

The cyber risk playbook

For the AI threat era

Thank you for your interest
in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies through NASPO and a wide range of other contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with Optro, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.



To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com



Explore More Resources:
carah.io/OptroResources



Join Events & Webinars:
carah.io/OptroEvents



Discover Technology Solutions:
carah.io/OptroSolutions



Learn About Procurement:
carah.io/OptroContracts



Connect With Our Team:
Optro@carahsoft.com
844-445-5688

The cyber risk playbook

for the AI threat era




Table of contents

Key takeaways04

5 steps to strengthen cyber resilience in the AI threat era05

1. Establish clear AI governance frameworks06

2. Improve visibility into third-party risk07

3. Develop mature cyber risk quantification practices08

4. Invest in AI security training and skill building10

5. Strengthen defenses against AI-enabled social engineering and ransomware11

Conclusion12

About Optro13

Methodology14

Executive summary

More than 80% of organizations have reported an increase in AI-enabled attacks over the past year. As threat actors use AI to scale and refine their tactics, AI-driven social engineering has emerged as the top defensive priority for security teams, surpassing ransomware.

At the same time, our research reveals a paradox. While 85% of organizations report strong confidence in overall cyber resilience, meaningful operational gaps persist. For example, 79% of security and GRC professionals express confidence in their organization’s third-party risk visibility, but only 59% of CISOs agree.

This disparity signals a critical transition: The acceleration of AI threats demands a fundamental shift toward automated governance and improved fluency. AI security expertise is cited as the top resource constraint, meaning that organizations struggle to both adopt AI and neutralize AI-driven exploits in real time.

The challenge for the modern security leader is no longer just recognizing AI risk. It’s operationalizing the defense against it. Drawing on insights from 210 security and GRC professionals, this playbook provides the roadmap to move beyond reactive policy and toward a state of systemic, AI-driven resilience.

Research methodology

Who we surveyed	Organizational profile	Panel sourcing
210 decision-makers via a global online panel	\$100M+ annual revenue (USD)	Experience: C-suite 44%, Lead 50%, Manager 5%
Roles: Cyber security, IT, GRC	Minimum size: 250+ employees	Geo: U.S., Canada, Germany, UK

For a full breakdown of the research methodology, role quotas, and industry segments, see page 14 of this report.

Key takeaways

1

CISO confidence seems decoupled from the board:

While overall organizational confidence is high at 85%, CISOs express significantly less certainty about third-party visibility (59%) than non-security peers (79%).

2

Policy gaps are creating operational pressure:

A lack of clear AI governance is the primary concern for 42% of CISOs, who fear that permissive policies will accelerate impersonation and identity-based attacks.

3

Risk quantification maturity is fragmented:

Although 73% of organizations claim to quantify cyber risk, the UK and compliance sectors lag, relying more on manual than automated processes.

4

The AI talent gap is the primary bottleneck:

Despite widespread adoption of AI tools, the number one resource constraint remains a critical shortage of personnel with AI security expertise.

5

Social engineering has overtaken ransomware:


AI-enabled attacks are on the rise, with 82% of organizations reporting an increase. Social engineering is now the top threat priority for security professionals.

5 steps to strengthen cyber resilience in the AI threat era

Achieving cyber resilience in the AI era requires rethinking information security as a continuous, intelligence-driven capability rather than a static line of defense. As threat actors increasingly deploy AI to scale and refine their attacks, security leaders must respond with equally advanced capabilities.

This includes using AI to strengthen detection and response, while building the governance, processes, and expertise needed to manage emerging risks. Technology alone cannot close the gap. Organizations must combine advanced security tools with modern governance and operational discipline to adapt quickly and maintain trust in an increasingly automated environment.

Keep reading for five steps you can take to strengthen your organization's cyber resilience.



