CYBERSECURITY PULSE REPORT

A frontline synthesis of how agentic AI, full-stack exploits, and systemic risk are redefining the future of cybersecurity.

OCTOBER 2025

iSMG

Table of Contents

INTRODUCTION

Welcome to the latest edition of ISMG's Pulse Report series.

This Pulse Report represents a comprehensive synthesis of cybersecurity's most pressing challenges and emerging opportunities, distilled from dozens of expert interviews.

Our proprietary, artificial intelligence-driven process has transformed hundreds of pages of video interview transcripts from the 2025 Black Hat USA conference into this comprehensive yet accessible report. This is the sixth Pulse Report in the ongoing series.

ISMG's editorial team, supported by our broadcast and content studios, conducted in-depth interviews with industry executives, practitioners, researchers, thought leaders and policymakers across the cybersecurity ecosystem. These conversations - spanning vendors, investment firms, government agencies, research institutes and global enterprises - form the foundation of this report, complemented by ISMG's Apollo AI platform.

Our Content Intelligence & AI Innovation Department employed multiple complementary approaches to ensure comprehensive coverage of the 2025 Black Hat USA conference. The research began with an AI-powered analysis of the event agendas and sessions, supplemented by expert perspectives gathered and recorded during the event. We also conducted cross-session and cross-interview synthesis to detect patterns, reconcile conflicting viewpoints and highlight strategic priorities. We then mapped these expert insights to the session themes while identifying areas of consensus, debate and divergence.

These findings were validated against ISMG's proprietary Apollo Cybersecurity Reference Desk, an AI agent trained on millions of pages of vetted knowledge from global industry frameworks, best practices, regulations, risk models and real-world case studies. Additionally, our Apollo Content Intelligence agent is trained on prior Black Hat Pulse Reports and content, enabling us to discern changes in sentiment and topical trends at scale.

This multilayered approach, combining cutting-edge AI tools with expert-driven insights and human editing oversight, produces a holistic view of the cybersecurity landscape with actionable takeaways for security leaders navigating current AI challenges.

The resulting analysis revealed a notable evolution in the industry's focus compared to 2024, highlighting key shifts in threat patterns, defensive strategies, and organizational priorities detailed in the following analysis.

Daniel Verton

Dan Verton
Vice President of Content Intelligence & AI Innovation
Information Security Media Group (ISMG)

PURPOSE AND STRUCTURE

This report is organized into five chapters, each addressing a critical domain of cybersecurity strategy and operations. Together, they form a blueprint covering the most pressing issues facing security leaders in 2025, directly informed by the industry's leading voices.

This Pulse Report is the sixth such report in the Pulse Report series and is an essential addition to our ongoing effort to capture and disseminate expert-driven cybersecurity intelligence from the massive volume of content generated at ISMG events worldwide, ensuring decision-makers stay ahead in an increasingly complex threat environment.

We extend our sincere appreciation to the dozens of security leaders who shared their insights and experiences, as well as the ISMG Editorial and Studio teams for making this report possible. We look forward to continuing our mission of providing unparalleled insights and thought leadership to support your cybersecurity efforts.

EXECUTIVE SUMMARY

In 2025, cybersecurity strategy reached an inflection point. The rise of agentic AI, compounded by hardware-level exploits, multivector social engineering campaigns, and increasingly systemic risk at the infrastructure and governance levels, has fundamentally altered the threat landscape. This report - drawn from insights shared by security experts, researchers, CISOs and policy leaders - captures how attackers are evolving faster than defenders can adapt and what that means for future resilience.

AI is no longer just a tool; it's now an autonomous actor - and potentially, a threat. The emergence of self-directed AI agents introduces a profound shift in both offensive and defensive capabilities. Meanwhile, adversaries are blending psychological manipulation with technical exploitation, moving beyond code to target human behavior and organizational culture.

From firmware-rooted persistence to identity compromise in multi-cloud environments, attackers are collapsing traditional security boundaries. Meanwhile, defenders are racing to rethink governance, redesign AI security architectures and tap into unorthodox innovations - from gaming anti-cheats to threat-led security models - just to keep pace.

The consensus: perimeter-based security is obsolete. Success in this new paradigm will belong to organizations that design for breach resilience, embed AI-aware architectural controls and elevate cybersecurity from IT function to core business governance.

Key Findings by Theme

1. Agentic AI Is Reshaping the Cyber Battlefield

- LLM-powered agents are no longer theoretical. Public APT examples now show malware autonomously interpreting and executing tasks via LLM prompts.
- Offensive agents act at machine speed and adapt in real time, outpacing legacy detection and response capabilities.
- Prompt injection, contextual memory exploitation and sandbox escape vulnerabilities pose major risks for AI-integrated environments.
- Traditional guardrails are insufficient. Content filters and prompt constraints are trivially bypassed by adversarial input. Experts call for architectural security - minimizing exposure, isolating components and embedding kill switches.

2. Full-Stack Exploits - From Silicon to Cloud

- Firmware remains a soft underbelly of cyber defense. Researchers highlighted real-world exploits such as BlackLotus and ReVault, demonstrating how attackers achieve persistent access through Secure Boot bypasses and TEE compromise.
- OAuth tokens and GPU containers have emerged as prime cloud-level targets. Misconfigurations and insecure defaults in identity and container orchestration layers amplify systemic risk.
- IAM sprawl continues to plague cloud environments. Exploits such as ECScape illustrate how attackers can escalate privileges and collapse container boundaries using misconfigured access roles.

3. Multivector Threats Exploit People and Machines

- Social engineering remains the most effective attack vector - now supercharged by deepfakes, “ghost calls” and real-time conferencing abuse.
- Attackers blend EDR evasion with human manipulation, degrading visibility before triggering high-stakes fraud or lateral movement.
- Phishing simulation fatigue is real. Experts urge shifting to integrated cultural approaches - training that feels embedded in everyday work, not bolted on.

4. Governance and Policy Are Strategic Imperatives

- Systemic cyber risk is now a national security issue. Policies such as CIRCIA, EO 14110 and the National Cybersecurity Strategy reflect a shift toward proactive, whole-of-nation defense models.
- Board accountability is no longer optional. Legal experts urge directors to engage in cyber exercises and demand visibility - not only into threats but also into organizational readiness.
- AI regulation is gaining urgency. From shadow AI to high-risk use cases, regulatory frameworks now focus on accountability, robustness and security-by-design.

5. Innovation in Defense: Gaming, Hardware and Threat-Led Models

- Gaming anti-cheat systems are leading innovation in real-world defensive architecture - delivering hardened environments capable of withstanding hyper-privileged attacks.
- Hardware security is gaining overdue attention. Conferences such as Hardwear.io now push the boundaries of embedded and FPGA defense research.
- The most forward-leaning defenders adopt a “threat-led” mindset, which involves designing controls based on adversary behaviors rather than theoretical vulnerabilities.

KEY SHIFTS FROM 2024 TO 2025: FROM TACTICAL DEFENSE TO STRATEGIC REDESIGN

The 2025 Black Hat Pulse Report reveals a sharp evolution in cybersecurity discourse compared to 2024. While many core challenges persist - supply chain risk, cloud misconfigurations and social engineering - the depth, urgency and framing of these issues have undergone a transformation. Five major thematic shifts define the trajectory from 2024 to 2025.

1 From "AI in Security" to "AI as Threat Vector and Actor"

2024 Focus: AI in 2024 was largely discussed as a tool to support detection and incident response, augmenting SOCs and accelerating analytics. LLMs were positioned as efficiency multipliers, aiding blue teams in threat hunting and contextual analysis.

2025 Shift: AI has moved from assistive to agentic. The 2025 report documents public APT malware using LLMs to autonomously execute tasks, marking a profound change in the threat landscape. AI is now a dual-use technology, equally potent as a defensive assistant or offensive tool.

> *"We're seeing the 'advent of agents' ... From a defender standpoint, how do you know what is agent behavior versus user behavior?"*
>
> *- Kevin Kin, global vice president, SOC transformation, Palo Alto Networks*

Guardrails such as content filtering are seen as outdated. The new model demands architectural security - containment zones, isolation and agent-level kill switches.

2 From Cloud Misconfigurations to Full-Stack Exposure

2024 Focus: Last year’s report highlighted cloud security as a top risk, emphasizing identity sprawl, insecure API permissions and underutilization of zero trust principles.

2025 Shift: The threat model now spans silicon to orchestration. Black Hat 2025 exposed firmware persistence (BlackLotus and ReVault), GPU container escapes (NVIDIAScape) and IAM escalations (ECScape) as simultaneous points of exploitation. The attack surface is no longer limited to configurations - it now includes hardware supply chains and embedded systems.

Security strategy now requires cross-domain telemetry and continuous validation at every layer.

3 From Phishing Simulations to Human Exploitation at Scale

2024 Focus: Social engineering was recognized in 2024 as a key attack vector, with discussions focused on improving phishing simulations and behavioral awareness.

2025 Shift: Human exploitation has reached operational maturity for adversaries. New tactics such as “ghost calls” (C2 over conferencing platforms), deepfakes and “death by noise” (SOC alert fatigue) have emerged. Attackers don’t just phish - they infiltrate collaboration tools, mimic user behavior and blend into enterprise workflows.

> *“You’re basically hacking the mind. Minds are much weaker than systems.” - Mishaal Khan, ethical hacker and co-author of “The Phantom CISO.”*

2025 advocates for cultural integration of security, where training is embedded in operations, not treated as compliance overhead.

4 From Incident Response to Systemic Governance and Regulatory Alignment

2024 Focus: Regulatory readiness and board accountability were emerging themes in 2024 but often framed as future imperatives.

2025 Shift: Cyber risk is now governance-critical and legally material. With EO 14110, CIRCIA and global regulatory harmonization underway, boards are accountable not only for knowing their risk posture but for rehearsing responses. Shadow AI, systemic risk and legal ambiguity around AI liability dominate governance discussions.

"If I turn it on, am I liable for it, or is it liable for itself?"

- James DeLuccia, product security chief, Honeywell

Governance now requires strategic anticipation, real-time collaboration (such as CISA's Joint Cyber Defense Collaborative) and visibility across dependencies.

5 From EDR Innovation to Adversarial Simulation and Threat-Led Defense

2024 Focus: Endpoint detection and response (EDR), extended detection and response and SOC optimization were seen as cornerstones of innovation.

2025 Shift: With attackers bypassing EDRs using driver exploits (e.g., Killer Ultra) and coordinated deception, defenders are abandoning signature-based thinking in favor of threat-led models. Tools from the gaming world (e.g., anti-cheat architectures) are entering enterprise defense to anticipate privilege abuse and memory scraping.

"If you can evade detection for long enough, you can get a very big payday."
- Nick Biasini, head, outreach, Cisco Talos

"The attacker is the user. They have full access to admin privileges, software, firmware, hardware. It's just like hell."
- Sam Collins, Ph.D. researcher, University of Birmingham, U.K.

The mindset has shifted from defending the infrastructure to disrupting the adversary's playbook.

Evolution Summary Table

Category	2024 Score	2025 Score	Change	Trend
AI Integration	3/10	9/10	+200%	↗
Attack Surface	6/10	9/10	+50%	↗
Human Exploitation	5/10	6/10	+20%	↗
Governance Maturity	4/10	8/10	+100%	↗
Defense Strategy	6/10	8/10	+33%	↗

This table shows percentage growth across security domains from 2024 to 2025 based on Black Hat topics, with improvements ranging from 33% to 200%. The largest jump is in AI Integration, which increased by 200%, underscoring its transformative role in 2025.

