IBM Institute for Business Value | Research Brief

Banking in the AI era

The risk management of AI and with AI

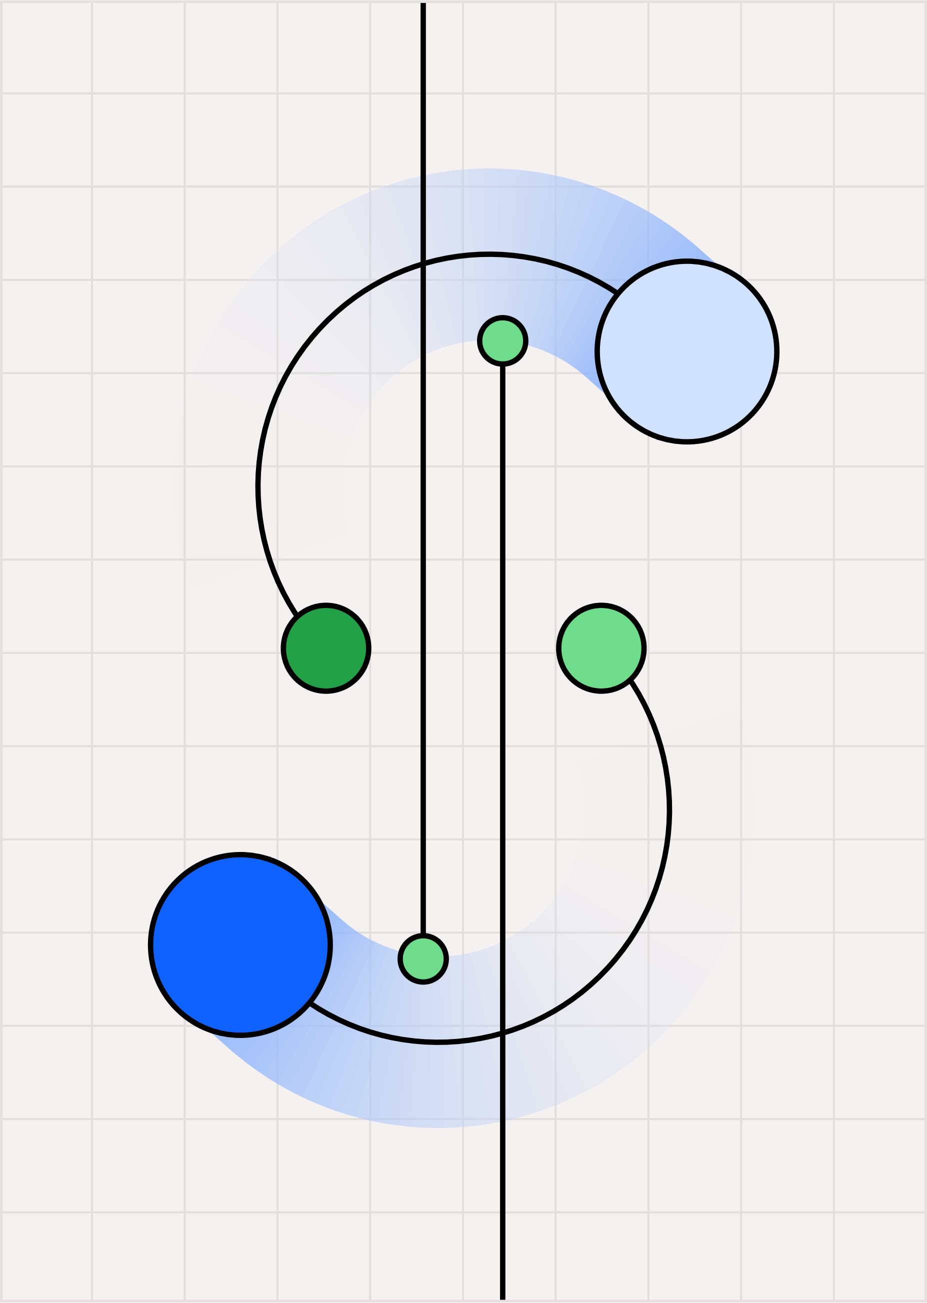

Foreword

Banking and financial markets have arrived at a pivotal juncture that is shaped by the accelerating pace of technological innovation. Also, intensified global trade tensions and shifting macroeconomics dynamics have created an uncertain business environment, requiring financial institutions to reevaluate business and technology strategies.

The rapid advancement of artificial intelligence (AI) is led by generative AI models and agentic AI. They stand out as instrumental forces to help institutions gain a competitive edge with personalized client engagement while building operational efficiencies. However, innovating with AI-powered tools must be balanced with a prudent approach that meets rigorous regulatory requirements. A reflection on risk and compliance practices in the era of AI becomes a prerequisite to scale the power AI enterprise-wide. These steps include strengthening validation capabilities and upskilling the workforce in a financial services world where every banker must be an AI risk manager.

Our research brief looks ahead at AI's role in addressing critical aspects of risk and compliance such as Know Your Customer (KYC), Anti-Money Laundering (AML) and fraud detection. It enriches a growing debate about validating AI models while accurately meeting demand for timely market deployments. And it also informs decisions about managing the risks of AI while successfully adapting operating models and adopting enhanced governance platforms. We trust that you will benefit from a set of actions to help your organization decode how to scale AI enterprise-wide.

Shanker Ramamurthy
Managing Partner, Global Banking
and Financial Markets
IBM Consulting

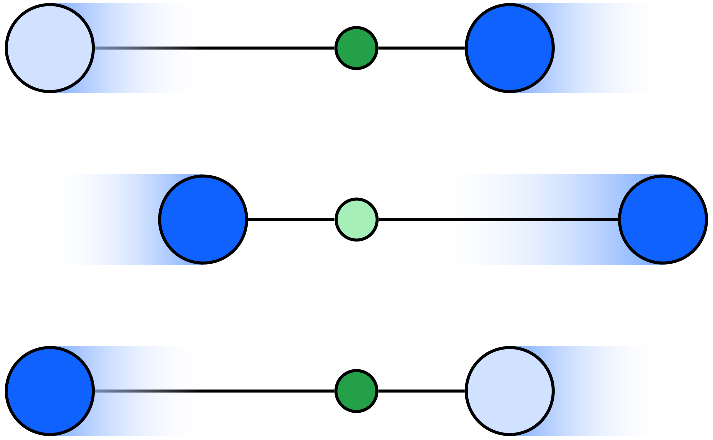

Key takeaways

Risk, compliance, and validation officers recognize AI’s potential to transform practices and provide competitive advantage—but also know they must strengthen their operations as a precondition for scaling AI enterprise-wide.

