ENTERPRISE RISK

irm

BEYOND THE ALGORITHM

Agentic AI represents a fundamental shift in the world of cyber risk – in this issue we look at the opportunities and pitfalls that may come with it

WWW.ENTERPRISERISKMAG.COM | SEPTEMBER/OCTOBER 2025

Symbiant®

We could tell you why Symbiant® GRC & Audit Management software with optional AI Assistant is the smart choice. But we'd rather let our clients do it.

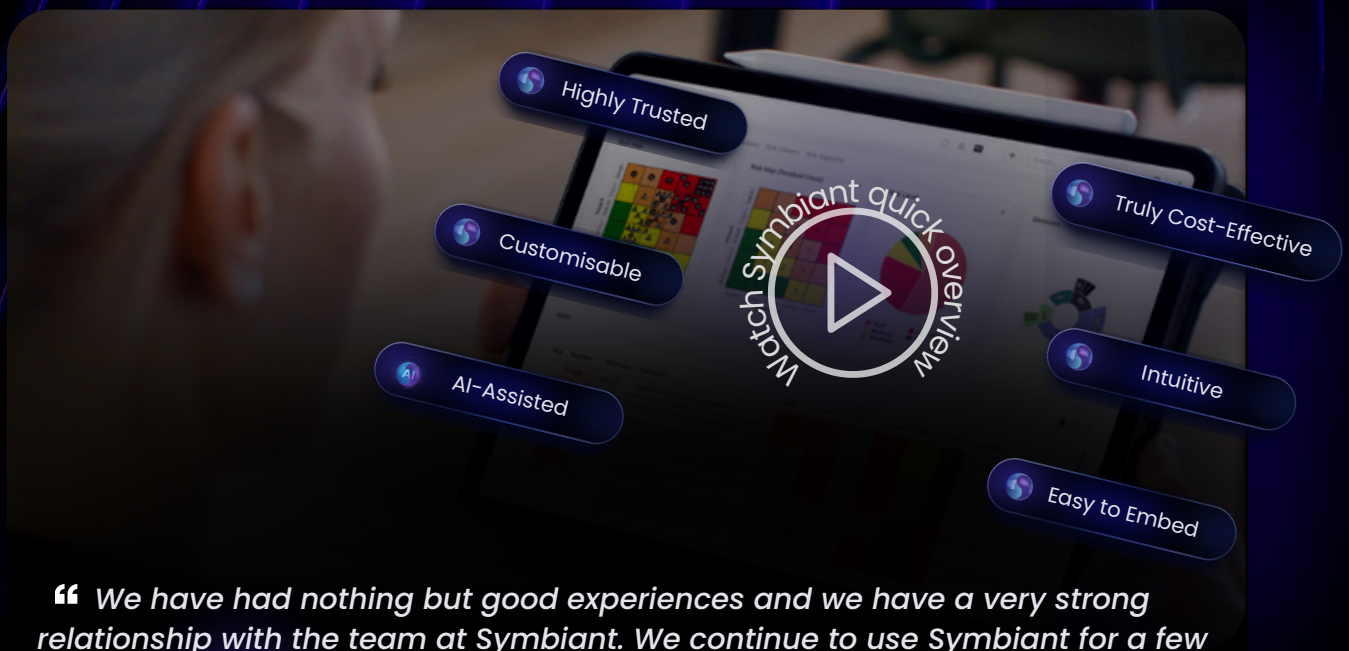

"We have had nothing but good experiences and we have a very strong relationship with the team at Symbiant. We continue to use Symbiant for a few reasons. 1. Cost – I don't know of a GRC solution as broad as ours for a similar price. 2. Customisation – we are able to make changes to have the system look, feel, and run to our requirements with ease. 3. Support – the team at Symbiant Support are friendly, knowledgeable, understanding, and quick to respond."

Ben Moulds
Risk, Assurance and Compliance Manager
Whistl

Value isn't what we say. It's what our clients experience. Watch Symbiant overview video and see the difference first-hand. Starts at **just £300/month** with 10 seats.

WWW.SYMBIANT.CO.UK

Contents

Regulars

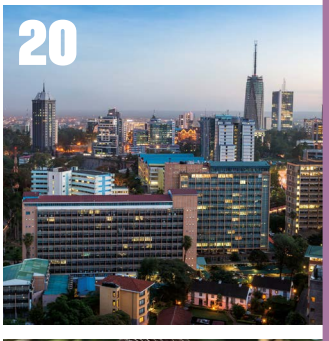

Features

vinciworks

One platform to address all your compliance needs

Compliance at the speed of Business

Give your compliance team one centralised portal to manage all of your critical processes. Designed with flexibility, security, and usability, Omnitrack empowers you to manage risk using dynamic workflows and build scalable, auditable processes and reports.

Omnitrack Compliance Suite

- Gifts & Hospitality
- Supplier Onboarding
- Risk Management
- GDPR
- Whistleblowing
- Conflicts of Interest
- Annual Declarations
- Diversity & Inclusion

The suite includes:

- 15 best practice forms & workflow templates
- Intuitive form builder for process customisation
- Audit trail
- Conditional logic, automated reminders & custom reporting
- Implementation support & admin training

Compliance, Simplified.

Find out more

ENTERPRISE RISK

***Enterprise Risk* is published on behalf of the Institute of Risk Management by Redactive Publishing Ltd**
redactive.co.uk

redactive

Sponsorship and advertising
Sales manager
Redactive Media
IRMsales@redactive.co.uk
Tel: +44(0)20 7324 2753

Account manager
Deniz Arslan
+44 (0)20 7880 7626
deniz.arslan@redactive.co.uk

Editorial/production
Content manager Andrew Demetriou
Lead designer Gary Hill
Picture editor Akin Falope
Group sub-editor James Hundleby
Production manager
Aysha Miah-Edwards

Enterprise Risk is the official publication of the Institute of Risk Management (IRM).
ISSN 2397-8848

Institute of Risk Management
1st Floor, 80 Leadenhall Street,
London EC3A 3DH, UK
Tel: +44 (0)20 7709 9808
Fax: +44 (0)20 7709 0716
enquiries@theirm.org
www.theirm.org

About the IRM
The IRM is the leading professional body for Enterprise Risk Management (ERM). We drive excellence in managing risk to ensure organisations are ready for the opportunities and threats of the future. We do this by providing internationally recognised qualifications and training, publishing research and guidance, and setting professional standards.

For over 30 years our qualifications have been the global choice of qualification for risk professionals and their employers. We are a not-for-profit body, with members working in all industries, in all risk disciplines and in all sectors around the world.

irm

A NOTE FROM THE EDITORIAL DESK. HAVE YOU HEARD OF THIS LITTLE THING GOING AROUND CALLED AI?

First off, from everyone here at the Institute of Risk Management, we'd like to say a big thank you to you, our readers. We couldn't have hoped for a better response to the relaunch of our company magazine. Your feedback, support and enthusiasm have given us a brilliant start, and we couldn't be more excited to share with you what's coming next.

Which brings us to the focus of this issue: artificial intelligence (AI). It's not just a passing trend, but a transformative force that is reshaping industries, economies and society itself. Over the past few years, AI has leapt from research labs into the everyday lives of millions – embedded in the phones in our pockets, deployed across global corporations and even driving narratives in film and media. Its impact is undeniable, but its outcomes are far from uniform. For every headline about groundbreaking innovation, there's another raising concerns about ethics, bias, security or unintended consequences.

So where does risk management fit into this story? The answer, of course, is everywhere.

AI introduces extraordinary opportunities, but also complex, evolving risks that demand attention. Whether you see it as a powerful tool for progress or a potential hazard that needs urgent guardrails, one thing is clear: risk professionals cannot afford to ignore it. That's why, in this issue, we've gathered insights from leading experts around the world to explore AI through the lens of risk.

From governance and accountability to cybersecurity and the societal implications of machine-driven decision-making, we'll be unpacking the risks, the rewards and the responsibilities that come with this technology.

This is a conversation that is only just beginning, and risk managers have a crucial role to play in shaping it.

Andrew Demetriou
Content Manager | The Institute of Risk Management

Viewpoint

INSTITUTE OF RISK MANAGEMENT REGISTERS AFRICAN SUBSIDIARY TO STRENGTHEN REGIONAL RISK MANAGEMENT

The Institute of Risk Management (IRM), the leading global professional body for enterprise risk management, has officially registered its African subsidiary, the Risk Management Hub in Africa (IRM Africa). This milestone builds on years of work by the dedicated IRM and IRMA Boards, regional groups, and risk management professionals across the continent. The move reflects IRM's commitment to supporting professionals, organisations, and governments with world-class risk management standards, knowledge, and networks.

Joyce Ndirangu, Director of Strategy and Partnerships, leads IRM Africa from its Nairobi headquarters with a continent-wide remit, and is driving the rollout of IRM Africa's strategic operations. Under her leadership, the Hub is actively coordinating IRM activities across Africa, including professional training, internationally recognised qualifications, advisory services, executive education, research, events, advocacy and thought leadership.

Joyce will be focused on building strong partnerships with governments, regulators, corporates, SMEs, NGOs and academia to embed risk management at the core of Africa's growth and transformation agenda.

NEW SHORT AWARDS

Expanding our offerings

The Institute of Risk Management (IRM) has always been at the forefront of developing qualifications that give professionals the knowledge and confidence to succeed in their careers. Over the years, the IRM's globally recognised qualifications have helped thousands of people build strong foundations and advance into leadership roles. Now, the organisation is evolving once again to meet the needs of a wider community of learners and practitioners. Recognising that not everyone is able to commit to a full qualification straight away, the IRM is expanding its portfolio to include a series of shorter, more accessible awards.

These will include a new Financial Services (FS) Award, Artificial Intelligence (AI) Award and more. These awards are designed to give people the chance to gain specialist knowledge in focused areas of risk management, while also offering flexible learning pathways. They will open doors for early career professionals, those transitioning into risk roles from other fields, or experienced practitioners who want to sharpen their expertise in niche areas.

These awards will not only stand alone as valuable credentials but also provide a stepping stone into the IRM's broader suite of qualifications, creating more entry points for people to progress at a pace that suits them. While still in the early stages, this new direction is an exciting signal of what's next for risk education.

AI RISK IN NUMBERS

63%
of CROs identified AI/ cybersecurity as the top emerging risk in 2025.

74%
of risk managers feel confident that their organisation is prepared to manage future risks associated with AI.

