
AXA
XL Insurance
Building Resilient AI
Managing AI risk through
governance, security and resilience
Produced in partnership with S-RM

Contents

Foreword	3
Executive summary	4
AI has entered the enterprise	5
The AI risk landscape	6
Navigating current and future regulation	12
From risk to resilience	15
Building secure AI by design	16
The AXA XL approach	17
Final reflections	18
Disclaimer	19

Foreword

Trust is the thread that runs through every successful business. It underpins relationships, enables innovation and gives organizations the confidence to embrace change. Today, that trust is increasingly being placed in Artificial Intelligence. However, the governance needed to support it is struggling to keep pace.

AI has moved well beyond isolated pilots or productivity tools. It is fast becoming part of the systems and processes organizations rely on daily to make decisions, to serve customers, to improve productivity and to drive innovation.

The rate at which developments are now moving has created significant opportunities but has also introduced an entirely new category of risk for enterprises. **As AI becomes woven into critical business functions, data governance, protecting the integrity of AI models, third-party risk and regulation are becoming very difficult to manage.** This shift is being felt everywhere, with AI now recognized as one of the world's most significant business challenges.

While organizations continue to expand investment in the capabilities AI offers, it is clear the decisions being made today will shape resilience, customer confidence and competitive positioning for years to come.

For boards and executive teams, AI security needs to be viewed as more than just a technology issue. It demands strong leadership, clear accountability, robust governance and absolute focus on business resilience. Those organizations set to gain the greatest value from AI investments will be those that understand where it is being used, establish clear oversight throughout its lifecycle, and build the necessary safeguards from the outset.

This report explores the evolving AI risk landscape and the practical steps that can be taken to adopt AI with confidence today. Drawing on insights from AXA XL, S-RM and leading industry research, it defines how companies can scale AI in a trusted and responsible way.

Jonathan Salter
Head of Risk Consulting

Rebiah Bardot-Girard
Head of Cyber Risk
Consulting Services

Executive summary

The AI security landscape is entering a new phase. Conversation is expanding far beyond individual models to include the systems, governance and dependencies that shape how AI is deployed across enterprises today. At the same time, rapid advances in agentic AI, evolving regulation and more sophisticated cyber threats are impacting how organizations think about risk.

Five developments in particular are shaping what business leaders must now prioritize:

New categories of cyber risk are emerging

AI is something of a double-edged sword for cyber teams. On the one hand it helps strengthen defenses yet simultaneously provide attackers with new capabilities. Companies are also grappling with risks unique to AI itself: data leakage, poisoned or manipulated models, prompt-based attacks, model outputs that downstream systems trust without validation, retrieval systems that expose information beyond a user's normal access rights, and the security of autonomous agents acting with their own credentials. These risks are no longer theoretical. Synthetic voice and video are already being used to authorize fraudulent payments, and organizations have been held to statements made by their own customer-facing AI systems.

Governance is becoming a competitive advantage

With AI now embedded across business operations, governance is moving beyond compliance to become a strategic business capability. Organizations with clear accountability and strong oversight are far more likely to scale technology successfully and keep pace with emerging risks, in turn building greater trust with audiences. Oversight must also cover the AI that organizations do not know they are running. Employee use of public AI tools, and AI features switched on inside existing software, are today the most widespread source of exposure.

AI needs to be managed across its lifecycle

Managing AI risk requires more than securing individual models. Visibility across the full lifecycle, from data collection and model development through to deployment and monitoring, will be needed. That also includes third-party providers recognizing that AI ecosystems are interconnected, and that accountability cannot be outsourced. Lifecycle management also means being able to reconstruct what happened. Without retained records of prompts, responses and tool calls, incidents involving AI systems cannot be investigated with the same confidence as traditional IT events.

Certainty on mitigating and insuring against AI risks is needed

Companies adopting AI will need support to understand the vulnerabilities it may unintentionally create. Risk management and insurance have an important role to play in helping businesses have clarity surrounding exposure, build resilience before incidents occur, and help with effective recovery when disruption cannot be avoided. Clarity is also needed on how AI-related losses interact with existing lines of cover. A fraudulent payment authorized through a synthetic video call, a defective AI-generated output causing third-party loss, and an intrusion enabled by an AI agent each engage different policies.

AI regulation is accelerating globally

Governments and regulators are moving quickly to establish concrete expectations for responsible development and use. Although approaches differ by region, one thing is universal: businesses will need to demonstrate that AI is being deployed responsibly.

Key Learning

AI is creating new opportunities alongside new forms of risk. Those that treat AI governance as a strategic business priority, rather than a compliance exercise, will be best positioned to help their organization scale AI securely, adapt to changing regulation and realize its long-term value.

AI has entered the enterprise

Across industries, organizations are embedding AI into core business processes, using it to support customer interactions, streamline operations, analyze data and inform decision-making.