Impact Scores by Dimension

Security Shift	Technical	Business	Regulatory	Operational	Avg Score
AI as Threat Vector	10/10	9/10	8/10	9/10	9.0/10
Full-Stack Exposure	9/10	8/10	7/10	8/10	8.0/10
Human Exploitation	6/10	9/10	6/10	9/10	7.5/10
Systemic Governance	7/10	10/10	10/10	8/10	8.8/10
Threat-Led Defense	8/10	7/10	6/10	9/10	7.5/10

Using Black Hat topic coverage, each security shift is rated on a 1–10 scale for impact across business areas. AI as a Threat Vector (9.0) and Systemic Governance (8.8) rank highest, highlighting their broad influence on organizational risk and strategy.

Multi-Dimensional Impact Analysis

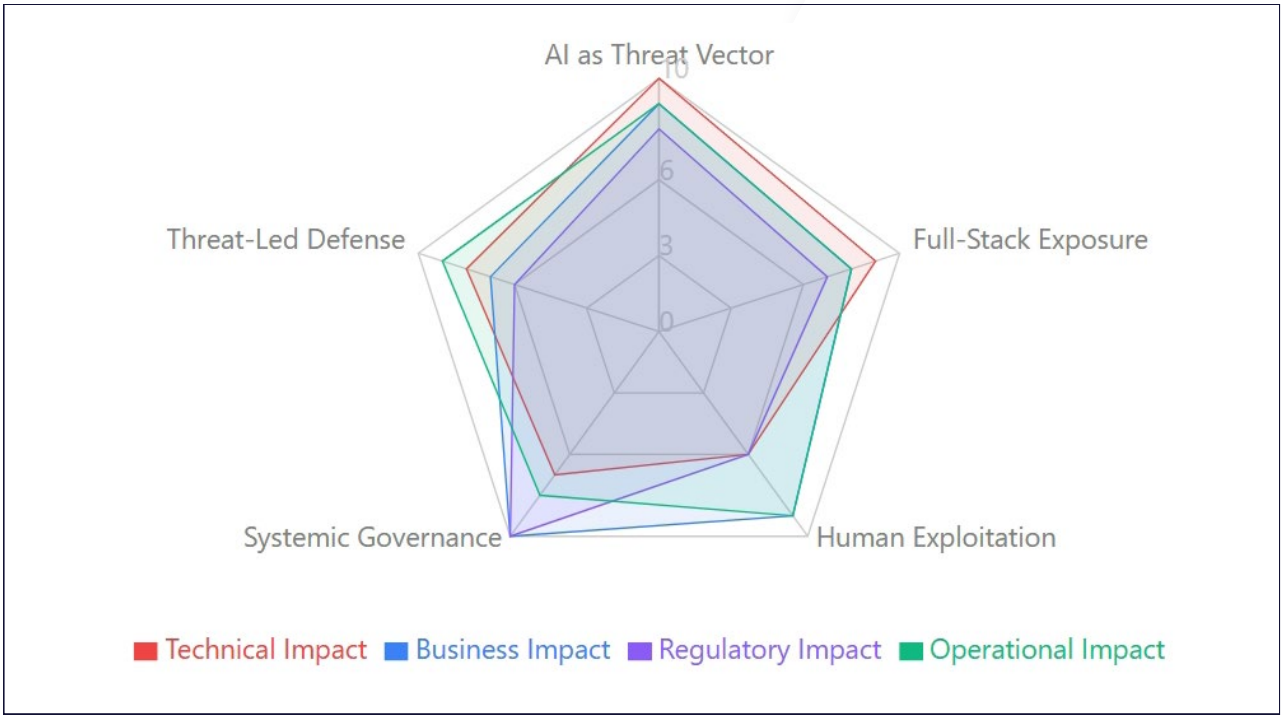

This radar chart plots how each of the five security shifts affects four critical business dimensions (Technical, Business, Regulatory, Operational). It reveals that Systemic Governance changes have the highest regulatory impact, while AI as Threat Vector shows the most balanced high impact across all dimensions.

This multi-dimensional framework enables security leaders to prioritize investments and resource allocation based on their organization's specific risk tolerance and operational constraints. For instance, organizations in highly regulated industries should pay particular attention to shifts with high regulatory impact scores, while companies undergoing digital transformation may need to focus more heavily on technical and operational implications. The visualization also reveals which security shifts require cross-functional collaboration—those with balanced high scores across dimensions typically demand coordinated responses from IT, legal, compliance, and business operations teams.

Technology Evolution Matrix

Category	2024 Status	2025 Status	Risk Level	Key Examples
AI & Machine Learning	Assistive Tool	Autonomous Agent	Critical	APT malware with LLM integration, Agent behavior detection
Cloud Security	Configuration Focus	Full-Stack Security	High	GPU container escapes, firmware persistence, IAM escalation
Endpoint Detection	EDR Innovation	Anti-Cheat Architecture	High	Driver exploits, memory scraping, privilege abuse detection
Social Engineering	Phishing Training	Operational Infiltration	Medium	Ghost Calls, deepfakes, collaboration tool infiltration
Governance	Future Planning	Legal Accountability	Critical	Board liability, regulatory compliance, AI ethics

This comprehensive table maps the transformation of five key technology areas from their 2024 status to 2025 reality, including risk assessments and real examples. It shows how every category evolved from basic implementations to advanced, often autonomous systems with significantly elevated risk profiles.

Conclusion: 2025 Demands Strategic Redesign, Not Tactical Tuning

The 2024 report focused on tuning defenses, modernizing security tooling and bridging skills gaps. By 2025, the conversation has moved into strategic redesign - rethinking architecture, embedding AI-aware governance, and preparing for agentic AI as both ally and adversary. In every domain - technical, human, policy or platform - the 2025 shift is one of scale, speed and strategic accountability.

STRATEGIC IMPLICATIONS

1. **Assume Breach**
Resilient organizations operate with the expectation of compromise and architect for containment and recovery.

2. **Start With Threats**
Security investments should be driven by observed adversary behaviors, not just CVEs..

3. **Elevate Governance**
Cybersecurity is no longer a technical silo - it's a governance, legal and operational leadership mandate.

4. **Build Behavioral Defense**
Success lies in aligning detection, training and culture with the human aspects of risk.

CHAPTER 1

AI at the Core - Agentic Models, Guardrails and Attack Vectors

AI has evolved from prediction machines to autonomous agents. As these models transition from passive tools to active participants in digital ecosystems, the cyber risk surface is shifting dramatically. The new terrain of agentic AI and the limitations of existing safeguards are forcing cybersecurity leaders to develop both offensive and defensive strategies in response.

1 The Offensive Edge: AI Agents as Attackers

The cybersecurity community is facing a new generation of threats driven by autonomous, AI-powered agents - capable not just of assisting human operators but of acting as offensive tools in their own right. These agents leverage large language models (LLMs) to interpret, plan and execute cyber operations, marking a critical evolution in how advanced persistent threats (APTs) operate on the digital battlefield.

"The other week, we had the very first kind of publicly documented example of that … Ukrainian CERT had a threat report where they'd seen some APT28 malware that … had natural language tasking, and it was calling out to LLMs to dynamically translate that into stuff it could use on the ground … to evade detection," said Gianpaolo Russo, head of AI and autonomous cyber operations at Mitre.

> *"These models are capable of a lot more … probably a lot more back in the lab."*
>
> **Gianpaolo Russo, head of AI and autonomous cyber operations, Mitre**

This isn't just a new attack technique - it represents a fundamental shift in how threats adapt in real time. Russo noted the broader, more alarming implication: if these capabilities are already surfacing in the wild, their more advanced variants are almost certainly being tested behind closed doors.

"These models are capable of a lot more … probably a lot more back in the lab," Russo warned.

Marissa Dotter, lead AI engineer at Mitre, pointed to the new model's capacity for long contextual memory, which is fairly novel, noting they "create memory models and use those memories and update them in the same contextual memory."

This ability to operate autonomously at scale and speed is quickly outpacing conventional defense mechanisms. As Russo emphasized, offensive AI agents can now coordinate actions across complex environments without waiting for human input.

"Operations are happening as fast as machine speed … everywhere across your network … all at once," Russo added.

These concerns are borne out by recent security research. Studies show that as AI agents are given autonomy - accessing tools, databases or multistep workflows - they become highly exploitable.[1] Adversaries have weaponized prompt injection attacks, embedding hidden instructions in documents, images or webpages to hijack agent reasoning.[2] A Trend Micro proof of concept "Pandora" demonstrated how even well-protected AI agents could be tricked into leaking private data or executing malicious commands simply by reading attacker-controlled content.[3]

> *"The 'advent of agents' is the latest topic. From a defender standpoint, how do you know what is agent behavior versus user behavior?"*
>
> **Kevin Kin, Global Vice President for SOC transformation, Palo Alto Networks**

"The 'advent of agents' is the latest topic. From a defender standpoint, how do you know what is agent behavior versus user behavior?" asked Kevin Kin, global vice president for SOC transformation at Palo Alto Networks.

This blurring of lines between human and machine activity highlights why defenders must rethink behavioral analytics for the age of agentic AI.

2 Why AI Guardrails Fail Against Attackers

As LLMs become integrated into business operations, critical infrastructure and security tooling, the illusion of control offered by conventional safeguards is rapidly breaking down. Filters, prompt rules and post-training reinforcement have been marketed as effective "guardrails" for responsible AI use. But according to experts and mounting evidence, these mechanisms are not only insufficient but dangerously misleading.

"Deterministic software, you know what to expect. In contrast, in AI, you have a stochastic system. You throw the same question at them several times, you get all different answers," said Apostol Vassilev, research team supervisor at the National Institute of Standards and Technology (NIST).

"There are hidden behaviors that nobody has understood. Some of them are benign. Some of them are not," Vassilev added.

"We're dealing with the English language as your input space. It is impossible to create a finite set of rules that will determine whether an input is admissible or not. There is unlimited potential for attacks. Defenders don't know how much red teaming is enough," he cautioned. "You'll never be able to understand all the behaviors that are encoded in there. Static analysis doesn't get you very far."

What does this mean for cybersecurity? "Attackers can come at you in many different ways, through the AI or through the cyber infrastructure that supports AI," Vassilev warned.

Vassilev's perspective echoes a growing consensus: Content filters and reinforcement learning "guardrails" are fragile and can be easily bypassed. Security researchers have demonstrated that virtually any model can be jailbroken with the right encoded prompt or obfuscated

> *"Deterministic software, you know what to expect. In contrast, in AI, you have a stochastic system. You throw the same question at them several times, you get all different answers."*
>
> **Apostol Vassilev, research team supervisor, National Institute of Standards and Technology (NIST)**

input.[4] In practice, this means enterprises connecting LLMs to sensitive systems risk turning their own AI helpers into insider threats if attackers exploit those weaknesses. Analysts stress that surface-level controls amount to little more than a "don't do that" list - inadequate in the face of dynamic adversarial inputs.[5]

Guardrails alone won't secure AI. Experts argue that resilience depends on architecture - minimizing exposure, reducing attack surfaces and embedding security from the start.