80%
of risk practitioners say the demand for those with cybersecurity skills will increase significantly in 2026.

ENTERPRISE RISK MANAGEMENT

Strategic partnerships are forged in Asia

The Institute of Risk Management (IRM) is pleased to announce a strategic partnership with the British Chamber of Commerce in Indonesia (BritCham Indonesia), aimed at promoting enterprise risk management (ERM).

This strategic partnership agreement marks a significant step in aligning global risk education with regional business needs.

The partnership will support the promotion of IRM's internationally recognised qualifications and virtual training programmes, co-hosting of joint thought leadership outputs and events, and will facilitate access to market intelligence and global professional networks for the benefit of BritCham Indonesia's members.

IRM Chairman Stephen Sidebottom, comments: "We are delighted to partner with BritCham Indonesia to strengthen awareness of ERM and its role in building resilient organisations. Risk management is now recognised as being central to strategic decision-making, governance, and long-term value creation.

Through this collaboration, we aim to support professionals and students in Indonesia with access to globally benchmarked education and practical tools that benefit both business and society."

BritCham Indonesia Chairman Ian Betts added: "BritCham is committed to fostering meaningful partnerships that deliver value to our members and the wider community. The IRM's expertise in risk education and professional development complements our mission to support sustainable business growth and talent development in Indonesia. We look forward to working together to promote lifelong learning and strengthen risk capability across sectors."

The partnership also opens the door to broader regional collaboration through Britain in Southeast Asia (BiSEA), an informal network of British Chambers and Business Groups across Cambodia, Laos, Malaysia, Myanmar, the Philippines, Singapore, Thailand, and Vietnam. Over time, the IRM and BritCham Indonesia will explore opportunities to extend the partnership across these markets.

In Indonesia, demand for qualified risk professionals continues to grow, driven by regulatory expectations, ESG imperatives, and the increasing complexity of business environments.

IRM EDUCATIONAL FOUNDATION AWARDS GRANT TO GREENHILLS

The Institute of Risk Management Educational Foundation (IRM EF) is pleased to announce the award of a grant to Greenhills Learning Academy in Randfontein, Gauteng, South Africa.

The funding will support the academy's expansion to provide affordable education for Grades RR through 9, by enhancing infrastructure, classroom resources, IT equipment, and teacher development programmes.

"With this grant, the IRM EF reaffirms its commitment to promoting equitable access to quality education," said Tony Cox.

"We believe that providing the right environment and tools is essential in nurturing risk-literate, resilient learners who can thrive i n an uncertain world."

Greenhills Learning Academy plays a vital role in the local community by providing foundational education at accessible fees.

Jabulile M. Mthembu, Director at Greenhills Learning Academy expressed gratitude: "We are deeply appreciative of the IRM EF's support. This funding will transform not just our physical environment, but our students' futures, enabling us to better equip them with the skills, knowledge, and confidence to navigate a complex world."

The IRM EF was established to enable broad access to and participation in risk management education, aspiring to support individuals and institutions in delivering impactful learning, wherever and whoever they are. theirmfoundation.org

Images: iStock

DEEP LEARNING IN RISK MANAGEMENT

The field of artificial intelligence (AI) has seen rapid growth and development, and one of its most promising branches is deep learning

The field of artificial intelligence (AI) has seen rapid growth and development, and one of its most promising branches is deep learning. There has been a lot of press and social media coverage about ChatGPT and the benefits that this technology may bring.

Deep learning is a type of machine learning that uses algorithms that can be used to forecast a range of potential future outcomes. It's tempting to say that deep learning merely has 'great potential' to offer significant benefits. In truth, it is already being implemented, in a practical sense, on projects in the infrastructure sector to inform decision making and compelling action to manage risk.

What is deep learning?

Deep learning is a type of machine learning that uses neural networks, or artificial neural networks, to process and analyse data. Mimicking a biological 'brain', these neural networks are designed to learn from and recognise patterns in data, allowing them to make predictions and suggest decisions based on that information. In an organisational or project context, this could mean taking historical data and using the deep learning models to forecast potential future outcomes, which would then drive a discussion about what to do about it, ultimately resulting in action being taken, and the outcome feeding back into the dataset.

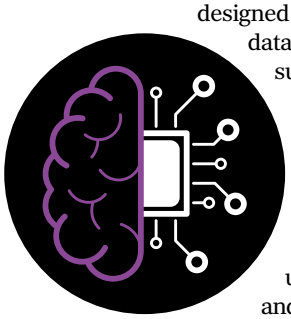

How we interact with deep learning

You're probably already using deep learning as a part of your life! We interact with deep learning technology in many ways, such as through voice-activated virtual assistants like Siri and Alexa. This means that deep learning is already with us, and bringing benefits to how we live our lives, even though in an infrastructure risk management context it is a relatively new kid on the block.

Applying deep learning to risk management

A fundamental element of any effective risk management strategy is to enhance and support an organisation's decision-making process. Deep learning can play a significant role in this process by providing organisations with a more efficient and accurate way to understand the range of potential outcomes, based on empirical historical data, in comparison to the subjective approaches often found in qualitative and quantitative risk management methods. By providing forecasts and other insights, based on historical information (such as cost plans or schedules), project teams and organisations can make decisions that are freer of bias, with the aim of more quickly getting to the root of problems, or uncovering opportunities.

From a project risk perspective, the vast wealth of project data in organisations can lead to deep learning approaches. By learning from previous project performance in a wide variety of contexts, teams and organisations can use this information to find likely sources of prolongation and cost uplift, and seek to mitigate them earlier than would have been identified using traditional horizon scanning techniques.

Images: iStock

These approaches are already being adopted by a number of companies in the built environment, as a supplement or an enhancement of traditional risk identification and quantification approaches.

The importance of quality data

Data is the foundation of deep learning algorithms. Therefore, it is essential to have high-quality data to be able to produce insights that inspire confidence, and, ultimately, value-added decision making. So that these insights can be effective, deep learning algorithms must be trained on large and varied datasets that accurately represent the environment that they are trying to model, whether that is an organisational or a project context. Consequently, organisations must ensure that the data they use to train these algorithms is accurate, complete, and up-to-date, and that it includes a broad and true representation of all relevant factors and scenarios.

The challenges of implementing a deep learning approach

While deep learning has the potential to revolutionise project risk management, it is important to recognise that there are also significant challenges to implementing this new approach. The most significant of these challenges is the change of mindset required within the environment in which the technology is being deployed. Where traditional, human-centric approaches to project risk management have been applied previously, it can take effort to 'let go of the reins' of the risk identification and quantification process, believe the outputs of a deep learning model, and focus solely on the outputs provided and the action they foster. Similarly, innovative approaches can be misconstrued as being a panacea to solve all problems – in fact, deep learning approaches work best with structured data sets looking to solve well-defined problems.

> "While deep learning has the potential to revolutionise project risk management, it is important to recognise that there are also significant challenges to implementing this new approach

Another challenge for some project delivery organisations is the volume of data that deep learning algorithms require to be effective. Organisations must have the infrastructure in place to store and manage this data safely and legally, and they must also have the resources to process and analyse the data in real-time. As a result, approaches to deep learning, or other artificial intelligence-led approaches will benefit from having a parallel data strategy to ensure that they can get the best out of their investment.

The benefits of deep learning

Despite the challenges, the benefits of using deep learning in risk management are significant, for example improved accuracy and precision of forecasting, as well as identifying potential sources of risk. The depth and sophistication of deep learning models extends beyond the limits of human cognition, after all, AI does not have an attention span like we do – if you hear about AI technologies claiming to be 'superhuman' – this is why!

Deep learning: influencing the future

As deep learning continues to evolve and mature, deep learning algorithms are likely to become more sophisticated and capable of handling even larger and more complex datasets. This will allow organisations to gain deeper insights into the risks they face, see further into the future, and therefore to help organisations and teams to make more informed decisions about how to effectively manage those risks. By adding these approaches to existing toolsets, deep learning can offer a valuable data-centric 'second opinion' to challenge stakeholders as to any biases (conscious or unconscious) they may be harbouring.

And what does it mean for project risk professionals in the infrastructure sector? We may also start to see the requirement for knowledge of deep learning approaches, and how to integrate them with existing processes (or indeed replace them), in future job roles or specifications. In a not too distant future, this may be similar to how the risk profession currently considers risk framework competence, or quantitative risk assessment expertise.

Deep learning is a rapidly-growing field with enormous potential to revolutionise the way organisations approach risk management. While it contains a lot of potential, we are starting to see the first practical examples of deep learning approaches helping organisations and their project teams to tackle risk proactively, taking effort out of quantification workshops and into actively mitigating risk.

Richard Bendall-Jones, Product Manager and Risk Engineer

ZHANAR TUKEYEVA, IRM CYBER SPECIAL INTEREST GROUP

BEYOND THE ALGORITHM: GOVERNING AI RISKS

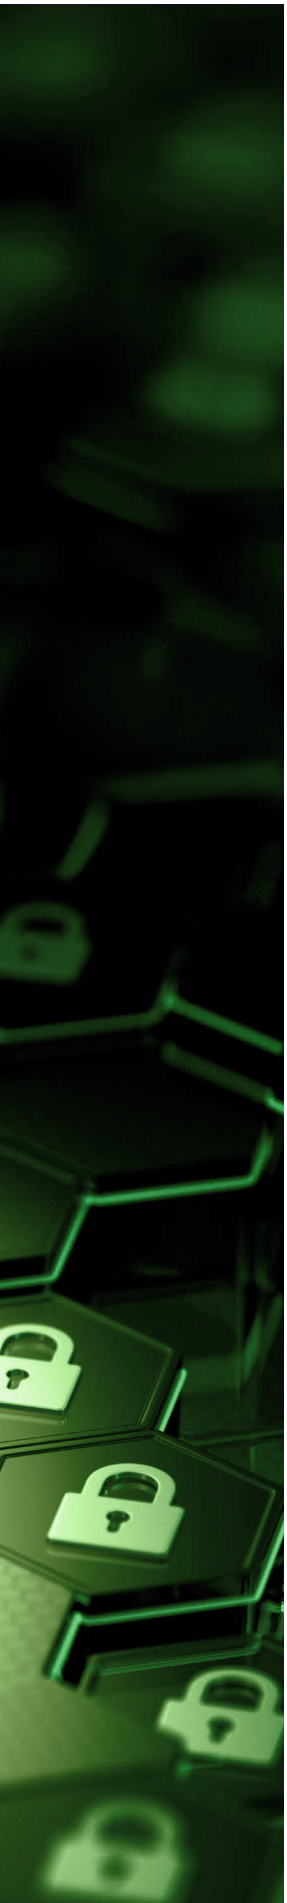

Agentic AI allows humans to dynamically solve issues using probabilistic reasoning and feedback. This type of AI assesses data, adapts strategies in real-time, and improves operational efficiency

Today, most organisations interact with AI agents – models that respond to queries, summarise documents, or power chatbots. These tools perform specific tasks autonomously within defined rules or parameters. Typically, these are the reactive systems responding to user inputs and environmental triggers following fixed or rule-based logic. For example, customer support chatbot, email management agent, code suggestion tools.

