencing the most significant structural transformation in decades, driven by compressed development cycles, unprecedented capital concentration and the emergence of entirely new AI security categories at velocities never before seen in the sector.

- **WHAT THE INTERVIEWS REVEAL**

  The numbers tell a story of fundamental change. Israeli cybersecurity companies alone raised $5.1 billion in 2025 - the highest figure since 2021 - with average seed rounds reaching $9.5 million, according to Ofer Schreiber, senior partner at YL Ventures. Globally, cybersecurity M&A deals hit over 400 transactions in 2025, a staggering 270% year-over-year increase from 2024's $46.1 billion. Strategic buyers dominated the landscape, deploying 92% of all M&A capital in 2025.

> **"A year ago, there was no category on the cyberscape at all for AI security - and that was a big gaping hole."**
>
> - Eric McAlpine

"A year ago, there was no category on the cyberscape at all for AI security - and that was a big gaping hole," said Eric McAlpine, founder and CEO of Momentum Cyber. "We now track over 400 AI security companies" out of roughly 4,200 total cybersecurity firms, representing the creation of 12 new market categories for AI security, where none existed previously.

The transformation extends beyond deal

volume to fundamental changes in how capital flows. Traditional three-phase investment cycles have compressed into single mega-rounds for AI-native companies with proven founders. This shift reflects both private market maturation and the reality that AI-enabled development allows companies to ship code in six weeks rather than the six to 12 months required by previous generations of security startups.

Private equity firms are struggling to adapt as they are caught between established SaaS valuation models and the realities of AI-native company performance. “PE is a little bit entrenched with this new tectonic shift in AI,” McAlpine said. “It’s happened at such a rapid pace that it’s going to take time to reprice the old SaaS model to a more AI-native model.”

time to reprice the old SaaS model to a more AI-native model.”

WHERE EXPERTS ALIGN - AND DIVERGE

Universal agreement exists on the scale and speed of structural transformation. Every investor interviewed confirmed that traditional cybersecurity investment frameworks are becoming obsolete, replaced by new models built around AI-native development velocity and accelerated category formation.

The strongest consensus centers on strategic buyer dominance. Multiple sources confirmed that large incumbents - CrowdStrike, Check Point, Palo Alto Networks and Cisco - are in an active AI-capability acquisition race directly linked to their product road maps. Umesh Padval, managing partner at Seligman Ventures, said that incumbents face an “Innovator’s Dilemma” between bolting AI onto existing architectures versus building AI-native platforms that could cannibalize current revenue streams.

Divergence emerges around timeline and sustainability. Some investors predict market stabilization within 1 to 3 years, while others suggest 3 to 5 years before new valuation frameworks and competitive dynamics settle. The sustainability question remains particularly contested - whether the hundreds of AI-native companies funded in just two years represent genuine innovation or an overfunding bubble that will produce significant losses.

SUPPORTING PERSPECTIVES

Global tier-1 venture capital firms have fundamentally altered their approach to the Israeli ecosystem, shifting from growth-stage to seed-stage participation. “The Israeli ecosystem has proven that it produces category-defining or category-leading companies on an annual basis - one, two or three multi-billion-dollar companies in security are being founded in Israel every year,” Schreiber said.

This recognition has accelerated competitive dynamics within categories. “Almost every category in security becomes very competitive and very crowded, very fast - much faster than before,” Schreiber said. “The race to who emerges as the category winner is very aggressive and much faster than before.”

“The Israeli ecosystem has proven that it produces category-defining or category-leading companies on an annual basis - one, two or three multi-billion-dollar companies in security are being founded in Israel every year.”

- Ofer Schreiber

The speed of change is forcing early strategic responses. Cato Networks’ acquisition of Aim Security demonstrates how incumbents are using early-stage M&A to enter new AI categories before competitive positions

solidify. “What we’re witnessing now is the major technology shift in the history of software,” said Ori Barzilay, partner at Team8. “There’s no other way to bet on AI-native [software], because they’re actually designing their products in a much more efficient way.”

