



In financial services, solving GenAI's challenges starts with trust

Introduction

Financial services companies are expanding the application scope of their current generative artificial intelligence (GenAI) initiatives. But almost two years after the arrival of Chat GPT-3, the industry is still struggling to deliver transparency and explainability while maintaining robust data security.¹ And, as companies struggle to integrate GenAI into existing systems and processes, industry leaders are asking uncomfortable but necessary questions about improving speed to market and increasing return on investment.

Focus isn't the problem. Of the industry's executives surveyed in 2024 by KPMG LLP (KPMG), 97 percent said they are investing in GenAI over the next year.¹ Risk management shortcomings aren't an issue, either: 91 percent conduct regular audits to help ensure ethical use, supplementing that effort with education and training, human oversight, and ethical frameworks.²

Trust, it seems, is what's missing. Though institutions believe they

have the infrastructure to handle and monitor GenAI, most remain highly cautious about the ways these powerful algorithms are used. For a highly regulated industry with strict data privacy restrictions, getting past these concerns is proving onerous: 77 percent of financial services respondents cited data privacy and security risk as the top concern when working with partners on GenAI.³

Given high expectations around the transformative power of GenAI, the eagerness to deliver faster is warranted. Fortunately, there are practical ways to establish the necessary trust and build momentum for GenAI programs: With the right governance, people, and processes, institutions can move their GenAI programs forward with confidence.

77%

of financial services respondents cited data privacy and security risk as the top concern when working with partners on GenAI

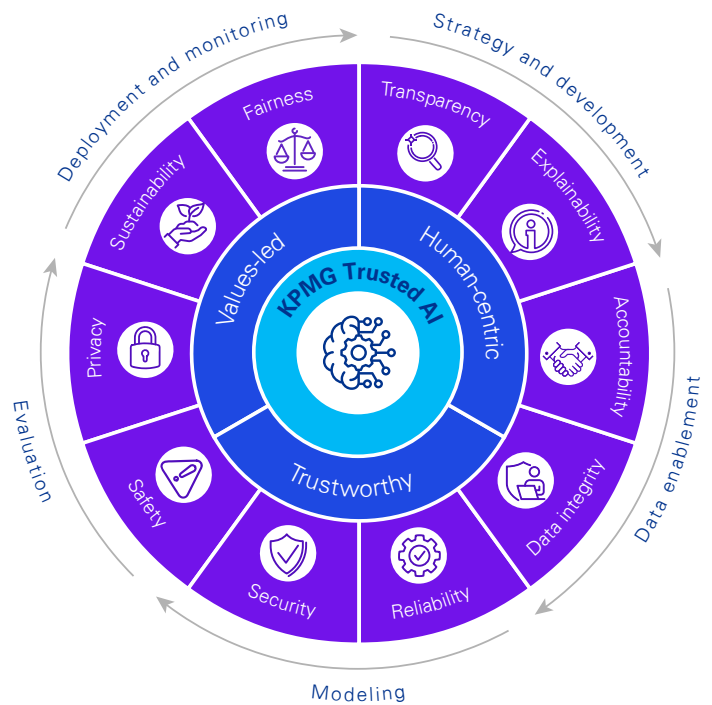
¹ "2024 KPMG Generative AI Survey," August 2024

² Ibid.

³ Ibid.

Strengthen AI foundations with a strong governance framework

Start by establishing a strong AI steering committee (SteerCo). The committee is an essential strategy, policy, and oversight body that directs GenAI efforts across the enterprise. The members of the committee help ensure that the vision and execution of your GenAI strategy is conceived, aligned, and implemented to your institution's risk appetite level. It also helps balance innovation agility, risk management, and governance, enabling policies that are comprehensive and flexible.



The SteerCo creates a company-wide AI charter, an ethical north star for trusted and responsible AI use, that guides and streamlines the GenAI lifecycle. Operations, people, technology, reputation, legal and regulatory, and financial factors are important considerations. For this reason, it's essential that the SteerCo includes stakeholders representing diverse perspectives from across the enterprise. Here, the risk function plays a crucial role, establishing the governance framework, policies, and processes for building trust in GenAI.

Addressing equity concerns around AI is essential.

These run the gamut from data bias to lingering workforce fears that AI will make human jobs obsolete. Transparent and explainable inputs, methods, and outputs are essential—especially when the technology is used in consumer-facing applications, like assisting with lending decisions. Companies that ignore this issue in their framework risk reinforcing historical discrimination and could face legal and financial consequences.