Establish clear AI governance frameworks

AI touches every aspect and every department of an organization. This means that both its benefits and its risks are a shared responsibility, not just across audit, risk, and compliance, but also for cyber risk teams. AI governance can help organizations ensure compliance, manage risk, and drive responsible, ethical AI innovation at scale. For the most part, it begins at the top with the board of directors.

However, our research found that a lack of clear AI governance is the primary concern for 42% of CISOs. Among the issues that create the most day-to-day pressure for security leadership today, “board expectations vs. operational reality” ranked last among security and GRC professionals at 12%. Still, it ties for third among risk management professionals at 18%.

In addition, when asked how board-level expectations are aligned with the technical reality of cybersecurity risk today, only 27% of total security and GRC leaders answered that they were “highly aligned.” CISO confidence is largely decoupled from the board: While overall organizational confidence is high at 85%, CISOs are significantly less certain about third-party visibility (59%) than their non-security peers (79%).


The board plays a critical role in underscoring that AI is a crucial area that requires a strategy. While management bears responsibility for developing the strategy and monitoring related risks, the board must stay informed to ensure proactive balance between innovation and risk management. Three areas of collaboration between security and GRC leaders and the board for developing and overseeing clear AI governance policies include:

- 1. Ensuring strategic priorities:** This entails working with the board to define how AI can best advance the company’s mission and competitive positioning. The board must ensure that policies reflect the organization’s values and the risk appetite that they’ve agreed to.
- 2. Conducting ongoing evaluation:** The board must continually monitor the implementation of AI policies and understand the key performance indicators (KPIs) reported by management.
- 3. Understanding current policies:** It’s important to ascertain which policies the organization can embed AI into and where the organization must create new policies to mitigate other risks. There should be alignment on who will lead this process to ensure accountability.

Consulting firm Deloitte recently released an [AI Governance Roadmap](#) that provides additional guidance for infosec professionals to better align with board-level expectations. It defines both the board’s and management’s roles in specific AI areas and questions the board might ask management in each.

TECH TIP

Confidence gaps often occur when data is siloed across different systems and teams. Bring your data (e.g., risks, controls, and policies) into a [central hub](#) to ensure everyone, from IT to compliance to the board, sees the same live cyber risk picture.



Improve visibility into AI vendor third-party risk

Hiring third-party AI vendors and partners can inherently introduce serious threats to your business:

- 34% of security and GRC leaders reported a third-party or supply chain security incident in the past 12 months.
- Twice that amount (69%) experienced an increase in these types of incidents compared to the prior 12-month period. That number jumps to 90% for CISOs.
- When asked about their confidence in their organization's visibility into third-party cyber risk, including risks introduced by vendor AI capabilities, just 35% of total security and GRC professionals answered "very confident."

The good news: Organizations are broadly taking a multi-pronged approach to AI-driven cyber threats, including enhancing traditional security controls, according to 55% of security and risk professionals. These traditional security controls can include:

- Identity and access management (IAM), which we'll cover more below
- Endpoint protection, sometimes using AI itself
- Network monitoring, as the historical "perimeter" is gone, necessitates smarter systems
- Employee security training, which 56% of security and GRC leaders are doing to address threats such as phishing and deepfakes

TECH TIP

Organizations are confronting increasingly sophisticated phishing, impersonation, and manipulation campaigns. Effective cyber risk management requires technology that scales, adapts, and grows to operationalize against expanding risks.

Infosec professionals must also strengthen third-party and supply chain risk management, which 51% of security and GRC professionals are doing.

Third-party risk management (TPRM) helps identify, assess, and reduce risks associated with external vendors, suppliers, or other typical outsourced entities not managed in house. This includes fourth- or "nth"-party dependencies, such as a supplier's supplier. To proactively manage vendor risks, you need a solution that:

- Automatically monitors vendors through integrations with security and vulnerability tools
- Leverages real-time risk insights into vendor security and compliance across your ecosystem



Develop mature cyber risk practices

By adopting practices such as cyber risk quantification (CRQ) and vulnerability management, organizations can transform their security posture into a strategic, value-driven operation.