- ## STRATEGIC IMPLICATIONS

The transformation creates both opportunities and risks that extend beyond traditional investment considerations. The compression of investment cycles concentrates capital around proven founders, creating winner-take-most dynamics where early movers capture majority funding. Organizations beginning AI-native security platform migrations now will likely hold advantages when the market matures over the next 3 to 5 years.

For incumbents, the 92% strategic buyer dominance in M&A suggests that major platform companies are treating AI security acquisition as core to competitive survival rather than opportunistic expansion. The PE repricing gap creates a temporary acquisition window for operating companies before new valuation frameworks stabilize.

But the landscape of over 400 AI security companies creates selection complexity for CISOs and evaluation fatigue for enterprises that cannot practically assess 10 new AI security tools daily. The speed-versus-security tension inherent in six-week development cycles raises questions about security-by-design practices during this rapid market formation period.

- ## YEAR-OVER-YEAR SHIFT

The shift from zero AI security categories in early 2025 to 12 distinct categories by March 2026 represents the fastest market segment formation in cybersecurity history. Investment patterns have evolved from traditional growth-stage focus to seed-stage concentration as global VCs recognize the reliability of the Israeli ecosystem’s category-defining company production.

M&A patterns show a decisive shift toward strategic consolidation, with the 92% strategic buyer dominance in 2025 representing a fundamental change from historically more balanced strategic-PE competition. This indicates that cybersecurity has become too strategically critical for large platforms to rely on traditional financial buyer intermediation.

- ## CONFIDENCE ASSESSMENT

High. The structural transformation is supported by quantitative data from multiple authoritative sources, including proprietary M&A datasets and verified funding figures. The directional shift toward AI-native companies, compressed investment cycles and strategic buyer dominance is consistent across all investor perspectives and backed by concrete transaction evidence.

The timeline predictions and specific outcome forecasts carry moderate confidence, as they represent investor thesis development rather than validated market outcomes. But the core judgment of structural transformation is supported by unprecedented data points across multiple market dimensions.

“What we’re witnessing now is the major technology shift in the history of software. There’s no other way to bet on AI-native [software], because they’re actually designing their products in a much more efficient way.”

- Ori Barzilay

INTERVIEW SOURCES (IN ORDER OF APPEARANCE)