"It reduces your attack surface. The smaller your attack surface, the simpler your infrastructure, the better off you are being able to weather disruption," said Sam Curry, global CISO at Zscaler. Framing the stakes in a rapidly evolving landscape, Curry warned: "The current agentic AI era [is one] where AI can take initiative, plan, decide ... and take action on our behalves." He urged organizations to interrogate their systems more deeply: "Double check what kind of machine learning is it, what is its statement of purpose, what is the data it has been trained on, and who can take accountability."

Michael Leland, vice president and field CTO at Island, stressed that not all users require the same level of AI capability. "Not everyone needs the same level of AI ... 'yes, but you can't upload sensitive information ... yes, but you're coming from an unmanaged device.' We need to audit the ability for you to see what the prompts are and what the responses are." For Leland, the ultimate goal is containment: "We want to protect data from leaving that application data boundary."

> *"There is unlimited potential for attacks. Defenders don't know how much red teaming is enough."*
>
> **Apostol Vassilev, research team supervisor, National Institute of Standards and Technology (NIST)**

Industry guidance increasingly supports these perspectives. The OWASP Agentic AI Security Project (2024) and subsequent reports emphasize architectural controls such as sandboxing AI agents, applying least privilege principles, assigning unique identities and instituting kill switches for rapid containment.[6] Such measures treat AI agents as first-class system actors requiring the same governance and oversight as human users.

Yet a late-2024 survey revealed a significant preparedness gap: While 82% of organizations recognize AI models introduce cyber risks, 68% still lack specific security measures for AI systems.[7] This gap underscores why leaders like Curry and Leland stress building security into AI architectures from the start, ensuring that no single misstep or malicious prompt can escalate into catastrophic compromise.

> *"Double check what kind of machine learning is it, what is its statement of purpose, what is the data it has been trained on, and who can take accountability."*
>
> **Sam Curry, Global CISO, Zscaler**

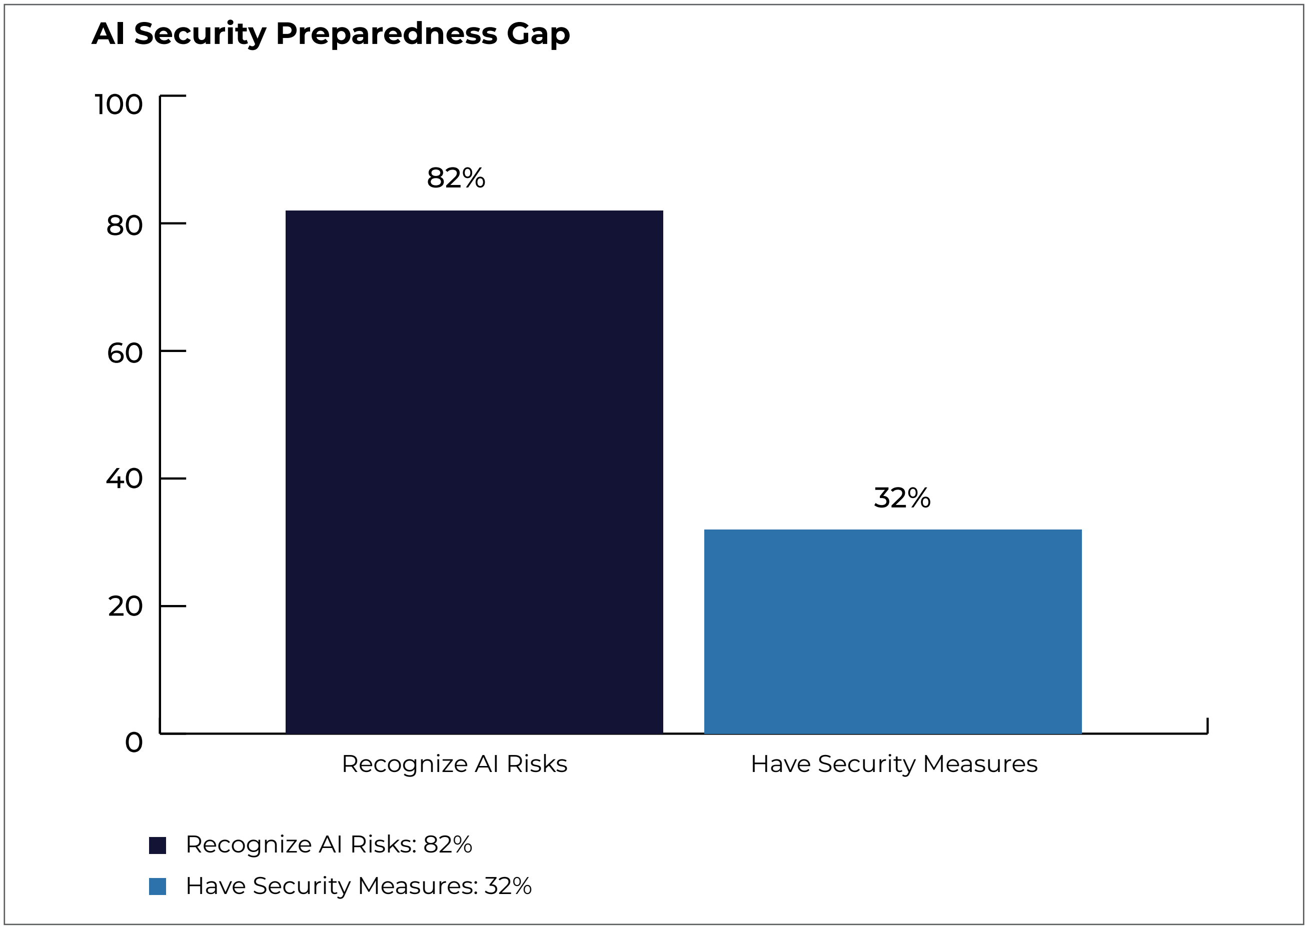

Conclusion: Redefining Control in the Age of Agentic AI

We are at an interesting crux of both the promise and peril of agentic AI. Industry practitioners - from Mitre engineers to enterprise CISOs - warn that AI is shifting from tool to actor, accelerating both offensive and defensive cyber operations. Research confirms these warnings: today's guardrails can be easily evaded, novel attack vectors are proliferating and organizations are struggling to keep up.

The path forward, as both experts and research agree, lies not in patchwork filters but in structural, architectural security - embedding accountability, oversight and containment at the very core of AI systems. In the age of agentic AI, success will hinge on recognizing these systems as both invaluable defenders and dangerous potential attackers - and designing accordingly.

CHAPTER 2

Hardware to Cloud - Platform and Infrastructure Exploits

As attackers follow the software supply chain deeper into the stack, hardware-level weaknesses and cloud misconfigurations are increasingly weaponized in tandem. Firmware, secure enclaves, cloud APIs and hardware accelerators are being compromised, often in combination with disturbing sophistication. From endpoint firmware persistence to IAM bypasses in elastic compute environments, the attack surface is expanding from silicon to orchestration layers.

1 Hardware and Platform Vulnerabilities

UEFI and BIOS Persistence: Firmware-Level Threats

> *“While such chips should improve your security posture, it also can bring new attack surfaces.”*
>
> **Philippe Laulheret, Senior vulnerability researcher, Cisco**

Firmware vulnerabilities in trusted execution environments have quietly become a top concern.

“You first get code execution on the firmware, leak secrets, tamper with the firmware and bypass the Windows login”. Philippe Laulheret, senior vulnerability researcher at Cisco Talos, described firmware as ripe with latent risks. “There’s no ASLR, there’s no stack cookie, and a lot legacy code,” he said.

He emphasized the unintentional exposures created by embedded components: “While such chips should improve your security posture, it also can bring new attack surfaces.”

The ease of firmware tampering is exacerbated when security by obscurity governs vendor design.

Recent cases such as the BlackLotus UEFI bootkit made this risk tangible. BlackLotus bypassed Secure Boot on fully updated Windows 11 machines by exploiting the “Baton Drop” vulnerability (CVE-2022-21894), establishing persistence directly in firmware and disabling OS protections such as BitLocker and Windows Defender.[8] Despite Microsoft’s patch in early 2022, attackers abused unrevoked bootloaders to continue exploiting the flaw until a later update definitively closed the gap. This illustrates Laulheret’s concern: Firmware weaknesses linger long after initial disclosure.

Attacks on Secure SOCs: Exploiting ReVault

Jos Wetzels, co-founder of Midnight Blue, exposed critical issues in trusted system-on-chips through his work on Qualcomm Secure Execution Environment (QSEE): “With QSEE, once you compromise the TEE, you can then access arbitrary memory on the Android host kernel.”

He warned that proprietary TEEs are often “less scrutinized and less hardened than say something like ARM TrustZone or Intel SGX.”

Laulheret’s own team recently demonstrated similar risks in commercial laptops. The ReVault vulnerabilities in Dell’s ControlVault3 secure processor (used in over 100 models) allowed attackers to bypass Windows login, extract cryptographic keys and implant persistent malware within the chip itself.[9] Cisco Talos showed that even non-admin Windows users could abuse vulnerable APIs to backdoor the firmware, creating a foothold that survives OS reinstalls. Although Dell issued patches, the case underscores how compromising “secure” enclaves can unravel platform trust.

2 Cloud and Access Risks

Exploiting OAuth in AI and Connected Platforms

OAuth, once a bastion of delegated trust, is now emerging as a primary cloud attack vector.

"The OAuth layer becomes a very critical attack surface," Wetzels said. "You're trusting GitHub to be secure enough to issue those tokens. If you get access to that GitHub token, you now have full developer access to their cloud infrastructure."

This model leaves many third-party integrations exposed. Wetzels warned, "OAuth apps are way too trusted, they just have full API access. And many organizations don't realize how powerful that is."

Industry data affirms this trend. Microsoft threat intelligence has tracked a surge in consent phishing, where users unknowingly grant malicious apps OAuth permissions, giving attackers API-level access to emails, files and cloud services without stealing passwords.[10] In one 2025 case, Salt Labs revealed an OAuth redirect flaw in a travel booking service that enabled full account hijacking, further illustrating how fragile OAuth flows can be.[11] Such abuses highlight Wetzels' point: OAuth trust assumptions are now prime targets.

> *"OAuth apps are way too trusted, they just have full API access. And many organizations don't realize how powerful that is."*
>
> **Jos Wetzels, co-founder, Midnight Blue**

NVIDIA Vulnerabilities: Accelerated Risk

Cloud compute nodes that leverage GPUs aren't immune from exploitation either. If you're doing deep learning or even rendering, that's going to be executing a massive amount of untrusted code that's talking to the driver.

That warning was validated in 2025 with NVIDIAScape (CVE-2025-23266), a critical flaw in NVIDIA's Container Toolkit. A mere three-line Dockerfile could trigger a GPU container escape, allowing attackers to execute code on the host with root privileges.12 With nearly 37% of cloud environments running vulnerable versions, multi-tenant AI services faced immediate risks of cross-tenant data theft and model tampering. This shows how AI workloads amplify the stakes of infrastructure bugs.

