Office of the CISO

Perspectives on Security for the Board

November 2025 – Edition 9

Google Cloud

November 2025

For more information, visit our Board of Directors Insights Hub

Table of Contents

Foreword

The board-level conversation on AI and security is evolving almost as quickly as AI technology itself. The central theme we hear in our boardroom discussions has shifted from debating whether directors should encourage spending on AI and security to evaluating how to measure the strategic return on that capital investment—and ensuring it can fuel long-term growth.

Our experts at Google Cloud confirm that a robust security posture can be a powerful enabler of trust and a crucial competitive asset. Perspectives on Security for the Board provides actionable insights that can help boards determine the best approaches to cybersecurity and risk for their organizations. In this edition, our final publication for 2025 and our ninth overall, we focus on three critical, interlinked themes based on our latest analyses and conversations:

Governing the agentic AI roadmap, from vision to value	Taking cyber risk governance from cost center to growth engine	Securing the unrelenting perimeter threat
We analyze how to oversee the agentic AI era, and demonstrate how early adopters of AI governance policies have achieved measurable ROI.	Organizations are successfully reframing security investment from operational cost toward the strategic value of risk reduction.	Boards who encourage investment in securing their public-facing corporate infrastructure as a critical risk mitigation strategy can reduce the likelihood of breaches.

When the Office of the CISO first started this report, our goal was to initiate a dialogue with the boardroom community. Since then, the conversation has fundamentally evolved—moving from discussing general AI potential to the strategic governance of agentic AI systems and the accelerated pace of external threats. This commitment to evolving insights sets the stage as we look forward to continuing this vital conversation with you throughout 2026.

Nick Godfrey, Senior Director, Office of the CISO, Google Cloud

From Vision to Value: Governing the Agentic AI Roadmap

AI evolution has progressed rapidly from predictive models to AI agents, sophisticated systems that combine advanced models with access to tools. They can independently execute tasks and make decisions—ideally under human oversight.

This change could provide a significant, long-term advantage for the company and its security posture. AI agents are more than just tools; they allow our security team to stop threats before they cause damage, not just after. This means security incidents are handled faster, and our top analysts can spend their time hunting for the most serious risks and to accomplish strategic tasks that would otherwise be deprioritized, such as developing proactive defenses and security architecture improvements.

Realizing measurable returns

Companies that pioneered the adoption of AI agents are already realizing significant, measurable returns on investment (ROI). Our recent report, *The ROI of AI in security*, highlights that 88% of agentic AI early adopters are now seeing a positive ROI on at least one generative AI use case. This value can help justify initial investments and also can help secure executive alignment on broader, scalable AI strategies.

Agentic AI Early Adopters Report Greater Security Improvements from Gen AI:

Unlocking success with board oversight in the agentic era

For boards charting their organization's AI strategy, the focus shifts to ensuring effective governance that maximizes strategic results while proactively managing emerging risks.

The report affirms the importance of executive support as a defining success factor: 78% of executives surveyed whose organizations have comprehensive C-level sponsorship report seeing ROI on gen AI now. 37% of executives report that data privacy and security is their top consideration for their organization when considering LLM providers.

The key to board support for secure AI lies in understanding the AI maturity roadmap, and validating that early successes can be converted into a repeatable, enterprise-wide scaling model.

Putting this into action

To maximize returns while governing the strategic deployment of AI, boards may consider the following three actions:

1. **Formalize executive sponsorship:** Boards should review the governance structure to ensure C-level sponsorship for AI initiatives is clearly defined and formalized. This oversight can help align the strategic AI vision with the resources required to maximize ROI.
2. **Affirm foundational governance:** Boards are well-suited to encourage the business to prioritize the responsible use of AI by confirming that data privacy and security are paramount considerations when adopting and scaling AI agents. Boards can help ensure that CISOs and business leaders have clear protocols for governing data throughout the AI lifecycle.
3. **Validate scaling and repeatability:** Boards can help their organizations convert early, measurable returns (such as employee productivity or direct revenue impact) into justified requests for broader, enterprise-wide agent deployment with a repeatable scaling model.

From Cost Center to Growth Engine: Governing Cyber Risk

Whatever else the discussion topic may be, boards of directors always return their focus to creating business value and reducing risk—not the technical cost-efficiency of security tools. That's the singular, key theme that has emerged from our Google Cloud executive roundtables and board interactions. Beyond ticking compliance boxes, boards want to know how cyber resilience has advanced the business strategy.

Boards should work to shift the mindset of business and security leaders so that cybersecurity is seen as an enabler and protector of the business, said Andreas Wuchner, CISO and board advisor, on an episode of the Cyber Savvy Boardroom podcast.

> *"Meeting compliance requirements isn't the same as strong cybersecurity and privacy practices that create the culture and conditions required for us to innovate—and the potential to differentiate a brand,"* he said.

This shift allows organizations to use smart cyber investments to enhance brand trust and unlock new revenue streams. The optimal outcome for boardroom dialogue involves confirming how the organization is positioning itself to treat cyber risk with the same rigor and strategic outlook as any other significant business risk.