- **Ed Skoudis,** *President, SANS Technology Institute*
- **Sandra Joyce,** *Vice President, Google Threat Intelligence*
- **Ian Thiel,** *Co-Founder and COO, Sublime Security*
- **Matt Olney,** *Director, Threat Intelligence and Interdiction, Cisco Talos*
- **Scott Taylor,** *Director, Cyber Resilience, Field Solution Architects, Everpure*
- **Rick Orloff,** *CISO, Everpure*
- **Anthony Cusimano,** *Chief Evangelist and Director, Solutions Marketing, Object First*
- **Steven Gerry,** *Vice President, Sales, Tidal Cyber*
- **Tim Pappa,** *Incident Response Engineer, Cyber Deception Strategy, Walmart Global Tech*
- **Despina Spanou,** *Director General, Communications Networks and Technology, European Commission*
- **Dave DeWalt,** *CEO, NightDragon*
- **Akshay Joshi,** *Head, Centre for Cybersecurity, World Economic Forum*
- **Paul Foster,** *Head, National Cyber Crime Unit, U.K. National Crime Agency*
- **Nadir Izrael,** *CTO and Co-Founder, Armis*
- **Phil Venables,** *Partner, Ballistic Ventures*
- **Kris Burkhardt,** *CISO, Accenture*
- **Michael Sikorski,** *CTO and Vice President, Engineering, Unit 42, Palo Alto Networks*
- **Apostol Vassilev,** *Research Team Supervisor, NIST*
- **Tom Leighton,** *Co-Founder and CEO, Akamai*
- **Sanjay Beri,** *Co-Founder and CEO, Netskope*
- **Nicolas Dupont,** *Founder and CEO, Cyborg*
- **Rafael Narezzi,** *CEO, Centrii*
- **Kamel Ghali,** *Vice President, Car Hacking Village*
- **Brian Deitch,** *Chief Technology Evangelist, Zscaler*
- **Ely Kahn,** *Chief Product Officer, Okta*
- **Jim DuBois,** *Former CIO and CISO, Microsoft*
- **Matt Caulfield,** *Vice President, Identity and Duo Security, Cisco*
- **Art Gilliland,** *CEO, Delinea*
- **Jeremy Grant,** *Managing Director, Technology Business Strategy, Venable*
- **Landen Brown,** *Field CTO, MIND*
- **Brian Vecci,** *Field CTO, Varonis*
- **Yaki Faitelson,** *CEO and Co-Founder, Varonis*
- **Kristie Chon Flynn,** *Data Protection Officer, Privacy, Safety and Security Engineering, Google*
- **Valerie Lyons,** *Chief Operations Officer and Senior Privacy Specialist, BH Consulting*
- **Srikanth Venkat,** *Senior Director, Databolt Product, Capital One Software*
- **Sebastien Cano,** *Senior Vice President, Cyber Security Products, Thales*
- **Ravi Monga,** *Healthcare CISO, Zscaler*
- **Eric Wenger,** *Senior Director, Technology Policy, Cisco*
- **Hans de Vries,** *Chief Cybersecurity and Operations Officer, ENISA*
- **David Cass,** *CISO, Keyrock*
- **Chad Alessi,** *Managing Director, Cybersecurity, CTG*
- **Rob T. Lee,** *Chief AI Officer and Chief of Research, SANS Institute*
- **Sumedh Thakar,** *President and CEO, Qualys*
- **Evan Peña,** *Co-Founder and Chief Offensive Security Officer, Armadin*
- **John Morgan,** *Senior Vice President and General Manager, Security, Splunk*
- **Lior Div,** *Co-Founder and CEO, 7AI*
- **Michael Nichols,** *General Manager, Elastic*
- **Nick Schneider,** *President and CEO, Arctic Wolf*
- **Steve Januario,** *CIO, Bill.com*
- **Rohit Ghai,** *CEO, Barracuda Networks*
- **Ofer Schreiber,** *Senior Partner, YL Ventures*
- **Eric McAlpine,** *Founder and CEO, Momentum Cyber*
- **Umesh Padval,** *Managing Partner, Seligman Ventures*
- **Ori Barzilay,** *Partner, Team8*

About ISMG

IInformation Security Media Group (ISMG) is the world's largest media organization devoted solely to cybersecurity, information technology, artificial intelligence and operational technology. Each of its 38 media properties provides education, research and news that is specifically tailored to key vertical sectors including banking, healthcare and the public sector; geographies from North America to Southeast Asia; and topics such as data breach prevention, cyber risk assessment, OT security, AI and fraud. Its annual global summit series connects senior security professionals with industry thought leaders to find actionable solutions for pressing cybersecurity challenges.

Contact

(800)944-0401 · sales@ismg.io

ISMG

BANK INFO SECURITY · CAREERS INFO SECURITY · HEALTHCARE INFO SECURITY · GOV INFO SECURITY · infoRisk TODAY

CU Just for Credit Unions INFO SECURITY · Data Breach TODAY · AIToday.io · OT.today · PaymentSecurity.io · FraudToday.io · DeviceSecurity.io

CIO.inc · CyberEd.io · CyberEdBoard · CXO ADVISOR · CYBER THEORY · Xtra mile LIFECYCLE MARKETING · ATHENA

QC MEDIA · SUMMIT · custom events · roundtables · NULLCON · hardwear.io · ISMG Studio

902 Carnegie Center · Princeton, NJ · 08540 · ismg.studio