AI’s greatest promise is revealed in fraud detection, cybersecurity, and KYC.

61% of executives say fraud risk detection will provide the biggest boost to business value, with cybersecurity close behind at 52%. 45% of these executives also believe that AI will significantly transform AML and KYC processes.

Validation and risk control present a critical gap in people and skills.

61% of executives identify validation as a critical area for investing in people and skills, with risk control ranking second at 46%.

Enabling stress testing of AI models is the highest risk and compliance priority to support scaling AI enterprise-wide.

63% of executives recognize that stress test simulations are vital. They need to address AI system reliability and model deficiencies before deployment across the enterprise. Real-time risk controls are the next priority, according to 48% of respondents.

How exponential technology is impacting bankers and their clients

The financial services industry exemplifies relentless adaptation of exponential technologies such as hybrid cloud, AI, and generative AI (see Figure 1).

Initially, the emergence of online banking augmented traditional branch-centric models by enhancing accessibility for customers. This shift was eclipsed by the widespread smartphone adoption which catalyzed the ascent of digital banking as the preeminent channel for customer interaction.

However, first-generation digital interfaces have become inadequate for managing the intricacies of core banking services. This constrains the ability of financial institutions to deepen client engagement and offer new products and services on digital platforms.

By integrating cloud-based solutions, banks have markedly improved omnichannel delivery and provided granular insights into customer preferences. Yet the abundance of data alone has not been sufficient to fully transform client relationships to digital engagement.

FIGURE 1

Evolving financial services with exponential technologies

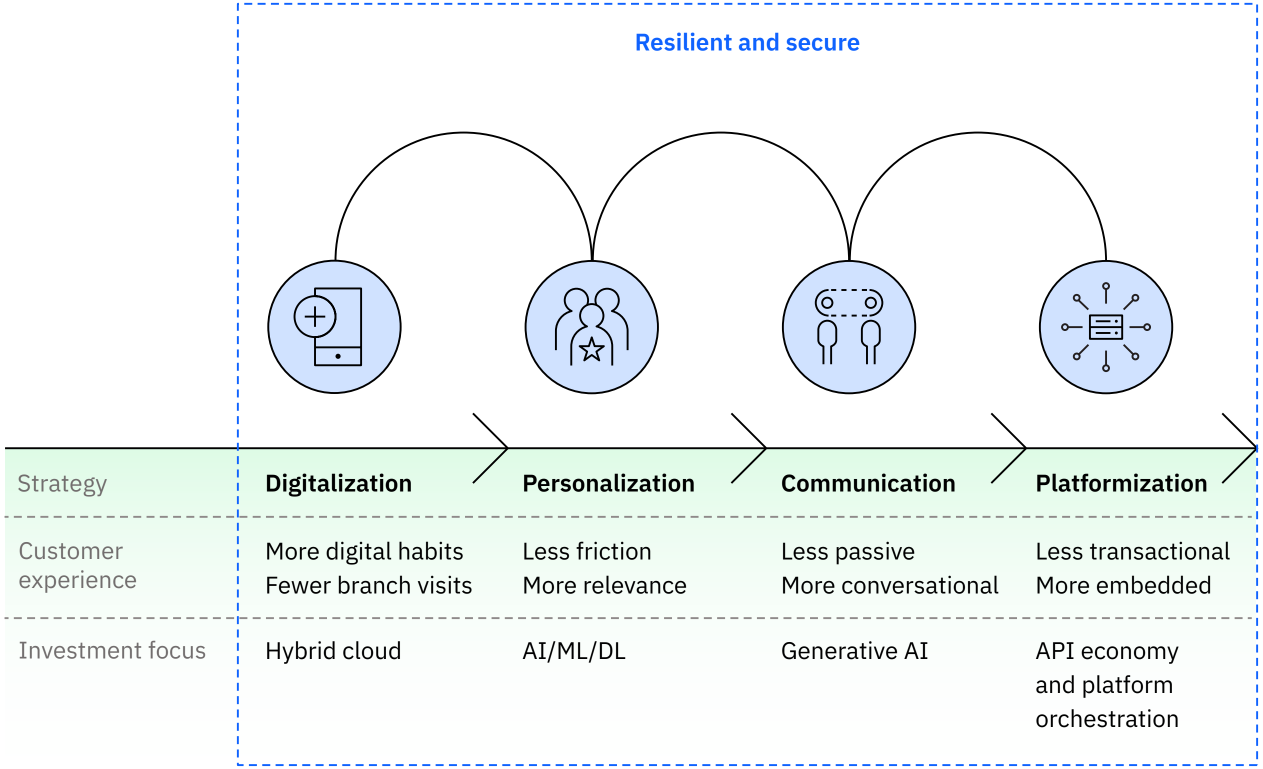

Strategy	Digitalization	Personalization	Communication	Platformization
Customer experience	More digital habits Fewer branch visits	Less friction More relevance	Less passive More conversational	Less transactional More embedded
Investment focus	Hybrid cloud	AI/ML/DL	Generative AI	API economy and platform orchestration

While mobile platforms adeptly serve self-directed users, personalized advisory services—a growing source of banking revenue—still rely heavily on human expertise. Without advanced conversational mechanisms, digital-only banks risk disengaging clients who value tailored financial counsel.

As AI evolves from machine learning to generative and agentic AI, it progressively bridges this divide to facilitate more substantive and individualized digital interactions.

Consumers increasingly use digital devices to shop for and purchase goods and services across industries. When they tap their banking applications, bank customers expect transactions to seamlessly weave into the fabric of their daily experiences. The notion of "banking everywhere" has evolved from convenience to fundamental expectation, characterized by real-time, context-sensitive financial solutions.

Nevertheless, this drive for innovation in the delivery of digital banking services must be judiciously balanced with an uncompromising focus on security and operational resilience. The digital frontier, while rich with potential, also introduces vulnerabilities that require rigorous oversight to safeguard institutional integrity and client trust. Within this risk-aware transformative context, a recent poll of banking CEOs indicated that over 60% recognize the necessity of embracing substantial risks to leverage automation's competitive benenfits.[1] This acknowledgment highlights the critical need for robust, scalable AI frameworks and a pervasive risk management culture where every banker is also an AI risk manager.

To elucidate this shift in bank culture, and its underlying processes, we surveyed 100 risk, compliance, and validation (RCV) officers from financial institutions with total assets exceeding $10 billion. Conducted across six countries with English proficiency in their banking sectors—the US, UK, Australia, India, Singapore, and Germany—the survey's findings are considered broadly representative of global trends. These insights affirm a dual perspective in the AI era of banking—the risk management with AI, and the enterprise-wide management of AI risks.

> "Innovation comes with new risks and refreshed compliance."