McKinsey's latest State of AI survey found that 88% of organizations now report regularly using AI in at least one business function, up from 78% a year ago, demonstrating how rapidly AI has become part of the technology organizations rely on daily.[1]

It is also becoming more autonomous, with businesses moving beyond tools that generate content to systems that retrieve information, interact with other applications and take action with limited human input. While this can bring significant gains in efficiency it also increases the potential impact of security failures and operational mistakes.

[1] McKinsey & Company, The State of AI, 2026

With AI now integral to many business operations, failures can spread rapidly across systems and be harder to detect before they have consequences.

Success in the next phase of AI adoption will depend on how well organizations manage the wider ecosystem surrounding their models. Understanding where AI is used, how it interacts with business processes and who is accountable for its outcomes, is becoming fundamental to building resilience surrounding its use.

As companies become more dependent on AI-driven ecosystems that create new dependencies across people, technology and third-party providers, these developments signal a fundamental shift that is changing both the scale of opportunity and the consequences of failure.

88%

of organizations regularly use AI in at least one business function

- McKinsey 2026

Key Learning

AI is becoming part of the operational fabric of modern organizations, woven into increasingly interconnected technology ecosystems. Businesses are rapidly inheriting a level of complexity few governance models were designed to manage. Security now needs to extend beyond protecting individual AI models to governing the data, systems, suppliers and business processes that enable AI to operate safely at scale.

The AI risk landscape

Organizations are inheriting an amplified version of the traditional cyber risks they once faced. These risks are not isolated technical issues. Instead, they have the power to influence governance, security and resilience simultaneously. Meanwhile, existing challenges around identity, data, supply chains and operational resilience are becoming increasingly complex.

The risks are not new. Data theft, malicious insiders and software vulnerabilities have long been challenges. However, what AI changes is scale and complexity. Today's systems process more information, interact with more applications and make or influence more decisions with greater autonomy than ever before.

Enterprises are starting to recognize the need to address this new world proactively. The World Economic Forum's latest Global Cybersecurity Outlook reveals *64% of organizations now have processes in place to assess the security of AI tools before deployment*[2], up from 37% the previous year. For business leaders, the implication is clear: AI security needs to be considered from the outset rather than after systems are already in production, from a global governance perspective rather than only a localized technical one. The first step here is about understanding how the AI risk landscape is changing, where new vulnerabilities are emerging, and what they mean for business.

64%

of organizations now have processes in place to assess the security of AI tools before deployment

[2] World Economic Forum, Global Cybersecurity Outlook, 2026

"From our work with organizations responding to cyber incidents and strengthening resilience, one lesson is becoming increasingly clear: AI risk is not emerging in isolation. It is being layered onto existing weaknesses in identity management, data governance, supplier oversight and incident readiness. For many organizations, AI is accelerating exposure faster than operating models are adapting."

Joani Green,
Global Technical Director,
Cyber Security at S-RM

Combatting data leakage and unauthorized access

AI systems rely on access to large volumes of data to generate accurate outputs, often drawing on commercially sensitive information and intellectual property. As AI is embedded more deeply into the everyday, that data becomes even more valuable and, in many cases, more widely accessible than ever before.

Without strong governance, sensitive information can easily find its way into prompts, outputs or third-party applications, creating risks that go far beyond cybersecurity. As AI gives more systems and applications access to sensitive data, controlling who and what can access that information becomes increasingly important. Attackers are also targeting the identities and access privileges that provide a route to that data.

IBM's 2025 X-Force Threat Intelligence Index found that more than 30% of incidents investigated involved the theft or misuse of credentials, reflecting a shift from traditional software exploitation toward identity-based attacks.[3]

[3] IBM, X-Force Threat Intelligence Index, 2025

Companies need more than ever a very clear understanding of where data is used, how it moves between systems and who can access it. That visibility is essential for managing risk effectively, in particular if sensitive or personal data are involved.

From an insurance perspective, organizations that can show strong data governance, robust access controls and clear oversight of AI systems are far better positioned to reduce exposure. For business leaders, protecting data is no longer simply an information security issue. It is becoming fundamental to maintaining customer trust and ensuring AI-generated decisions remain reliable.

Key Learning

Effective AI governance starts with understanding what data AI systems can access, where it resides and how it is protected, wherever it is used.

Data poisoning

In the architecture of AI, data is one of the core parts of the model. Data poisoning occurs during the training phase when malicious actors deliberately inject corrupted fragments into training sets to steer outputs and shape the model's behavior. It aims to degrade overall performance or unlock hidden behaviors (backdoors) under certain conditions.

When AI outputs are manipulated from the beginning, it is often hard to detect after deployment and it is costly to reverse without full retraining on clean data. The impact goes beyond performance: it erodes trust, undermines safety and governance, and can introduce regulatory or business risks.

Organizations need to ensure that rigorous AI data provenance tracking is done from the early stage of the project, that robust training techniques are defined and implemented, with continuous monitoring of the whole training process to catch signs of early manipulation, including human monitoring steps.

Key Learning

Embed good data governance practices from the first stage of an AI project and adapt them as needs evolve throughout the project lifecycle.

