



How can
AI build
resilience

for the risks
we are yet
to imagine?

Global Risk
Transformation Series



The better the prompt.
The better the answer.
The better the world works.



Shape the future
with confidence

REIMAGINE RISK

Foreword

Welcome to the third installment of the EY Global Risk Transformation Study. In our first article, we showed how the risk environment has become NAVI (nonlinear, accelerated, volatile and interconnected) and how “Risk Strategist” organizations are better adapted for this new world of risk. Our second article examined how companies thrive in the NAVI world by reimagining the risk operating model to make it strategy-first, trigger-based, and governance-forward.

But that’s easier said than done. How do you define triggers and establish governance for a vast and expanding universe of tail risks? How do you do this for risks nobody has even imagined yet? Building the new operating model is all-but-impossible with manual processes and human constraints. The speed, scale and insight of AI are indispensable game-changers.

Yet, many risk functions are relatively slow in adopting AI. Action lags ambition. This chapter focuses on the barriers to adoption, and how risk leaders can overcome them with an approach that is both AI-native and future-ready. Instead of automating existing processes, achieving the full value opportunity from AI requires reimagining process in an AI-first way. Value Blueprints provide customizable frameworks that address multiple constraints simultaneously, while realizing value that compounds, rather than plateaus.

So far, the crises of the NAVI operating environment have seen most companies caught unprepared, leaving them to scramble after the fact to respond. So far, this hasn’t been a source of competitive disadvantage, since everyone was in the same predicament.

But that’s changing. Some companies are using AI to fundamentally transform their Risk functions, and pairing that shift with the telemetry and oversight needed to know their agents and stand behind every AI-driven decision. In future crises, these firms may be prepared for unfolding impacts ahead of time and quicker to respond than their competitors.

How can you use the speed, scale and insight of AI to build the future of risk management – instead of automating its past?



Kapish Vanvaria

EY Global Risk Consulting Leader

Contents

	Introduction	01	
01	Reimagining the risk function for the AI era	03	
02	“Wait and see” is no longer good enough	07	
03	AI is accelerating, with unexpected pivot points	15	
04	Building an AI-native, future-ready risk function	21	
	Summary	31	

Research methodology

About the EY Risk Transformation series

60+

interviews

28

senior risk management executives

15+

leaders from the EY Risk Consulting practice

1,200

survey respondents

This article is part of the EY Risk Transformation series. This series explores the symbiosis of risk and strategy by identifying and learning from the behaviors and mindsets of Risk Strategist firms, a group of companies that approach risk management in more strategic ways than “Risk Traditionalists.” These organizations are better prepared for the NAVI risk environment, allowing them to achieve better risk and business outcomes.

For more information, read the foundational article.

The article is based on a combination of quantitative and qualitative research.

Quantitative research for this article includes a survey of 1,200 risk professionals conducted by the global EY organization in April and May 2025. Respondents span 21 sectors and 12 countries and represent organizations with over US\$1 billion in annual revenue. To participate in the survey, risk professionals were required to have responsibility for risk management at their organization with a role in their organization’s C-suite, head of business unit, board member, or part of audit, risk, compliance, or nominating and governance committees.

Qualitative research for this article included a dozen interviews with risk professionals, building on more than 50 interviews conducted for the previous article in this series. Some of these interviews form the basis for guest perspectives by Risk Strategists that are featured in this article.

Defining Risk Strategists and Risk Traditionalists

To help us identify companies that closely link risk with strategy and the business, we captured the extent to which organizations:

- Align the risk management function with overall business strategy
- Make business decisions based on a risk-informed strategy
- Incentivize risk management professionals using business metrics
- View risk management as an enabler of business growth

After running latent class segmentation to understand how the responses to these questions were grouped, we created an index to reflect these relationships. Each question was asked on a five-point scale, and we scored each response based on level of agreement or the extent the statement describes their organization. Aggregating these scores and normalizing the scale from 0 to 100 formed the basis of the index. Through statistical testing, we identified Risk Strategists as the top 32% of index scores and Risk Traditionalists as the bottom 25%.


AI's speed, scale, and insight are game changers for driving resilient growth in today's complex risk environment.

As we write this article, Risk leaders everywhere are grappling with two external shocks: the Iran War and subsequent closure of the Strait of Hormuz, as well as the unprecedented capabilities of frontier AI models to identify cyber vulnerabilities at massive scale.

This piece offers no advice on responses to either crisis. Because how you react to these challenges is not the issue – the issue is whether you are still reacting, and how much longer you plan to remain in reactive mode. The question is not just these disruptions; it's disruptions like these. By the time you read this, these crises may well have faded, but they will almost certainly have been supplanted by other systemic shocks.

These two disruptions highlight a central tension confronting Risk, Strategy, and Technology leaders. The Hormuz Crisis exemplifies a shift EY calls the NAVI world – a post-pandemic environment in which risks are increasingly nonlinear, accelerated, volatile and interconnected. Surviving and thriving in this climate requires a fundamentally different approach to risk management, which is all but impossible without the speed, scale, and insight of AI. Yet, as illustrated by frontier AI's cyber breakthroughs, the evolution of AI is itself NAVI. New frontier models and capabilities can emerge unexpectedly, upending companies' AI adoption plans and the assumptions on which they were based – uncertainty that can make leaders hesitant to fully embrace AI.

How you react to these challenges is not the issue – the issue is whether you are still reacting, and how much longer you plan to remain in reactive mode.



Reimagining the risk function for the AI era

Realizing value from AI's speed, scale and insight requires rethinking risk with an AI-native, built-in approach.

01

"The post-pandemic risk environment has become increasingly complex," says Bill Diaz, CEO, Archer. "Risks now show up everywhere, at any time, often triggering chain reactions. The pace is faster, the effects larger. AI is vital for managing risk in this environment; it's critical for handling today's increased volume and complexity, as well as for responding to risks in real time."

As companies have been buffeted by the volatility of recent years, resilience has become a mantra. But the NAVI risk environment didn't just bring more risks; it also fundamentally reshaped the nature of risk itself. By the same token, companies need not just more resilience, but different resilience – they need to rethink and reframe what resilience means in a new world of risk.

If the traditional approach to resilience was reactive and limited to safeguarding business continuity, the new approach is equally about being proactive and driving strategic growth. Like any good game plan, this includes both offense and defense. Offense ensures your strategic planning is built on a comprehensive understanding of emerging risks and their strategic impacts. Defense ensures you can maintain fundamental promises in the face of NAVI's disruptions. The end result? A symbiotic relationship in which strategy becomes resilient, and resilience drives strategic growth.