Maria Cristina Arrastia Uribe
former Business Vice President
Bancolombia[2]

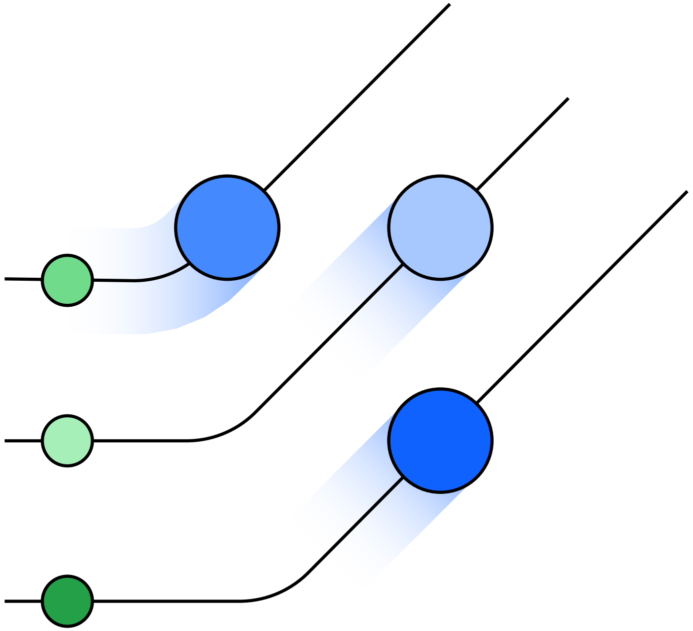

Perspective

Response to tariff-related tensions

Imposing tariffs on international trade is accelerating growing divergence in financial performance across economies and jurisdictions. Financial institutions are facing increased uncertainty with ultimate financial impacts depending on prevailing business models, geographical footprints, diversification, and adaptation of new technologies.

Investment banking and capital markets

Increased volatility in interest rates, foreign exchange, and equity markets may provide new opportunities for trading floors to earn additional fee income from increased trading and hedging activity. Technology-driven capabilities to expand trading operations while taming computing costs will also fuel financial performance.

Corporate banking and trade finance

Disruption of international supply chains may reduce demand for trade finance. Expected deterioration in the credit risk of large corporate clients is likely to impact risk-adjusted profitability in the short term. Technology-driven efficiency will be a differentiator in financial performance, and allows institutions to understand the impact on working capital across the expanded supply-chain of their clients.

Retail and commercial banking

Higher inflation and interest rates may heighten credit risks in the short term. When rebalancing lending portfolios, using technology to refresh datasets and recalibrate risk models will define impact on profitability.

Wealth and investment management

Capital market uncertainty may affect investor appetite for risk taking, potentially reducing fee income due to lower Assets Under Management (AUM). In response, banks can weather market turmoil by intensifying advisory activity—technology will be a differentiator in their capabilitity to advise at scale, provided that they can manage the risk of applying generative AI.

Managing risk and compliance with AI

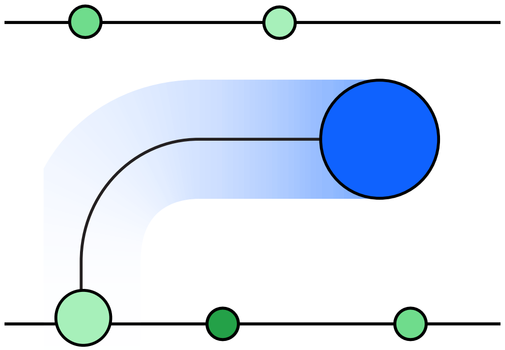

The survey reveals a telling snapshot of where RCV officers see AI's greatest promise—and its most daunting challenges.

The findings are as illuminating as they are provocative: 61% of respondents zero in on fraud risk detection as the area where they could harvest the most significant boost to business value within their functions (see Figure 2).

FIGURE 2

Potential impact of AI on risk and compliance business value

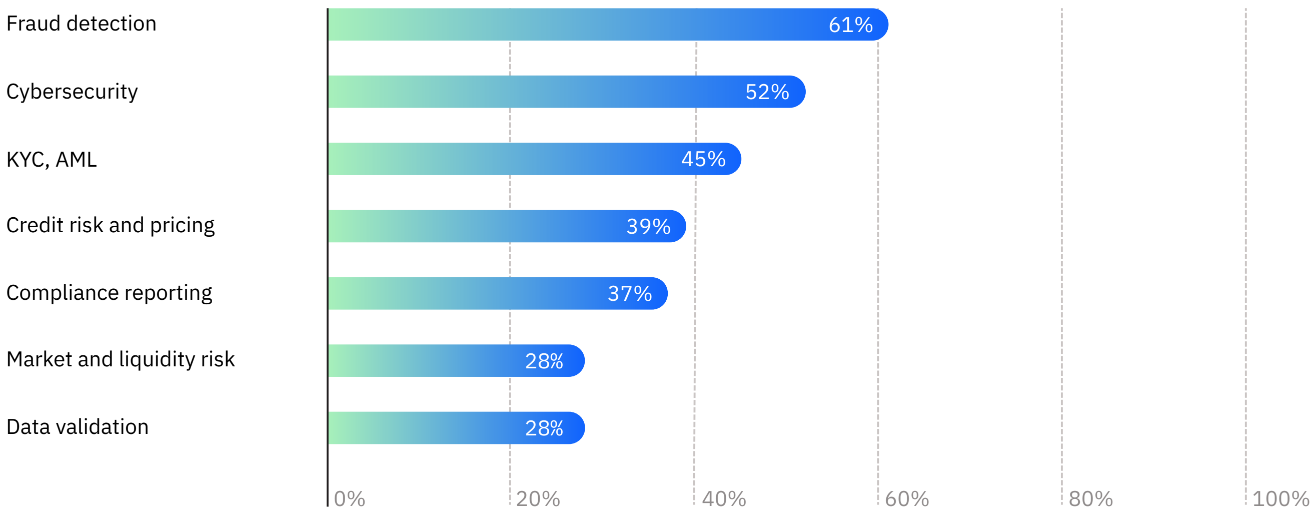

This is hardly surprising. In an era where fraudsters are as sophisticated as they are relentless, the ability to detect and prevent illicit activity in real time is nothing short of mission-critical. Close behind, 52% of respondents highlight cybersecurity, another domain where the margin for error is razor thin.

Rounding out the top three, 45% of these executives point to KYC and AML processes. These areas have long been the Achilles' heel of operation cost structures, bogged down by manual checks, outdated systems, and an ever-growing burden of regulatory requirements.