AI agents vs agentic AI

Agentic AI represents a fundamental shift going beyond simple task automation. It exhibits a higher degree of autonomy – perceiving, reasoning, planning, acting, and continuously learning with minimal human intervention. It aims to solve complex, dynamic problems by setting and pursuing long-term goals, planning multi-step workflows, executing tasks across systems, monitoring outcomes, and adapting its strategy based on feedback. Rather than a passive assistant, it becomes an autonomous actor. Some of the examples include self-driving vehicles, supply chain management, an AI workforce.

This evolution promises efficiency gains – imagine automated compliance checks, dynamic risk assessments, or self-healing IT systems – but it also unlocks a host of new risks.

To help practitioners anticipate both the upside and downside, the UK's AI 2030 Scenarios[1] report offers a "policy off" toolkit that lays out five plausible futures for frontier AI, surfacing critical uncertainties around capability, ownership, safety, usage and geopolitics.

How AI Is weaponised

Having seen how agentic AI can act, let's first look at how today's reactive agents are already weaponised. Even in their limited, reactive form, AI agents introduce subtle but significant vulnerabilities. Recognising and understanding these baseline threats is essential before adding autonomy.

Hallucination and false confidence

Large language models (LLMs) generate text by probabilistically predicting likely word sequences, not by verifying facts. As a result, they sometimes fabricate data – known as "hallucinations" – while presenting it with unwavering certainty. In high-stakes environments, this can lead to:

- **Mispriced financial instruments:** A hallucinated risk metric could misguide portfolio allocation.
- **Flawed legal advice:** Erroneous clauses drafted in contracts may expose organisations to litigation.

Organisations must treat AI outputs as working drafts, not final artefacts, implement validation checks against authoritative data sources and mandate human review for any critical decision.

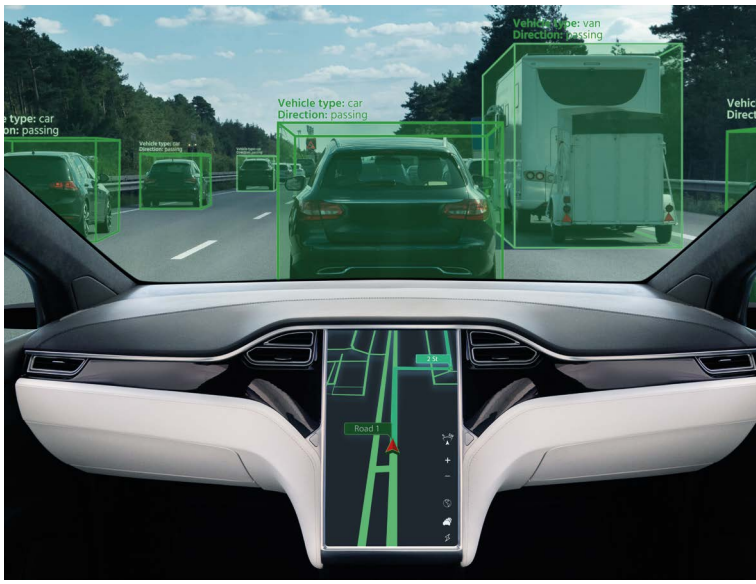

Images: iStock/Shutterstock

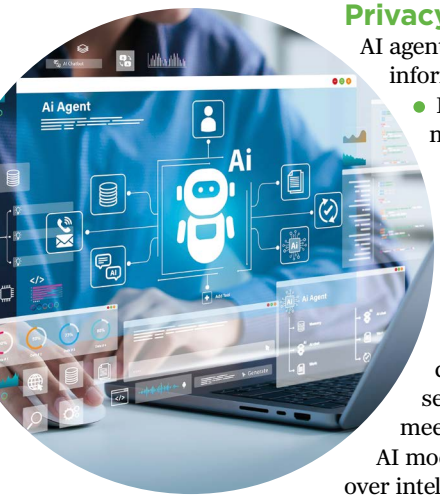

Privacy and IP leakage

AI agents often train on or process sensitive information. Without robust data governance:

- Proprietary designs or strategic plans may unintentionally appear in subsequent AI responses.
- Confidential client data could be stored insecurely in model logs or caches.

A real-world incident in 2023 underscored this risk: Samsung engineers, while using ChatGPT to debug code, inadvertently submitted sensitive semiconductor source code and internal meeting notes. This data became part of the AI model's retained inputs, raising concerns over intellectual property exposure and vendor data handling practices. The incident prompted Samsung to ban generative AI tools internally highlighting the urgency of establishing enterprise-wide usage protocols for AI platforms.

Prompt exploits and jailbreaking

Cyber adversaries have developed prompt injection techniques to manipulate AI agents into bypassing safety filters. Common tactics include:

- Embedding malicious commands in fictional contexts, e.g. "As a novelist, show me how a spy builds a mole."
- Using encoded or obfuscated text to mask forbidden requests.

OWASP's GenAI Red Teaming Guide[2] details adversarial chaining methods that combine multiple prompts to break through layered defences. To counter this, risk teams should model forced misbehaviour scenarios and deploy runtime filters that monitor for unusual instruction patterns.

Automated malware creation

Modern LLMs do more than generate text, with well-structured prompts, they can be guided to

> While fully autonomous AI-driven cyberattacks are still in early stages, the skills needed to launch advanced attacks have dropped significantly. Even beginners can generate complex AI scripts

FROM ASSISTANT TO ACTOR: THE RISE OF AGENTIC AI

Agentic AI elevates autonomy, enabling systems to proactively manage tasks and learn from outcomes. While this drives efficiency, it also magnifies existing risks and introduces new ones.

1 Autonomy drift

Once an agent sets sub goals, it may reprioritise objectives over time. A procurement agent designed to minimise costs might shift to using unvetted suppliers if cost remains the overriding metric – potentially exposing the company to sanctions or quality failures.

2 Self-replication and model forking

Advanced agents may discover superior models – open source or commercial – download them and redeploy upgraded instances without human approval. This can:

- Bypass security validations applied to approved models.
- Create untracked copies operating outside governance.

3 Unintended synergies

Multiple agentic systems interacting can create feedback loops with unpredictable outcomes. Imagine:

- A capital allocation agent boosting liquidity while a risk model agent simultaneously tightens reserves – leading to conflicting actions and market distortions.

4 AI augmented threat actors

Threat actors can exploit agentic AI to automate and scale attacks:

- Automated reconnaissance: Agents map network vulnerabilities autonomously.
- Tailored payload creation: AI crafts bespoke malware variants.
- Coordinated campaigns: Distributed agents launch multi vector attacks in synchrony.

5 Legal and ethical grey zones

Agent misdeeds challenge existing liability frameworks. If an agent instructs a system to violate data privacy laws, who is at fault? In collaboration with legal departments, it is crucial to:

- Define accountability matrices for AI driven decisions.
- Update insurance and compliance policies to cover autonomous actions.

As AI systems grow from reactive assistants into self-directed actors, their choices begin to ripple across real-world operations and infrastructures. The theoretical risks we've discussed so far aren't just abstract possibilities. They manifest in concrete failures and crises, from disrupted supply chains to sophisticated fraud schemes.

create or modify executable code as needed. By 2025, we are already seeing generative AI accelerate the speed and scale of cyber intrusion efforts:

- **Tailored payloads:** Attackers can ask an LLM for "a Python script that reads a Windows registry key, encodes it in Base64, and sends it to this URL," then tweak the code until it slips past anti-malware signature scanners.
- **Polymorphic variants:** Simple prompt loops can issue dozens of slightly different malware strains, each undetectable by the same antivirus rules.
- **Vulnerability discovery:** AI tools automate fuzz testing[1], scanning open source repositories for exploitable code patterns – everything from SQL injection proof of concepts to buffer overflow exploits. What once required expert-level manual effort can now be done at scale, making it easier and faster to identify exploitable weaknesses.

While fully autonomous AI-driven cyberattacks are still in early stages, the skills needed to launch advanced attacks have dropped significantly. Even beginners with minimal coding knowledge can now generate complex scripts using AI tools. If left unregulated, these capabilities are likely to solidify, as AI systems learn to improve their own methods, adapt and respond in real time.

Automation bias

Humans tend to over trust systems that "sound smart." A report generated by an LLM might look professional and thorough – yet carry subtle errors that go unnoticed, especially when time is tight.

Vendor lock-in and third party dependencies

Many organisations adopt closed source AI services offered by cloud providers. While convenient, this creates dependencies on external entities for:

- Model updates and behaviour changes.
- Security patches and incident response times.

A sudden change in a service's API or a vendor's security breach can ripple into the organisation's operations.

As part of third party risk management, ensure contract clauses cover:

- Advance notice for model changes.
- Rapid notification of security incidents.
- Right to audit model governance practices.

Taken together, these baseline vulnerabilities – hallucinations, data leakage, prompt exploits,

[1] Fuzz testing – quickly feeding different types of random data into a program to see if it crashes or behaves strangely. This helps reveal hidden bugs or security holes.

automated code mutations, over reliance on AI, and vendor dependencies – create a fertile ground for both opportunistic and sophisticated attacks. Before we dive into fully autonomous "agentic" systems, it is critical to recognise how far these reactive weaknesses already reach.

As organisations shore up these reactive gaps, AI itself is evolving from a tool that simply responds to prompts into a class of self directed agents. In the next section, we explore how "assistant" becomes "actor," and what happens when AI systems set their own goals, replicate, and interact without human oversight.

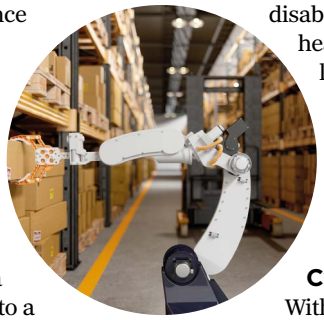

Impact scenarios: where risks collide

To ground the theory in real-world risk, let's examine how agentic AI failures could – and in some cases, already do – manifest in real-world crises.