To achieve this vision, risk management needs to change in three critical ways:

1. Companies need the ability to sense and respond to rapidly changing developments in real time.
2. Risk functions need to model vast numbers of assessments and scenarios at massive scale as the number of risks, and interconnections between risks, proliferate – and as tail risks once considered low-probability merit more serious consideration.
3. Firms need to make failures of imagination a thing of the past at a time when companies are repeatedly caught unawares by external shocks and nonlinear tipping points – as well as the unexpected downstream impacts of interconnected, cascading risks.

These shifts require moving beyond the limitations of human capability; achieving a NAVI-ready Risk function with manual processes alone is all but impossible. So, emerging technologies, and especially AI, become indispensable. AI can reinvent risk management with three game-changing capabilities:

1. **Speed** replaces slow-and-periodic manual processes with rapid-response, real-time automated ones.
2. **Scale** analyzes variables and scenarios in quantities many orders of magnitude beyond human capacity.
3. **Insight** meets the complexity of the NAVI risk environment with advanced analytical capabilities, as well as the potential for overcoming human behavioral biases and blind spots – from confirmation bias and desensitization to cognitive overload and decision paralysis.

But realizing value from the full potential of these capabilities is not a given. Instead, the value you realize depends on how you deploy AI.

AI is the indispensable game changer – but only if you use it to build the future of risk management; not automate its past.

Initially, many companies have pursued the proverbial low hanging fruit of automating their existing processes for efficiency gains – such as by adopting governance, risk and compliance (GRC) platforms. This can deliver tangible and near-term results, from standardizing taxonomies and automating manual processes to increasing coordination. Leading GRC platforms are increasingly embedding agentic capabilities to evolve static workflows into more adaptive, AI-enabled execution.

The incremental approach of automating existing processes has been a good starting point. It can demonstrate proof-of-concept and deliver near-term return-on-investment. It is also the least disruptive to legacy operating structures and the easiest lift in terms of implementation.

But it’s not where the biggest gains lie. “Most companies have felt compelled to adopt AI,” says Raul Villar Jr., CEO, Optro. “Every board and C-suite has been focused on how they can leverage this technology to their benefit. And so everyone’s invested in AI, but most have yet to see the return they were expecting on those investments.



Leading companies are moving away from making AI accretive to a process, and are instead looking at using AI to fundamentally reinvent processes.

Dan Diasio
EY Global Consulting AI Leader

Our acquisition of Midship, an agentic AI platform, was driven by a desire to accelerate value realization and a conviction that the future of GRC isn’t faster audits, it’s fundamentally different ones.”

Indeed, realizing the full value of AI comes from structural transformation: not automating existing processes, but reimagining them from the ground up in an AI-native way.

“Leading companies are moving away from making AI accretive to a process, and are instead looking at using AI to fundamentally reinvent processes,” says Dan Diasio, EY Global Consulting AI Leader. “They are transforming ways of working. This means challenging whether a particular process is still needed, and what the future, AI-native process should instead be. The real value realization opportunity is not from using AI as a bolt-on but making AI built-in to processes and functions.”

This “built-in” approach to AI opens the door to fundamentally rethink existing processes and functions. Examples of such shifts are provided in the accompanying table.

	Bolt-on AI	Built-in AI
Risk Appetite	AI supports existing risk appetite processes (e.g., summarizes exposures, drafts committee materials, compares appetite statements to recent incidents)	AI dynamically reassesses risk appetite, in alignment with strategic goals (e.g., continuously synthesizes internal + market data and aligns with strategic goals to adjust risk appetite in real time)
Risk Identification and Assessment	AI automates manual processes (e.g., implements surveys and interprets free-text responses, edits/standardizes risk descriptions, calculates assessments based on traditional likelihood x impact metrics)	Scenario-generation engine identifies and assesses risks at massive scale, eliminating failures of imagination (e.g., simulates thousands of scenarios to surface tail risks and cascading impacts, assesses risks based on future impact on strategic goals, moves from sampling to population analysis)
Risk Mitigation	AI makes existing mitigation workflows more efficient (e.g., recommends controls, reviews mitigation responses, standardizes remediation actions inside GRC platforms)	AI predefines strategic mitigation at scale (e.g., ownership/escalation are pre-wired and embedded into agentic workflows, scenario generation at scale enables mitigation of huge universe of risks)
Monitoring and Detection	AI assists episodic, calendar-driven approach (e.g., writes/edits periodic reports and audits focused on compliance metrics and driven by audit/reporting calendar)	Multimodal scanning agents monitor risks in real time across the enterprise (e.g., agents embedded in the first line monitor diverse data streams to continuously reassess probabilities of emerging risks, world model of company operations continually assesses impact on operations and strategic goals)
Risk Response	AI assists after-the-fact response to enable business continuity (e.g., automated alerts require humans-in-the-loop to review, deliberate and implement the response)	Agents and humans-on-the-loop use predefined playbooks for agile, real-time response (e.g., dynamic dashboards update human managers about emerging risks, crossed thresholds trigger protocols by agents and notification of humans-on-the-loop at relevant decision points, protocols can adapt as needed to changing circumstances)

“Wait and see” is no longer good enough

The future is already here —
are you ready for it?

02

AI adoption – and particularly, using a built-in AI approach to drive risk transformation – is essential for resilience in the NAVI risk environment.

Yet, much of the Risk world has been slow to adopt AI and other emerging technologies, all the more so when it comes to deploying tech in truly transformative ways.

Large majorities of Risk Strategists – which frame their Risk functions as enabler of strategic growth – see the disruptive potential of emerging technology:

- Seven in ten (70%) of Risk Strategists agree that AI will fundamentally transform the operating model of their Risk function, compared to only 40% of Risk Traditionalists.
- Similarly, 67% of Strategists say that emerging technologies have the potential to fundamentally change how they approach risk management to be better aligned with the NAVI risk climate, compared to only 41% of Traditionalists.

But while robust majorities of Risk Strategists see the potential, adoption numbers are considerably lower, and the gaps between Strategists and Traditionalists are smaller still. Action lags ambition.

The two technologies with the greatest adoption are traditional AI and natural language processing. These are also the oldest technologies covered in the survey, both pre-dating the mainstream emergence of generative AI. For the other technologies in our survey – which have greater potential for driving risk transformation – adoption rates, even by Strategists, are much lower.


Risk functions are applying emerging technologies across both incremental and transformative use cases to a similar extent. However, Traditionalists lag Strategists by almost twice as much in transformative use cases (8.2 percentage point gap) than for incremental use cases (4.5 percentage point gap).

Risk Strategists are much more likely to see the transformative potential of emerging technologies than Risk Traditionalists




Percentage of respondents who agree or strongly agree with each of the statements


Many companies have not yet adopted emerging technologies for risk management...




