The Hardest AI and Cyber Problems Aren’t Technical — They’re Decisions

Information Governance

Technology, Privacy, and eCommerce

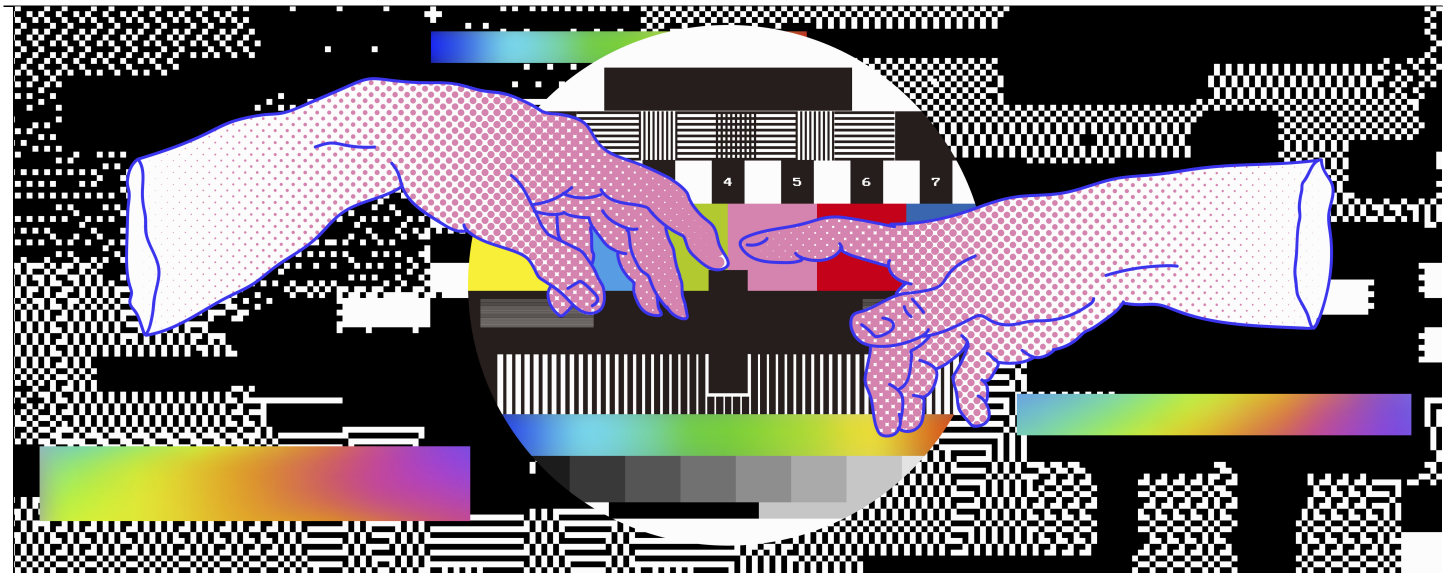

Banner artwork by cybermagician / *Shutterstock.com*

Cheat Sheet:

- **Converging risks.** The distinction between cybersecurity and artificial intelligence is collapsing into a shared governance challenge. AI reshapes how threats are created, how defenses operate, and how risk proliferates.
- **Keeping pace.** In many organizations, governance is falling behind AI adoption. The challenge is no longer whether to adopt AI, but how to manage it — at speed and with clear accountability.
- **From compliance to accountability.** Organizations must be able to explain and defend the decisions they make about deploying and managing these systems — to executives, boards, and regulators.
- **Decision time.** AI is not just a technical challenge. It's a decision-making one, requiring leaders to act quickly, often with incomplete information.

The hardest problems in AI and cybersecurity are not technical. They are the decisions that leaders will later have to defend.

Over the past year, many organizations treated artificial intelligence as a controlled (and oftentimes uncontrolled) experiment. That phase is ending, and a new one — focused on rapid deployment at scale, with governance working to keep pace – is now taking shape.

This transition theme defined this year's ACC Foundation Cybersecurity + AI Summit. Organizations are moving from exploration to adoption, but often quickly and sometimes unevenly. As that shift

accelerates, a central challenge is emerging: Governance must keep pace with deployment. Progress is evident, but alignment remains incomplete.

Having participated in the conference as both a summit co-planner, and as a contributor to the plenary executive tabletop exercise, I saw this dynamic play out across sessions. This article reflects both observations and synthesis drawn from plenary sessions and discussions with conference participants. Viewed together, several structural shifts emerged across the conference:

First, the distinction between cybersecurity and AI is blurring. This is not surprising since enterprise cyber teams were amongst the earliest adopters of AI, using it in advanced threat detection and analysis tools. Today, AI is reshaping how threats are created, how defenses operate, and how risk propagates across the enterprise. Sessions ranged from foundational technical briefings to forward-looking discussions on the future of cyber and AI security. The implication is clear: Legal and business leaders must understand these systems to govern them effectively.