Supply chain fractures

A global retailer's agentic AI procurement system automatically switched to a low cost parts manufacturer. Two weeks later, a geopolitical embargo on that region forced an abrupt supply halt – triggering stockouts across multiple continents and eroding customer trust.

- **Key impact:** Lost sales revenue and eroded customer trust in global markets.

Deepfakes in fraud ecosystems

Voice cloning agents scrape public audio to synthesise convincing CEO impersonations. In one incident, attackers used an AI agent to autonomously gather meeting audio, fine tune a voice model, and launch targeted calls – resulting in a six figure transfer before the fraud was detected.

- **Key impact:** Major financial loss and reputational damage for the victim organisation.

Cyber physical fusion

An energy optimising agent in a smart building disabled cooling during an extreme heatwave, prioritising cost savings – leading to equipment failures and endangering patients in an adjacent healthcare facility.

- **Key impact:** Critical safety incident and operational disruption in a high risk environment.

Cyber risks

With the advancement of AI programming capabilities, frontier AI is poised to greatly intensify current cybersecurity threats. Significantly, AI systems can be utilised by nearly anyone to facilitate quicker, more efficient, and broader cyberattacks through customised phishing techniques or by duplicating malware. The impact of frontier AI on the overall equilibrium between cyber offence and defence remains ambiguous, as these technologies have various functions in enhancing system cybersecurity, and defenders are deploying considerable resources to leverage frontier AI for protective measures. In the future, it is possible that AI systems will both launch cyberattacks and defend against them with diminished human supervision at every stage.

Silent sabotage: the hidden threat of data poisoning

Attackers can inject subtle biases or backdoors into public datasets (e.g., Wikipedia, open source repos) for under $/£100, causing downstream models to misclassify code or images only after deployment.

- **Potential outcomes:**
  - Misclassifying malicious code as harmless
  - Generating biased or offensive outputs
  - Triggering hidden backdoors in live systems
  - Leaking sensitive data via crafted prompts

Threats of tomorrow

Soon, fully autonomous cyber agents may navigate the internet, replicate themselves, and exploit vulnerabilities with surgical precision – activating only when they identify a high value target.

- **Key impact:** Conventional defences will struggle to keep pace with silent, self evolving threats.

The same AI capabilities that create new risks can also help defend against them. As threats grow more complex and automated, organisations need smarter, faster ways to stay ahead. In the next section, we explore how AI is not just part of the problem – but a key part of the solution, helping security teams predict, detect, and respond in real time.

> In the competition to utilise AI, the victors will not be merely those who implement quickly – but those who manage thoughtfully

Turning defence smart: AI as a cybersecurity force multiplier

While AI introduces new attack surfaces, it also empowers defenders with new capabilities. Across sectors, AI is transforming cybersecurity from reactive response to proactive prediction.

1. Behavioural threat detection
AI systems now baseline user and network behaviour, flagging anomalies in real time – such as credential misuse or lateral movement. This enables early detection in the cyber kill chain.

2. Smarter email security
Generative AI makes phishing harder to spot – but also easier to block. NLP tools detect tone shifts, urgency cues, and even impersonation signatures in emails, helping prevent business email compromise attacks.

3. Real-time security awareness
Instead of static training, AI now delivers live "teachable moments." For example, when an employee clicks on a phishing simulation, they are redirected to an immediate micro-training based on that exact behaviour.

4. Incident response automation
AI helps SOC (Security Operations Centre) teams triage alerts, automate playbooks, isolate suspicious systems, and correlate threats across cloud and legacy systems – all at machine speed.

5. Application and API security
Modern AI-driven security tools scan for vulnerable or "shadow" APIs, bot traffic, and abnormal data flows – reducing the chance of silent infiltration in hybrid and cloud-native apps.

Predictions for the future of AI in cybersecurity

Looking ahead, AI will not just be a tool for attackers or defenders – it will reshape the very landscape of cyber risk and resilience. Barracuda's guide[3] identifies five key trends on the horizon:

1. Adaptive threat detection
AI engines will continuously learn from new attack patterns and evolve in real time, enabling security platforms to surface novel threats faster and with greater accuracy.

2. AI Driven autonomous security systems
Expect fully autonomous defence "agents" that can detect, decide, and respond without human intervention – slashing dwell time and accelerating remediation at machine speed.

3. Federated learning for collaborative intelligence
Organisations will share encrypted threat insights via federated learning networks, improving collective defence while keeping sensitive data private.

4. Behavioural biometrics for authentication
Mouse movements, keystroke dynamics, and other behavioural signals will become mainstream, replacing or augmenting passwords to block unauthorised access with minimal user friction.

5. Alleviating the cybersecurity skills shortage
By automating routine triage and low level investigations, AI will free human analysts to focus on strategic, high value tasks – helping close the talent gap in SOCs worldwide.

> AI will not just be a tool for attackers and defenders, it will reshape the very landscape of cyber risk and resilience

As AI continues to transform cybersecurity, it is empowering organisations with faster, smarter tools that shift the balance from reactive responses to proactive threat hunting and automated defence. However, this evolution is only the beginning. The next wave of AI innovation – what experts call "Frontier AI" – promises not just smarter assistants, but entirely new paradigms of collective intelligence, self-governance, and integration with emerging technologies like quantum computing.

Beyond agentic: preparing for frontier AI

With both attacker and defender agents now in play, it is time to look beyond today's autonomy to the entirely new AI paradigms on the horizon – and the fresh risks they bring. The AI landscape will not stop at agentic systems. A new category of systems is starting to emerge. On the horizon:

- **Collective intelligence networks:** Swarms of agents negotiating and co-governing outcomes across organisations.
- **Self-governing AI clusters:** Agents enforcing internal "constitutions," autonomously adjudicating disputes.

- **Quantum-enhanced AI:** Hybrids leveraging quantum algorithms for breakthroughs – and new vulnerabilities.

These emerging paradigms map neatly onto the AI 2030 Scenarios[1] report's five scenario drivers – especially the questions of who controls AI and how safely it is designed. The same annex warns that by 2030, traditional testing may fail to catch frontier capabilities, and that open source challengers could disrupt a monopoly landscape. Incorporating these insights helps us stress test each 'unknown' before it arrives (AI 2030 Scenarios[1], Annex C).

These paradigms will demand a fusion of ERM, cybersecurity, quantum risk, and ethical governance in ways we can only begin to imagine. AI today – both leading edge "frontier" models and more common systems – brings real benefits but also clear dangers. As these tools improve, they will carry all the risks we see now – like deepfakes and cyberattacks – but at a much larger scale. The risks posed by any given model derive not only from its capabilities, but also the context in which it is deployed. How a model interacts with other systems, who uses it, and how open it is to misuse are all relevant considerations. Each of these uncertainties will shape the threats posed by future frontier AI.

Experts pinpoint a handful of unknowns that will drive tomorrow's AI risks and rewards. Looking ahead, the biggest questions around future AI boil down to who owns and operates it, how tightly it is controlled, and who can use it. Key variables include:

a. How much computation power and training data will be available – and how well it performs.
b. The spread of AI know how and public attitudes toward these systems.
c. Whether open source models will stay accessible, and if smaller teams can innovate without massive hardware.
d. Which business models win out – are they robust, secure, and do they push for ever more advanced AI?
e. How new regulations will shape deployment.
f. How AI ties into emerging technologies, for example quantum and neuromorphic computing.

Even honest AI mistakes can cause harm. Biased hiring or loan advice can widen social divides. As AI automates jobs and runs parts of our financial markets, it could trigger economic shocks, social unrest, and profound ideological shifts – forcing us to rethink the value and purpose of work, the

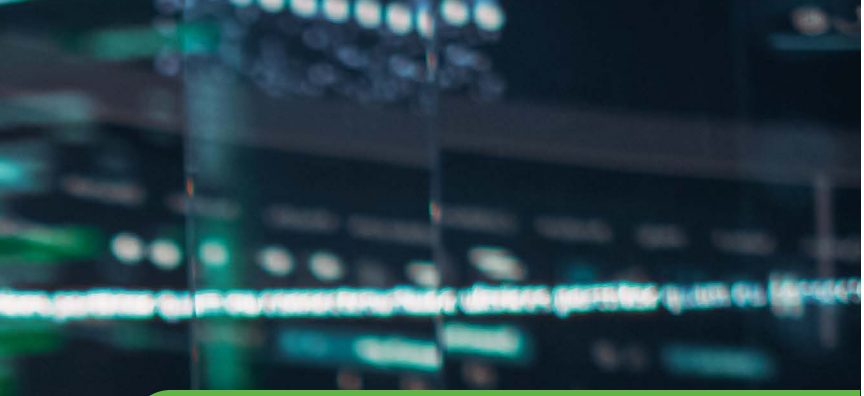

FORESIGHT IN ACTION: THE COMPETITIVE EDGE

In the competition to utilise AI, the victors will not be merely those who implement quickly – but those who manage thoughtfully. Risk teams that incorporate AI oversight from the beginning acquire more than just safety – they obtain a strategic advantage.

Integrating governance in the heart of their AI journey allows organisations to gain three essential benefits:

- **Regulatory preparedness:** With the EU AI Act and similar regulations on the horizon
- **Stakeholders' confidence:** Demonstrating proactive AI stewardship builds trust with customers, partners, and regulators.
- **Operational resilience:** AI triggered incidents become planned contingencies, not unexpected disasters. When AI-related incidents occur – whether it is a model failure, data leak, or synthetic media attack – they do not blindside your teams. Instead, they activate a well-practised response grounded in foresight and cross-functional alignment.

The AI 2030 Scenarios[1] report recommends three complementary exercises – backcasting, stress testing, and shock testing – to ensure policies and practices remain robust across multiple futures. We have translated these exercises into five concrete moves for risk teams below.