..and the gaps between Risk Strategists and Risk Traditionalists are small



Percentage of respondents who currently use each emerging technology within the risk function for each stated purpose or use case



Risk Traditionalists lag Risk Strategists in using AI, particularly in transformative use cases



Percentage of respondents who plan to use or currently use each of the emerging technologies for risk management across the enterprise

The top AI adoption barrier cited by survey respondents (45%) is low prioritization of risk versus other use cases. This raises a corollary question: why is AI adoption for risk transformation not being prioritized? From our interviews with risk leaders, including at many Risk Strategist organizations, an oft-cited sentiment was that they were in “wait-and-see” mode until the technology is more proven to justify wide-scale adoption.

What’s left unsaid is that assuming a “wait-and-see” approach to technology adoption also means that a company remains in wait-and-see mode on its ability to respond to emerging risks. So far, the crises of the NAVI operating environment have seen most companies caught unprepared, leaving them to scramble after the fact to respond. So far, this hasn’t been a source of competitive disadvantage, since everyone was in the same predicament.



But that’s changing. The accompanying guest perspective, by the Global Chief Risk Officer of a leading European automotive manufacturer, illustrates how one company is using AI and embedding AI agents in its first line to transform its approach to risk management. The company has already achieved proof of concept, and says its agentic system gave it advance warning about the impact of helium shortages arising from the Strait of Hormuz closure.

In future crises, companies that are using built-in AI to transform their Risk functions may be prepared for unfolding impacts ahead of time and, consequently, well before their competitors. At this point, “wait-and-see” will no longer be good enough. If you’re waiting until after a crisis to figure out the knock-on implications that could impact your business, it’s already too late. You will be competing against organizations that have identified the scenario, mapped out cascading risks and impacts on their business, prepared response plans, and assigned ownership and oversight – all ahead of time. The game will be over before it’s started.

Global Chief Risk Officer of a leading European automotive manufacturer



How a large European manufacturer is using AI to reimagine risk management

We started building a risk management technology platform 3-4 years ago, to replace what had previously been a siloed and manual approach. AI’s “ChatGPT moment” had not yet happened, but the world of risk had already changed. We were in a more uncertain environment – what EY calls the “NAVI world” – which demanded a nimbler approach.

We have global footprint, with a large number of legal entities around the world. Five years ago, we had a Risk Coordinator within each function and entity, who was responsible for reporting their top risks, but these were manual and fragmented processes. We were mostly identifying risks that were already well known. The real need was to instead identify the emerging risks that could hinder our strategy.

The rise of generative AI allowed us to think bigger, and implement a future-ready vision for risk management. Instead of using AI to just improve our existing processes (e.g., to edit or rewrite descriptions of risks) we used it to fundamentally rethink our entire approach.

We are reimagining risk management from the ground up with a multi-agent system consisting of about 15 AI agents. These agents monitor developments in the outside world and analyze their potential impact on our business and our strategy. We have a geopolitical agent focusing on the US, another on China, and so on. The inputs from these agents are run through an orchestrator agent that connects dots and removes redundancy.

Initially, foundational LLMs were not doing a good job of this kind of analysis, but the models continue to improve, and we are now getting the quality and consistency we can rely on. Next, we’re working on an implementation using an AI-powered tool that will allow board members to consume outputs in their preferred mode of communication, whether that’s a slide deck or a podcast they can listen to during their commute. We’re also working on a scenario-building agent that would simulate hypothetical, low-probability shocks and assess potential impacts on our company.

The system is not yet live, but we’ve achieved proof of concept and are already seeing results. It identified the potential impact of a Strait of Hormuz closure on the supply of helium and its impact on our manufacturing processes – well before this topic was being explored in the media.

Our AI-driven approach is also helping us better align Strategy with Risk Management. Risk can provide an unbiased view of different strategic options and keep strategy on track in a move volatile world. AI does this in real time, enabling a more agile strategy and organization.

For compliance reasons, we will continue to perform some activities, such as conducting quarterly meetings or reporting high-probability risks – but this is not where the value lies. In the NAVI world, the real “value add” is using AI to give the company the agility to respond to shocks, and the board the foresight to steer us in the right direction.

“

[Our AI-based system] identified the potential impact of a Strait of Hormuz closure on the supply of helium and its impact on our manufacturing processes — well before this topic was being explored in the media.

Global Chief Risk Officer
European automotive manufacturer

AI is accelerating, with unexpected pivot points

Is your tech strategy
future-ready for the next
frontier model disruption?

03

In recent months, the approach to GRC implementation has been complicated by breakthrough capabilities in autonomous/agent AI and “vibe coding.”

These breakthroughs are disrupting the software-as-a-service (SaaS) model and reframing the traditional build-vs-buy decision. There are perceived benefits from either approach – such as leveraging higher levels of support and expertise from a SaaS provider, versus increased control over enterprise data and customizability from going it alone. But, in this fast-changing space, even those assumptions aren’t standing still. Even as some organizations are bringing development in-house to gain tighter control, major GRC and SaaS providers are accelerating their own AI-native architectures – offering a “buy” option that delivers similar customizable, agentic capabilities without the heavy internal engineering burden. (For an example of a company that is going the “build” route, see the accompanying guest perspective from Adam Frank and Ramesh Raju of Uber).

Adam Frank, Deputy Head of Financial Risk Management, and Ramesh Raju, Head of Data, Automation and AI – FRM Compliance Tech, at Uber



How AI's expanding capabilities are enabling Uber to rethink its approach to tech adoption

The Finance Risk Management (FRM) team serves as Uber's second line of defense, managing certain compliance programs, including Sarbanes Oxley (SOX).

Historically, our approach to technology adoption was anchored in traditional robotic process automation to automate discrete tasks. However, following Uber's IPO in 2019, rapid growth created a continuously changing risk and operating environment. Increased volatility and velocity made it clear that off-the-shelf automation alone would not provide the control, adaptability, or transparency that was needed so we increased our investment toward in-house tech capabilities.

Initially, this took the form of traditional software development, with our team building a dedicated platform to support controls testing. As AI capabilities matured, our focus evolved again – moving from point automation toward AI-enabled solutions specifically designed for FRM use cases. Rather than simply accelerating existing manual processes, we use AI to fundamentally change how information is consumed, interpreted, and applied across the compliance lifecycle. Instead of attempting to automate every human step, we use AI to synthesize information, surface patterns, and generate insights at a scale beyond human capacity.

Much of our progress has come through modular, purpose-built solutions – tools that solve individual problems but often exist in silos. The next phase of value creation lies in connecting these capabilities. The objective is to provide control owners and FRM specialists a single platform that brings data, insights, and workflows together in one place.

Advances in AI are also reshaping how we build. Our team is increasingly using AI to assist with writing and understanding code, lowering the barrier between technical and non-technical roles. Risk managers and internal control professionals no longer need to rely exclusively on engineers to translate complex logic into plain language. AI can help bridge this gap, enabling shared understanding and faster iteration between risk, compliance, and engineering teams.