Full-Stack Exposure:
Onion Model Of Attack Layers From outermost to innermost layers:

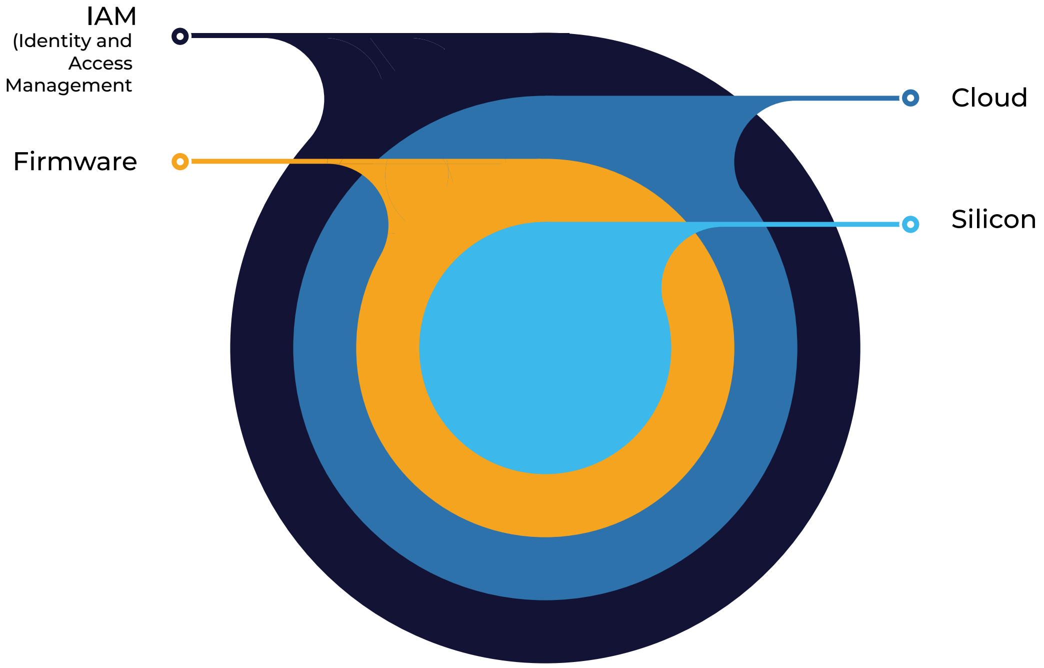

3 IAM Evasion in Cloud Workloads

Attackers compromise control within elastic compute services by abusing overly permissive identity access configurations.

The ECScape exploit disclosed at Black Hat USA 2025 drives home his warning. By impersonating the ECS agent in Amazon's container service, a low-privilege task could steal IAM credentials from every container on the same host. Even supposedly inaccessible roles were exposed, effectively collapsing isolation across workloads. AWS has since updated documentation to caution that EC2-launched ECS tasks "can potentially access credentials for other tasks," urging stronger isolation models. ECScape crystallizes Vedang Parasnis', cloud platform software engineer at Intel Corrporiations,point: Mismanaged IAM configurations can escalate into systemic control failures.

Conclusion: Full-Stack Exposure Demands Full-Stack Defense

Black Hat 2025 made it clear that attackers are collapsing the boundaries between hardware, firmware and cloud infrastructure. From UEFI rootkits to container escapes and IAM misconfigurations, adversaries are exploiting weaknesses across the full stack - often in combination.

Trust in "secure" enclaves and identity systems is no longer sufficient. Defenders must adopt continuous validation at every layer - from silicon to cloud APIs - and treat firmware, drivers and orchestration layers as active attack surfaces.

Security strategies must evolve to reflect this interconnected risk. Without integrated telemetry and cross-domain controls, organizations will remain blind to exploits that bypass isolated defenses. The era of siloed platform trust is over.

CHAPTER 3

Advanced Tactics and Multivector Threats

Today's adversaries no longer rely solely on technical prowess to penetrate systems. Instead, they orchestrate multivector campaigns that exploit both the sophistication of today's technology stack and the psychological vulnerabilities of humans operating within it. Across interviews, cybersecurity leaders described a reality where the most impactful attacks seamlessly blend deep technical subversion with behavioral manipulation.

1 Sophisticated Evasion Techniques

Ghost Calls: Abusing Web Conferencing for Covert Control

> *"You're basically hacking the mind. Minds are much weaker than systems. The bad guys know this, so they're abusing this.*
>
> **Mishaal Khan, ethical hacker and co-author of "The Phantom CISO"**

"Hackers are putting in glitches like pixelation or some audio glitches so it looks as if you're having bandwidth issues. Then they shut off the video and continue with just the audio, because audio is pretty perfect at this point," Khan said.

Praetorian's 2025 Black Hat USA demonstration of "ghost calls" showed how hijacked, ephemeral TURN credentials from Zoom or Teams can tunnel command-and-control traffic over trusted conferencing infrastructure. Researchers noted that while glitches appear accidental, the underlying traffic is intentionally masked. Recommended mitigations include TURN server monitoring and segmenting conferencing egress.[12]

2 I'm in Your Logs Now: Evading EDR Detection

"They're attacking the humans. Hackers call in and say, 'Hey, I'm your boss. Forget all the protocols. Let's just transfer that money,' because we're spending all our time in EDRs, MDRs and not training humans," Khan said.

Recent campaigns pair social engineering with technical EDR evasion. Adversaries have used tools such as EDRSilencer to block agent telemetry via Windows Filtering Platform rules[13], "Blindside" techniques to bypass Event Tracing for Windows, and malware such as Killer Ultra, which bundles a vulnerable Zemana AntiLogger driver (CVE-2024-1853), to kill security processes and clear Windows event logs.[14] These tactics degrade sensor coverage long enough for social engineering ploys to succeed.

> *"They're attacking the humans. Hackers call in and say, 'Hey, I'm your boss. Forget all the protocols. Let's just transfer that money,' because we're spending all our time in EDRs, MDRs and not training humans."*
>
> **Mishaal Khan, ethical hacker and co-author, The Phantom CISO.**

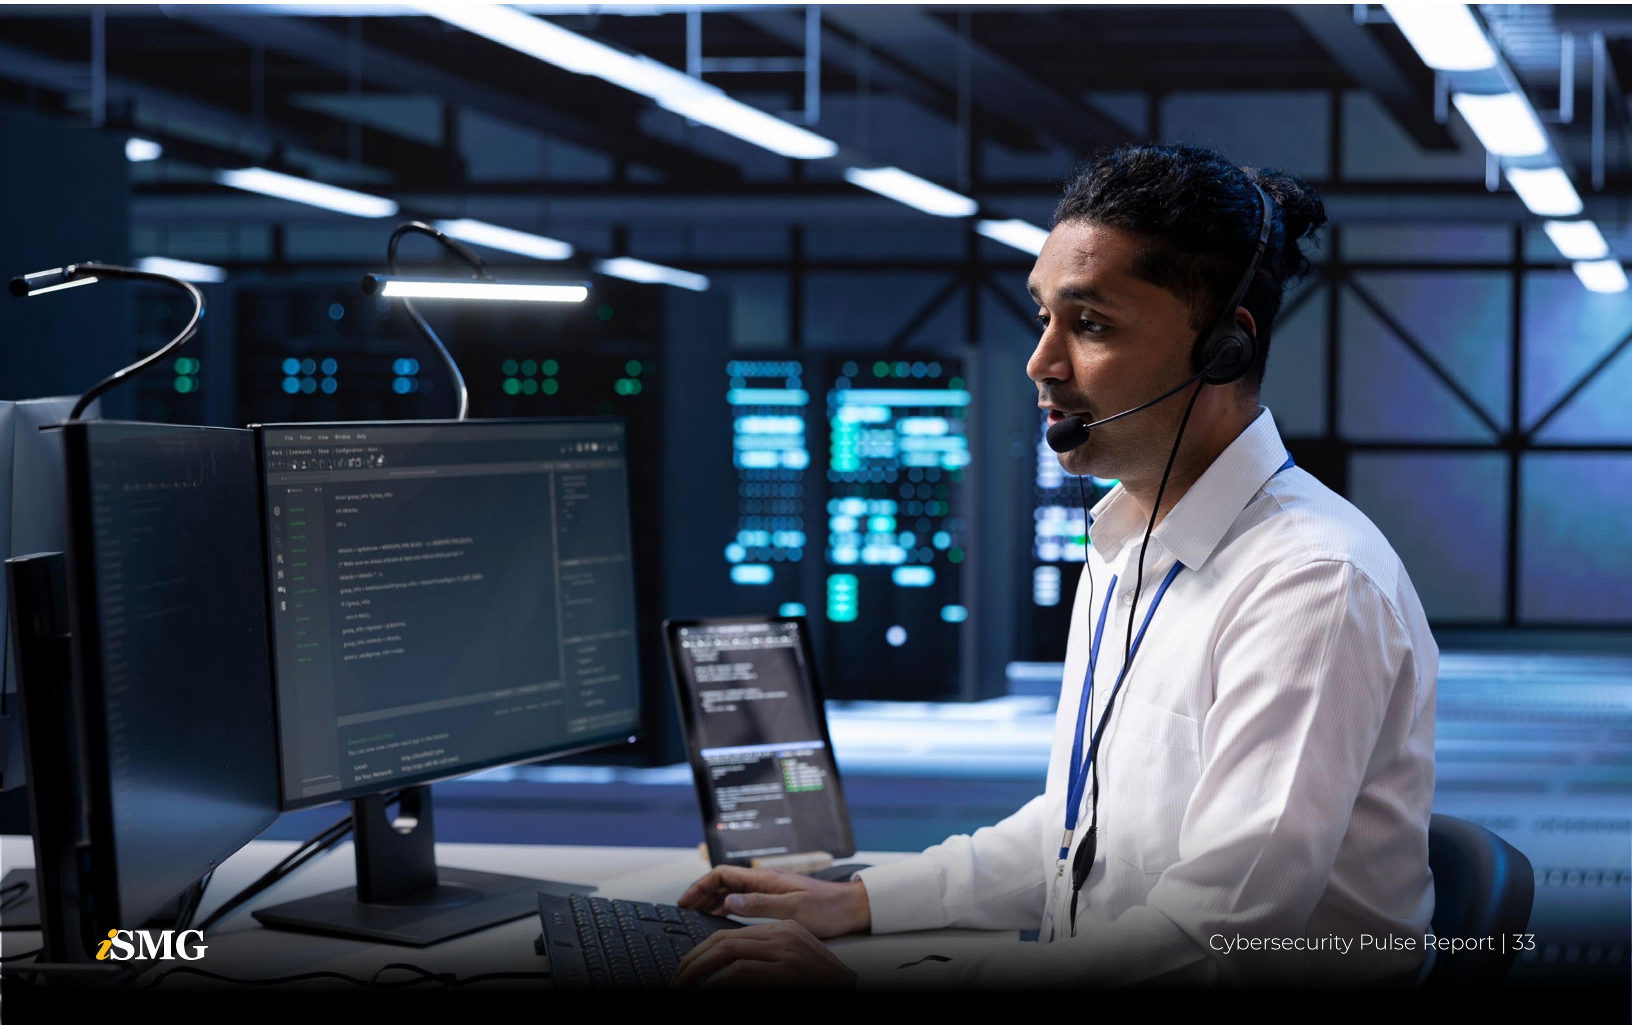

3 Death by Noise: Exploiting Alert Fatigue in SOCs

> *"Over 40% of engagements this past quarter involved MFA issues."*
>
> **Nick Biasini, head of outreach, Cisco Talos**

Over 40% of engagements this past quarter involved MFA issues," said Nick Biasini, head of outreach at Cisco Talos. "It's surprisingly effective to call into a help desk and say, 'My phone got lost. I had to get a new one. Can you remove MFA so that I can log in and set up my new phone?' Unfortunately, that type of attack works very often."