But here's where the narrative takes an intriguing twist. Despite the clear focus for using AI in these cost-intensive areas, fewer than 40% of respondents see more business-oriented processes as a major source of value. Specifically, only 39% highlight credit risk management and pricing, and a mere 28% focus on market and liquidity risk management.

This muted response suggests that while AI is poised to reshape certain aspects of RCV activities, traditional quantitative methods still hold the reins in others. When it comes to the core of banking risk calculus, the industry seems to be more cautious on betting the farm on AI. Nevertheless, it is the application of machine learning to operational data, such as cash flows, that enables banks to better address growth on digital channels and serve key client segments including small and medium enterprises.

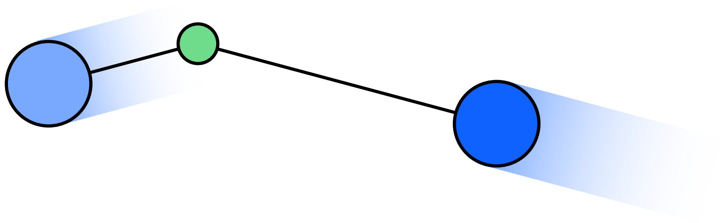

> "Improving the accuracy ratio through machine learning has been the enabler that makes our risk management approach fit for a digital multichannel world."

Davide Alfonsi
Group Chief Risk Officer
Banca Intesa Sanpaolo[3]

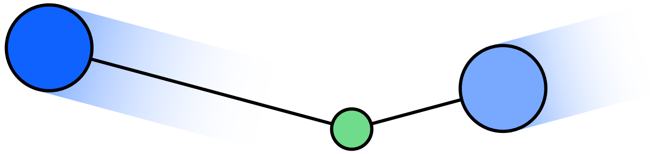

The two toughest transformation tasks for AI

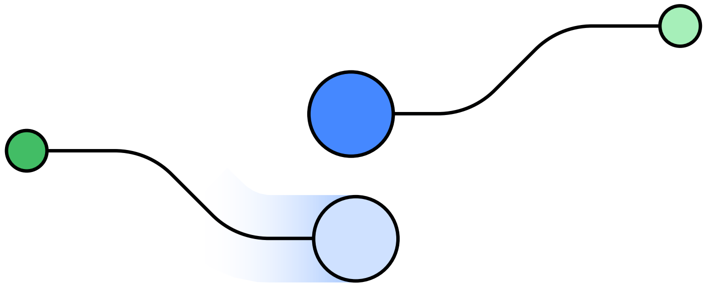

If fraud detection and cybersecurity come to mind first when reaping the benefits of AI, KYC and AML represent the Mount Everest of challenges (see Figure 3).

When asked about the complexity of using AI for key RCV processes, 43% of executives identify KYC and AML as the most daunting tasks to transform with AI.

FIGURE 3

The complexity conundrum in risk and compliance

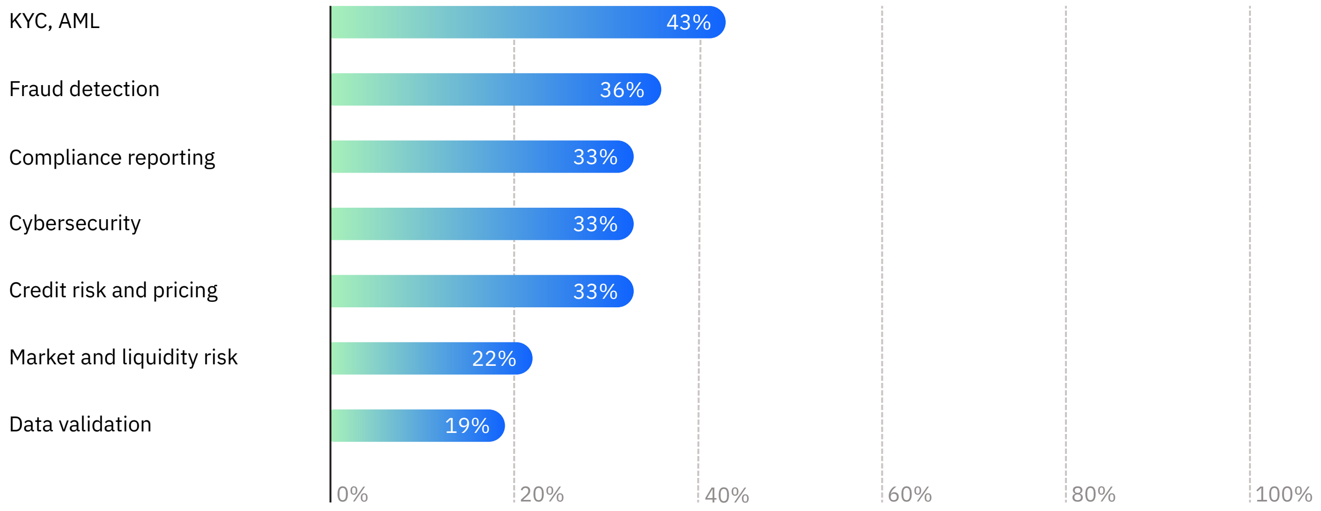

And it's no wonder why. These processes are labyrinthian by nature. They require navigating a maze of global regulations and verifying vast amounts of data with pinpoint accuracy. It's not just about crunching numbers; it's about understanding context, interpreting behavior, and making judgment calls—all in the shortest time possible.

This is where the journey takes a new turn. Enter agentic AI, the latest frontier in artificial intelligence. Unlike traditional AI, which follows more specific task definitions, agentic AI autonomously orchestrates a set of sub-agents that learn and adapt on the fly.

Imagine an AI system that doesn't just flag suspicious transactions but actively investigates them as it cross-references data points and analyzes patterns. The potential is tantalizing: a world where compliance is not just a box to check but a dynamic, proactive shield against financial crime and a constantly updated tool for regulatory compliance (see Perspective: "Know your customer with agentic AI").

Know your customer with agentic AI

Customer Due Diligence (CDD) constitutes a critical component of KYC frameworks. CDD serves as a regulatory mechanism to combat financial crimes, including money laundering and terrorist financing. This process requires financial institutions to exercise refined professional judgment to address escalating regulatory complexity and conduct periodic reviews.

The challenge is compounded by the nonstandardized and often ambiguous nature of the data involved, which renders CDD tasks manual, resource-intensive, and vulnerable to inconsistencies stemming from human interpretation. Traditionally, the process unfolds sequentially, involving multiple compliance officers collecting identity documents, verifying their authenticity, and assessing risks—a workflow that frequently spans days or weeks.

Agentic AI offers a groundbreaking solution to enhance CDD operations. By employing autonomous agents operating in parallel, this technology substantially reduces processing times while improving both efficiency and transparency, thereby bolstering trust in the system.