Board oversight: Defining a strategic security program

To assess investment effectiveness, board focus may center on confirming the cybersecurity strategy is "fit for purpose," protecting crown jewels against the most relevant threats. There are three key topics boards members should discuss with management:

- **Defining risk accountability:** Boards can confirm that risk ownership is explicitly accepted across all business unit leaders, ensuring security is integrated into operational decisions.
- **Tracking program health:** Boards can review reporting that demonstrates measurable progress to link security controls with business outcomes, including reducing fraud and improving uptime. The output should allow the board to clearly assess whether current security spend is adequate to manage risk in the defined tolerance.
- **Prioritizing resilience:** Boards can seek clear reporting on the organization's capability for rapid recovery and adaptation following an incident, and emphasize leading resilience indicators as core metrics.

Productive dialogue, effective oversight

Fostering productive dialogue and effective oversight involves clarity on two fronts: translating cybersecurity value and governing emerging technologies.

Translating cybersecurity value

Boards want to know that technology investments are strategically integrated with growth objectives.

- **Business impact, not technical metrics:** Guide the CISO to use established frameworks to translate security efforts into clear business impact and trends to support informed capital allocation.
- **Building a trusted source:** Cultivate the CISO as a trusted source of information through consistent communication and integrity in reporting on attack scenarios and control maturity.

Governing emerging technologies

New technologies introduce new forms of risk that necessitate strategic foresight—especially AI.

- **Risk-first guidance:** Affirm the necessity of a risk-first mindset for all AI adoption, and guide the organization to deploy new technology securely and in line with corporate values.
- **Strategic oversight:** Stay on top of overviews of new control maturity and systemic weaknesses, especially for AI-driven technologies.

Putting this into action

To advance the conversation and ensure cyber risk is managed as a major business risk, boards may consider the following three actions:

1. **Reframe the investment conversation:** Transition discussions away from the technical cost of security tools and toward the strategic value of business risk reduction and resilience. Encourage the CISO to present security investments to the board using a business-centered framework that clearly links expenditure to the protection of key assets and competitive advantage.
2. **Verify accountability and clarity:** Review reporting to confirm that risk ownership is explicitly established across the business. Ensure that all security reporting is translated into clear, actionable business impact metrics, and avoid technical jargon.
3. **Integrate AI risk into governance:** Initiate a focused dialogue on the organization's AI strategy, affirming the need for a risk-first mindset. Request management to present a clear plan for how new AI deployments will be monitored, secured, and governed to maintain business continuity and trust.

Exploitation of Edge Infrastructure Continues to Drive Breaches

Exploitation of public-facing enterprise infrastructure continues to play a role as one of the most significant ways organizations are being breached today. Threat actors are seeking vulnerabilities to exploit in routers, virtual private networks (VPNs), firewalls, and email gateways—the security and networking infrastructure that makes up much of modern enterprise computing and network perimeters.

Crucially, traditional endpoint detection and response (EDR) is not able to protect these technologies. Because these vulnerabilities are so valuable to threat actors looking for a toehold on corporate systems, they've created an environment that has required defenders to develop a holistic security strategy to account for blind spots in traditional security.

For boards, proactive cyber defense should not be seen as a tactical IT expense. Instead, it's a strategic cost-avoidance measure essential for mitigating systemic financial risk and safeguarding core business resilience.

With the edge increasingly under attack, boards should work with the business to prioritize a security investment strategy that focuses on detection and mitigation, as well as prompt vulnerability patching. Three key areas of effort that boards can help push for are:

- **Patching quickly** to mitigate the risk of mass exploitation once vulnerabilities are discovered in edge infrastructure.
- **Implementing increased detection** through logging in other parts of the environment, to increase the chances of identifying successful intrusions if the actors are able successfully make it past the perimeter.
- **Taking steps to harden** critical assets and parts of the network, such as virtualization environments, through segmentation and controls to help make breaches less impactful when they occur.

While the breadth and range of initial intrusion vectors network defenders need to be concerned with has grown, in recent years this avenue of compromise has remained a constant issue. Data from the 2025 M-Trends report shows that exploits as an initial intrusion vector are roughly one-third of the breaches Google Cloud Mandiant has responded to in each of the last three years.