Prompt injection attacks and adversarial activities

The ability to manipulate the instructions AI systems receive via prompts or to deliberately mislead the model into wrong predictions through crafted malicious input data (images, text, audio) is another area companies must be mindful of. Influencing behavior or bypassing intended safeguards to produce unintended or harmful outputs can prove very damaging.

When AI-generated outputs are manipulated, ramifications can range from exposing confidential information to influencing business decisions to disrupting customer-facing services, particularly when AI is woven into everyday operations.

Organizations need to ensure that data and system security testing, monitoring and human oversight evolve alongside the use of AI-enabled technology. This means treating AI systems with the same discipline as any other critical business system.

Because AI takes a more active role in customer interactions, organizations cannot afford to treat prompt injection as an isolated technical vulnerability. It is becoming another enterprise risk that requires constant oversight.

Key Learning

AI outputs should be verified, monitored and supported by appropriate governance, rather than assumed to be accurate or secure by default.

Model manipulation and model theft

For many companies, AI models themselves are becoming valuable business assets, reflecting a significant investment and containing valuable proprietary knowledge. As this importance grows, so too does their attractiveness to cyber criminals. Protecting AI models needs to become as strategically important as customer data or critical infrastructure.

When AI models are manipulated, the integrity of AI-driven processes can also be compromised. If stolen, attackers may gain access to valuable intellectual property or be able to replicate capabilities that took years to develop. This can severely undermine confidence in AI-enabled services and cause lasting damage to innovation.

Protection here means looking beyond technical controls alone. Organizations should understand where models are deployed, who can access them, how changes are managed and what safeguards are in place to detect unauthorized activity.

Key Learning

AI models should be protected and governed as critical business assets.

Hallucinations and unreliable outputs

Generative AI is capable of extraordinary things but can be confidently wrong. Hallucinations stem from gaps in training data, ambiguous prompts or limitations within the model itself. Recent research from Stanford University's 2026 AI Index found that across 26 leading AI models, hallucination rates ranged from 22% to 94%.[4]

A single inaccurate answer might not matter in isolation. But when AI is helping to shape business decisions, support customers or run operational processes, the stakes are much higher. Errors can spread quickly, leading to regulatory scrutiny, financial loss and lasting damage to trust.

Rather than assuming AI outputs are inherently reliable, organizations must establish appropriate validation processes applied at regular frequency and not as a one-time action and define where human visibility is required and how this evolves over time.

Companies need confidence not simply in what AI can produce, but in knowing when its outputs can be trusted, and when human judgment should prevail.

Key Learning

Trust in AI depends on verifying outputs, not assuming accuracy.

[4] Stanford University, AI Index Report, 2026

Dependency on unverified models or third-party models

Relying on subcontractors or external vendors for AI models introduces supply-chain and governance challenges. Limited visibility into how the model was trained, what data was used, and how updates are managed can increase undesired behaviors and policy non-compliance. As for any other IT supplier, dependency on an AI vendor creates bottlenecks during outages, roadmap shifts, or price changes, and can constrain strategic alignment with internal standards.

On top of that, unverified or poorly vetted models can expose organizations to potential backdoors or unsafe behaviors from the model. Policy, regulatory or ethical misalignment is another major concern, especially when models drift after deployment or are updated without thorough testing. Licensing ambiguities and lack of transparency about training data can also complicate compliance and auditability.

Rather than assuming that external models are safe by default, a robust governance and due diligence for all models and data must be implemented. Clear provenance and auditable activity must be required, security and functional validations must be enforced before deployment, and strong contractual protections (SLAs, transparency clauses, disclosure obligations for updates) must be set.

Key Learning

AI vendors should be governed as critical third parties and both the vendor, and the AI model should be processed through due diligence controls.

Combatting shadow AI

Shadow AI refers to situations where employees or business units use AI tools or services that have not undergone formal security, privacy, and governance reviews.

These tools may process sensitive data or interact with corporate systems, often bypassing procurement processes and risk assessments. The risks include data leakage or exfiltration, and the introduction of insecure configurations or hidden backdoors that expand the organization's attack surface.

Operationally, shadow AI can create uncontrolled data flows, licensing and cost problems, inconsistent outputs, and auditing challenges. Strategically, it undermines risk management, increases vendor dependence, and complicates incident response and business continuity.

To address this, companies should block unauthorized tools by default and establish a clear policy and approval workflow for AI tools, as well as applying data classification and least-privilege access principles. Regular training of employees and regular cybersecurity controls on the attack-surface can also be key to avoiding such situations.

Key Learning

More than ever, shadow AI mitigation is key.

"The most difficult risks are often not the most visible ones. Businesses may have formal governance around strategic AI programs, while unmanaged AI use continues through public tools, embedded software features or third-party services. This creates a gap between perceived control and actual exposure. Closing that gap starts with a practical inventory of where AI is used, what data it can access, and where it can take or influence action."