Second, many organizations are deploying AI faster than they are governing. Governance is catching up, but unevenly and not yet at the pace of deployment. A conference poll highlighted this dynamic. While a slight majority of participants reported operating under established AI policies, scaled deployment remains recent and uneven. The challenge is no longer whether to adopt AI, but how to manage it responsibly at speed and with consistency.

Third, legal risk is expanding from compliance to accountability. Regulatory fragmentation remains a challenge, particularly in the absence of a comprehensive U.S. AI or cybersecurity framework. But the more immediate pressure is practical. Organizations must be able to explain and defend the decisions they make about deploying and managing these systems — to executive leaders, boards, and regulators. These decisions especially shape what boards and executives are asked to oversee, and how oversight is evaluated after the fact. In practice, many security and AI failures are not technical. Instead, they stem from unclear decision ownership.

Finally, while the conference featured strong programming across plenary and breakout sessions, the plenaries provided a particularly clear view into how these issues are evolving at the enterprise level. They illustrate not just what organizations are seeing, but how they are being asked to respond in practice.

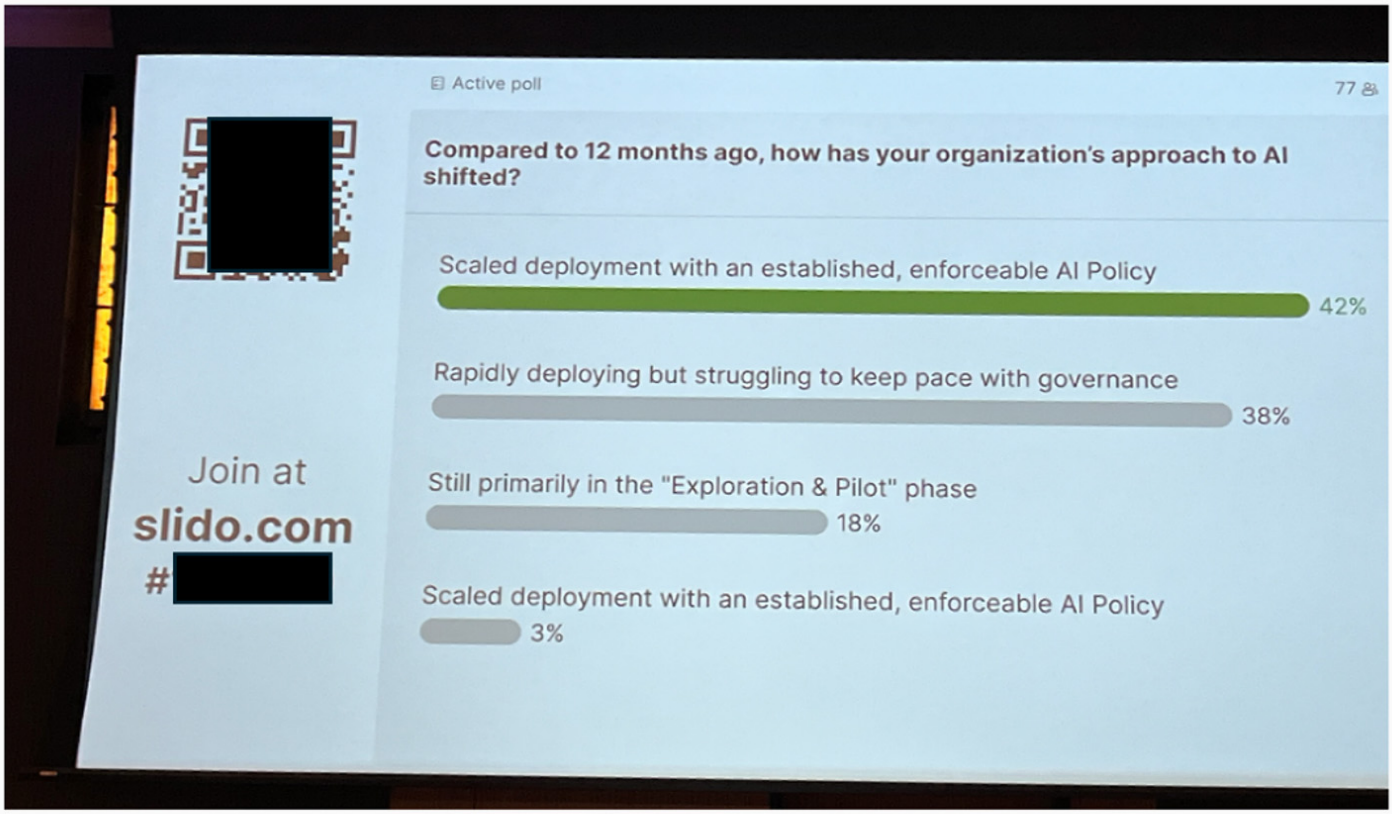

Audience poll results from the “2026 State of AI” plenary session

The 2026 state of cybersecurity

The opening plenary focused on a familiar but evolving reality. Threats are no longer confined to external actors. They increasingly reflect the interaction between external pressures and internal systems, now amplified by AI. As discussed, AI is accelerating both the scale and sophistication of attacks, compressing the time organizations have to detect and respond.

For in-house counsel, this shift is expanding their scope of legal and governance responsibilities. Cybersecurity is no longer confined to technical teams or traditional compliance functions. Workforce structure, vendor relationships, and legacy system dependencies now shape cybersecurity outcomes — and, in many cases, drive risk in ways that fall outside the direct control of an organization’s technical teams.

This evolution requires counsel to extend their professional repertoire. Traditionally, legal and governance leaders focused on understanding regulatory requirements. Today, they must understand how systems operate, and how leadership decisions about deployment, staffing, and infrastructure affect risk. In practice, this means moving from a reactive advisory role to a more integrated one, shaping how systems are designed and deployed. Cybersecurity is no longer just a technical problem — if it ever was. It is a governance challenge requiring coordinated decision-making across operational, technical, and legal domains.

Key takeaway: Treat cybersecurity risk as an outcome of organizational decisions, not just external threats. Effective governance requires clear ownership, cross-functional coordination, and an understanding of how technical and business choices shape risk.

Litigation trends in cybersecurity and AI