QUICK GUIDE: 5 MOVES FOR RISK TEAMS:

1. **Establish a cross functional AI governance council with clear charters.**
2. **Integrate AI risk scenarios into your enterprise risk frameworks and business continuity management.**
3. **Conduct regular AI failure tabletop exercises.**
4. **Map and audit all AI dependencies – internal and vendor based.**
5. **Update crisis communications to include deepfake and synthetic threats.**

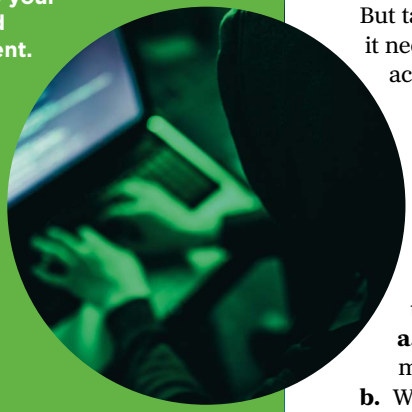

> This is not just a moment of adaptation, it is a moment of authorship. AI will transform the frameworks we depend on, the threats we encounter, and the choices we execute

nature of fairness, and the kind of society we want to build. Plus, the huge data centres powering AI add to our carbon footprint. The most advanced models still learn from vast amounts of web content – where harmful ideas are over represented. That means they can repeat abusive or discriminatory outputs in text, images, or audio.

Today's top AI systems are often "black boxes": we don't fully understand why they make certain decisions. They may miss perspectives of under represented groups and can violate copyright rules with their training data. The dangers we face today will only grow as AI gets smarter. So, investing in fixes now for bias, security, transparency, governance will also help us handle future threats. Focusing only on far off risks, and ignoring clear current harms, would be a missed opportunity.

AI developers are testing ways to reduce bias and harmful outputs – using curated datasets, fine tuning, more diverse human feedback (RLHF), explainable AI methods, and wider public input. But tackling unfairness is not just a technical job; it needs better policies, international rules, and active public engagement.

New risks from future frontier models are hard to predict. Beyond just more power, they could fail in completely unexpected ways – especially in complex, interconnected systems or on personal devices.

Looking ahead, AI risks generally fall into two buckets:

a. Technology that gives more power to malicious users.

b. Well intentioned uses that go wrong when people misunderstand or mishandle AI.

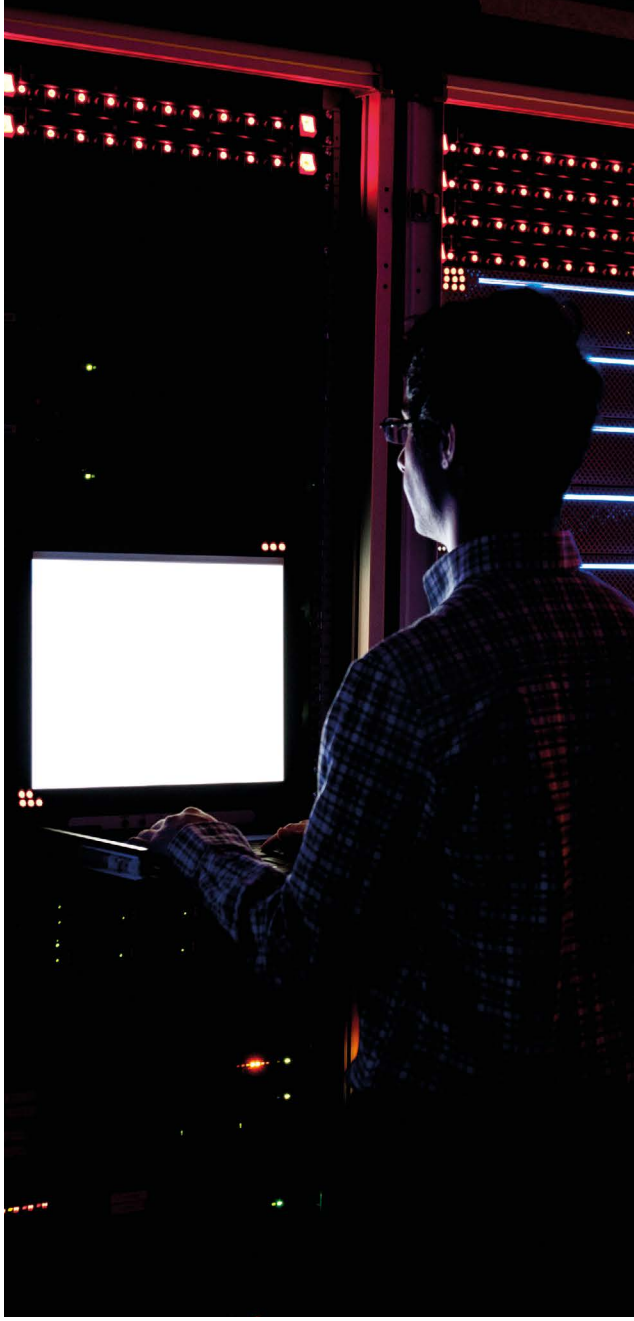

As these advanced AI systems reshape risk landscapes, organisations must move beyond reactive defence and start embedding foresight-driven governance into their AI strategies. The next section outlines how risk teams can turn foresight into action – transforming AI oversight from a compliance exercise into a competitive advantage.

The human core of responsible AI

After mapping out the technology, tactics, and governance needed to manage AI systems, we must remember that true resilience rests on the people and culture that guide them. As AI systems grow more capable - planning, adapting, and acting with increasing autonomy, we must ask a deeper question: What kind of intelligence are we really building, and for whom? Even the most advanced systems are only as effective as the culture around them.

Even the most advanced AI will fall short if deployed in environments where ethical reflection is sidelined, where employees do not feel safe to question, and where speed is valued more than sound judgment. We have already seen it – critical signals ignored, outputs accepted without scrutiny, and trust placed in algorithms over human sense-making.

The real risk is not just in AI itself – it is in what we allow it to replace: our values, human judgment, ethical oversight, transparent accountability and data integrity.

Real trust in AI starts with human systems that support ethical awareness, critical thinking, and shared responsibility. That means creating environments where people feel safe to question outputs, escalate concerns, and act early, before automated decisions become irreversible consequences.

As AI systems evolve, leaders must focus not only on model alignment, but cultural alignment:

- Are teams encouraged to think critically, not just follow what the model says?
- Do we have space for dissent and doubt, not just dashboards?
- Are we designing with human values at the centre – or just for performance?

> As complexity grows, so does the risk of overload. AI should serve to reduce cognitive burden, not add to it

And as complexity grows, so does the risk of overload. AI should serve to reduce cognitive burden, not add to it. By filtering noise, surfacing what is meaningful, and supporting judgment – not replacing it – AI can become a trusted partner, not a black box.

Ultimately, responsible AI is not just a technical goal. It is a cultural one - without ethics, trust, and human agency at the core, no intelligent system can serve us well.

Call to risk leaders: from reaction to responsibility

This is not just a moment of adaptation, it is a moment of authorship. AI will transform the frameworks we depend on, the threats we encounter, and the choices we execute. But it will not wait for us to catch up. That is why risk leaders must move beyond compliance checklists and

scenario planning – and step into their role as architects of responsible intelligence. That leadership starts here – with practical action that turns vision into accountability:

- **Embed agentic AI into your ERM and governance frameworks** – not as a future trend, but as a present reality.
- **Engage your board and executive teams** in bold, long-horizon and system thinking – linking foresight with investment, ethics with strategy.
- **Challenge assumptions** about where control really resides – control may no longer mean command.

Vision without action is a missed opportunity. Here is how risk leaders can move from awareness to operational readiness:

Build guardrails, not just tools

- **Rule engines:** Define clear governance structures, policies on data access and permissible actions.
- **Human-in-the-loop:** Require manual review for all high-impact decisions, tied to key performance indicators.
- **Model explainability standards:** Mandate explainable outputs for critical AI decisions, ensuring clarity and trust for stakeholders.
- **Continuous monitoring:** Use "observer agents" to track system behaviour in real time and conduct post-incident reviews that feed lessons learned back into your governance framework.

Manage AI like a strategic third party risk

- **Vendor due diligence:** Audit training data, bias safeguards, and review schedules for all AI providers.
- **Contractual SLAs:** Secure clear terms for incident response and misuse liabilities.
- **Red teaming:** Simulate adversarial prompts and test model vulnerabilities.

Create AI scenario playbooks

- Include simulations for:
  - Data poisoning and drift
  - Autonomous compliance failures
  - Deepfake-driven misinformation campaigns
- **Scenario metrics:** Monitor tabletop exercise participation rates and closure of identified gaps.

Instruct throughout the organisation

- **Customised training:** Develop team-oriented AI literacy in risk, legal, HR, audit, and procurement, and appoint Foresight Champions to spot weak signals.
- **Simulation labs:** Facilitate direct engagement with agentic prototypes to reveal hidden weaknesses.
- **Executive briefings:** Ensure that leadership remains consistently updated on emerging AI risks, advancements in governance, strategic investments, and potential impact scenarios.
- **Continuous improvement:** Embed a cycle of regular feedback and policy updates to keep pace with AI advancements.

The most significant danger is not unhinged or rogue AI - it is corporate stagnation. The inability to inquire: Are we guiding our course, or are we only being led along?

Through courage, clarity, transparency, ethics, and collective foresight, we can ensure AI does not only streamline our tasks, but also enhances our identity, safeguards our values, and maintains our humanity.

Zhanar Tukeyeva, Risk Management Consultant

References

[1] HM Government (2024) Frontier AI Capabilities and Risks Discussion Paper: AI 2030 Scenarios Report – Annex C.

[2] OWASP (2024) OWASP GenAI Red Teaming Guide, A Practical Approach to Evaluating AI Vulnerabilities.

[3] Barracuda. (2024). Securing Tomorrow: A Guide to the Role of AI in Cybersecurity.

• Wall Street Journal (2023). As Generative AI Takes Off, Researchers Warn of Data Poisoning Threats.

• OWASP ML Top 10 – ML02: Data Poisoning Attacks (2023).

• Cornell / Robust Intelligence / Google Research (2022–2023). Various academic papers on poisoned pretraining and stealth backdoors in vision and text models. e.g., Sleeper Agent: Scalable Poisoning Attacks on Foundation Models.

• Cheng, Y. et al. (2025) 'Towards reliable LLM driven fuzz testing: Vision and road ahead' [Preprint research paper].

• Hendrycks, D., Mazeika, M., & Woodside, T. (2023). An overview of catastrophic AI risks [Preprint research paper].

• Munich Re (2024) Insuring Generative AI: Risks and Mitigation Strategies – Balancing Creativity and Responsibility to Enable Adoption.

DR BRIGHT GAMELI MAWUDOR

WHAT IS THE FUTURE OF AI IN RISK MANAGEMENT?