These developments have also influenced our "build vs. buy" approach. While we previously focused on purchasing commercial platforms and building integrations around them, we now prefer building native solutions tailored to our internal tools. This enables tighter integration, faster evolution, and better team alignment.

Yet, we remain pragmatic. Buying will still be a viable option for static, well-defined tasks, or when a commercial product is highly cost-effective and requires no ongoing customization. Modern risk management requires technology to be embedded across every process, team, and decision. With rapid advances in AI's capabilities, we need to continually challenge old assumptions – from how we develop code to how we approach the build vs. buy decision.

“

Rather than simply accelerating existing manual processes, we use AI to fundamentally change how information is consumed, interpreted, and applied across the compliance lifecycle.

Adam Frank and Ramesh Raju
Uber

The decision about where, and to what extent, to build vs buy will depend on the circumstances and goals of individual companies. But the fact that these decisions are being reconsidered at all illustrates a larger truth. Emerging technologies aren't just a solution to managing risk in a NAVI world – their evolution is itself nonlinear, accelerating, volatile and interconnected.


Consider how AI models have been developing new capabilities at an accelerating pace. Or consider how frontier models achieve breakthrough capabilities in a nonlinear manner, creating volatility and repeatedly catching the business world by surprise.

The initial 2022 launch of ChatGPT was itself a nonlinear moment, for which most of the business world was unprepared. Most recently, new frontier models with the capability to expose software vulnerabilities at unprecedented scale have blindsided the cybersecurity world, challenging long-standing assumptions about threats, capabilities, and cybersecurity approaches

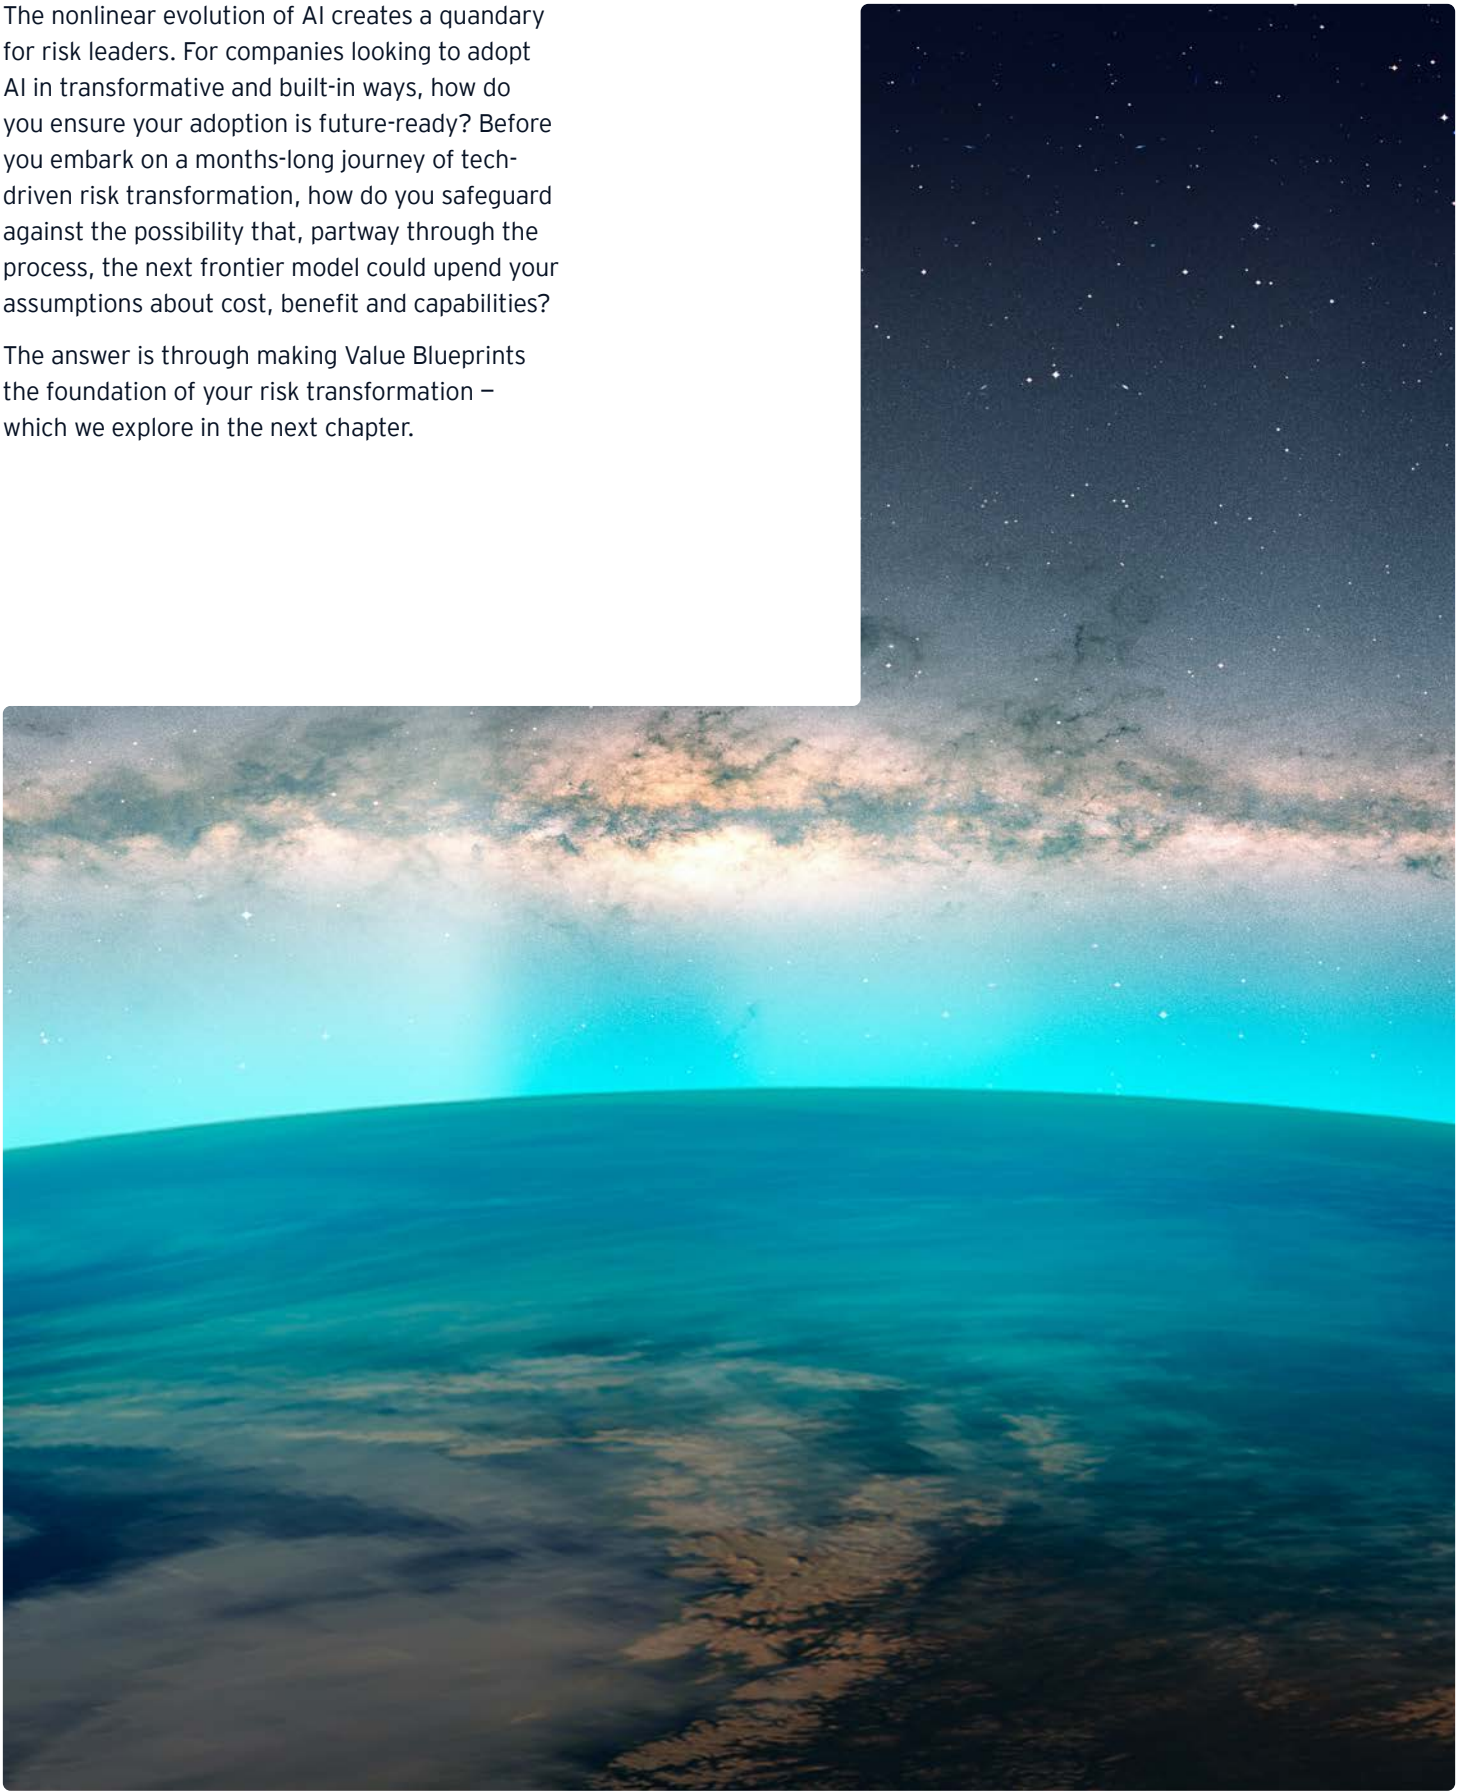
The nonlinear evolution of AI creates a quandary for risk leaders. For companies looking to adopt AI in transformative and built-in ways, how do you ensure your adoption is future-ready? Before you embark on a months-long journey of tech-driven risk transformation, how do you safeguard against the possibility that, partway through the process, the next frontier model could upend your assumptions about cost, benefit and capabilities?