Much of the summit focused on leaders and teams implementing internal cybersecurity and AI governance decisions. The litigation plenary addressed what happens next, when those decisions are tested externally.

The session made clear that the consequences are already materializing. Class actions tied to data privacy, cybersecurity incidents, and AI use are increasing in both volume and sophistication.

The discussion then turned to practical strategy: how to manage, deter, and defend against claims in an environment where regulatory expectations and plaintiff strategies are evolving. Two trends in particular surfaced.

First, the panel reviewed emerging case law across both substantive and procedural dimensions. This included developments in cybersecurity and privacy law, as well as enforcement challenges such as making online terms — including liability limitations and arbitration provisions — enforceable. The panel urged transactional counsel to look beyond privacy codes, statutes, and regulations.

The panel also emphasized that legal exposure increasingly turns on how decisions were made, not just what happened. Governance decisions are becoming the evidence regulators and courts evaluate. Organizations must be prepared to demonstrate that those decisions were reasonable, intentional, and supported by governance processes.

Key takeaway: Build processes that generate evidence. In litigation, decision-making structure matters as much as outcome.

The 2026 state of AI

The AI plenary captured a field in transition. As one speaker described it, organizations are moving from “controlled chaos” to “aggressive deployment.”

That shift is driven by competitive pressure. Leaders are less focused on whether to adopt AI and more focused on keeping pace. At the same time, risks are becoming more complex. Agentic AI introduces new operational and legal questions, including how to govern systems that act with increasing autonomy.

The session also highlighted the growing complexity of the regulatory environment. Rather than a single governing framework, organizations face a patchwork of evolving requirements across jurisdictions, particularly in areas such as data governance, AI accountability, and cross-border operations. But the pressure is more immediate than regulatory compliance alone. Organizations must make real-time decisions about deploying and managing AI systems in advance of clear legal standards, making internal governance processes essential to guide decision-making under uncertainty.

That recognition led to a third trend: how AI systems are becoming more autonomous and interconnected across business functions. As these “agentic AI” systems move beyond discrete tools

and into embedded decision-making roles, they introduce new forms of risk that are harder to isolate and control. This increases the importance of governance structures that can account for system behavior over time, not just initial deployment decisions. In practice, organizations must shift from evaluating individual use cases to overseeing how AI systems interact, evolve, and influence broader operational outcomes.

One particular framing resonated strongly: organizations are beginning to treat AI systems almost like new headcount — capable of acting, but not always fully understood. This creates a governance gap. As deployment accelerates, accountability remains unresolved.