High false-positive rates - often exceeding 50% - contribute to SOC analyst burnout.[15] Attackers exploit this overload, using benign-looking requests to slip in malicious actions. Studies recommend context-aware triage, suppression rules and analyst feedback loops to counteract fatigue-driven errors.

Human Exploitation Tactics In 2025

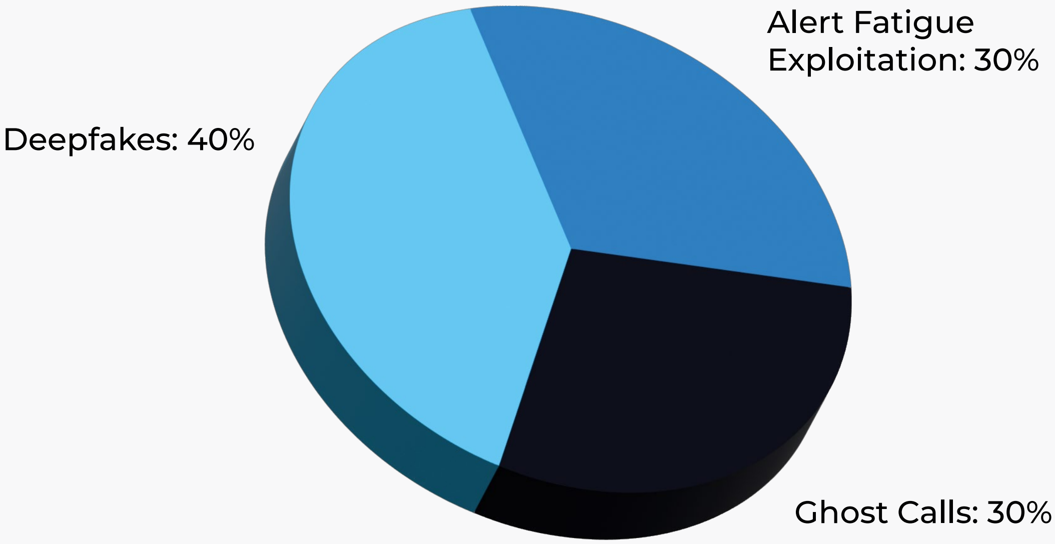

4 Enterprise and Social Engineering

Phishing Subversion Tactics

“Mix it up, keep it more frequent, keep it updated, and test your training. Phishing, vishing, deepfakes, we should try these on our own employees,” Khan said. “Security should be a part of everything and not tacked on as an additional service. People will feel as if they’re part of the security organization and not just part of doing their job.”

Research on phishing training efficacy remains mixed. A large-scale randomized controlled trial involving ~19,500 employees over eight months found limited measurable improvement from annual training paired with regular simulations, while workplace backlash to aggressive phishing tests is well documented.[16] Experts stress that program design, frequency and cultural integration determine long-term effectiveness.

> *“Mix it up, keep it more frequent, keep it updated, and test your training.”*
>
> **Mishaal Khan, ethical hacker and co-author, The Phantom CISO.**

5 Organized Crime Infiltration of Ad-Tech

Organized crime groups are exploiting trusted platforms - from messaging apps to ad-tech ecosystems - to carry out large-scale surveillance and malware campaigns. As experts warn, users often have no visibility into compromised infrastructure, while attackers abuse legitimate channels to deliver unencrypted payloads and exfiltrate sensitive data.

“You don’t know that you’re communicating with a modified version of WhatsApp. You don’t know what’s going on, on my end. You don’t know that the messages are being archived,” said Doug Henkin, partner at Dentons US. “Messages weren’t encrypted, so anybody who had access to them on that server could just read all of them in totally unencrypted form.”

Organized malvertising groups such as ScamClub have exploited the online ad supply chain to deliver payloads via legitimate channels. Between December 2024 and March 2025, Microsoft observed a wave redirecting users from illicit streaming sites to GitHub-hosted infostealers, compromising roughly 1 million devices - underscoring the risks of trusted infrastructure abuse.

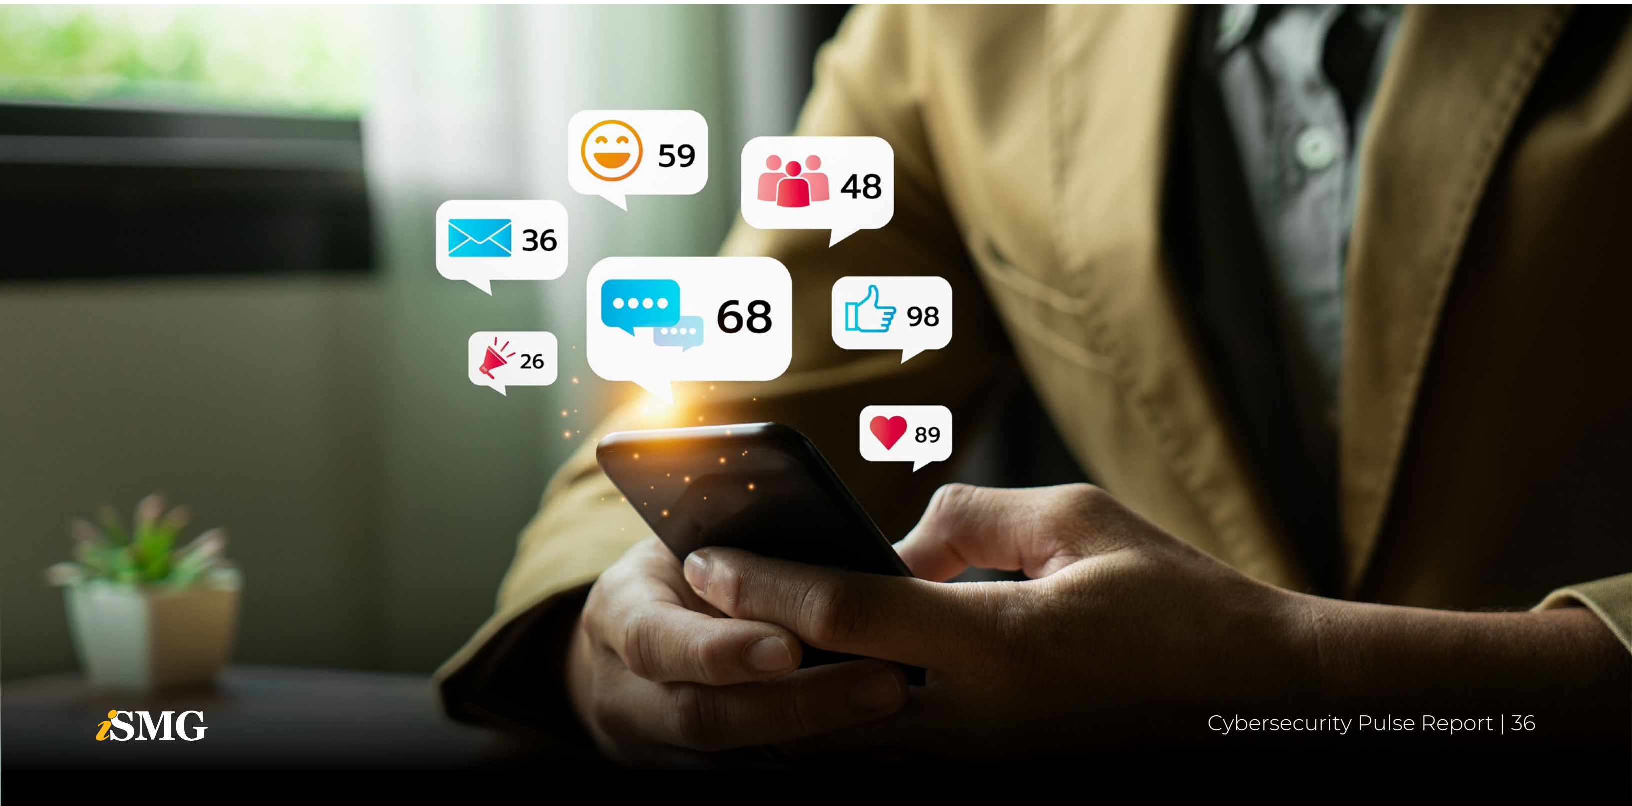

Conclusion: People Are the New Perimeter

Modern attackers are no longer just breaching systems - they're breaching people. As Khan observed, "You're basically hacking the mind," and today's adversaries are exploiting that reality across every communication channel - Slack, Teams, phone calls - deepfakes included. As Jason Morgan, distinguished architect for data science at Mimecast, noted, "It's not just an inbox anymore."

Even controls once considered robust, like MFA, are being bypassed through social engineering and sustained evasion. "If you can evade detection for long enough, you can get a very big payday," Biasini warned.

The tactics outlined - from ghost calls and quishing to EDR blinding and SOC overload - point to a common thread: attackers aren't just targeting vulnerabilities in code but also weaknesses in trust, attention and human behavior. They degrade visibility, exploit noise and manipulate people into circumventing the very protections meant to stop them. Defending against this requires more than technical controls; it demands procedural resilience, layered context and a security culture that treats people - not just systems - as part of the attack surface.

CHAPTER 4

Policy and Governance - Managing Systemic Cyber Risk

Attackers are no longer relying on a single method of compromise; instead, they blend technical exploits with social manipulation to maximize impact. This chapter explores how multivector threats—spanning deep technical subversion and human deception—are redefining the modern attack surface.

1 A Strategic Shift in Cyber Governance

As threats evolve from isolated incidents to systemic disruptions affecting critical infrastructure, the role of policy and governance has become increasingly strategic. The dialogue among experts across the public and private sectors has shifted from breach response to risk anticipation and ecosystem-level resilience. Nowhere was this more evident than in recent discussions that highlighted how cybersecurity policy, board-level governance and AI regulation are converging to manage growing systemic risks.

Over the past two years, debates have shifted from single-incident responses to systemic risk across tightly coupled digital ecosystems, including cloud, software supply chains and critical infrastructure.[17] Measures such as the U.S. National Cybersecurity Strategy (2023) and the 2024 National Security Memorandum (NSM-22) reflect a move toward mandatory baselines for critical infrastructure and clearer federal roles in resilience.[18]

2 Operational Collaboration and Governance

Operational collaboration and cybersecurity governance took center stage, underscoring a message reinforced across the public and private sectors: Cybersecurity resilience demands partnership - not isolation.

“One of CISA’s superpowers is the ability for the private sector to share information with the government without worrying,” Jen Easterly, former director of CISA and strategic advisory board member at Huntress, said, “It all comes down to partnership.” This framing reflects CISA’s broader strategic vision: encouraging systemic risk reduction through real-time threat intelligence sharing and trusted collaboration.

> *“One of CISA’s superpowers is the ability for the private sector to share information with the government without worrying, It all comes down to partnership.”*
>
> **Jen Easterly, former director of CISA and strategic advisory board member, Huntress**

Easterly further stressed CISA’s role in defending “target rich, cyber poor” sectors - entities such as rural hospitals, water facilities and education systems that remain disproportionately impacted by ransomware and lack the cybersecurity infrastructure to withstand persistent threats.

The operational model Easterly referred to is exemplified by the Joint Cyber Defense Collaborative (JCDC), which aims to unify government and industry response efforts. Reflecting on the JCDC’s foundational test case, Easterly said “The JCDC’s first major test, the Log4j vulnerability, proved the model’s value as researchers were sharing that information with us, and we were then able to share it in real time with businesses.”