The answer is through making Value Blueprints the foundation of your risk transformation – which we explore in the next chapter.

Artificial intelligence is approaching or exceeding human performance on a range of cognitive tasks



Source: Stanford AI index report 2026




Building an AI-native, future-ready risk function

How can you address multiple adoption barriers to move ahead with confidence?

04

Strategic prioritization and data challenges are the biggest barriers to emerging technology adoption



Percentage of respondents ranking each factor among the top three biggest barriers to adopting emerging technologies for the organization's risk management function across the enterprise


The chart above shows the barriers Risk functions face in adopting AI. These include the low prioritization of risk use cases, integration challenges, data constraints, security concerns, and issues related to talent, budget and cost.

By now, this is a familiar list. Numerous surveys have surfaced similar concerns with respect to AI adoption. The question is what you can do to overcome them, particularly when you simultaneously face several, or all, of these constraints – as many organizations do.

The nub of the matter is that these constraints are not independent variables. Instead, they are deeply interconnected, exacerbating and compounding each other. Low prioritization hampers investment, which intensifies budget and talent constraints. Talent and data gaps deepen integration challenges, which in turn delay data readiness. Such linkages create a self-reinforcing vicious cycle that undermines the thesis for AI adoption.

To break out of this doom loop, Risk functions need to address adoption constraints with a cohesive and interconnected approach that tackles all of them. This is exactly what the EY.ai Value Blueprints framework provides, and why it should be the foundation of your approach to AI adoption. The seven layers of framework are interlinked; they feed off each other, creating value that grows exponentially instead of plateauing quickly. (For more, see the perspective by Dan Diasio on the next page).

Dan Diasio, Global Consulting AI Leader, EY



How EY is helping clients realize value from AI with EY.ai Value Blueprints

“

When organizations add AI use case-by-use case, value rises incrementally, then plateaus. When they instead shift to a blueprint-by-blueprint approach... effort decreases with each subsequent blueprint, while value compounds exponentially.

Dan Diasio, Global Consulting AI Leader
EY

Realizing the full potential value of AI requires an approach in which AI is built-in, not bolt-on: reimagining processes with AI at the core, rather than merely automating the past. To help organizations make the shift, we have developed EY.ai Value Blueprints. These are layered, ready-made, and customizable frameworks that guide transformation blueprint by blueprint – from strategy to execution. Each layer focuses on a critical dimension of enterprise value and builds on the last, compounding value and making the next one easier to achieve:


- 01 **Systems of record:** reliable information that connects humans and AI agents
- 02 **AI-native foundation:** enterprise infrastructure that enables AI to operate safely and at scale
- 03 **Intelligence:** unified organizational memory and knowledge that enables reasoning in the context of the enterprise and the individual
- 04 **Trust:** new governance frameworks that maintain appropriate oversight while enabling AI autonomy
- 05 **Processes:** process redesign to eliminate manual handoffs, create continuous flows, and enable workflows that can adapt and optimize themselves

- 06 **Workforce:** upskilling initiatives to enable more effective collaboration with AI agents, especially as the role of human experts shifts to human-on-the-loop oversight
- 07 **Customer:** enabling new offerings and business models to create new value

When organizations add AI use case-by-use case, value rises incrementally, but then plateaus. When they instead shift to a blueprint-by-blueprint approach, the opposite happens. Since Value Blueprints are built on reusable capabilities and cross-functional targets, effort decreases with each subsequent blueprint, while value compounds exponentially.