Joani Green,
Global Technical Director,
Cyber Security at S-RM

Autonomous AI Agents

Rather than simply generating content or responding to prompts, agentic AI systems are now able to plan tasks and interact with applications, carrying out actions with limited human intervention. The potential here is huge, allowing companies to automate complex workflows.

But with the increased level of autonomy comes a set of new governance challenges. AI systems require clearly defined permissions and appropriate safeguards to ensure they operate within the right boundaries. Without effective oversight, organizations run the risk of the wrong data being accessed or interactions with connected systems that inadvertently create new vulnerabilities.

McKinsey's survey found that 23% of organizations are already scaling agentic AI in at least one business function, while a further 39% are experimenting with AI agents[5]

The pace of adoption here highlights just how quickly these systems are moving into the day-to-day. Companies need to balance the benefits of automation with appropriate levels of accountability and human control. The more autonomy organizations delegate to AI, the greater the importance of clearly defining where responsibility sits for its actions.

Key Learning

The more autonomy businesses give AI, the greater the need for effective governance, clearly defined permissions and meaningful human oversight.

Viewed individually, data leakage and poisoning, shadow AI, prompt injection, model manipulation and theft, hallucinations and autonomous agents can appear as separate technical challenges. But taken together, they point toward a much bigger shift in the enterprise risk landscape – a need for governance that reflects the entire AI ecosystem rather than addressing individual risks in isolation.

However, understanding risks is just one part of the challenge. As adoption accelerates, organizations must also navigate a regulatory environment that is redefining expectations. Managing risk now requires companies to take a more integrated approach than ever, one that considers how technology, people, processes and third-party providers all influence each other.

[5] McKinsey & Company, The State of AI 2026

Navigating current and future regulation

The global regulatory environment for AI is evolving quickly as policymakers move toward defining the guardrails that will determine how it is governed in practice. The OECD[6] has highlighted this change, noting that governments are moving well beyond broad AI strategies, and toward clearer rules with greater accountability to support safe adoption. While individual approaches differ, the direction of travel is becoming clear. Organizations are expected to understand where AI is being deployed, establish appropriate governance and demonstrate accountability across the entire lifecycle.

Rather than responding to each new regulation in isolation, organizations should instead focus on building governance frameworks that are flexible enough to adapt as requirements evolve.