The SteerCo and charter create a robust governance framework that sits at the core of any successful GenAI risk management program and is essential for evaluating categories of risk. It must be airtight, built on trusted AI principles and girded by security, privacy, and ethics measures. Importantly, it reinforces an organization's risk tolerance levels, and prepares a ranking and acceptance process to accelerate approvals and speed to market.

As with all systems and procedures in financial services, calibrating this framework to meet stringent regulatory, privacy, and security requirements is critical. With privacy a growing concern among oversight bodies, organizations that fail to adequately secure sensitive customer data may face increasingly harsh legal and reputational consequences.

A large regional bank tackles new GenAI risks

A large regional bank established a mature governance program for models alongside an established AI risk management framework. However, this framework hadn't evolved to consider the unique risks presented by GenAI, particularly around the creation of new content.

While the bank incorporated some foundational responsible AI elements, it hadn't aligned with leading practice, including the National Institute of Standards and Technology GenAI risk management framework; nor did the company's risk taxonomy reflect the potential hazards posed by generative AI. It had no AI acceptable use policy and no externally facing AI charter to declare its ethical position on AI use.

The solution:

KPMG helped the company incorporate GenAI risks into the bank's responsible AI framework, extend its risk taxonomy, finalize and publish an AI charter, and implement trusted AI pillars consistent with leading practice.

Create an AI center of excellence

An AI center of excellence (COE), led by a chief AI officer and supported in an advisory role by appropriate subject matter experts, can help drive cultural alignment, education, and adoption of AI. The COE will also advise the business on how to use established governance pathways—built by defining risks and risk tolerance levels—to speed assessment, determine mitigation steps, and gain approvals.

This AI COE plays other key roles, particularly around evaluation and

readiness preparation as vendors introduce new AI features in their products. The COE inventories the enterprise's GenAI "wish list" and monitors the AI deployments and third-party solutions in consideration or in use. It also helps organizations re-evaluate third-party risk exposure by unearthing potential "gotchas" in contracts, including identifying misaligned guidelines and principles for responsible and trusted use.

An AI COE helps organizations re-evaluate third-party risk exposure by unearthing potential "gotchas" in contracts, including identifying misaligned guidelines and principles for responsible and trusted use.

Rethink the ways you onboard new GenAI models

When it comes to new GenAI models, even the largest financial services institutions are likely to buy or license rather than build: 48 percent of respondents in the KPMG August 2024 survey will buy or license and another 30 percent say they will mix build, buying, and partnering.⁴ Not all GenAI partners are created equally, so save time in risk identification and approval by redesigning your intake process for new GenAI models.

Intake should include guidance on suitability and security and a thorough review and approval process before development or implementation. Evaluating providers up front helps avoid introducing new risks that may not be a part of your current vendor onboarding process.

It is not necessary to reinvent your risk framework or your policies. Instead, review your existing policies to assess their suitability for GenAI. Where they are insufficient, adapt them to account for any gaps.

Creating one COE for a large regional bank

A large regional bank had already formed an AI SteerCo and a responsible AI subcommittee to oversee implementation of AI programs at the bank. Efforts across business units, however, largely operated independently of one another.

With information housed across disconnected systems, the bank was struggling to fully inventory its AI projects and risks. Its nascent AI COE primarily validated and supported use cases and solutions, instead of assuming leadership for enterprise-wide inventorying, education, and coordination. Because the bank lacked a process to evaluate and categorize AI systems at intake, it wasn't prepared to unify its GenAI technology in a single, enterprise-wide model.

The solution:

KPMG helped the client coordinate and enhance its existing COEs, evolving them into one core team empowered to centralize efforts around innovation and risk. A central team provided better line of sight into AI efforts, including new vendor intake and inventorying processes, and eliminated potential duplication and inconsistencies in GenAI project lifecycles.



⁴ Ibid.



Addressing model and nonmodel risk

Since most organizations already have robust governance in place for models (per supervisory guidance SR-11-7), treating model and nonmodel GenAI risk separately is logical. Many, though, have not accounted for GenAI in their nonmodel risk assessments.

For nonmodel risk management, one size rarely fits all. For one, existing frameworks built for traditional risks most likely don't account for risks from GenAI. For example, a chatbot that responds to customers' frequently asked questions has a different risk posture than a nonmodel that uses customer data to assign new loan interest rates. In addition, some GenAI risks don't carry the same level of institutional peril as others—leaked client data, for example, is probably far more serious than poor output from a generative AI-supported editing assistant.

Existing frameworks built for traditional risks most likely don't account for GenAI.

By classifying, tiering, and aligning risks to your organization's risk tolerance practices, you can better allocate resources to the most pressing issues. You can also fast-track the approval of projects that are less likely to create serious problems, such as an internal chatbot created to increase employee activity. That will help your organization build the trust it needs to proceed with go-live decisions.