The process starts with a value-led diagnostic, working with client teams to quickly define the ambition, reimagine the highest-impact processes, and design the architecture and roadmap to deliver them. This provides a practical, prioritized playbook and establishes a self-funding path forward. With strategic intent defined, we reimagine processes and the workforce together, designing with AI at the center. Finally, we develop the data, technology, and governance foundations so future-ready processes can run at scale – fueling transformation, rather than just automation.

Use the Value Blueprints framework to address your constraints – data, governance, talent, etc. – in an interconnected way, so they build on each other to create a multiplier effect. Approach every step of the process with a focus that is both AI-native and future-ready:



- **Being AI-native.**
Think boldly to fully leverage the capabilities of today’s frontier models. Instead of carrying forward legacy processes and seeking to automate them, reimagine what processes would be best suited for an AI-native Risk function.
- **Becoming future-ready.**
Anticipate and prepare for future inflection points. This is feasible, despite the seeming unpredictability of nonlinear change. While we may not know when the next nonlinear shifts will arrive, it’s much easier to see what they are likely to be – they are being worked on, discussed, and even incorporated in some newer models. Think through how emerging capabilities would fit into your approach, and be ready to pivot when the next inflection point arrives. Reframe adoption barriers, transforming them from justifications for remaining in wait-and-see mode, to potentially temporary challenges that may resolve with AI’s continuing evolution.

01

Define your ambition and value opportunity

Start with defining the vision and ambition for your Risk function. What does the AI-native, built-in version of your Risk function look like? This includes exploring several questions with respect to both today’s AI and the likely capabilities of future frontier models.

Being AI-native:

- Starting from a blank sheet (i.e., not limited by legacy processes) what is your AI-native vision for your risk processes and operating model?
- Which risk management practices will become obsolete, and which new ones will need to emerge, for this vision to be realized?
- Where do the biggest value pools lie? What will become commoditized and plentiful in this future state, and what will become the new drivers of value?

Becoming future-ready:

- How will potential inflection points and relevant emerging capabilities (e.g., agentic and autonomous AI, on-device/edge AI, small language models, multimodal AI, reasoning AI, world models/spatial intelligence, persistent memory) shift the frontier of capabilities for risk management in your organization?
- What limitations today are temporary barriers that might be overcome in the near future as capabilities improve? For instance, might synthetic data generation help address data availability, small language models lower high compute costs, and no-code/low-code AI mitigate talent shortages?
- How can you layer different models/capabilities based on where they are most suited? For instance, small language models may be ideally suited for developing domain-specific AI agents embedded in the first line, while Reasoning AI and World Models could take a scenario generation engine to the next level.

Reimagine your processes and workforce

Next, reimagine your processes and workforce. These are inextricably linked in any AI-native risk transformation; they should be reimaged together as a single design problem, rather than as separate transformation tracks.

Processes

Reimagine processes to leverage the core capabilities and strengths of your AI and human workers. This could include creating continuous flows so AI can conduct real-time monitoring and response, or develop scenarios and simulations at massive scale. The goal should be to design workflows that can adapt and optimize in real time, while humans direct strategic outcomes.

Challenge legacy structures and silos. For instance, the traditional approach of classifying risks into discrete categories – financial risk, operational risk, regulatory/compliance risk, cyber risk, and so on – may no longer be optimal in an AI-native organization optimized for the NAVI risk environment. Instead of discrete risk types, how should you restructure your processes and organizational structure based on interconnected and systemic risks? Where should you intentionally blur boundaries within the risk operating model – such as in the three lines model?

Workforce

Realizing value from redesigned processes requires developing skills and redefining roles to empower human employees to work effectively alongside AI.

Rethink job definitions and categories. Every job is a bundle of skills, responsibilities, and tasks. Do the ways in which these have been bundled in the past still make sense? If not, how should they be unbundled and repackaged? Which parts of these legacy bundles can be done best by AI and other technologies, and which ones should humans specialize in?

With the arrival of agentic systems increasingly capable of operating autonomously, a future-ready approach also requires preparing the workforce for the transition from human-in-the-loop to human-on-the-loop oversight.

“As AI matures and takes on more of our work, it will increasingly look to humans to make decisions requiring judgment,” says Sinclair Schuller, EY Americas Responsible AI Leader. “A model might perform vast amounts of work, then turn to a human to decide between Option A and Option B. Both options are valid and correct, so human judgment is required to pick between them. This is an example of human-on-the-loop oversight – as opposed to a human-in-the-loop approach, in which people directly intervene at a much finer grain and smaller scope. Keeping humans in the loop is certainly appropriate in some situations, but as technology matures, we will increasingly move to human-on-the-loop.”

Risk functions, in particular, may have a natural instinct to retain humans in the loop, especially if this is assumed to be the lower risk option. But, to maximize value from agentic AI and autonomous systems, risk leaders should carefully assess where an on-the-loop approach might be feasible. They should also be willing to challenge the assumption that in-the-loop necessarily carries lower risk. Being involved in every detailed decision can lead to a numbing effect for workers; an approach that selectively escalates decisions for human oversight only when appropriate might well provide more effective governance.

Anticipating and preparing for the shift to human-on-the-loop systems requires deliberate action and investment. This includes investing in training and upskilling to empower your workforce to function effectively in this new capacity. It includes quantifying and articulating thresholds that trigger human oversight, escalation pathways to route decisions to appropriate individuals, and dashboards to enable monitoring by human overseers.

“

Keeping humans in the loop is certainly appropriate in some situations, but as technology matures, we will increasingly move to human-on-the-loop.

Sinclair Schuller

EY Americas Responsible AI Leader

Develop the supporting foundation of data, governance, and technology

The third step is to focus on technology, data, and governance. AI cannot effectively scale unless these are designed together as a unified and integrated supporting foundation.

Data

Data forms the basis for decision making in the AI-native Risk function. Unify and connect disparate data streams, and embed data directly into workflows. Agentic AI, in particular, requires investing in semantic and structured data, so agents can accurately interpret meaning and relationships in context. Developing real-time risk monitoring and decision-making systems requires investing in and integrating continuous data feeds from within and outside the organization.


To be future-ready, identify and anticipate the data readiness requirements of emerging AI capabilities – for instance, edge AI and small language models will require a decentralized data strategy, while multimodal AI requires diverse signal integration.

Governance

Robust governance is a scaling enabler – creating the trust essential for AI to gain adoption. Agentic AI systems need governance embedded in execution – such as real-time controls, guardrail agents, and kill switches. As discussed above, it also requires assessing where governance can effectively move from a human-in-the-loop to a human-on-the-loop approach.