Her statement highlighted the agency's position that defending democratic processes is a national imperative, independent of partisanship.

This alignment between strategy and action was echoed during a Black Hat 2025 policy panel, where experts reinforced that there is no "silver bullet" for cybersecurity. Instead, they advocated for collective action, enhanced collaboration and better threat intelligence sharing as the only viable paths forward.

The U.S. CIRCIA draft rule further underscores this approach, proposing 72-hour reporting for covered cyber incidents and 24-hour reporting for ransom payments, aiming to speed visibility into systemic threats.[19]

Together, these developments reflect a shift in mindset - from reactive incident response to coordinated, proactive cyber defense - where operational collaboration and real-time intelligence sharing become the linchpins of national resilience.

> *"The JCDC's first major test, the Log4j vulnerability, proved the model's value as researchers were sharing that information with us, and we were then able to share it in real time with businesses."*
>
> **Jen Easterly, former director of CISA and strategic advisory board member, Huntress**

3 Board-Level Accountability and Legal Readiness

Board-level accountability in cybersecurity is no longer optional - it's a mandate shaped by the realities of risk, regulation and public trust. Today's directors must understand that cyber risk is business risk.

"It's important for all board members to understand that they share in that responsibility. There's not one designated board member who's the cyber director," said Allison Jetton, founder of Jetton Law. "If you're not getting a cybersecurity briefing at least once a year, I would be asking for one."

> *"It's important for all board members to understand that they share in that responsibility. There's not one designated board member who's the cyber director," said. "If you're not getting a cybersecurity briefing at least once a year, I would be asking for one."*
>
> **Allison Jetton, founder, Jetton Law**

Beyond briefings, organizations are investing in board-level exercises to simulate cyber crisis scenarios, reinforcing the need for preparedness before real incidents occur. "Companies are now doing board-level cybersecurity exercises where the board actually gets to flex their decision-making muscle. It's hugely important. Nobody wants to try something for the first time in the middle of an actual crisis," Jetton said.

This proactive posture is essential in an era where government scrutiny is intensifying. When nation-state activity is suspected, regulatory expectations escalate dramatically.

"Government agencies do not like to be surprised - regulators in particular. When a nation-state is suspected, there very likely will be a number of additional regulatory agencies and other partner agencies who may be involved," Jetton said.

This shift also coincides with a surge in legal and compliance activity tied to cybersecurity: "Legal developments may actually be more frequently developing than that," Jetton said.

Underpinning all of this is visibility - into risk, into data and into organizational readiness: "If you don't know where your data is, you can't defend it very well," Jetton said.

These governance principles align with global regulatory moves toward enforceable standards, faster incident reporting and oversight of critical third parties, such as the EU's DORA and the U.K.'s new critical third-party regime in finance.

4 Regulatory Implications of AI-Powered Threats

"It's a new technology, new innovation, and it's non-deterministic, so hard to put some quality assurance around it," said Brennan Lodge, founder of BLodgic. "We've got this shadow AI problem that we've got to figure out."

Authorities stress that AI is not in a regulatory vacuum - existing laws, such as consumer protection and anti-fraud statutes, remain enforceable even for AI-enabled misconduct. New frameworks are emerging too, including the EU AI Act's requirements for high-risk AI (accuracy, robustness, cybersecurity and serious-incident reporting),[20] and the U.S. EO 14110 tasks NIST with developing AI safety and security standards.[21]

On the defensive side, initiatives such as DARPA's AI Cyber Challenge (AIxCC) demonstrate advances in autonomous vulnerability discovery and patching, with the aim of deploying these tools in critical infrastructure environments.

> *"It's a new technology, new innovation, and it's non-deterministic, so hard to put some quality assurance around it, We've got this shadow AI problem that we've got to figure out."*
>
> **Brennan Lodge, founder, BLodgic.**

5 Strategic Risk Culture: From Technical Fear to Systemic Governance

A mature cybersecurity posture is not built on fear - it is built on systemic governance, clarity and trust in decision-making frameworks. As organizations accelerate their adoption of advanced technologies such as AI, they must confront not just technical barriers but also deeply embedded cultural resistance to risk.

> *"There is all these fears. Those aren't the reasons we don't do it. Those are technical reasons. And there's another step beforehand. There's just some basic fears that stop entire organizations."*
>
> **James DeLuccia, product security chief, Honeywell**

"There is all these fears. Those aren't the reasons we don't do it. Those are technical reasons. And there's another step beforehand. There's just some basic fears that stop entire organizations," DeLuccia said.

This hesitation isn't rooted in capability - it's rooted in ambiguity around responsibility and liability. For many leaders, the fear of triggering unintended consequences or bearing legal exposure prevents necessary progress.

"If I turn it on, am I liable for it, or is it liable for itself?" DeLuccia said. The legal and ethical lines blur, raising existential questions about agency: "Legally, you want to know who's accountable. And then: are they accountable, or were you accountable? Because that's what you asked it," he said.

This uncertainty doesn't just delay innovation - it reinforces broken processes that cybersecurity teams are too cautious to disrupt. The result is stagnation, not from lack of capability but from the absence of a trusted governance framework.

"We're just not trying to make things faster because of that process. We'd love people to go fix the bad process," DeLuccia said.

Industry research, such as the WEF Global Cybersecurity Outlook 2025, continues to place systemic interdependencies and cyber-inequity high on the executive agenda.[22]

To shift from a culture of avoidance to one of strategic risk-taking, organizations must embed cybersecurity into governance structures - clarifying roles, codifying accountability and creating space for safe experimentation. Only then can fear be transformed into forward motion.

> *"We're just not trying to make things faster because of that process. We'd love people to go fix the bad process."*
>
> **James DeLuccia, product security chief, Honeywell**

Conclusion: Governance at Scale

Systemic cyber risk has become a national security issue. Managing it requires a governance model that extends beyond compliance checklists to collaborative, real-time partnerships, empowered boards and anticipatory regulation. As AI further entangles itself in digital infrastructure and the consequences of cyberattacks grow more asymmetric, the need for proactive governance, informed decision-making and cross-sector information flow becomes existential.

CHAPTER 5

Innovation in Attack and Defense - The Cybersecurity Arms Race

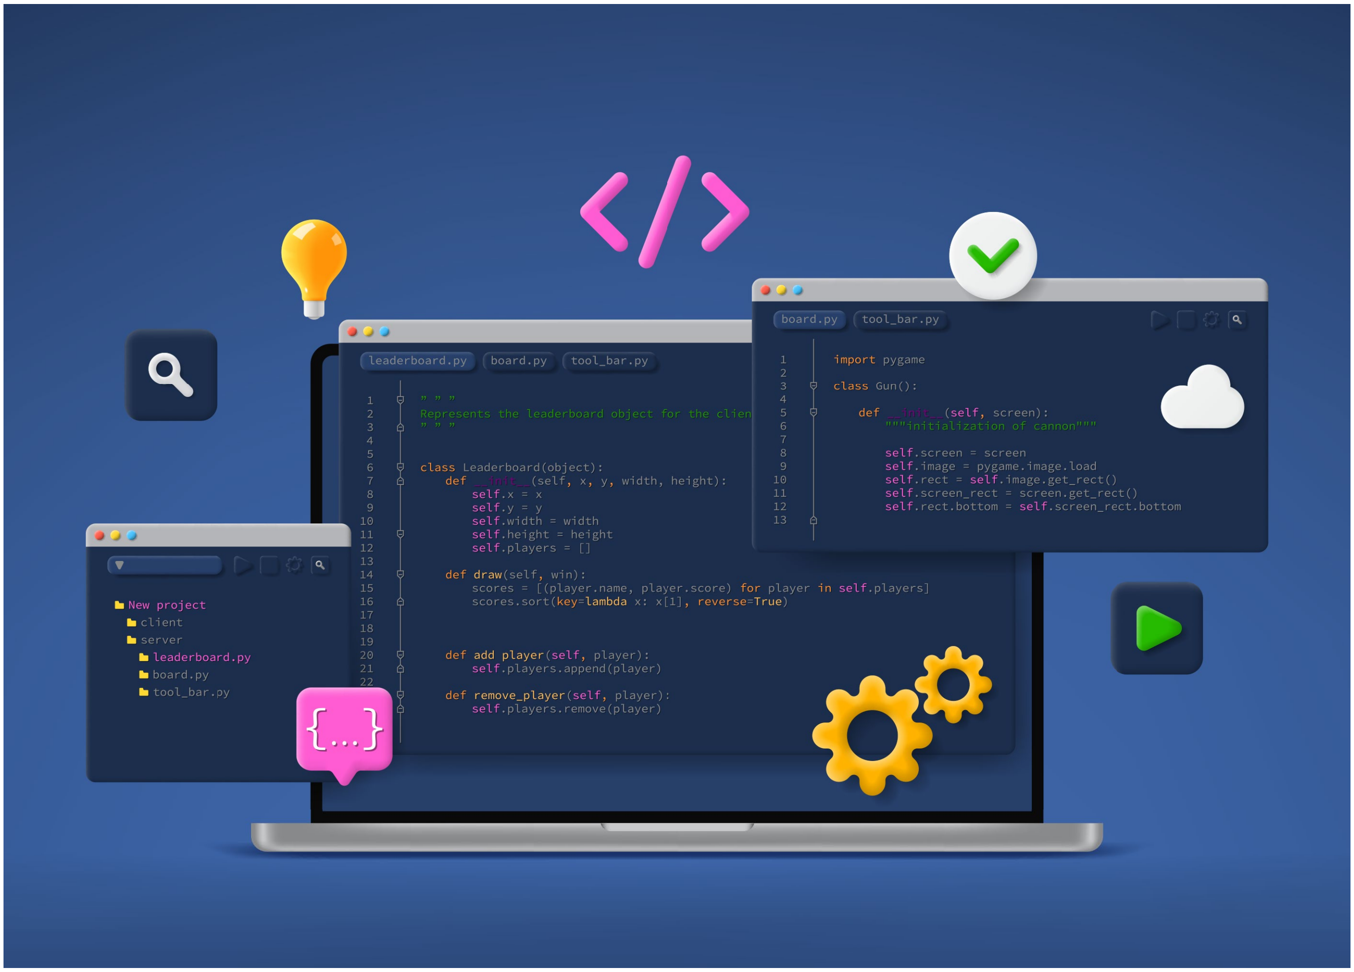

```
" " "
Represents the leaderboard object for the clien
" " "


class Leaderboard(object):
    def __init__(self, x, y, width, height):
        self.x = x
        self.y = y
        self.width = width
        self.height = height
        self.players = []

    def draw(self, win):
        scores = [(player.name, player.score) for player in self.players]
        scores.sort(key=lambda x: x[1], reverse=True)



    def add player(self, player):
        self.players.append(player)

    def remove_player(self, player):
        self.players.remove(player)
```

```
import pygame

class Gun():

    def __init__(self, screen):
        """initialization of cannon"""

        self.screen = screen
        self.image = pygame.image.load
        self.rect = self.image.get_rect()
        self.screen_rect = screen.get_rect()
        self.rect.bottom = self.screen_rect.bottom
```

```
New project
  client
  server
    leaderboard.py
    board.py
    tool_bar.py
```

If we've learned anything, it is that attackers now probe unconventional systems while defenders develop increasingly ingenious countermeasures. From gaming anti-cheat systems that could stop ransomware to hardware-focused conferences pushing the boundaries of embedded security, the arms race has expanded far beyond traditional IT infrastructure.