The work doesn't stop there. Solutions incorporating GenAI must be continuously monitored. Testing for issues such as model drift and hallucinations is essential.

Going faster with good governance

The bank feels pressure to innovate and stay ahead of the competition, but its risk and governance functions are having a hard time keeping up with risk assessments and approvals. Its current AI lifecycle is creating challenges to innovate quickly, especially during research and development. Although it already has governance in place to account for the model risks of GenAI, its nonmodel risk governance practices have significant gaps, including for GenAI.

The solution:

KPMG helped the bank revise its AI lifecycle to balance speed and governance, while building a thorough risk evaluation process for nonmodel risk. The bank also implemented a responsible AI platform to support effective governance.

Support GenAI at scale with durable, repeatable governance

With a consistent framework, a thorough understanding of risks, and better organizational alignment around reviews, policies, and processes, organizations can deliver AI projects at scale.

The goal: A streamlined, understandable, and repeatable governance process that supports the rapid and safe deployment of GenAI solutions.



How KPMG can help

From strategy to design to operations, KPMG combines our deep industry experience, technical skills, and robust partner ecosystem to help business leaders create GenAI solutions they can trust. We help clients rapidly assess existing GenAI frameworks, perform maturity and benchmark analyses, and implement a GenAI governance process from intake to production.

Authors



Vivek Mehta
*Partner,
Financial Services*

Vivek Mehta is the Head of Trusted AI for Financial Services at KPMG, where he is responsible for the firm's Trusted AI strategy. In this role, Vivek spearheads the development and implementation of a comprehensive framework to ensure the responsible and ethical use of AI in the financial industry. Vivek collaborates with regulators to align the firm with new AI rules and regulations, ensuring compliance and promoting transparency in AI systems. He also works with cross-functional teams to drive the Trusted AI strategy forward.



Bryan McGowan
*Principal,
US Trusted AI Leader*

Bryan McGowan is a Principal in the KPMG Advisory practice and the leader of US Trusted AI for Consulting. Bryan leads the Trusted AI go-to-market effort for Risk Services, focusing on business process improvement, control design and automation, and managing risks associated with emerging technologies. Bryan has over 20 years' experience supporting clients on their automation and analytics journeys, including AI governance and framework development, regulatory compliance, and governance technology implementation.



Rob Fisher
*Partner, Advisory,
Sustainability Hub*

Rob Fisher is the US leader for the Financial Services Consulting business at KPMG. Previously, Rob served as the US firm's Sustainability leader and Global Executive Sponsor of the Trusted market proposition, including Trusted AI. Rob has focused on helping organizations build trust with customers, employees, investors, regulators, and business partners in their communities. Rob has more than 25 years of experience across various domains with a background in risk and regulatory within Financial Services.



Binoy Palakkal
*Principal, Advisory,
C&O Financial Services*

Binoy currently serves as the Financial Services AI Lead for KPMG. With more than 16 years of experience in the consulting space, he focuses on creating innovative solutions for financial services clients that use machine learning, NLP, statistical modeling and low code workflow automation. Binoy has more than 15 years of financial services experience, with specialties in intelligent automation, digitization, AI, managed services, business transformation, business process management, client onboarding, and client experience.

The authors wish to thank Dasanj Aberdeen, Michael Bender, Dhruvi Bhagat, Lisa Bigelow, Cynthia Bridges, Megan Marco, Daniel Miller, Brennan Morris, Rachel Napear, Kalpana Ramakrishna, Sandeep Sinha, and Lara Volpe for their contributions to this paper.

For more information, contact us:


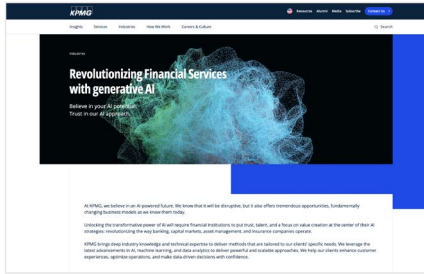
Vivek Mehta
Partner,
Financial Services
vivekmehta@kpmg.com

Bryan McGowan
Principal,
US Trusted AI Leader
bmcgowan@kpmg.com

Rob Fisher
Partner, Advisory,
ESG Hub
rfisher@kpmg.com

Binoy Palakkal
Principal, Advisory,
C&O Financial Services
bpalakkal@kpmg.com

Recommended reading:



Revolutionizing Financial Services with generative AI

The impact of artificial intelligence on the insurance industry

Artificial intelligence and model risk management

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Please visit us:  [kpmg.com](https://www.kpmg.com)

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.

DASD-2024-16013