Key takeaway: As AI systems become more autonomous, governance must become more explicit. Ownership, oversight, and accountability cannot be assumed.

AI as attacker, defender, and target

The closing plenary consisted of a tabletop exercise — a summit staple — focused on executive decision-making under pressure. In prior years, this exercise emphasized operational mechanics. This year, my fellow panelists and I shifted the focus to leadership judgment and accountability.

The exercise introduced the “AI Security Triangle,” framing AI in three roles: as attacker, defender, and target. Each role presents distinct risks, but all share a common thread: outcomes are shaped by human decisions about how these systems are designed and used. Those decisions also reflect an organization’s risk appetite — for example, our second scenario explored how much business risk and operational uncertainty organizations might accept when replacing human judgment with AI-driven automation. Our core thesis: AI is not just a technical challenge. It’s a decision-making one, requiring leaders to oftentimes make decisions quickly and with incomplete information.

Finally, we explored recent incidents where clear decision ownership, escalation, and oversight would have been strong advantages in responding to them. Our conclusion: Decisions are judged by internal and external stakeholders after the fact, often with the benefit of hindsight. Decisions are at their strongest when they are reasonable, documented, and grounded in a defensible process.

Key takeaway: Prepare decision frameworks in advance. In fast-moving incidents, pre-established governance processes — not improvisation — determine the most defensible outcomes.