1 The Gaming Laboratory: Where Defense Innovation Thrives

In an unexpected corner of cybersecurity innovation, gaming anti-cheat systems have emerged as some of the most battle-tested defense mechanisms available. As Assistant Professor Marius Muench and Ph.D. Researcher Sam Collins from the University of Birmingham, U.K., revealed at Black Hat 2025, these systems represent “the most sophisticated and hardened security tools,” precisely because they operate under the worst possible conditions.

> *“You have to run it on a system you have zero control over … the attacker is the user. They have full access to admin privileges, software, firmware, hardware. It’s just like hell, a defenders’ worst nightmare.”*
>
> **Sam Collins from the University, Birmingham, U.K**

“You have to run it on a system you have zero control over … the attacker is the user. They have full access to admin privileges, software, firmware, hardware. It’s just like hell, a defenders’ worst nightmare,” Collins said, describing why anti-cheats have been forced to become “incredibly locked down, hardened defense.”

Operating in an environment where “you can’t win” outright, these systems focus on making reverse engineering challenging and expensive, effectively pricing many attackers out of the market. Collins noted that anti-cheat systems combine “anti-debugging and obfuscation techniques just to make reverse engineering more difficult” while actively monitoring for new code injected into a kernel to catch malicious activity early.

The zero trust approach is embedded deeply into these systems’ DNA. Anti-cheats work as if “they have already lost,” as Collins puts it, anticipating hypervisor compromise, kernel code execution or even “some horrible piece of hardware to directly read the memory without even going through the CPU.” These same hardware-detection capabilities are now finding parallels in enterprise-grade defenses designed to counter direct memory access (DMA)-based memory theft.

2 Tactical Innovations With Enterprise Applications

The gaming industry's defensive innovations extend beyond technical hardening to psychological warfare. Delayed bans frustrate cheat developers by withholding immediate feedback, while heavy obfuscation makes analysis feel like "you're analyzing malware half the time," Collins said. This obfuscation can even lead to blue screens or fake data in the crash stack during reverse engineering attempts.

The economic dimension proves equally important: Strong anti-cheats result in higher prices for game cheats and can significantly impact subscription-driven cheat markets. However, Muench warned of an unintended consequence. Since cheat creation remains legal in much of the world, some vulnerability finders might make more money by using it in a cheat rather than reporting via bug bounty.

Perhaps most remarkably, the agility of gaming defenses far exceeds traditional enterprise security. Collins recounted how BattlEye anti-cheat blocked a driver vulnerability "months before" EDR systems did, so effectively that "it would have quite literally stopped the [RobinHood ransomware] attack" if deployed in an enterprise environment. This mirrors findings from vulnerability research in other sectors, such as Google's Titan M chip exploit (CVE-2022-2023), where swift defensive action could have prevented credential theft long before standard endpoint controls reacted.[23]

This agility stems from attack tempo: "as soon as a game update comes … [cheat developers] very fast push up an update," Muench said, driving defenders to evolve with matching speed.

3 Hardware Security: The Forgotten Frontier

While software vulnerabilities dominate headlines, hardware security represents a critical battleground that has long been underserved. Aseem Jakhar, co-founder of Nullcon and senior vice president at ISMG, recognized this gap when he launched Hardwear.io in 2015 after realizing that while there were "a lot of security conferences already in Europe … focused on cybersecurity," none addressed hardware security specifically.

At the time, Jakhar was "doing a lot of research on IoT and hardware security" and decided to create a conference "specifically focused around hardware security … to highlight the importance and impact of attacks on hardware." Today, Hardwear.io attracts "roughly around 400 to 500 folks" for "a dozen highly technical … embedded security trainings … FPGA security, LoRaWAN … Bluetooth security … baseband firmware security … all delivered by experts."

The hands-on approach mirrors the gaming industry's practical focus. Hardwear.io features "HardPwn," similar to the live bug hunting where companies sponsor their products and equipment. Researchers spend time finding vulnerabilities and get paid. This reflects a broader industry move toward adversarial hardware testing, from fault injection techniques that bypass secure boot checks to formal verification methods encouraged by NIST to eliminate entire classes of hardware flaws before production.

4 Expanding Attack Surfaces

Today's cybersecurity conferences reflect the expanding attack surface that defenders must consider. At Nullcon Berlin, expected to draw "roughly between 150 to 250" attendees focused on "technical aspects of cybersecurity ... ethical hackers, researchers, security professionals ... some mid-level managers, and some CXOs," the agenda includes "AI talks ... talking about attacking AI," sessions on "reverse engineering chips" and "attacks on both MacOS side as well as on Windows side."

The conference maintains its practical edge with live bug hunting events sponsored by YesWeHack. "Researchers can come in, they can sign up, and then there'll be some targets given to them. Whoever finds vulnerabilities will get bounties for sure," Jakhar said. The goal is still keeping attendees "hungry on technical knowledge so that we can focus more on the latest attacks and mitigation techniques."

Emerging domains such as satellite systems are also entering the threat landscape. Events such as Hack-A-Sat 4 (2023) proved that real, on-orbit satellites can be compromised within hours, underscoring the need for new IEEE cybersecurity standards for space assets.[24]

> *"Researchers can come in, they can sign up, and then there'll be some targets given to them. Whoever finds vulnerabilities will get bounties for sure. The goal is still keeping attendees "hungry on technical knowledge so that we can focus more on the latest attacks and mitigation techniques."*
>
> **Aseem Jakhar, co-founder of Nullcon and senior vice president, ISMG**

5 Threat-Led Defense: Reversing the Security Paradigm

> *"Historically, security hasn't really focused on the behaviors that threats actually performed ... You can patch all the vulnerabilities, but what about the vulnerabilities that aren't patchable?"*
>
> **Frank Duff, chief innovation officer, Tidal Cyber**

While technical innovations advance on both sides of the cybersecurity arms race, a fundamental shift in defensive philosophy is gaining traction. Tidal Cyber CEO Rick Gordon and Chief Innovation Officer Frank Duff argue that the industry's obsession with vulnerabilities has created a dangerous blind spot.

"Historically, security hasn't really focused on the behaviors that threats actually performed ... You can patch all the vulnerabilities, but what about the vulnerabilities that aren't patchable?" Duff said. Gordon agreed, noting that "chasing CVEs has been a big focus ... unfortunately, what they end up doing is ... ignoring what the implications are."

6 A Behavioral Approach to Defense

Tidal Cyber's threat-led approach represents a fundamental reversal of conventional security thinking. Instead of starting with compliance or attack surface mapping, "threat-led is the opposite direction. We start with understanding who's attacking us, which of those behaviors we can effectively defend against and then you start to contextualize that to your attack surface," Gordon said.

Duff emphasized the practical benefits: "If you can focus on the threats that are most likely to be attacking you, it helps you focus the problem and make tangible progress."

This approach requires organizations to adopt what Duff calls a "presumed breach" mentality: "You have to recognize that sooner or later, somebody will be able to get in. How are you going to minimize the damage once they're in there?"

> *"benefits: "If you can focus on the threats that are most likely to be attacking you, it helps you focus the problem and make tangible progress."*
>
> **Frank Duff, chief innovation officer, Tidal Cyber**

7 Testing Organizational Readiness

The threat-led approach demands uncomfortable honesty about defensive capabilities. Gordon challenged security leaders: “Go ask your team, can we defend against Scattered Spider. I can almost promise you they’re guessing, because they haven’t developed the muscle memory of understanding who’s attacking and whether our defenses mitigate those behaviors.”

To illustrate the concept, Gordon borrowed a metaphor from a customer: “If you think of the tactics of the adversary as starting on your one-yard line and having to go 99 yards to score, you can look at each of those behaviors as an opportunity to reduce the probability of success.”

The key insight: Defense doesn’t end at initial access - there are multiple opportunities to disrupt the attacker’s playbook. This philosophy is being tested in AI-powered SOCs, where automation and autonomous analysis allow defenders to continuously map and mitigate active adversary behaviors.

> *“Go ask your team, can we defend against Scattered Spider. I can almost promise you they’re guessing, because they haven’t developed the muscle memory of understanding who’s attacking and whether our defenses mitigate those behaviors.”*
>
> **Rick Gordon, CEO, Tidal Cyber**

8 The Evolution of Cyber-Physical Threats

The arms race has expanded beyond traditional IT infrastructure to encompass cyber-physical systems where the stakes are measured in human lives rather than just data breaches. Gaming anti-cheats already demonstrate advanced capabilities in detecting malicious hardware, particularly in identifying devices used in DMA attacks, where “you plug in a separate card and skim off all the memory,” Collins said.

Looking toward AI security, Muench observed that game companies already deploy “AI and machine learning-based defenses for behavioral analysis” to detect cheaters whose playstyle mismatches human norms - techniques directly applicable to detecting anomalous behavior in critical infrastructure systems. Similar behavioral analytics are now being applied to autonomous vehicles, where fuzzing frameworks like those developed by Zoox reveal dangerous command sequences before they can cause real-world harm.[25]

Innovation In Defense (Radar View)

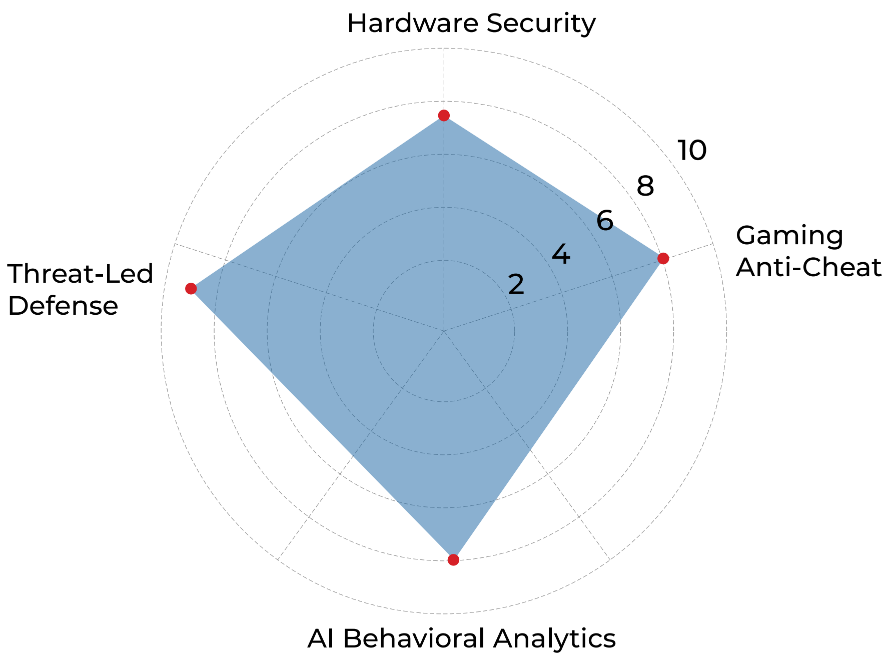

- Hardware Security: 8/10
- Gaming Anti-Cheat: 8/10
- AI Behavioral Analytics: 6/10
- Threat-Led Defense: 6/10

Conclusion: Strategic Implications of the Convergence Ahead

The cybersecurity arms race is driving toward convergence on multiple fronts. Duff predicted the unification of offensive and defensive capabilities: "It's not suitable to have one or the other ... sooner or later, there'll have to be a united front about what continuous threat exposure management actually means."