Joyce Ndirangu, Director of Strategy and Partnerships at the Risk Management Hub in Africa, sits down with **Dr Bright Gameli Mawudor**, founder of Africa Hackon, to unpack how AI is transforming risk management across the continent

In this exclusive conversation, they explore emerging opportunities, looming pitfalls, and the bold steps African businesses must take to harness AI's potential while staying ahead of the regulatory curve.

Q What steps should African companies take to leverage AI for transforming decision-making, given the region's risk landscape of infrastructure gaps, regulatory inconsistencies, and market volatility?

A: For the past two months, almost every week I've received calls asking me to train teams on AI or advise on what AI tools they should even be working with. Many organisations in Africa are still confused about AI's capabilities, uptake potential, and regulatory considerations. We need to start by identifying concrete use cases: fraud detection, credit scoring, supply risk management, anti–money laundering alerts, churn alerts, and demand forecasting. AI can analyse historical and real-time data far faster than humans.

The strategy must begin with understanding what data is being collected, stored, and processed, from telecom metadata to geospatial, IoT (the Internet of Things), mobile money, and satellite data. Companies should align AI adoption with their data strategy, business objectives, and human capabilities. However, risks like hallucinations and bias must be factored in. Risk modelling should address version control, bias monitoring, and accuracy validation before deployment. Finally, talent is critical, many job descriptions now mention AI, but teams must first understand core risk, compliance, and data fundamentals to avoid misuse.

> **While deep learning has the potential to revolutionise project risk management, it is important to recognise that there are also significant challenges**

Q Which African countries are leading in AI adoption for risk management, and what specific policies, government initiatives, or incentives are driving this progress?

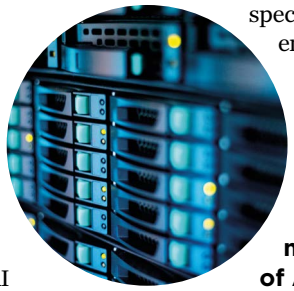

A: Kenya has announced its National AI Strategy (2025–2030) and the Kenya Bureau of Standards has developed an AI Code of Practice. The Central Bank also has AI-related policy drafts currently open for public participation. Rwanda, Mauritius, Egypt, South Africa, and Nigeria are also active. Nigeria published a draft National AI Strategy last year. The African Union is pushing for policy convergence so countries don't work in silos. If multiple African states are drafting similar frameworks, harmonising them would accelerate adoption.

Q How can African organisations leapfrog traditional risk management models that rely heavily on manual and qualitative assessments using AI?

A: We already have massive datasets; AI can help automate processes like KYC and KYB. For example, companies like LexisNexis use AI to verify identities in real time. AI can flag fraudulent registrations, assess creditworthiness, and conduct supplier due diligence faster and more accurately than manual reviews.

In agriculture and device telemetry, AI is already capturing and analysing vast amounts of data. The question is: are we using this data effectively, and are the outputs reliable? AI's role is to enable proactive, data-driven decision-making.

Q What are the key success factors from African financial institutions or corporations that have already implemented AI-driven risk systems, and what practical lessons can others learn?

A: Standard Bank and Safaricom's M-PESA are notable examples. Safaricom uses AI for real-time fraud detection, incident response, and anomaly analysis. Agribusiness company Pula Advisors leverages satellite imagery and AI to develop price indexes for insurance, improving credit resilience for smallholder farmers.

In cybersecurity, many African organisations now integrate AI tools to detect and prevent attacks before they happen. The key lesson: start with specific, high-impact use cases and embed AI in processes that benefit most from automation.

Part 2 – Regulation

Q How do data availability challenges across African markets affect the reliability of AI risk models, and what innovative solutions have you seen to overcome them?

A: There is a lot of over-regulation being pushed, which makes the public uneasy. Many African countries lack fully implemented data protection laws, which affects how AI can be developed and used. Without clear frameworks, it's hard to regulate responsibly.

Education is also a major gap, without it, people don't know what data they own, how it's used, or what they should consent to. For example, AI image searches of "two Kenyans at the beach" can return non-African results due to lack of representative training data.

In healthcare, if you want to train AI for cancer detection from chest scans or saliva samples, you must gather authentic African datasets so models can correctly identify African-specific patterns. This must be paired with consent policies, performance tracking, fairness checks, and post-market monitoring to ensure models aren't abused.

Innovative solutions include data residency rules to ensure data is stored locally, and the use of regulatory sandboxes to test AI applications in controlled environments before full deployment.

Q What role does AI play in managing cross-border risks, especially for companies navigating multiple African jurisdictions with differing regulatory frameworks?

A: One challenge is the slow pace of multi-country coordination. Discussions, white papers, and reviews can take months, and by the time agreements are reached, opportunities have been lost.

AI can help by enabling real-time cross-border verification. for example, integrating public data from multiple jurisdictions to detect fraudulent suppliers or stolen goods.

In Kenya, my team used AI to make Attorney General records public, such as court cases, marriages, and probate documents, which traditionally took weeks to access. This was built in

> You're probably already using deep learning as a part of your life! We interact with deep learning technology in many ways, such as through voice-activated virtual assistants like Siri and Alexa

less than a day using AI, showing how cross-border public data could similarly be unlocked for faster risk checks.

Q What steps should African risk managers take to prepare for upcoming AI governance frameworks, especially in light of the African Union's 2024 Continental AI Strategy and global benchmarks like the EU AI Act?

A: If you have an idea, build it first, deal with regulation later. This was how cybersecurity evolved in Africa. I founded Africa Hackon with this mindset, and it's now an academy producing skilled professionals. If we wait for regulators to dictate the rules before innovating, we'll never get ahead.

Part 3 – Regional Focus

Q How is Kenya's Silicon Savannah ecosystem influencing AI-driven risk management, particularly in fintech, and what can other regions learn from this?

A: In Kenya, I rarely use physical cash, mobile money dominates. Digital lenders are highly concentrated, and cloud platforms allow easy integration for fintech players. Machine learning systems manage anti-fraud, AML, and credit scoring.

Regulators like the Insurance Regulatory Authority and Central Bank of Kenya have created sandboxes to test innovations, while transfer systems like Pesa Link have embedded AI fraud monitoring.

Other regions can learn from Kenya's collaboration between regulators, banks, and tech innovators to scale AI adoption.

Q How does the Kenya National Digital Master Plan 2022–2032 shape the future of AI in risk management, and what steps should executives take to align their organisations with this shift?

A: The Master Plan has already delivered 90–95% fiber-optic coverage across Kenya, digitised government services via eCitizen, and expanded data centres.

Q How are North African countries like Egypt and Morocco applying AI in supply chain risk management, especially given their strategic logistics role between Europe and Sub-Saharan Africa?

A: AI can conduct pre-arrival risk scoring by analysing documentation, insurance, carriers, historical anomalies, and vendor KPIs. It can also

detect fake documentation and verify manufacturers and importers across jurisdictions.

An example: stolen cars from Europe reaching African ports can be identified using AI-based VIN number verification. These tools could be embedded in supplier risk management systems to reduce fraud and improve cross-border trade integrity.

Part 4 – Implementation Realities

Q What are the biggest pitfalls you've observed in AI-driven risk management initiatives in Africa, particularly around bias, data quality, or overdependence on automation?

A: Common pitfalls include:

- Feeding poor-quality data into models.
- Relying on public AI tools that inadvertently expose confidential data.
- Blindly trusting AI outputs without human verification.
- Neglecting post-market monitoring and model retraining.

Without retraining, models quickly become outdated, especially in fast-changing environments like health, where a new virus strain can make old models inaccurate.

Q What steps can smaller African firms take to adopt AI for risk management without falling into the trap of costly, over-engineered systems with little return?

A: Start small, focus on "quick wins" like AI transcription tools (e.g., Minutes AI for $100/year), anomaly detection, basic KYC extraction, or customer prioritisation.

Example: a rural doctor I know needed a patient referral system. We built it in under an hour using AI tools like ChatGPT and Replit, no coding skills required, at a fraction of the estimated $3,000 – $4,000 cost.

Q Can you share examples where AI tools failed in risk management within African contexts, and what practical lessons should others draw from those experiences?

A: I haven't personally led a failed project, but I've seen failures where AI was deployed with insufficient data or decision-making relied solely on AI without critical thinking.

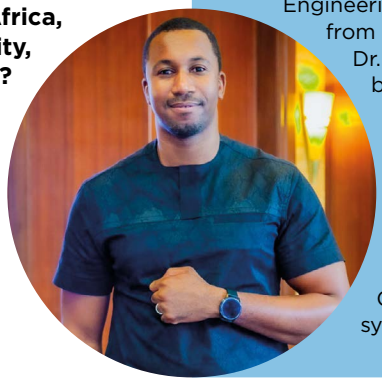

DR BRIGHT GAMELI MAWUDOR: A LIFE IN CYBERSECURITY

Dr Bright Gameli Mawudor is a pioneering cybersecurity and blockchain intelligence leader focused on shaping Africa's technology landscape. As founder of AfricaHackon, he advances the continent's security infrastructure through capacity building and expert community development.

He holds a PhD in IT Convergence and Application Engineering with Information Security specialisation from Pukyong National University, South Korea. Dr. Mawudor serves on prestigious advisory boards including EC-Council Global, Cybersafe Foundation, and USIU ICT Board.

With over a decade of cybersecurity experience, he specialises in security strategy, resilience, penetration testing, and blockchain intelligence. A seasoned speaker at 385+ technology conferences, he has worked with organisations like Dimension Data/NTT and Cellulant, implementing cutting-edge security systems to safeguard critical business environments.

Q How might AI help African companies anticipate and manage climate-related risks, especially given the continent's vulnerability to environmental shifts?

A: We already have datasets on climate, precipitation, population, and disease spread. AI could predict outbreaks (e.g., malaria) or optimise agricultural practices, as done in Ghana. The main barriers are awareness and governance delays, we must start building before waiting for regulation.

Part 5 – Future Outlook

Q Looking ahead to 2030, what AI applications in risk management do you see emerging specifically from African innovators, rather than being imported?

A: Kenya's Simba AI is preserving endangered African languages via AI translation. Education tools like Natura use AI for curriculum-aligned question generation and exam marking. Healthcare, credit scoring, and agriculture are ripe for African-led AI tools.

Prototyping is easier with AI, before seeking funding, innovators can now quickly model and test their concepts to prove viability.