[6] OECD, Governing with Artificial Intelligence

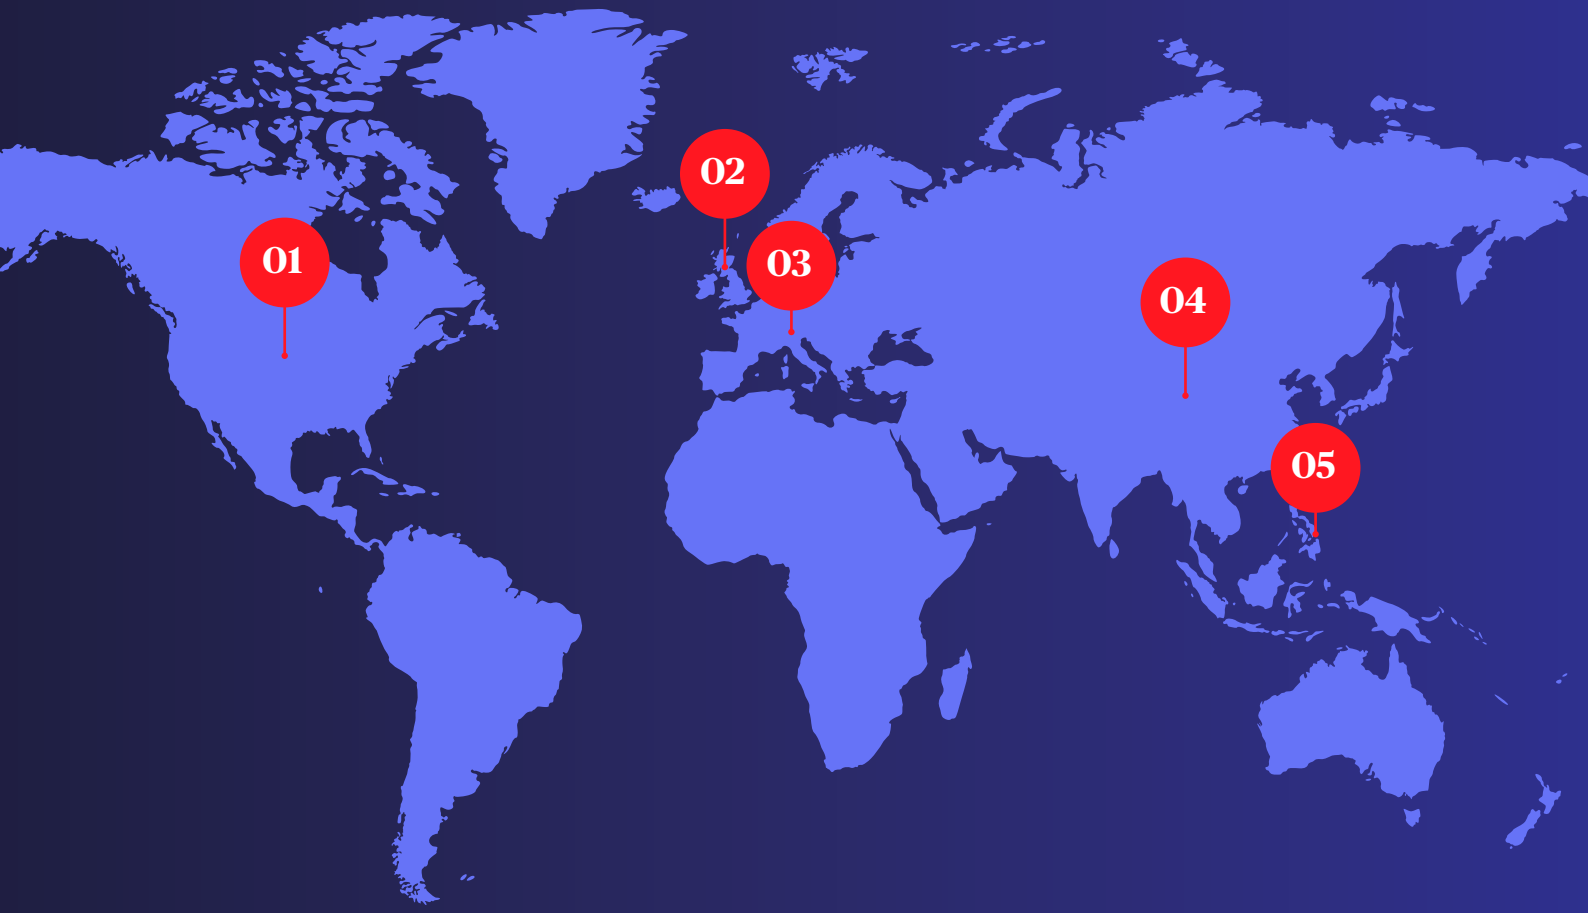

Global regulatory developments

Businesses operating internationally will need to navigate a patchwork of AI regulation. Although many governments share similar objectives, the pace, scope and focus of regulation will vary significantly by region:

01 **The United States** faces a more fragmented and uncertain regulatory picture, with AI policy continuing to evolve at both federal and state level. Differences between these approaches, and ongoing questions over how federal and state requirements will interact, mean businesses need to prepare for a landscape that could continue to shift significantly.

02 **The United Kingdom** has adopted a principles-based approach, seeking to encourage innovation while applying existing regulatory frameworks to AI. The focus is on demonstrating responsible governance, transparency and accountability rather than introducing entirely new rules.

03 **Europe** continues to set the pace through the EU AI Act, which categorizes AI systems according to the potential risks they pose to individuals, including their safety and fundamental rights. Higher-risk applications face greater requirements around oversight, transparency and accountability. Although the legislation applies within the EU, its influence is already extending further, encouraging businesses to strengthen how they govern AI across international operations.

04 **China** has developed its own regulatory framework for AI, with requirements spanning areas such as generative AI services, algorithms, data and content. Its approach continues to evolve alongside the country's rapidly growing AI sector, making developments in China increasingly important for businesses operating internationally or relying on Chinese AI technologies and providers.

05 **Across the wider Asia-Pacific region** regulatory approaches continue to mature at different speeds, with common themes emerging around responsible AI, secure data management and effective oversight. Organizations operating internationally will need an approach that can adapt to differing requirements as the regulatory landscape matures.

What AI regulation means for business

Although regulatory approaches vary, they converge around similar expectations. Organizations need to understand where they're using AI, identify high-risk applications, establish clear ownership and keep clear records of deployment throughout its lifecycle. Explainability and accountability are becoming just as important as technical performance, particularly where AI influences decisions that affect customers, employees or critical business operations. For business leaders, this reinforces the need to view AI governance as an enterprise responsibility.

Compliance as a business enabler

Organizations can view regulation as a constraint on innovation, but the reality is that effective governance can actually become a competitive advantage. Deloitte's latest State of AI in the Enterprise research[7] found that only one in five organizations has a mature governance model for autonomous AI agents, highlighting an opportunity for businesses that establish robust governance early. Organizations that can demonstrate clear ownership, consistent governance and transparency will be better placed to meet new regulatory requirements and earn the trust of customers and partners.

Knowing what regulators expect is one thing, but building the organizational capability to meet those expectations consistently is another. This is where governance frameworks become valuable, not only as compliance exercises, but as practical tools for embedding resilience into everyday operations. Organizations that invest in the process today are unlikely to see regulation as a burden tomorrow.

"AI also changes the evidence organizations need after an incident. Traditional logs may not be enough to understand what happened, particularly where prompts, outputs, model behavior, retrieval sources or agent actions are involved. Organizations should consider what information they would need to reconstruct an AI-related event, demonstrate accountability and support legal, regulatory or insurance processes."