Additionally, the incorporation of generative AI facilitates the contextual interpretation of complex regulatory requirements and customer data, streamlining the process for near-zero idle time. This advancement not only elevates the precision of regulatory compliance but also reduces the likelihood of disputes, given higher transparency on decision-making processes and related reporting actions.

An example of how it works

In a typical KYC review, an orchestrator layer oversees specialized AI agents that concurrently collect data from diverse sources, such as government databases and customer submissions. Leveraging Optical Character Recognition (OCR) and Natural Language Processing (NLP) technologies, these agents authenticate documents with embedded expertise.

This parallel processing capability can significantly accelerate review timelines, transforming a traditionally protracted procedure into a swift and efficient operation as dependencies among AI agent tasks are processed seamlessly.

Perspective

Know your customer with agentic AI

(continued)

However, maintaining agentic AI effectiveness demands continuous updates to align with evolving regulations and rigorous risk management of AI models within operational frameworks. Drawing on IBM's expertise, this process involves systematic checkpoints.[4] These include:

Explicit goal specification.

Define comprehensive specifications of the agent's objectives, ensuring alignment with business goals, regulatory requirements, and ethical standards.

Goal-oriented guardrails.

Rules and dynamic mechanisms that actively restrict/guide what the agent is allowed to do toward achieving intended objectives.

Continuous monitoring of agent behavior.

Real-time monitoring of agent's alignment, including goal adherence and completion rate.

Value learning mechanisms.

Enable agents to continually learn and refine their understanding of human values and organizational priorities through training/fine-tuning/feedback on data.

Evaluation benchmarks and frameworks.

Leverage evaluation benchmarks to validate application-specific agents—such as software development and conversational agents—against common tasks. Assess their planning/reasoning capabilities, including task decomposition, multi-step reasoning, and reflection/recovery capabilities.

As regulatory landscapes continuously evolve on scope and details, the flexibility and scalability of agentic AI will prove indispensable in upholding more robust and cost-efficient KYC frameworks.

Managing the AI proving ground

Beyond the buzz coming from high-impact, high-complexity areas of AI application, an equally significant trend is quietly unfolding. RCV officers report that investments in compliance reporting are not identified among the top use cases in areas of expected business value. At first glance, this may seem counterintuitive. The answer lies in the practice of risk management: not just managing banking risks but risks associated with innovation (see Figure 4).

Compliance reporting may be seen as a less volatile domain for AI adoption. It's the perfect proving ground—a place where institutions can build confidence in their AI capabilities, smooth out the kinks, and gather the data needed to scale up to more complex applications. In essence, it's a strategic foothold, allowing banks to dip their toes into AI waters without diving headfirst into the deep end.

FIGURE 4

AI investment priorities in risk and compliance

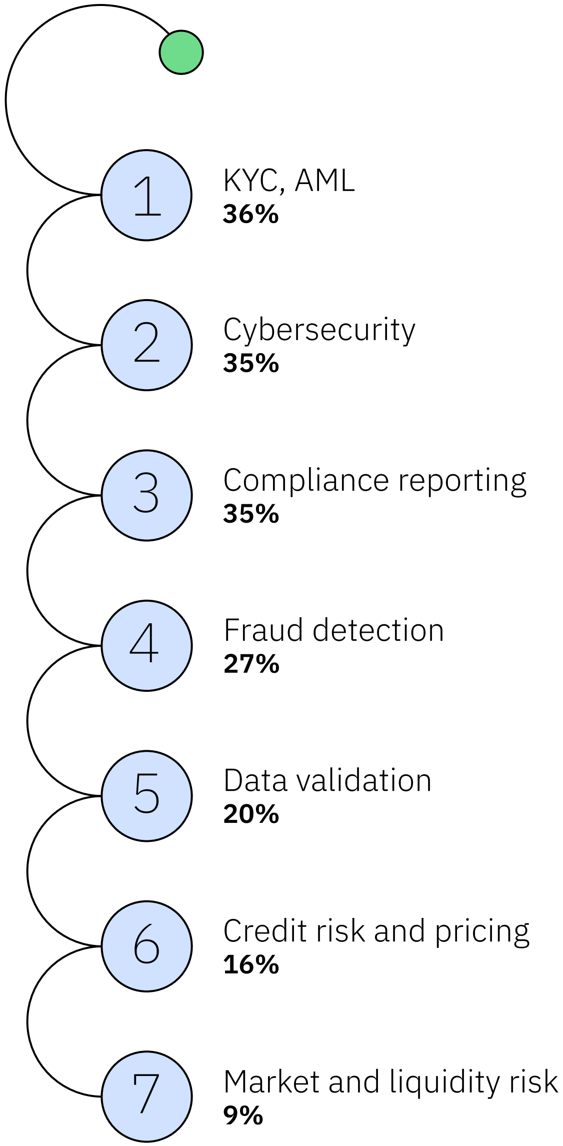

1	KYC, AML	**36%**
2	Cybersecurity	**35%**
3	Compliance reporting	**35%**
4	Fraud detection	**27%**
5	Data validation	**20%**
6	Credit risk and pricing	**16%**
7	Market and liquidity risk	**9%**

Addressing the AI talent gap

Given all the changes required to modernize risk and compliance in the AI era, accessing AI talent is a key consideration.

Significant expertise deficiencies are most apparent in the fields of model validation (61%) and risk control (46%). These disciplines are indispensable when helping to ensure the operational integrity and security of AI systems. Without skilled professionals in these areas, even the most sophisticated solutions may become liabilities rather than reliable tools, jeopardizing their intended strategic benefits (see Figure 5).