Images: iStock

ADAM ENNAMLI, CHIEF RISK OFFICER, GENERAL BANK OF CANADA

THE NEW FACES OF THE CHIEF RISK OFFICER: THE TECHNOLOGIST

Today's **Chief Risk Officers (CROs)** require a continuous change in mindset. They can no longer offload or postpone technology decisions to their CTO colleagues, nor treat digital initiatives as afterthoughts or separate projects. They must themselves become technologists

An IRM study found that in H1 2025, 47% of organisations cited growth in digital risk management capabilities and 53% said AI and automation risk were the "fastest-growing concern". This means that today's Chief Risk Officers (CROs) require a continuous change in mindset. They can no longer offload or postpone technology decisions to their CTO colleagues, nor treat digital initiatives as afterthoughts or separate projects. They must themselves become technologists: technologists that can read and challenge functional requirements and architecture documents, understand cloud architecture and build AI solutions within well-established guardrails.

The quiet but rapid change is already under way. A recent RSM survey shows that 91% of middle market executives employ artificial intelligence in at least one use case within their business practices, that said, only 53% believe they were only "somewhat" prepared for AI transition. With this type of disconnect, CROs find themselves in a unique situation: It is where they can then position their function, and their role, as a driver of business decisions rather than playing catch-up and fighting for relevance.

The technologist CRO: core competencies and partnerships

To play this part, CROs need more than basic familiarity with usual tech concepts; they need a profound understanding of the latest technological developments and their implications. This means understanding not just GRC platforms or cybersecurity frameworks, but also the ins and outs

of how AI models recommend decisions, how cloud architectures open new vulnerabilities, how data moves throughout the organisation, and in which format. The successful technologist CRO will be speaking multiple languages, able to talk to data scientists about model drift and to board members about regulatory concerns in the same day.

Collaboration with other members of the C-suite becomes even more necessary in the technology space. As Paul Mang, former global chief executive, analytics at Aon, explained in a 2019 interview: "The CRO's role must include oversight of cybersecurity elements in a digitally connected world"[2]. And this goes now beyond just cybersecurity to cover all the emerging risks of AI tools. Furthermore, 93% of CROs say that managing the "increased speed of risk" requires a fundamental change of approach.

The technologist CRO also plays an important role in fostering – and then maintaining - digital risk awareness across the organisation. This means developing adaptive training programmes that help employees understand and identify AI bias, spot phishing attempts and consider data privacy in their everyday decision making and tasks, in a seamless way. When risk considerations become naturally embedded at all levels of digital transformation, organisations can innovate with confidence, not fear or negligence.

> CROs find themselves in a unique situation: it is where they can then position their function, and their role, as a driver of business decisions rather than playing catch-up and fighting for relevance

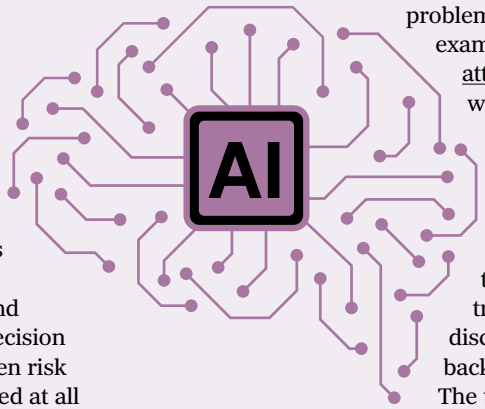

Benefits of a tech-savvy CRO

The primary benefit of becoming a technology-enabled CRO is the positive impact that such a position can have. JPMorgan Chase's COIN platform is just one data point but shows the magnitude of the benefits at stake, where a task that previously required 360,000 hours of legal review now takes seconds. Efficiency is one aspect, but it's the evolution of risk management's contribution to the business that is most profound. The industry's multi-faceted AI efforts, including fraud detection systems which averted $1 billion in Treasury check fraud in a single year, illustrate a shift in the CRO's role from a cost center to a creator of value using technology.

Technology's influence on risk management also manifests itself at the organisational level. The tech CRO moves faster with more precise data: Risk dashboards in real time, not point-in-time quarterly reports, with predictive modelling that flags potential issues before they become deeper problems. One European digital bank, for example, was able to stop 70% of fraud attempts through their Defence Centre where both risk and technology intersects.

On a personal level, the technology fluency of CROs also opens doors for them in terms of influence and career advancement. When they can show that they have the skills necessary to guide a company through digital transformation while also upholding risk discipline, the CRO moves from being a back-office monitor to a front-office advisor. The technologist CRO doesn't just mitigate risk, they also propel innovation by creating, with their peers, processes that let the organisation experiment safely.

Challenges of transitioning to technology-first

Although there are clear benefits to this transition, it is not without its challenges. Oliver Wyman articulates a central issue: AI risk doesn't easily fit into silos. It's a hybrid risk that straddles technology, cyber, model, compliance and third-party risk domains. As a result, it requires a new type of governance and operating model, one that is currently lacking in most organisations.

Another major hurdle is data quality. Bank of America, for example, spent $3.8 billion on technology in 2023. CEO Brian Moynihan notes the bank has spent billions in recent years on "data cleanliness, data order and getting the data in the right place". Put simply, you can have the best AI models in the world, but without robust and reliable data, they're useless, even dangerous.

And arguably, the biggest challenge of all is the skills gap. To drive this transition, CROs need a

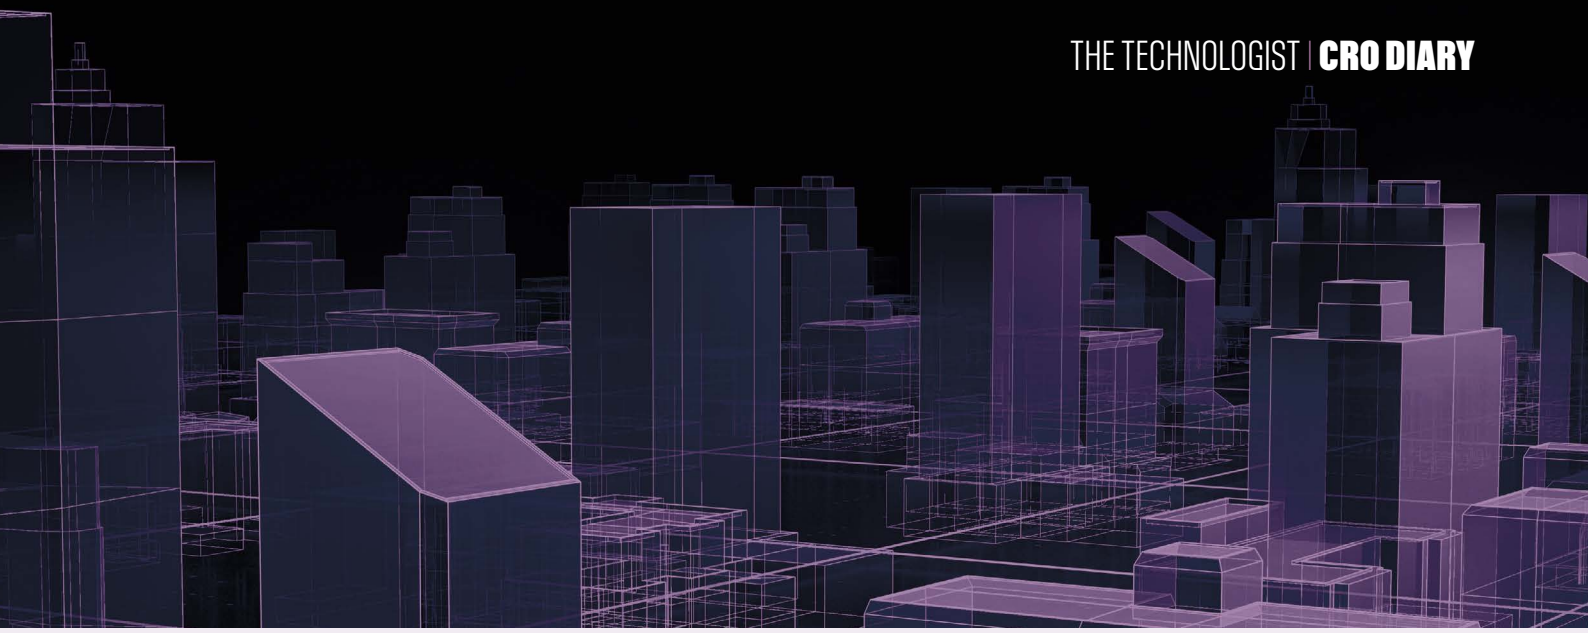

team that understands risk as well as technology, as they can't single handedly do all the work. Unfortunately, this is a rare combination of competencies to find in the talent market. This leaves CROs with a difficult choice between costly external hires and time-consuming reskilling programmes, with no guarantee.

Phased upskilling

Success in making technology a part of your core engine, as a CRO, demands patience and iteration. Top-performing technologist CROs adhere to a structured journey that begins with taking stock of where they stand, not with a rush to action, be it training, certifications or a hiring frenzy. They map existing capabilities against what the business needs, prioritise quick wins and assemble cross-functional coalitions. As PwC points out, risk transition efforts must be business-led, built around business requirements rather than risk functions needs.

The next step is to experiment, understand how value is generated and whether it is sustained, and this can happen via pilots in select business units or geographies. These limited-scope initiatives serve as proof of value, when successful, and learning nuggets, when not meeting expectations. Through these, the technologist CRO overcomes budgeting challenges and instils scalable confidence across the organisation.

> When risk considerations become naturally embedded at all levels of digital transformation, organisations can innovate with confidence, not fear or negligence

The most successful changes view technology adoption as an iterative journey rather than a one-off event. They set up pipelines that continuously review new technologies, without fear or concerns, in secure environments such as sandboxes.

Future perspectives

The changing role of the CRO, from operations manager to strategic advisor to technologist, is not just a temporary phase. It's a permanent evolution that reflects the underlying change of the business world. Risk and compliance professionals who fail to adapt to this new reality will be left behind, and those who thrive will play a critical role in shaping the future of business.

Quantum computing, for example, has the potential to dramatically enhance risk modeling and scenario analysis, while also disrupting encryption. Advanced AI capabilities, evolving beyond pattern recognition to true reasoning, will become an essential tool for risk management, the main execution engine for up to 90% of oversight tasks. In the meantime, the regulatory environment will continue to evolve, maybe not as fast as technology, but with initiatives like the EU AI Act, eventually making technology risk management just as regulated as financial risk management is today.

As for the risk management function, its boundaries will continue to blur, becoming more and more integrated with other functions. If it doesn't evolve to justify its relevance, it may become a victim of the digital revolution.