Joani Green,
Global Technical Director,
Cyber Security at S-RM

[7] Deloitte, State of AI in the Enterprise

From risk to resilience

As this technology becomes more and more embedded, organizations will need practical ways to manage risk consistently as systems evolve. This means moving beyond individual technical controls toward governance models that bring various teams within a business around a common approach.

No single framework can address every challenge associated with AI. However, established standards such as the NIST AI Risk Management Framework, ISO/IEC 42001, the OWASP Top 10 for Large Language Model Applications[8] and Top 10 for Agentic Applications[9] provide a valuable foundation. Rather than prescribing individual technical controls, they give organizations a consistent way to assess risk, assign ownership and govern AI. Together they support a more consistent approach to identifying high-risk AI applications, managing third-party dependencies, monitoring AI performance and responding to incidents as systems evolve over time.

[8] OWASP, Top 10 for Large Language Model Applications

[9] OWASP, Top 10 for Agentic Applications for 2026

Ultimately, resilience depends less on adopting a particular framework than on embedding good governance into everyday decisions. Maintaining visibility across AI systems and continuously reviewing performance will ensure companies are equipped to manage emerging risks while continuing to innovate with confidence.

Looking ahead

Governance frameworks provide direction, but they only create value when translated into everyday practice. The next challenge for businesses will be embedding those principles into how AI systems are designed, developed, deployed and managed – making security and resilience part of the AI lifecycle rather than an activity that happens after deployment.