By quantifying the potential financial losses from various cyber scenarios, **CRQ differs from a compliance-first approach by informing and prioritizing risk decisions and related investments.**

DEFINITION


Cyber risk quantification, or CRQ, is a data-driven approach to cyber risk management that determines the potential financial impact of cyber risk using real-time risk telemetry and historical data.

More organizations are pursuing CRQ and vulnerability management (VM) for several reasons:

- Increasing regulatory focus on demonstrating risk management maturity, not just compliance, is upping the ante for infosec and GRC leaders.
- The growing financial impact of breaches and heightened scrutiny of cyber-related spending make CRQ a higher priority.
- A continuing, complex, and chaotic operating environment heightens the need for data-driven decision-making.
- Organizations want to make more informed decisions about the type and amount of cyber insurance coverage they need to avoid being underinsured or overinsured.
- The growing volume of exploited vulnerabilities and the shrinking window between “patch release” and “active exploit” make rapid, risk-based VM a higher priority.
- A continuing, complex, and chaotic infrastructure, spanning cloud, containers, and remote endpoints, heightens the need for automated, high-fidelity visibility.

While nearly three-quarters (73%) of organizations quantify their cyber risk, the approaches vary. The plurality (42%) maintains a fully mature CRQ program integrated into regular business decision-making, but the second-most cited response is manual processes (e.g., spreadsheets, qualitative-to-quantitative heat maps) at 26%. This drops to 71% and 37% for CISOs, respectively.

The most significant challenges organizations face when developing a mature CRQ program:



Most organizations can begin risk quantification with their existing resources and data sets, rather than waiting for ideal timing, conditions, or resources. Why? Delaying the process leaves your IT assets vulnerable and prevents risk-informed decision-making. Instead, view cyber risk quantification as an evolutionary process that builds upon your existing qualitative risk data.

Security teams can use existing assessment data to quantify risks right away, rather than waiting for the full implementation of a “perfect” methodology. Delaying quantification in pursuit of a flawless methodology is itself a risk decision and rarely the right one.

Asset data, for example, provides a strong starting point. This information enables business leaders to assign value to systems that directly affect performance metrics. Existing asset inventories and compliance documentation from frameworks, such as ISO 27001, PCI DSS, NIST SP 800-53, and COBIT 5, also provide valuable input. A focused approach makes the CRQ process more manageable and quickly improves communication of risk to leadership.

The **Factor Analysis of Information Risk (FAIR)** model also provides a helpful framework for understanding the data needed for CRQ. However, treating full FAIR adoption as a prerequisite can create unnecessary hurdles. This perception often slows progress, since comprehensive FAIR implementation can take a year or more. Instead, organizations can begin quantifying and managing cyber risks today using the data they already have, while gradually preparing for FAIR or even operating effectively without formally implementing it.

TECH TIP

If you're still using manual heat maps, you're likely spending more time on data entry than on actual defense. By bringing together your IT asset data and risk scores, you can see exactly which systems are most vulnerable in real-time.

Invest in AI security training and skill building

With a constantly evolving technology like AI, it's easy to feel like you're already behind before you've even gotten started. Training and upskilling can feel like futile exercises. "How will we ever keep up when everything changes on a dime?"

When authority and accountability are misaligned, it limits what security professionals feel empowered to act on. For instance, our research revealed that nearly a third of organizations (31%) feel they have authority for some risks but are accountable for risks beyond their control. CISOs feel slightly differently: One-third say accountability and authority are roughly balanced, while another third say they are held accountable for third-party risk they can't completely control.

To manage AI risk comprehensively, businesses must balance **authority and accountability**. Well-defined governance structures help by:

- Establishing clear roles
- Fostering cross-functional collaboration
- Implementing transparent policies to innovate with AI while maintaining control, trust, and adherence to regulatory and ethical standards


This is why effective training and skill-building are an organization's best defense in the era of AI-driven cyber risk: They strengthen an overarching governance framework. Increasing employee awareness and training was the top answer (at 56%) when we asked security and GRC professionals, "In which of the following ways is your organization primarily defending against AI-driven cyber threats?"