FIGURE 5

AI talent gap in risk and compliance functions

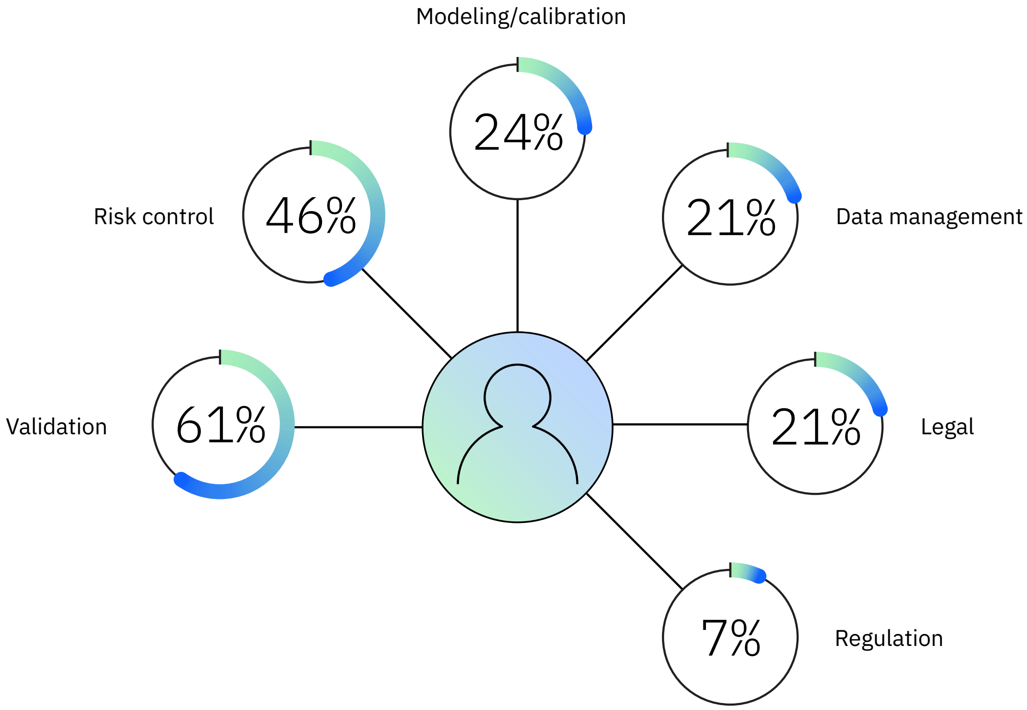

Scaling AI at the enterprise level

As RCV officers progressively adopt AI to transform business-critical operations, they also understand that managing these operations is essential for scaling AI across the enterprise.

In this context, the accelerated life cycle of technology hardly aligns with the more measured pace of validating and applying risk controls. Institutions must invest in strengthening risk and compliance functions, while helping RCV officers transform the way they work in a more automated and AI-driven financial services space.

We asked survey respondents about which initiatives they believe will empower RCV functions to improve the scaling of AI across their institutions, while managing associated risks (see Figure 6).

FIGURE 6

RCV initiatives to support scaling AI across the enterprise

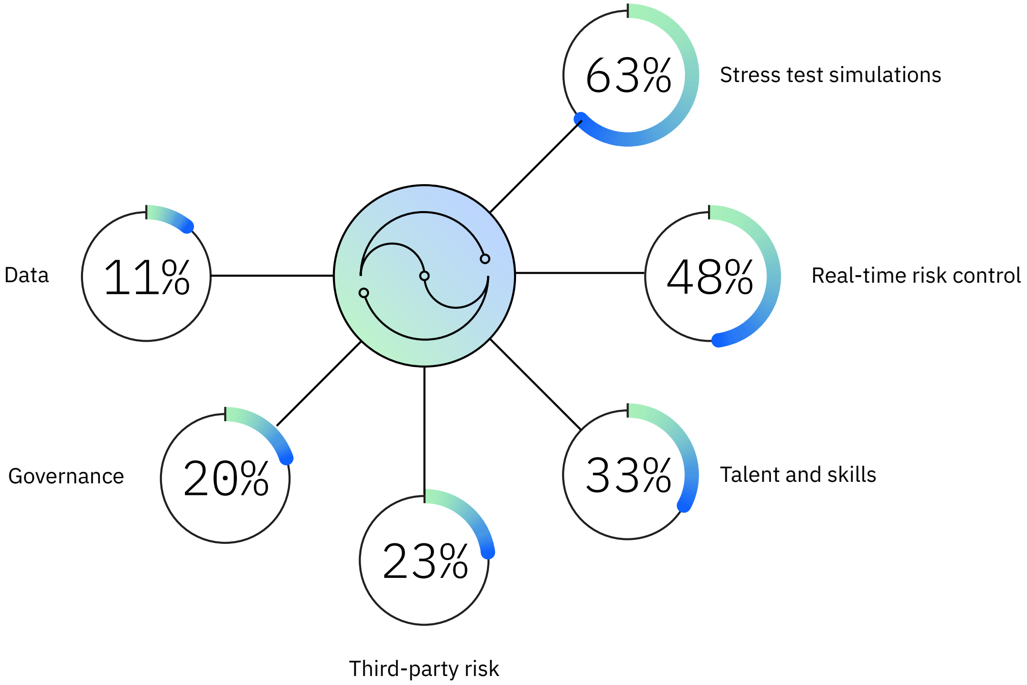

These findings underscore three pivotal initiatives:

- **Stress test simulations.** Selected by 63% of respondents as the foremost priority, simulations are vital to rigorously assess AI system viability and reliability. They proactively identify potential model weaknesses before operational deployment.

- **Real-time risk controls.** Endorsed by 48% of respondents, these mechanisms provide ongoing monitoring. They also have capacity to take immediate corrective measures to help ensure AI systems don't drift and hallucinate.

- **Talent and skill development.** Noted by 33% of respondents, this emphasizes the necessity of fostering workforce proficiency to oversee and enhance AI technologies.

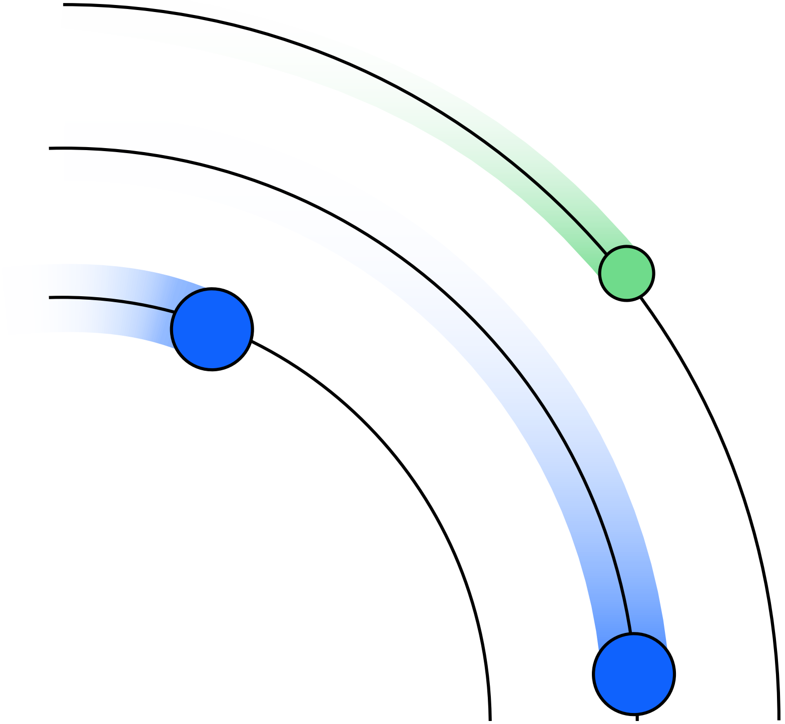

Together, these initiatives highlight the importance of comprehensive understanding, meticulous control, and adept risk management competency in successfully managing AI integration across the enterprise.