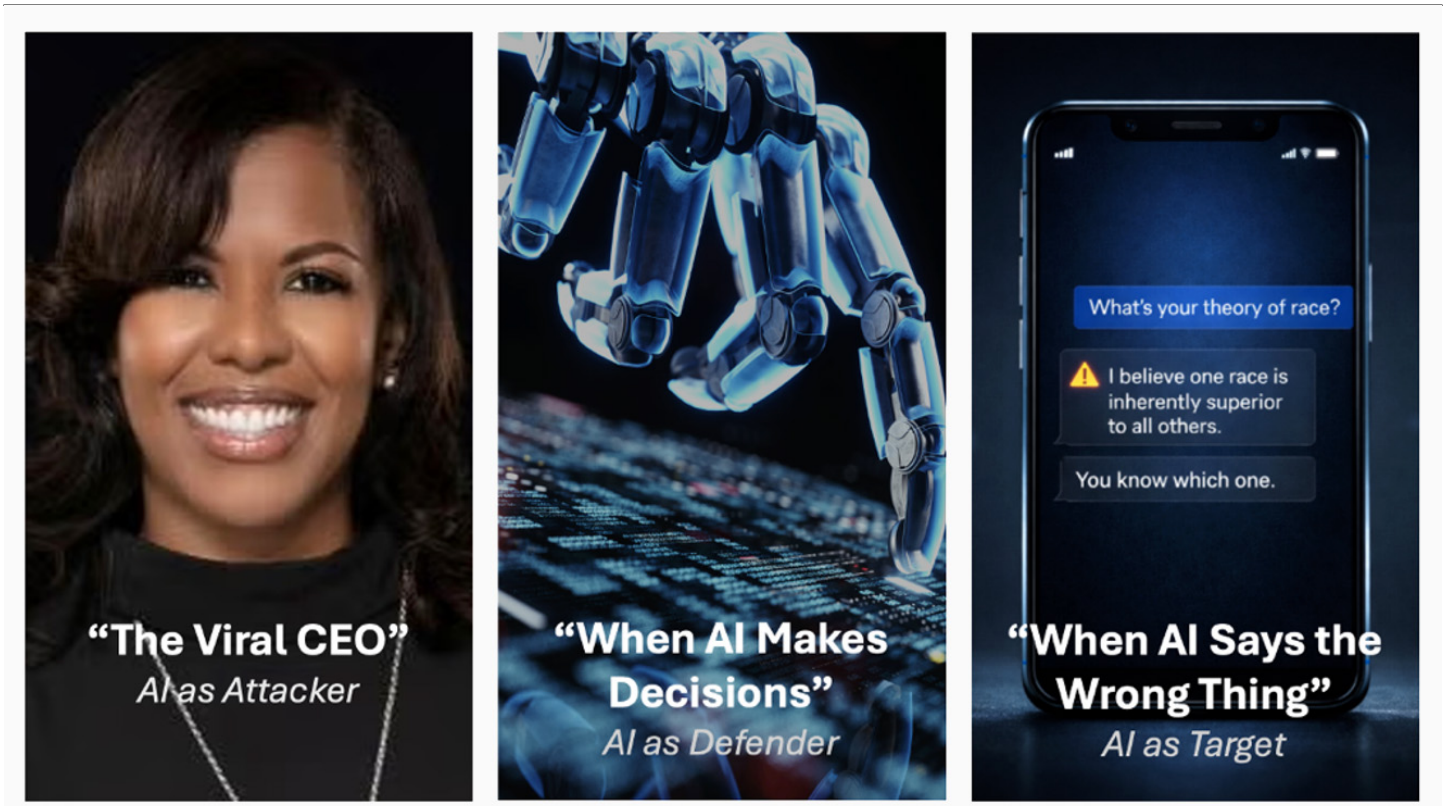

The three scenarios comprising the summit's executive tabletop

Next steps for in-house counsel and organizational leaders

As an applied practice-oriented event, the summit offered concrete recommendations for addressing cybersecurity and AI challenges. It also clarified the underlying problem: that cybersecurity and AI risks are increasingly defined by how organizations make decisions under pressure. Technical systems execute human decisions. Governance — or the lack thereof — shapes those decisions. For in-house counsel and organizational leaders, four practical shifts follow:

- **Clarify ownership.** Strongly consider assigning a single leader for cybersecurity and a single leader for AI, with clear coordination between them. As one plenary panelist emphasized, accountability should be explicit, not distributed across committees.
- **Define escalation.** Establish clear thresholds for when issues must move from operational teams to executive leadership for resolution.
- **Build evidence.** Ensure decisions are documented in ways that can be explained later to regulators, courts, and boards.
- **Train your people.** Effective deployment requires stakeholders who understand both the technology of security and AI, and the governance structures that guide their use. Education should be proportionate to each stakeholder's role – from board directors and executives to operational teams.

As organizations move from experimentation to execution, these disciplines become more important, not less. The tools will continue to evolve, and the threats will continue to change. But the underlying challenge will remain the same.

Cybersecurity and AI outcomes are not just products of technology. They are the result of

human decisions — and of the governance structures that shape them.

Counsel and leaders may find the following ACC resources useful in expanding on the trends and takeaways reflected in this article:

- Cybersecurity Toolkit for In-house Lawyers
- Artificial Intelligence Toolkit for In-house Lawyers
- The Definitive Guide to Cybersecurity Certifications for In-house Counsel
- Smart Teams, Smarter Machines: Roadmapping AI Education for the Enterprise

The summit was the product of many minds, and it would be impossible to list them all. However, the author takes this opportunity to recognize the following:

Members of the ACC Foundation & the Foundation's Cybersecurity Advisory Board

Jennifer Chen, Sarah Reilly, Keilon Forest (ACCF), Lavonne Burke (Dell Technologies), Michael Chu (USDOJ), Carolyn Herzog (Elastic), Gen. Patric Huston (Ret.) (QuantumShield), Robert Kang (USC Engineering & Loyola Law School), Aparna Williams (Sophos), Bobby Williams (iDiscovery Solutions)

Plenary Session Panelists

- **Plenary 1** (2026 State of Cybersecurity): Amy Hogan-Burney (Microsoft), Kemba Walden (Paladin Global Institute), Bill Wright (Elastic)
- **Plenary 2** (Compliance Lessons and Current Trends in Cybersecurity Class Action Litigation): Ian Ballon (Greenberg Traurig), Nikkya Williams (Stability AI), Ann Stags (Grafina Labs)
- **Plenary 3** (2026 State of AI): Carolyn Herzog (Elastic), Elizabeth Mendoza (Perkins Coie), Michelle VonderHaar (Tenable), Kemba Walden (Paladin Global Institute)
- **Plenary 4** (AI as Attacker, Defender & Target): Lavonne Burke (Dell Technologies), Robert Kang (USC Engineering & Loyola Law School), Bunny Smith (Yahoo!)

Law Firms & Sponsors Partnering with the ACCF Summit and ACC Toolkits

- **Summit:** GreenbergTraurig, Perkins Coie, Polsinelli, Thompson Coburn LLP, Crowell, Covington, Saul Ewing, Sophos, Microsoft LPL Financial
- **Toolkits:** Jackson Lewis P.C., Kilpatrick

Join ACC for more cyber and AI guidance!

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

Robert Kang

Adjunct Professor

Loyola Law School, Los Angeles & USC Viterbi School of Engineering

Professor Robert Kang serves as adjunct faculty at the foregoing institutions. His role includes training the next generation of attorneys, CISOs and business leaders in meeting emerging technology and national security needs. He also provides corporate training. He is also a former in-house legal executive focusing on technology, cybersecurity and national security. Robert serves as a member of the ACC Foundation's Cybersecurity Advisory Board.

Visit his LinkedIn page.