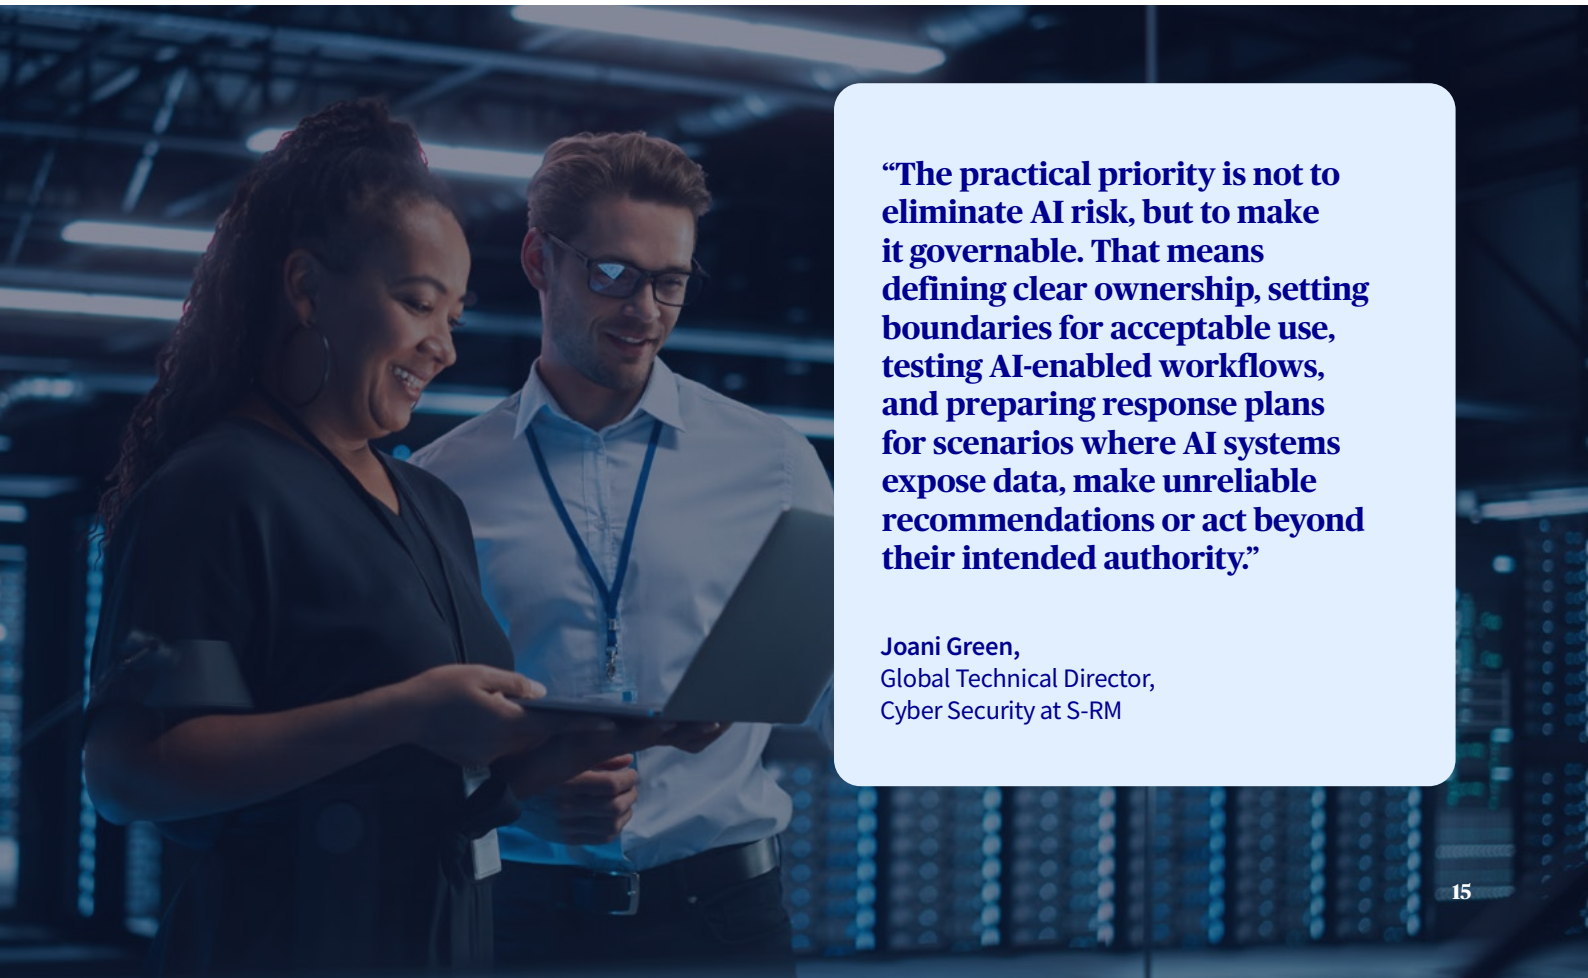

"The practical priority is not to eliminate AI risk, but to make it governable. That means defining clear ownership, setting boundaries for acceptable use, testing AI-enabled workflows, and preparing response plans for scenarios where AI systems expose data, make unreliable recommendations or act beyond their intended authority."

Joani Green,
Global Technical Director,
Cyber Security at S-RM

Building secure AI by design

Organizations that build security into AI from the outset will be on better ground to scale adoption and strengthen resilience over the long term, as regulations evolve. Technology alone cannot achieve this.

The following five priorities provide a practical foundation for identifying risks early and enabling AI to be adopted with confidence:

Better data governance

Strong AI governance begins with understanding the data that powers AI systems. Organizations should know what information AI can access, where it originates and how it is protected to reduce both regulatory and security risks.

Securing models and applications

AI models should be treated as critical business assets. Security testing, regular monitoring and effective change management can all help maintain confidence in AI-enabled services, and reduce the risk of manipulation or unauthorized access.

Building resilience across the AI ecosystem

Most organizations rely on a network of cloud providers, technology partners and AI vendors. Understanding these interdependencies helps organizations manage their network and identify any potential weaknesses early.

Strengthening identity and access controls

As systems start to act with more autonomy, identity becomes more important. Organizations should ensure both users and AI systems only have access to the data and applications necessary to perform their intended role, strengthening governance while limiting the impact of compromised accounts.

Continuously monitoring and adapting

AI is not static, with models evolving constantly, regulations changing and new threats emerging. Regular monitoring, governance reviews and testing can help companies detect issues early, demonstrate ongoing compliance and adapt as the technology and risks surrounding it evolve.

Security-by-design is ultimately about giving organizations greater confidence to innovate and scale AI securely and responsibly, with trust built into the process. That depends on embedding governance, security and resilience throughout adoption, supported by an organizational culture and, where needed, trusted partners that can help navigate an increasingly complex risk landscape.

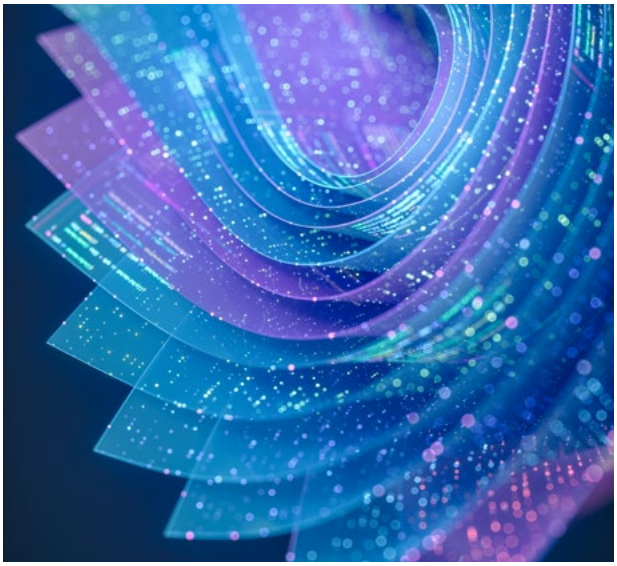

The AXA XL approach

As AI becomes embedded in core business processes, organizations need to understand how it changes their overall risk profile. AI should not be viewed only as a technology risk, but as an enterprise resilience issue connecting cyber security, data governance, third-party dependency, operational continuity, regulatory accountability and insurance.