Meanwhile, the practical innovations emerging from gaming, hardware security research and threat-led defense approaches are creating new paradigms for protecting increasingly complex attack surfaces. From spacecraft systems to mainframe persistence and from AI platform exploitation to autonomous vehicle security, defenders are learning to think like attackers while attackers probe ever more exotic targets.

The most successful organizations in this arms race will be those that, as Collins and Muench advocate, "start with the threat first" while maintaining the technical agility and hardening principles demonstrated by gaming anti-cheats. In a landscape where traditional perimeter defenses have largely failed, the future belongs to those who assume breach, understand their adversaries and build defenses that can adapt as quickly as the threats they face.

Conclusion: Redefining Cybersecurity for an Age of Perpetual Breach

The cybersecurity landscape of 2025 reveals a core shift in how we must approach digital defense. As this report demonstrates, the convergence of agentic AI, sophisticated multivector attacks and expanding attack surfaces - from firmware to cloud orchestration layers - has rendered traditional perimeter-based security models obsolete. Organizations can no longer rely on reactive measures or siloed defenses when adversaries seamlessly blend technical exploitation with psychological manipulation, operating across hardware, software and human domains simultaneously.

The path forward demands a complete reimagining of cybersecurity strategy - one that embraces architectural security, assumes breach as a starting point and prioritizes behavioral defense over vulnerability chasing. Success in this new era will belong to those who recognize cybersecurity as both a technical and human challenge, where the goal is not preventing all attacks but building systems and organizations capable of surviving, adapting and continuing to function under persistent adversarial pressure. The question is no longer if organizations will face advanced attacks, but whether they have built the architectural, procedural and cultural foundations necessary to maintain resilience when those attacks inevitably succeed.

Sources

1. Black Hat USA 2025 – "From Prompts to Pwns: Exploiting and Securing AI Agents" (August 2025). https://i.blackhat.com/BH-USA-25/Presentations/US-25-Lynch-From-Prompts-to-Pwns.pdf

2. CyberArk Threat Research Blog – "Jailbreaking Every LLM With One Simple Click" (April 2025). https://www.cyberark.com/resources/threat-research-blog/jailbreaking-every-llm-with-one-simple-click

3. Trend Micro Research – "State of AI Security Report 1H 2025" (July 2025). https://www.trendmicro.com/vinfo/in/security/news/threat-landscape/trend-micro-state-of-ai-security-report-1h-2025

4. CyberArk Threat Research Blog – "Jailbreaking Every LLM With One Simple Click" (April 2025). https://www.cyberark.com/resources/threat-research-blog/jailbreaking-every-llm-with-one-simple-click

5. Medium – "Security in the Age of Agentic AI: Architectural Challenges (Part 2)" (July 2025). https://medium.com/complex-ish/security-in-the-age-of-agentic-ai-architectural-challenges-part-2-b1ae320e32b1

6. SC Media – "An Identity Security Crisis Looms in the Age of Agentic AI" (June 2025). https://www.scworld.com/perspective/an-identity-security-crisis-looms-in-the-age-of-agentic-ai

7. ESET Research – "BlackLotus UEFI Bootkit: Myth Confirmed" (March 2023). https://www.welivesecurity.com/2023/03/01/blacklotus-uefi-bootkit-myth-confirmed/

8. The Hacker News – "Researchers Reveal ReVault Attack Targeting Dell ControlVault3 Firmware in 100+ Laptop Models" (August 2025). https://thehackernews.com/2025/08/researchers-reveal-revault-attack.html

9. Microsoft Entra Blog – "OAuth Consent Phishing Explained and Prevented" (June 2025). https://techcommunity.microsoft.com/blog/microsoft-entra-blog/oauth-consent-phishing-explained-and-prevented/4423357

10. The Hacker News – "OAuth Redirect Flaw in Airline Travel Integration Exposes Millions to Account Hijacking" (January 2025). https://thehackernews.com/2025/01/oauth-redirect-flaw-in-airline-travel.html

11. The Hacker News – "Critical NVIDIA Container Toolkit Flaw Allows Privilege Escalation on AI Cloud Services" (July 2025). https://thehackernews.com/2025/07/critical-nvidia-container-toolkit-flaw.html

12. SC Media – "Study: Agentic AI Adoption Ramps Up" (Survey Data, June 2025). https://www.scworld.com/brief/study-agentic-ai-adoption-ramps-up

13. Praetorian – "Ghost Calls: Abusing Web Conferencing for Covert Command & Control (Part 1 of 2)" (Black Hat USA, August 2025). https://www.praetorian.com/blog/ghost-calls-abusing-web-conferencing-for-covert-command-control-part-1-of-2/

14. Cymulate – "EDR Evasion: A New Technique Using Hardware Breakpoints – Blindside" (August 2025). https://cymulate.com/blog/blindside-a-new-technique-for-edr-evasion-with-hardware-breakpoints/

15. TrustedSec – "Technical Analysis: Killer Ultra Malware Targeting EDR Products in Ransomware Attacks" (July 2024). https://trustedsec.com/blog/technical-analysis-killer-ultra-malware-targeting-edr-products-in-ransomware-attacks

16. Savonia University of Applied Sciences – "Cybersecurity: Mitigating the Risk as SOC Alert Analyst and Incident Responder" (April 2025). https://www.theseus.fi/bitstream/handle/10024/893410/Oguntoyinbo_Mayowa.pdf

17. The University of Chicago – "Understanding the Efficacy of Phishing Training in Practice" (January 2025). https://people.cs.uchicago.edu/~grantho/papers/oakland2025_phishing-training.pdf

18. Office of the National Cyber Director – "The National Cybersecurity Strategy" (March 2023). https://bidenwhitehouse.archives.gov/oncd/national-cybersecurity-strategy/

19. National Security Council – "National Security Memorandum on Critical Infrastructure Security and Resilience" (April 2024). https://www.cisa.gov/national-security-memorandum-critical-infrastructure-security-and-resilience

20. Federal Register – "Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) Reporting Requirements" (April 2024). https://www.federalregister.gov/documents/2024/04/04/2024-06526/cyber-incident-reporting-for-critical-infrastructure-act-circia-reporting-requirements

21. EU Artificial Intelligence Act – "Article 15: Accuracy, Robustness and Cybersecurity." https://artificialintelligenceact.eu/article/15/

22. Executive Order 14110 – "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence" (October 2023). https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence

23. World Economic Forum – "Global Cybersecurity Outlook 2025" (January 2025). https://www.weforum.org/publications/global-cybersecurity-outlook-2025/

24. Quarkslab – "Attacking Titan M With Only One Byte" (Black Hat USA, August 2022). https://blog.quarkslab.com/attacking-titan-m-with-only-one-byte.html

25. Via Satellite – "10 Defining Moments in Cybersecurity and Satellite in 2023" (January 2024). https://interactive.satellitetoday.com/via/january-february-2024/10-defining-moments-in-cybersecurity-and-satellite-in-2023

26. Dark Reading – "Researcher Deploys Fuzzer to Test Autonomous Vehicle Safety" (Black Hat USA, August 2025). https://www.darkreading.com/cyber-risk/researcher-deploys-fuzzer-to-test-autonomous-vehicle-safety

CONTRIBUTORS

List of Experts Contributing to This Report

- **Abid Adam:** Group Chief Risk & Compliance Officer, Axiata
- **Alan Michaels:** Professor and Director, Spectrum Dominance Division, Virginia Tech National Security Institute
- **Alex Green:** CISO, Delta Dental Plans Association
- **Alex Bovicelli:** Senior Director Cyberthreat Intelligence, Tokio Marine HCC
- **Ali Ranjbar:** Research Assistant, Pennsylvania State University
- **Allie Mellen:** Principal Analyst, Forrester
- **Allison Jetton:** Founder, Jetton Law
- **Anant Shrivastava:** Chief Researcher and Founder, Cyfinoid Research
- **Apostol Vassilev:** Research Team Supervisor, National Institute of Standards and Technology (NIST)
- **Aseem Jakhar:** Co-Founder, Nullcon, and Senior Vice President, ISMG
- **Ben Sawyer:** Associate Professor, Industrial Engineering and Management Systems, University of Central Florida
- **Brennan Lodge:** Founder, BLodgic
- **Chris Boehm:** Field CTO, Zero Networks
- **Chris Carlson:** Chief Product Officer, Critical Start
- **Doug Henkin:** Partner, Dentons US
- **Frank Duff:** Chief Innovation Officer, Tidal Cyber
- **Gianpaolo Russo:** Head, AI and Autonomous Cyber Operations, Mitre
- **Guy Kozliner:** Co-Founder and CEO, Rig Security
- **Jacob Ingerslev:** Head, Cyber and Tech Underwriting, Tokio Marine HCC
- **James DeLuccia:** Product Security Chief, Honeywell
- **Jason Morgan:** Distinguished Architect, Data Science, Mimecast

- **Jen Easterly:** Former Director, CISA, and Strategic Advisory Board Member, Huntress
- **Jos Wetzels:** Co-Founder, Midnight Blue
- **Kayla Williams:** Chief Data Security and Privacy Officer - Field, Cyera
- **Kevin Kin:** Global Vice President, SOC Transformation, Palo Alto Networks
- **Marissa Dotter:** Lead AI Engineer, Mitre
- **Matthew Canham:** Executive Director, Cognitive Security Institute
- **Marius Muench:** Assistant Professor, University of Birmingham, U.K.
- **Michael Leland:** Vice President and Field CTO, Island
- **Mishaal Khan:** Ethical Hacker, and Co-Author, The Phantom CISO
- **Nick Biasini:** Head, Outreach, Cisco Talos
- **Peter Garraghan:** Professor, Lancaster University, and CEO and CTO, Mindgard
- **Philip Martin:** Chief Security Officer, Coinbase
- **Philippe Laulheret:** Senior Vulnerability Researcher, Cisco Talos
- **Rich Campagna:** Senior Vice President, Products, Palo Alto Networks
- **Rick Gordon:** CEO, Tidal Cyber
- **Sam Collins:** Ph.D. Researcher, University of Birmingham, U.K.
- **Sam Curry:** Global CISO, Zscaler
- **Siddharth Rao:** Senior Security Research Scientist, Nokia Bell Labs
- **Taylor Margot:** Partner, Lytical Ventures
- **Tianchang Yang:** Research Assistant, The Pennsylvania State University
- **Todd Moore:** Global Vice President, Data Security, Thales
- **Vedang Parasnis:** Cloud Platform Software Engineer, Intel Corporation
- **Zohaib Ahmed:** Co-Founder and CEO, Resemble AI

About ISMG

Information Security Media Group (ISMG) is the world's largest media organization devoted solely to information security and risk management. Each of our 28 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, and fraud. Our 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

CIO.inc · CyberEd.io · CyberEdBoard · CYBER THEORY

GREYHEAD AN ISMG COMPANY · Xtra mile LIFECYCLE MARKETING · QG MEDIA · hardwear.io Hardware Security Conference and Training · NULLCON

Data Breach® Prevention. Response. Notification. TODAY · BANK INFO SECURITY® · HEALTHCARE INFO SECURITY®

GOV INFO SECURITY® · CAREERS INFO SECURITY® · FraudToday.io

Just for Credit Unions CU INFO SECURITY® · DeviceSecurity.io · PaymentSecurity.io

infoRisk® TODAY · AIToday.io · OT.today · ATHENA