Technology

Technology is the core of the AI-native Risk function. As discussed above, harnessing its transformative power requires building AI-native infrastructure, not bolt-on tools. A cohesive approach, using the Value Blueprints framework to develop a vision for your AI-powered Risk function – while simultaneously transforming your processes, workforce, data, governance, and technology – will enable agentic execution at scale across the Risk function and enterprise.



Summary

A new world of risk needs new ways of managing risk. Manual processes and failures of imagination will no longer suffice.

AI is the indispensable game changer – but only if you use it to build the future of risk management; not automate its past. Implement AI that's built-in and future-ready. Use Value Blueprints to ease adoption constraints, find a faster and accelerating path to value – and, ultimately, shape an organization that's resilient in the turbulence of the NAVI world.

Authors



Scott McCowan

EY Global Consulting Risk Markets Leader



Gautam Jaggi

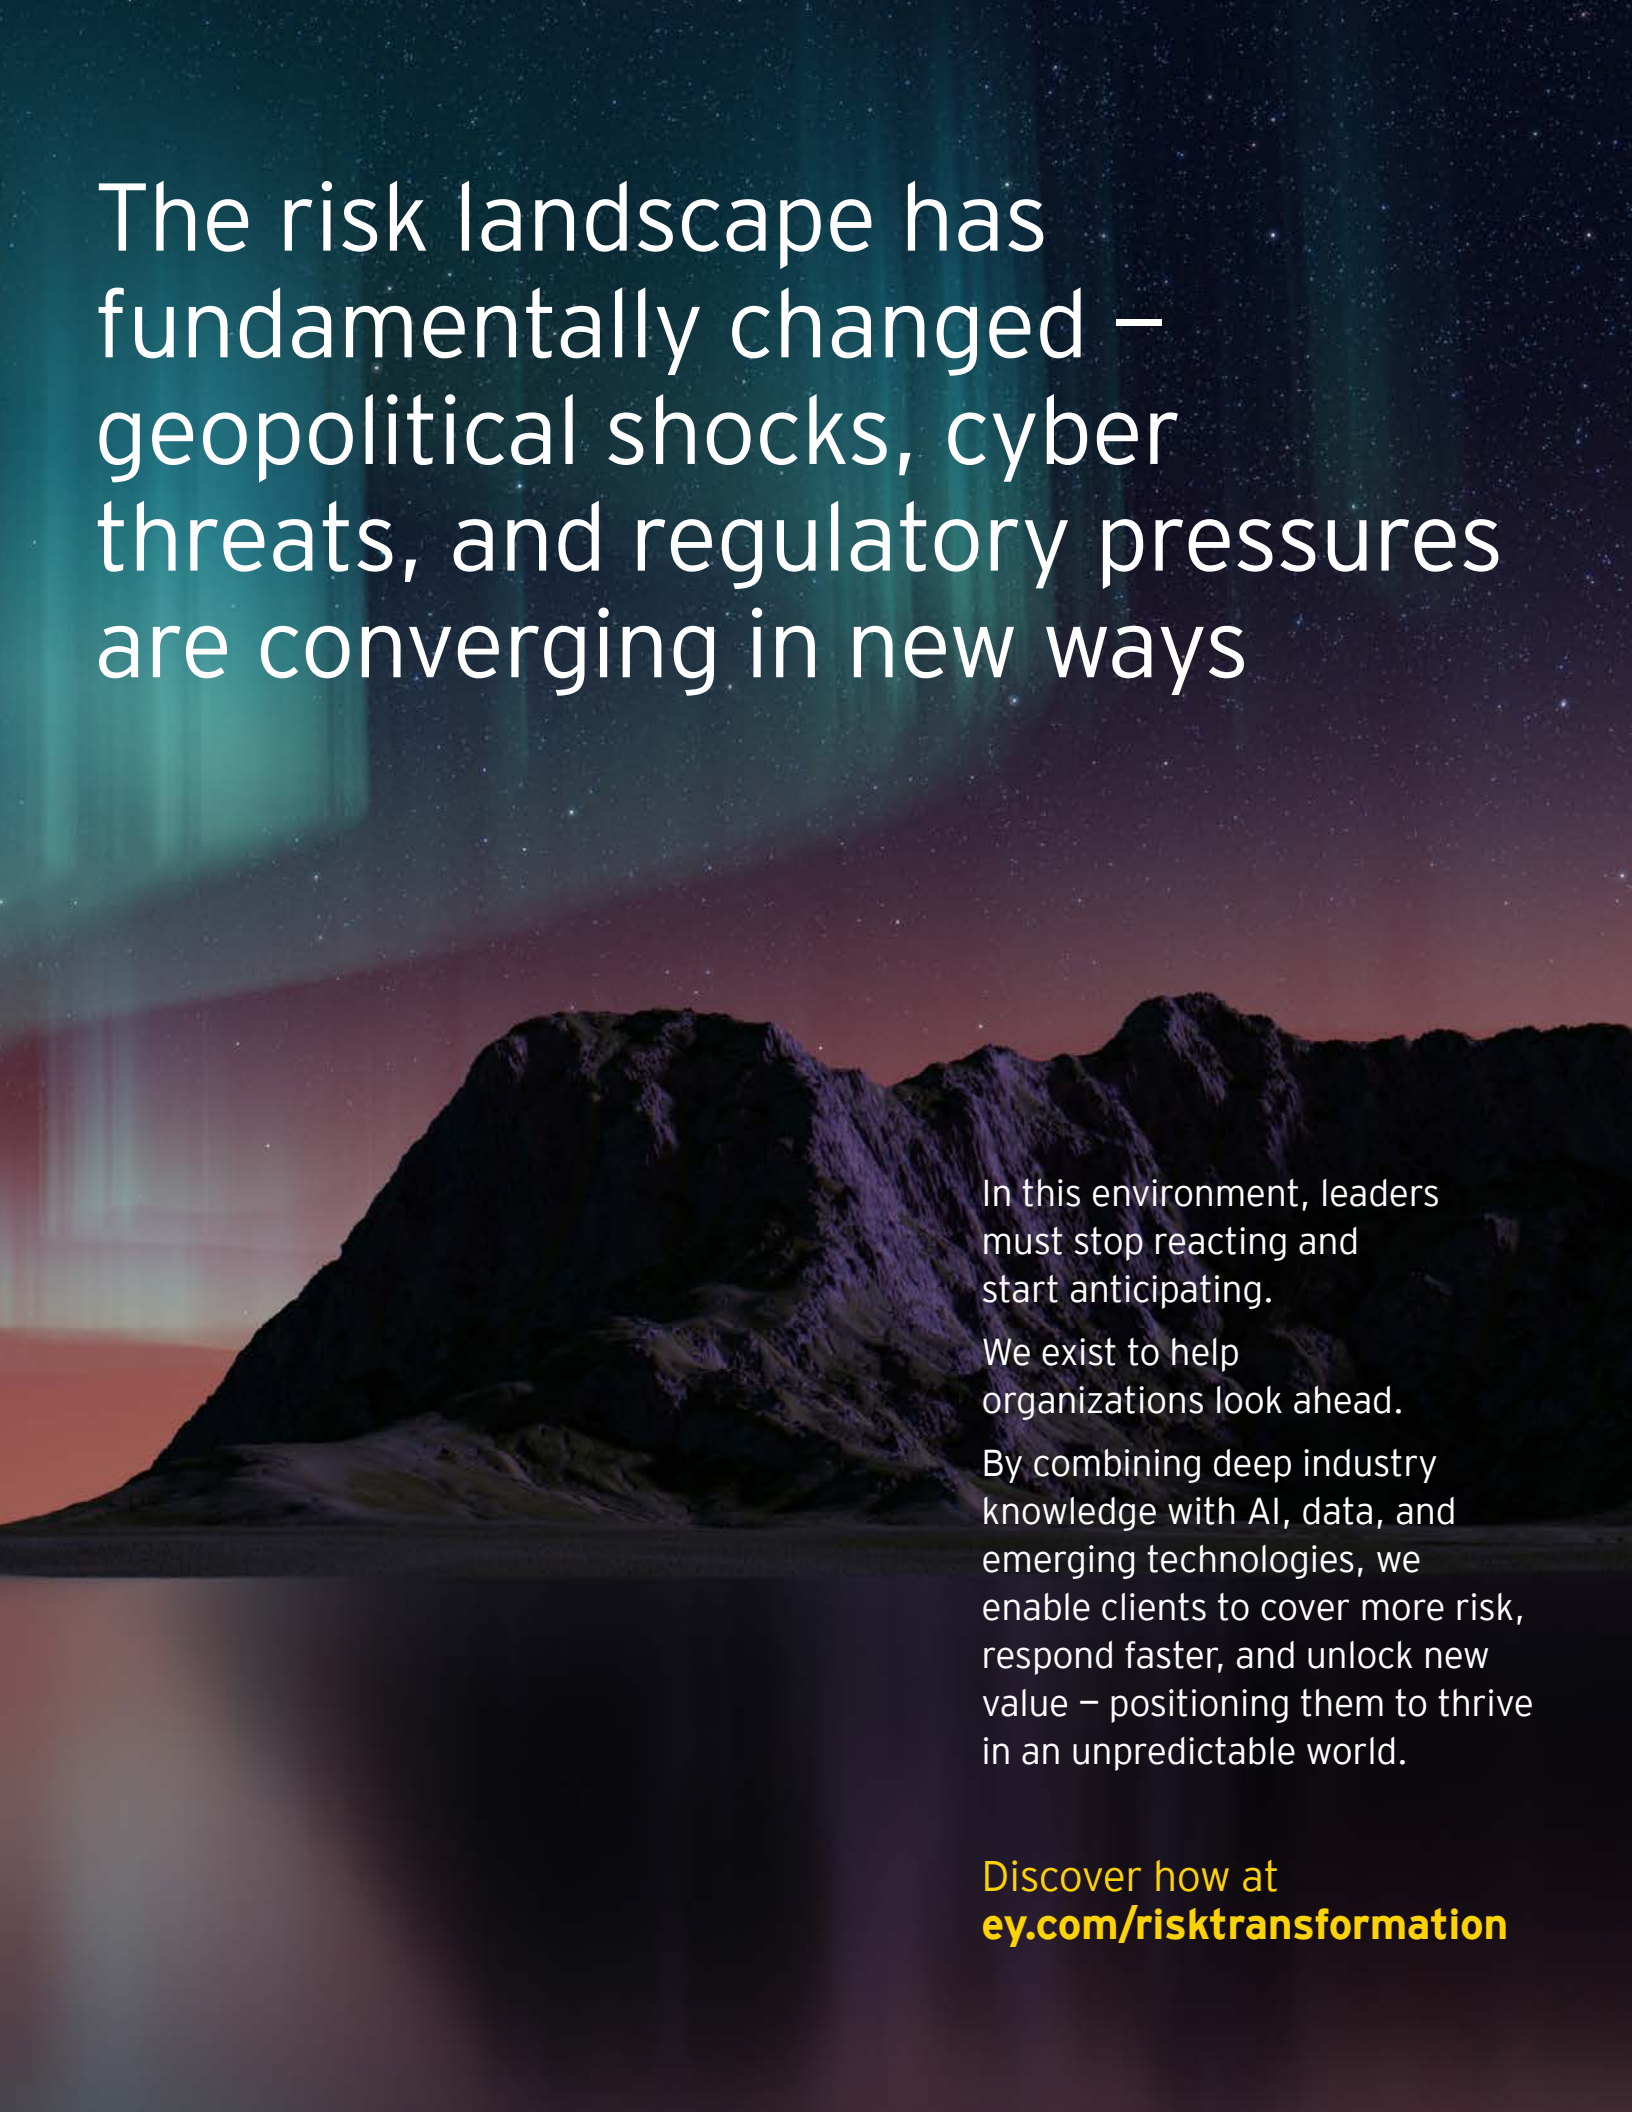
Director, Thematic Insights, Ernst & Young LLP



