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.

01. 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.

02. 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.

03. 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.

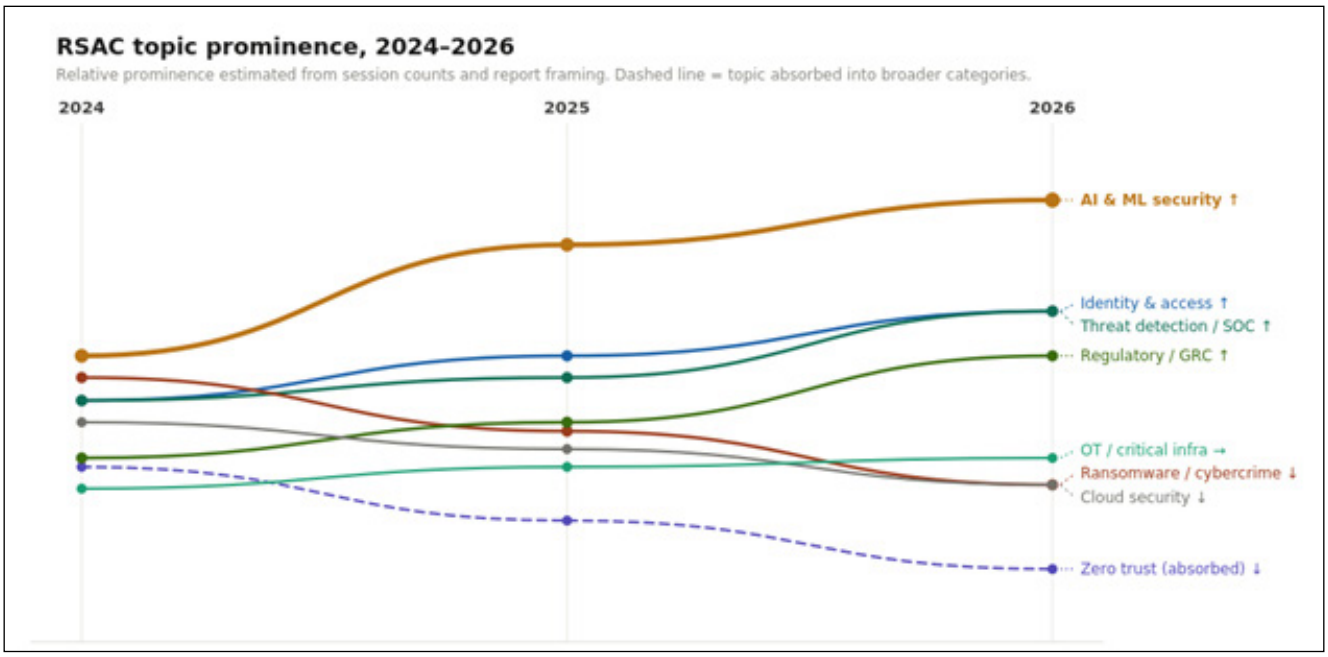

The slope chart above makes the structural argument visible. AI and machine learning security's trajectory is not a steep climb among other climbers - it separates from the field. Everything else either holds, declines modestly or gets absorbed. The lines that fall - ransomware, zero trust and cloud 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.

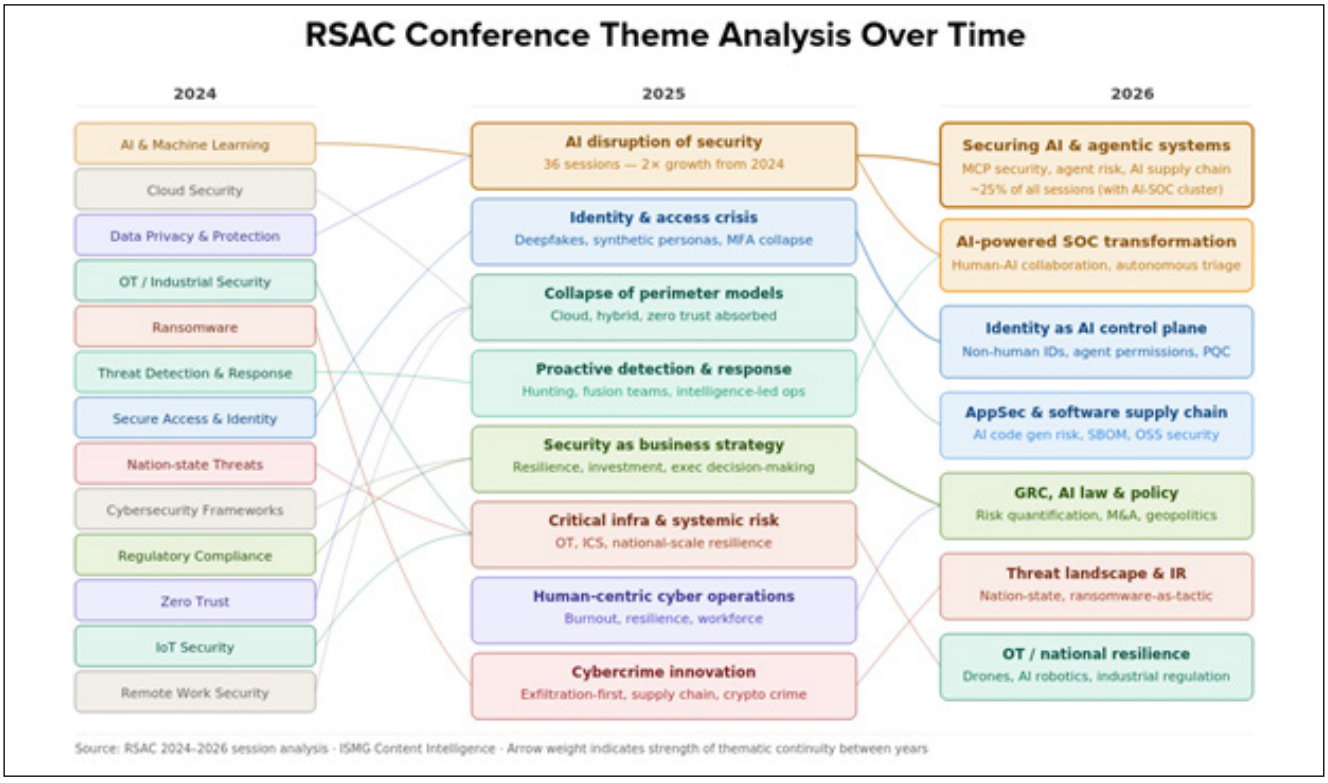

The flow diagram above shows where the 13 themes of 2024 actually went. Most did not disappear - they merged upward into larger strategic narratives in 2025 and then re-anchored around the AI and governance clusters that define 2026. The arrow weights reflect thematic continuity. The compression is not 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 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

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, 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.

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.