Skill development is an ongoing process that might involve practical exercises that test knowledge in controlled environments. Organizational initiatives can oversee and deliver skill-building efforts, or external firms or consultants with specialized expertise can supplement them. The objective in both cases is to develop the knowledge and practical competencies needed to identify, evaluate, and thwart AI-enabled fraud.

Skill-building can also scale with organizational maturity and risk exposure. The continuum can range from foundational awareness training for broad audiences to more advanced upskilling, specialized training, or formal certification for auditors with deeper responsibilities in AI-enabled fraud detection and deterrence.

TECH TIP

One of the biggest manual tasks for security experts is the collection of evidence. You can regain significant time by adding automation layers, such as gathering audit evidence and testing controls.



Strengthen defenses against AI-enabled social engineering and ransomware

AI gets a lot of the blame for the risks it poses, which makes plenty of sense. However, as organizations struggle to keep pace with technology, they often overlook that humans are responsible for both intended and unintended consequences of AI outputs. When we asked security and GRC professionals whether they had experienced any of the following in the past 12 months, phishing attacks (44%) and AI-enabled social engineering (42%), such as a deepfake voice or impersonation, were the top responses. What's more, 61% said AI-enabled social engineering has increased. However, these are also human-driven incidents that can be prevented or reduced with increased employee awareness and training, and a strong majority of CISOs (72%) plan to do just that.

To bolster organizational defenses against attacks like these, infosec needs a more modern IAM approach. By redefining their IAM approach to think like both attackers and end users, organizations can use the identity lifecycle to visualize user journeys, map identity-based threats to them, and leverage predictive AI, orchestration, and modernization to unify the lifecycle and achieve a holistic view. Here are the five identity lifecycle phases and how to upgrade your IAM approach in each.

1. **Onboarding and registration:** Look for technologies, such as deepfakes, liveness detection, identity verification, telemetry, and device-level checks, to identify known bad IPs or devices.
2. **Digital identity account creation:** Make sure you use detection technology, trigger authentication, and contextual signals to assess whether any account or address changes appear suspicious. Furthermore, conduct a secondary type of verification.
3. **Authentication and "ongoing access":** Balance user experience, security, and total cost of ownership such that, every time a user logs in, they're not proving who they are — they are just going through the authentication process.
4. **Account recovery stage:** Mitigate deepfakes, look for alternative authenticators, enable telemetry and bot detection, and monitor activity after the reset.
5. **Trust layer:** Focus on predictive AI, authentication, behavior, fraud, device trust, telemetry, and continuous authentication

Implementing these upgraded IAM recommendations enables continuous authentication and consistent device evaluation across these five phases, helping you pinpoint and address gaps to stay ahead of today's identity-based attacks and their future iterations.

TECH TIP

Centralize all your vendor info in one place so you can spot red flags across your entire supply chain. Using a system that automatically flags inconsistencies in vendor answers helps you catch identity-based threats before they become a major problem.

Conclusion

The rise of AI-driven threats is forcing organizations to rethink their approach to cyber resilience. As our research demonstrates, many security and GRC leaders are confident in their security posture, yet meaningful gaps persist. Closing them requires more than incremental improvements. Together, the five areas we've outlined above can help transform cybersecurity from a reactive function into a proactive, intelligence-driven capability that can keep pace with rapidly evolving threats in the AI era.

Ultimately, the organizations that will navigate this era most effectively are those that treat cyber resilience not as a security tech problem, but as an enterprise capability. That requires alignment across leadership, investment in the right expertise, and governance structures that can move as fast as the threats themselves.




Cyber resilience for the AI threat era

Move beyond manual tracking and centralize your risks, controls, and policies into a single, unified platform.