Zero-Days Exploited In-The-Wild by Year

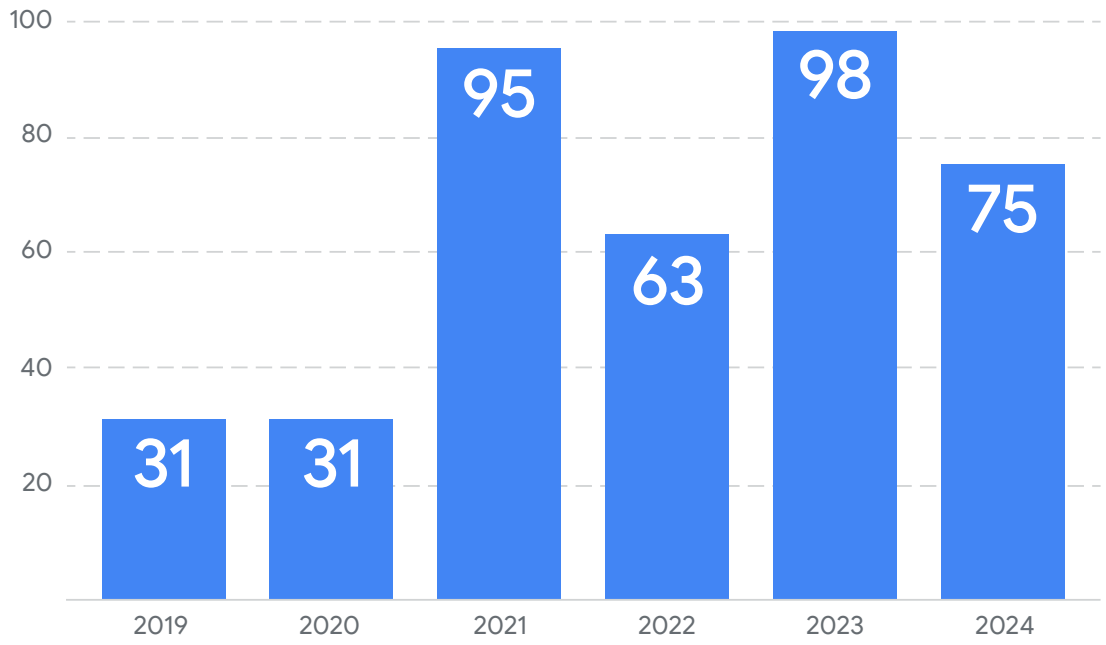

Exploitation of perimeter technologies occurs due to both known and unknown vulnerabilities. However, the continued rise of zero-day vulnerabilities (previously-unknown and unpatched security holes that threat actors use to launch their cyberattacks) in recent years has made this problem particularly acute.

One recent example of how this exploitation plays a role in a larger breach is the series of recently-disclosed intrusions comprising the BRICKSTORM malware campaign. This long-running, stealthy espionage campaign highlights a common pattern we have seen from China-nexus cyber espionage campaigns: leveraging edge appliances to gain network access.

The BRICKSTORM campaign by UNC5221 is part of a trend of zero-day exploitation driven by state-sponsored actors, especially from China-nexus APT groups. While it is easy to view these individual events and campaigns singularly, it is important to not lose sight that this is part of a larger trend that security teams must now contend with.

And it is not just edge device exploitation that is increasing the difficulty for defenders to detect intrusions and attacks today. It is important to

contextualize and understand this as part of the larger trend. Threat actors have been increasingly focusing their resources on evading detection. Similar to compromises of employee credentials or other attacks on identity—which is also notable this year—actors are seeking to carry out intrusions in a manner that evades traditional security controls and detections. Investments in EDR technologies, while important, have helped make endpoints more secure. The bad news is that they've pushed more threat actors to use stealthier intrusion methods and other techniques that are more difficult to detect.

As these devices increasingly become the initial vector for a successful intrusion, detection efforts have to prioritize elsewhere in the environment to catch the adversary as they move to other stages of the attack lifecycle.

Putting this into action

When organizations invest in holistic cyber defense, they're looking to ensure their long-term, future security through cost avoidance, increased resilience, and systemic risk reduction. A strategy that focuses future security investments around what threat actors are doing today can lead to tactical threat prevention.

Boards can help their organizations shift from reactive defense to proactive risk mitigation by proactively discussing the following strategic priorities with executive management and the CISO.

- **Prioritized patching:** Start by learning how the organization uses threat intelligence to help guide its vulnerability and patch management program. In addition to vendors, cross-sector information sharing can be helpful to this end. Prioritize resources on patching appliances that adversaries are actively exploiting—not just those with a highest criticality score. Ask executive management and the business to ensure the asset inventory is robust enough to cover all edge infrastructure.
- **Enhanced logging and detection:** Invest in detection capabilities to quickly identify follow-on lateral movement. Because threat actors' early and initial access to vulnerable devices on the perimeter is often difficult to detect, quickly finding the attacker once they move into other parts of the environment is key to reducing the impact from a breach.
- **Hardening critical assets:** Verify that organizational crown jewels—particularly attractive internal assets like virtualization environments—are sufficiently protected against cyberattacks. As they have become attractive targets for cybercriminals and state-sponsored actors alike, increased security here mitigates a range of threats. Because breach remediation can be extremely expensive, board action here can help safeguard the company's pocketbook.

How boards can learn more

To delve deeper into these critical discussions and explore all of our "Perspectives on Security for the Board" reports, including AI and cybersecurity, cyber risk oversight, cloud adoption risk mitigation, supply chain security, and insider risk, please visit our dedicated Board of Directors Insights hub.

Google Cloud Contributors

Alicja Cade, Director, Financial Services, Office of the CISO

Nick Godfrey, Sr. Director and Head of Office of the CISO

David Homovich, Advocacy Lead, Office of the CISO

Luke McNamara, Deputy Chief Analyst, Google Threat Intelligence Group

Seth Rosenblatt, Cybersecurity Editor

Google Cloud