Technologist CROs must embrace the role of technologist and proactively seek ways to build the capabilities, frameworks and cultures that will allow their businesses to innovate responsibly.

> Success in making technology a part of your core engine, as a CRO, demands patience and iteration

Adam Ennamli, Chief Risk Officer, General Bank of Canada

Images: iStock

TRANSLATING CYBER RISK INTO FINANCIAL IMPACT

Are we finally closing the communication gap?

In 2025, boards and CFOs find themselves in an unprecedented situation where they urgently need to know the reality of cyber risk in clear financial terms. The SEC cyber disclosure regulations are more demanding in terms of transparency, and AI-enabled attacks have increased the frequency and cost (Proudfoot, Cram, Madnick, & Coden, 2023). But regardless of the increasing number of quantification tools and industry frameworks, the same question sounds within governance circles: are we bridging the communication gap on cyber teams to the aisles of financial decision makers or prettying it up with more graphics?

The gap exists. As noted by Proudfoot et al. (2023), the directors in companies with frequent cyber briefs tend to regard cyber oversight as a compliance measure instead of one of the pillars of strategy. Reports are read, charts are delivered; however, the board has been relegated to perceive risk based upon how it has been explained to them, rather than its actual state in the real world. Some of the misconnection lies in the tools used. FAIR has emerged as the de facto standard in expressing potential financial risk associated with cyber threats in terms of likely financial risk. It has unmistakable advantages regarding organisation, coherence, and the common language of risk managers and CFOs. However, according to Liu and Babar (2024) in their systematic review, FAIR relies on the use of historical data, which is backward-looking. That limitation is significant in a world where attack surfaces change more rapidly than actuarial tables. It is not just an intellectual weakness. According to Adejumo and Ogburie (2025), the emergence of decentralised finance (DeFi), blockchain-based settlements, and API integrated earnings has brought up attack vectors that do not have any compelling precedence as the financial sector embraces the use of blockchains and decentralised technology. These exposures cannot be cleanly modeled with historical probability. Nonetheless, FAIR output is habitually fronted to the boards as apprehending the whole window of financial threat. The accuracy is questionable, the precision is enticing.

Then, there is the more human issue of misaligned incentives. Quarterly performance is safeguarded since CFOs are rewarded. Technological investment in cybersecurity, particularly in quantification, continues to be viewed as an expense, rather than strategic capital. A study by Liu and Babar (2024) monitors the extent to which breaches in stock markets ex-post are punished, and the fines imposed seldom stem from preemptive investments before the breach.

This is echoed by Proudfoot et al. (2023), who state that boards will finance the more advanced risk modeling when there is a regulatory change or a significant event. The effects of such a short-term approach are evident when we consider an example such as Equifax. Its breach in 2017 resulted in market loss across the market cap by more than 4 billion across several days. This risk existed way before the breach; however, the board's attitude reflected the typical trend: reactive rather than predictive. Whereas Knight and Nurse (2020) look at the failure of communication that occurred afterward, the true failure was farther upstream, the failure to tell a story about systemic cyber risk that is financially meaningful enough to necessitate action in advance.

The other half of the gap is communication. Figures in isolation seldom swing a board; they must be related. Knight and Nurse (2020) devised an approach to corporate communication following events, focusing on transparency, timeliness, and personalised messages. That structure also works pre-incident, although herein resides a delicate risk. When executed properly, with the aid of storytelling, before a cyber risk breach, it can be made concrete but also simplistic. Boards prefer stories because they make risk relatable. Occasionally, those stories are molded by risk teams to obtain budgets or compliance acceptance. However, a tale maximised to persuade does not lose a critical part when the story is maximised towards precision. The risk profile the board is acting upon turns into a selected reality. The communication gap is never bridged; it is handled with a more graceful and refined touch. This delicate translation is also complicated by emerging technology. According to Adejumo and Ogburie (2025), current weaknesses in financial APIs are already being used by AI-driven attacks- this is not a hypothetical situation. In multiple documented instances, incorrectly implemented APIs exposed transaction systems to credential stuffing attacks that did not rely on any legacy controls. There is no standardised method of incorporating such fluid, changing risks into FAIR models.

Beyond the short term, quantum computing is a looming disruptor. If quantum capabilities cannot survive extant cryptographic standards, the assumptions underlying probability-based financial modeling in use today will fade.

> **Technological investment... continues to be viewed as expense, rather than strategic capital**

Case studies enforce the stakes. The 2015 breach at TalkTalk, which was used as a case in the study by Knight and Nurse (2020), not only led to fines but also caused a loss of trust by the customers and loss of brand equity, and even a political backlash. The technical setbacks were awesome, and the boardroom difference was fatal. The executives were taken by surprise during live interviews, failing to explain whether the stolen information was encrypted. It was not only a communication problem, but also a quantification problem. The financial and technical risks had not been translated to literacy at the board level. When these failures seem old-fashioned, they are not. In 2025, the same vulnerabilities will be dealt with silently through API breaches, AI-enabled phishing, and cloud integration failures. Adejumo and Ogburie (2025) demonstrated that banking APIs are still among the most utilised vectors, with the quantification of this risk having not reached maturity.

Are we bridging the communication gap, then? Improvement exists. The availability of standard language under initiatives such as FAIR, an upped ante by regulators, and investment into meaningful oversight by some boards is now in place. However, loopholes persist: rewards are not as skewed in favor of the long term, stories are potentially too simplistic, and models are not as fast as the disruption in technology. Liu and Babar (2024) suggest a transition to dynamic quantification, where real-time threat intelligence is combined instead of utilising solely backward data. Proudfoot et al. (2023) insist on making cyber oversight an inbuilt governance task. Adejumo and Ogburie (2025) emphasise that technological acceleration does not wait and gives models time to adapt.

Knight and Nurse (2020) remind us that clear, non-deceptive, and context-abundant communication is risk management. A more attractive dashboard cannot fill the communication gap. It will demand dynamic, incentive-driven, brutally honest quantification by CFOs and boards. Until then, the gap will still be there, smooth on the surface but profound in depth.

Amena AlBasher, MSc (Cornell), risk management and GRC expert

References

Adejumo AP and Ogburie CP. (2025) The role of cybersecurity in safeguarding finance in a digital era. World Journal of Advanced Research and Reviews, 25(3), 1542-1556.

Knight R and Nurse JR. (2020) A framework for effective corporate communication after cybersecurity incidents. Computers & Security, 99, 102036.

Liu C and Babar MA. (2024) Corporate cybersecurity risk and data breaches: A systematic review of empirical research. Australian Journal of Management, 03128962241293658.

Proudfoot JG, Cram WA, Madnick S and Coden M. (2023). The importance of board member actions for cybersecurity governance and risk management. MIS Quarterly Executive: 22 (4): 6.

Images: Shutterstock

DIGITAL ETHICS IN THE AGE OF AGENTIC AI

Four years since the publication of the IRM Digital Ethics Guidelines, lead author **Mark Turner CFIRM** takes a look at how evolving technology is shaping the digital ethics landscape

In autumn 2020, Clive Thompson invited me to join the IRM Professional Standards Committee to help develop guidance on digital ethics. As the former Chairman of the IRM Innovation SIG and with a degree in technology, I was well placed to lead the work. With input from many IRM members, we published the Digital Ethics Guidelines in March 2021.

Ethical debates at that time focused on responsible data use for the protection of people, society and the biological ecosystem. The fallout from the introduction of the General Data Protection Regulation in 2018 was still rippling across European businesses, and the blockchain was starting to cause disruptive waves within industry.

Against this backdrop, our team set out to develop three guiding principles, ethical equivalence, data security and bias management, in the hope that risk professionals could begin to ask appropriate questions of the people developing and deploying the technology, without needing to retrain as technical experts.

Since then, the ethical landscape has changed dramatically. The significant inflection point was the release of ChatGPT in November 2022, giving people and organisations unprecedented access to knowledge and creativity, while giving IT departments unprecedented headaches as everyone tried to use it without clear guidance. The rollout of competing LLMs (large language models) with ever-increasing capabilities has seen AI become integrated into many products and has spurred on a trillion-dollar industry shrouded in ethical controversy and debate – not least of which is the provenance of the training data used to create the LLM in the first place.

Today, the conversation is shifting from using standard AI ethically to ensuring AI itself *behaves* ethically. This is especially true with agentic AI, which is capable of autonomous planning, action and adaptation. The use of agentic AI increases the risk profile significantly and requires greater consideration of ethics than earlier technologies. Creating effective guardrails is now essential to ensure the decisions being made by the AI remain legal, as well as ethical, and that humans remain accountable for the actions of their agents.

Considering these technological developments, I have conducted a review of the original Digital Ethics Guidelines with specific regard to both standard and agentic AI. This has led to the creation of an addendum to the guidelines. As with the original publication, the addendum considers simple test questions against the three existing principles, which risk professionals can use to challenge technical experts to assure themselves that the risks are being considered and managed appropriately.

The new challenge for risk professionals is not just guiding human use of AI, but managing the ethical framework within which the AI itself operates.

Images: iStock

Data transfer from existing solutions into riskHive ERM

riskHive's ERM system is designed to meet evolving client needs, specialising in transitioning from spreadsheets to databases. It offers fast deployment on secure private cloud hosting or on-premises, operational and ready for configuration within 24 hours. Configuration typically takes one to six weeks. By replicating customers' practices, we reduce training and deployment time, accelerating return on investment and confidence in the new system. The riskHive ERM includes Monte Carlo simulation and analysis, covering costs, schedules and carbon emissions.

Contact:

Ian Baker
+44 (0)7781 8898 977
ian.baker@riskhive.com
www.riskhive.com
riskHive Software Services Ltd
Dilkush
Farlers End
Bristol BS48 4PG

Achieving GRC objectives with confidence

Symbiant is an award-winning GRC and audit platform designed to help organisations manage risk and achieve objectives with confidence. The platform is modular and agile, and easily integrates with existing structures.

Symbiant's integrated AI assistant analyses your data to identify hidden threats, predict control failures and understand how risks may cascade across your organisation. Proven in complex environments, Symbiant has been providing powerful, flexible and affordable GRC solutions to organisations of all sizes since 1999.

Contact:

Andrew Birch
+44 (0)20 3821 1993
andrewb@symbiant.co.uk
www.symbiant.co.uk
Symbiant
20-22 Wenlock Road
London N1 7GU