> "Regulatory compliance needs to be there, and financial risks managed professionally. But that does not mean the traditional ways of working as a financial institution will be applied. Combining the two worlds fruitfully means everything from new skillsets to new ways of working to how teams are built, what culture is needed, how success is evaluated, and how people get incentivized."

Christoffer Malmer
Chief Financial Officer
SEB[5]

Risk-based tiering of AI use cases

While these essential steps are being taken, risk and compliance functions are tasked to adopt a proportional approach that balances prudent risk management with the speed of AI innovation.

The most sought-after approach by risk and compliance executives is to categorize use cases according to overall risk levels (see Figure 7).

This structured methodology helps ensure that applications with the most heightened risk concerns are scrutinized as a priority, facilitating a disciplined and effective allocation of resources. Interestingly, cybersecurity-based tiering receives limited emphasis, although robust cybersecurity should be the starting point for any AI strategy.

FIGURE 7

Key methods for a proportionate level of controls

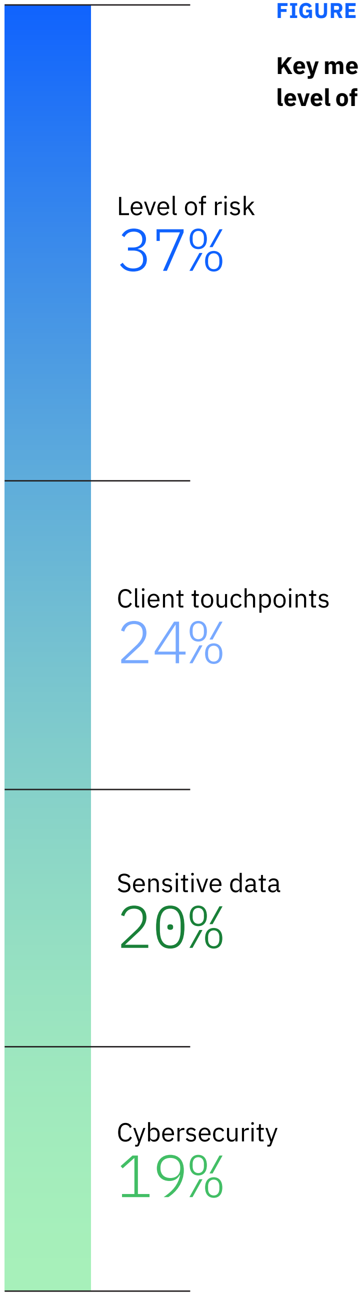

Addressing deficiencies in real-time monitoring

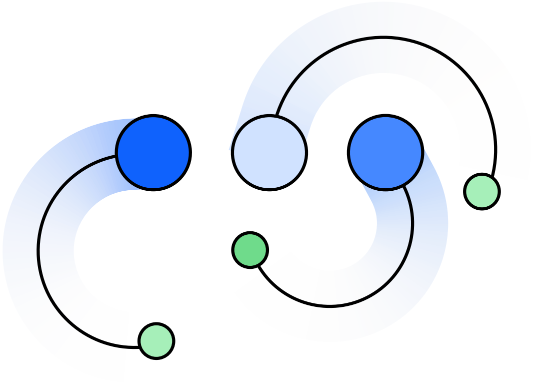

Real-time risk monitoring is recognized as the second most critical risk and compliance initiative for scaling AI across enterprises (see Figure 6).

FIGURE 8

Institutions that always apply these risk management practices to high-risk AI use cases

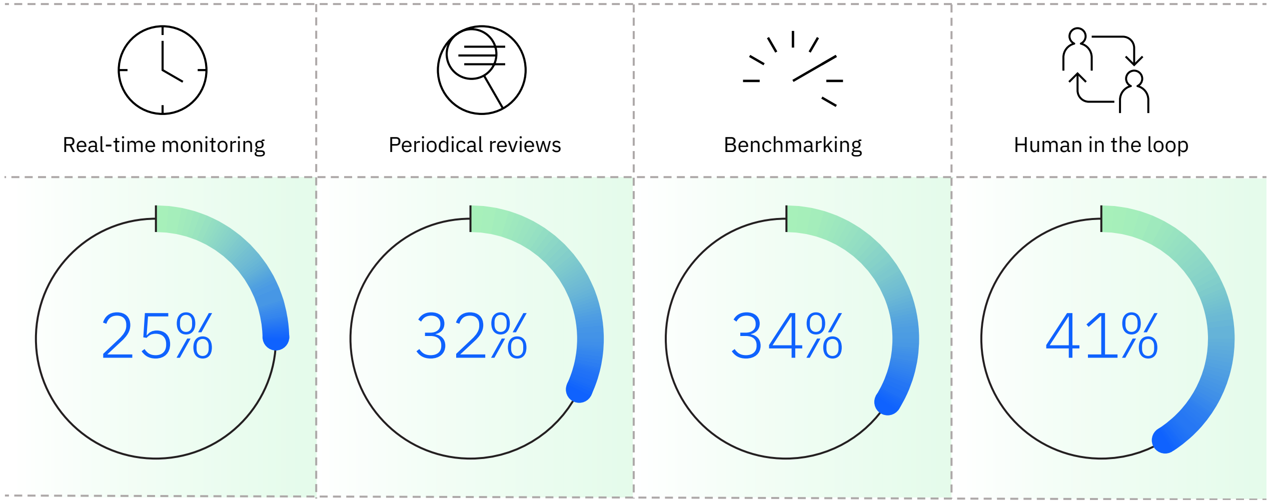

Despite the significance of this practice, only 25% of surveyed professionals report that they consistently apply it to high-risk use cases (see Figure 8).

Rectifying this discrepancy is essential to safeguard the reliability and efficacy of AI deployments. Banks require deliberate investments in platforms and capabilities in order to develop AI models and use cases with built-in and real-time monitoring tools.

Action guide

In a business environment of heightened economic uncertainty, AI is emerging as a powerful tool for financial institutions to build resilience, enhance KYC and AML, detect fraud, and boost cybersecurity. At the same time, adopting AI can also add elements of risk, with only 25% of banks maintaining real-time oversight of critical AI applications— an alarmingly low percentage.

Here are six strategic steps you can take now to strengthen how your institution manages risk and regulatory compliance in the era of AI.

Strengthen talent development.