[EXPLORE PLATFORM](#)

About Optro

Optro (formerly AuditBoard) helps enterprises transform risk into opportunity, redefining GRC through an agentic system of action. More than 50% of the Fortune 500 trust Optro to elevate audit, risk, and compliance in addressing a new era of risk. Optro is top-rated by customers on G2 and was named a Leader in the 2025 Gartner® Magic Quadrant™ for Governance, Risk and Compliance (GRC) Tools, Assurance Leaders. To learn more, visit: optro.ai or [get in touch with our team](#) today.



Methodology

RESEARCH METHODOLOGY

The survey included 210 respondents sourced from a leading global online panel provider. They were selected from the panel based on geographic and role-based quotas, as well as screening questions based on role in audit and compliance, decision-making role, company size, and how long they have been in their audit role. All participants were audit, GRC, or IT decision-makers and purchase influencers working at companies with annual revenue of at least \$100 million USD. Selected respondents were further screened based on self-reported audit and compliance knowledge and attentiveness to survey questions.

ROLE QUOTAS

The survey divided respondents into three broad roles: C-suite 44%, Lead 50%, Manager 5%. Respondents were asked to select which role – from a list of 31 options – most closely described their primary responsibility, even if none were quite right or even if they performed more than one of these roles. Answers were consolidated into those three broad roles.

GEOGRAPHIC QUOTAS

The survey included respondents from the U.S., Canada, Germany, and the UK.

INDUSTRY

Although no industry-level quotas were deployed, we monitored the data to ensure that no single industry was overrepresented in the data. The final breakdown of respondents by industry is as follows: Financial Services 16%, Retail / Ecommerce 9%, Industrial and Manufacturing 7%, Energy & Resources 4%, Transportation and Logistics (including supply chain) 1%, Life Sciences (including healthcare and pharmaceuticals) 6%, Insurance 7%, Technology 31%, Business / Professional Services 3%, Education 2%, Government / Public Sector 5%, Telecommunications 8%, and Marketing and Advertising 1%.

RESPONDENT SCREENS

Role: All respondents were required to indicate that they were responsible for or had influence in evaluating and/or selecting audit compliance solutions or software for their organization.

Company size: All respondents must self-report that their companies have a minimum of 250 employees. All potential respondents from smaller companies were excluded. In total, the survey includes 8% of respondents from companies with 250-499 employees, 29% from companies with 500-999 employees, 35% from companies with 1,000 to 4,999 employees, 10% from companies with 5,000 to 9,999 employees, 3% from companies with 10,000 to 24,999 employees, 3% from companies with 25,000 to 49,999 employees, and 12% from companies with 50,000 or more employees.

Time in IT: Respondents must have spent a minimum of 3 years managing, planning, or purchasing compliance and/or cyber risk management software services or infrastructure in order to qualify for the survey. In total, 39% of respondents have spent 3 to 5 years in this role, 44% have spent 6 to 10 years in this role, 10% have spent 11 to 15 years in this role, and 7% have spent 16 years or more in this role.

Information level: In our experience, it is possible to have “qualifying respondents” who nevertheless prove to have too little information or knowledge about the space to provide useful data from which to draw insights. We therefore apply an “information screen” to respondents as well. Specifically, we ask whether or not respondents could explain certain terms to their colleagues if asked to do so. In order to qualify for this survey, a respondent must say “yes” to this question for the term “GRC (Governance, Risk, and Compliance)”

“Attention” level: It is easy for respondents to speed through surveys or not pay enough attention to provide useful data. We make an effort to exclude these respondents as well, as they provide generally less useful data. In this survey, respondents were screened out for “attention” reasons if they said they could explain the made-up term “CRISM Framework” to a colleague in the same question used for the information screen noted above.

RESPONDENT SCREENS

It is technically impossible and improper to list a margin of error for a survey of this type. The respondents for this sample were drawn from an online panel with an unknown relationship to the total universe, about which we also do not know the true demographics. As such, the exact representativeness of this, or any similarly produced sample, is unknown.