Rohit Mathur

EY Global Risk Consulting, Head of Strategy and Operations

The authors would like to thank Kyle Lawless, Senior Manager - Risk Consulting, Ernst & Young LLP; Arun V. Tom, Assistant Director, Ernst & Young LLP; AnnMarie Pino, Associate Director, Ernst & Young LLP; Joe Morecroft, Associate Director, EYGS LLP; and William Reid, Assistant Director, Ernst & Young LLP, for their contributions to this article.

The background of the slide is a photograph of a rugged mountain range under a dark, starry night sky. A vibrant aurora borealis is visible on the left side, with green and purple light streaks. The mountain is dark and silhouetted against the lighter sky.

The risk landscape has
fundamentally changed –
geopolitical shocks, cyber
threats, and regulatory pressures
are converging in new ways

In this environment, leaders
must stop reacting and
start anticipating.

We exist to help
organizations look ahead.

By combining deep industry
knowledge with AI, data, and
emerging technologies, we
enable clients to cover more risk,
respond faster, and unlock new
value – positioning them to thrive
in an unpredictable world.

Discover how at
ey.com/risktransformation

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

© 2026 EYGM Limited.
All Rights Reserved.

SCORE 112589-26-GBL
ED None

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, legal or other professional advice. Please refer to your advisors for specific advice.

ey.com/risktransformation