A resilient approach starts with visibility. Organizations need to know where AI is being used, what data it can access, which decisions it influences and which third parties it depends on. This includes formal AI initiatives as well as less visible uses of public tools, embedded software features and autonomous agents.

AXA XL believes AI risk management should be grounded in practical business scenarios. Organizations need to consider what could happen if a model produces an unreliable output, an AI agent acts beyond its authority, sensitive data is exposed through a prompt, or a critical AI provider becomes unavailable. These scenarios help translate technical vulnerabilities into business consequences such as financial loss, operational disruption, liability, regulatory scrutiny and reputational harm.

They also support clearer conversations around resilience and insurance. AI-related losses may not fall neatly into one category of risk. A deepfake-enabled payment fraud, an AI-related data breach, a defective output causing third-party loss or an outage at a critical provider may each engage different parts of an organization's risk and insurance program. Understanding these scenarios early helps clarify where existing controls and coverages may respond, where uncertainty remains and where further mitigation may be needed.

AXA XL brings together risk consulting, cyber expertise and insurance insight to support this broader view of AI resilience. Working alongside specialist partners such as S-RM, the focus is on helping organizations identify exposure, strengthen governance, prepare for emerging incident scenarios and make more informed decisions about risk transfer.

The objective is not to slow innovation, but to help organizations adopt AI with confidence: understanding where risks sit, managing them through stronger governance and resilience, and assessing where they may be transferred or insured.

Final reflections

Adoption of AI shows no sign of slowing down, so the biggest challenge for businesses today is making sure it can be trusted at scale.

This means understanding how your company uses AI, embedding governance throughout its lifecycle, and ensuring security is evolving alongside technology. Organizations that build these foundations today will be better-equipped to adapt as threats evolve and regulations advance.

Building trusted AI also requires tight collaboration across leadership, risk, legal, security and operational functions, supported by partners who understand both the opportunities AI creates and the challenges it presents.

Those that embed governance, security and resilience in the way they operate will be the ones that unlock AI's true potential and adapt with confidence, as this powerful technology continues to evolve.

Disclaimer

This report is the result of analysis from a wide variety of sources including publicly available reports, insights and analysis, and proprietary analysis and insights by AXA XL, and S-RM based on their internal data. A non-exhaustive list of publicly available sources is included in the report but does not constitute the sole source of insights that made this report possible.

The analysis presented in this report is based on information available at the time of writing. When future evolution is expressed, it does not constitute prediction of the evolution of trends and threats, but the expression of the possible future context in which organizations may have to operate.

Global Asset Protection Services, LLC, XL Catlin Services SE and their affiliates ("AXA XL Risk Consulting") provide loss prevention and risk assessment reports and other risk consulting services, as requested. In this respect, our property loss prevention publications, services, and surveys do not address life safety or third-party liability issues. This document shall not be construed as indicating the existence or availability under any policy of coverage for any particular type of loss or damage. The provision of any service does not imply that every possible hazard has been identified at a facility or that no other hazards exist. AXA XL Risk Consulting does not assume, and shall have no liability for the control, correction, continuation or modification of any existing conditions or operations. We specifically disclaim any warranty or representation that compliance with any advice or recommendation in any document or other communication will make a facility or operation safe or healthful, or put it in compliance with any standard, code, law, rule or regulation. Save where expressly agreed in writing, AXA XL Risk Consulting and its related and affiliated companies disclaim all liability for loss or damage suffered by any party arising out of or in connection with our services, including indirect or consequential loss or damage, howsoever arising. Any party who chooses to rely in any way on the contents of this document does so at their own risk.

The information contained herein is intended for informational purposes only. Insurance coverage in any particular case will depend upon the type of policy in effect, the terms, conditions and exclusions in any such policy, and the facts of each unique situation. No representation is made that any specific insurance coverage would apply in the circumstances outlined herein. Please refer to the individual policy forms for specific coverage details. This summary does not constitute an offer, solicitation or advertisement in any jurisdiction, nor is it intended as a description of any products or services of AXA XL.

AXA XL is a division of AXA Group providing products and services through three business groups: AXA XL

Insurance, AXA XL Reinsurance and AXA XL Risk Consulting. In the US, the AXA XL insurance companies are: AXA XL Insurance Company Americas, Greenwich Insurance Company, Indian Harbor Insurance Company, XL Insurance America, Inc., XL Specialty Insurance Company and AXA XL Excess & Surplus Lines Insurance Company. In Canada, insurance coverages are underwritten by XL Specialty Insurance Company – Canadian Branch. In Bermuda, the insurance company is XL Bermuda Ltd. Coverages may also be underwritten by Lloyd's Syndicate #2003. Coverages underwritten by Lloyd's Syndicate #2003 are placed on behalf of the member of Syndicate #2003 by Catlin Canada Inc. Lloyd's ratings are independent of AXA Group. Not all of the insurers do business in all jurisdictions nor is coverage available in all jurisdictions.

axaxl.com