A comprehensive AI literacy program should cover technical skills such as programming and data analysis, as well as nontechnical areas including AI ethics, legal implications, and social impact. Take in interdisciplinary perspectives to understand AI's broader implications and foster critical thinking about bias and fairness. Customized training can align with specific organizational needs, such as improving customer-facing AI tools or optimizing internal processes. Promoting psychological safety helps ensure diverse voices contribute to identifying risks, align training with business goals, and prepare employees to navigate AI's complexities responsibly.

Build organizational resilience.

Agentic AI systems that autonomously manage critical operations also introduce complex dependencies, where a single failure could disrupt entire workflows. To counter this, implement fallback designs so essential functions can persist during outages. Dependency mapping identifies critical failure points, and resilience testing—such as stress tests or failure simulations—prepares systems for real-world challenges. Sandbox environments allow safe experimentation with AI updates and compartmentalization further enhances stability by isolating system components, limiting the scope of any failure.

Enhance real-time risk monitoring.

A unified, real-time monitoring framework should track key metrics—performance, latency, quality, and cost—across all AI applications, enabling early detection of issues such as model drift or anomalies. This approach provides actionable insights for improvement, reducing risks before they escalate. Centralized monitoring allows new AI deployments to tap into existing infrastructure rather than building anew. Automated controls enforce compliance with operational and ethical standards, embedding risk management into system design.

Embed cybersecurity practices.

AI expands vulnerabilities, so organizations must collaborate with cybersecurity and risk teams to update frameworks and address emerging threats specific to AI. Red team testing simulates attacks and exposes weaknesses before they're exploited. Integrating security insights into governance dashboards provides a holistic view of risks across on-premises and cloud-based AI systems. Cross-functional teamwork among developers, security experts, and leadership fosters a security-first mindset. A controls-as-code library standardizes security measures—such as encryption or access controls—across AI lifecycles, from development to deployment.

Manage third-party risk.

As organizations increasingly rely on external AI vendors, clear accountability frameworks define responsibilities across the supply chain. Cross-functional teams—including procurement, legal, and AI specialists—should craft vendor guidelines to address risks such as data mishandling or noncompliance. Contracts must mandate transparency, requiring vendors to disclose changes in AI functionality, security practices, or bias mitigation efforts. Regular audits and performance reviews can verify ongoing adherence to standards, while disclosure requirements around data use, such as copyrighted material and benchmark models, build trust.

Optimize the software development lifecycle (SDLC).

To balance innovation with risk management, clear guidelines should dictate when AI-generated code is permissible and when human oversight is mandatory. Developer training on prompt engineering, code validation, and bias detection helps ensure responsible AI use, supported by ethics and security playbooks. Automated code analysis catches vulnerabilities or compliance issues pre-deployment, while senior developer reviews verify quality. Rigorous testing supports robustness and maintaining records of AI versus human contributions aids accountability. An AI risk review board should oversee tool integration into the CI/CD pipeline.

Authors

Shanker Ramamurthy
Managing Partner, Global Banking and Financial Markets
IBM Consulting
sramamur@us.ibm.com
linkedin.com/in/shankerramamurthy

Shanker leads core banking modernization and payments. He's an IBM Acceleration Team member, recognized globally for his patents, white papers, and as one of Euromoney's top 50 most influential consultants.

Marc Haddad
Senior Partner, FSS Leader, EMEA
IBM Consulting
marc.haddad@ibm.com
linkedin.com/in/marcadad

Marc has guided banking and insurance leaders through strategic and operational transformations. He drives major modernization projects by introducing innovative practices in the financial sector. A former PwC and Accenture partner, he is an ESSEC graduate and Chevalier de l'ordre national du mérite.

Paolo Sironi
Global Research Leader, Banking and Financial Markets
IBM Institute for Business Value
paolo.sironi@de.ibm.com
linkedin.com/in/thepsironi

Paolo leads IBM IBV research in Banking and Financial Markets. He's a respected fintech voice and hosts The Bankers' Bookshelf podcast. He is a bestselling author on digital transformation and quantitative finance.

Prashant Jajodia
Managing Partner,
FS Sector Leader, UKI
IBM Consulting
prashant.jajodia@uk.ibm.com
linkedin.com/in/prashantjajodia

Prashant has a 25+ year record as a global financial services leader. He worked with Citi as head of technology for EMEA Treasury and built an electronic trading platform for the Indian stock market. Prashant has in-depth experience with AI, cloud, and digital transformation programs driving modernization for improved customer experience.

Rashmi Das
Managing Client Partner, US Banking and Financial Markets Industry Leader
IBM Consulting
Rashmi.Das@ibm.com
linkedin.com/in/rashmidas

Rashmi has over 27 years of experience working with commercial and investment banks and asset and wealth management firms on business transformation. She has leveraged exponential technologies for global and US financial institutions including agentic AI, hybrid cloud, RPA, blockchain and more recently, quantum, to deliver significant business impact.

Asanga Lokusooriya
Lead Client Partner and
FS Sector Leader
IBM Consulting
asanga.lokusooriya@ibm.com
linkedin.com/in/asangalokusooriya

With 25+ years of experience, Asanga drives business outcomes across industries through data and generative AI. Combining deep technical expertise with strategic vision, he bridges emerging technology and enterprise value, grounded in hands-on roles and scaled delivery expertise.

Authors

Liaquat Parkar
Executive Partner, Banking
IBM Consulting
liaquat.h.parkar@ae.ibm.com
linkedin.com/in/liaquatparkar

Liaquat partners with C-level banking executives to drive AI-enabled digital agendas that deliver real-time, seamless, and personalized experiences for customers and employees. With deep expertise in cloud, data, and AI, he helps banks reduce operational costs, boost revenue, and lead large-scale transformation initiatives.

Fabio Carvalho Pessoa
Vice President and Senior Partner,
Financial Services Sector Leader,
Latin America
IBM Consulting
pessoa@br.ibm.com
linkedin.com/in/fabio-carvalho-pessoa-12a7b452

Fabio is an experienced IT executive with 30 years of experience in consulting, digital transformation, AI, cloud migration, and outsourcing. He has proven leadership in sales, portfolio, and client management, specializing in the financial services industry and leading major transformation projects for banks throughout Brazil and Latin America.

Yuuji Sonku
Managing Partner and Sector Leader,
BKFM, Japan
IBM Consulting
sonku@jp.ibm.com
linkedin.com/in/yuji-sonku-b896501ab

Yuuji joined IBM in 1998 and has been helping Japanese financial institutions solve their challenges for more than 25 years. He is a recognized IT thought leader in the financial industry in Japan.

Contributors

Michal Chorev and *Marjolijn Herman*

How IBM can help

Modern financial institutions demand modularity, security, openness, AI-driven capabilities, and collaboration on a hybrid cloud. At IBM, we empower you to elevate customer experiences, modernize core banking infrastructures, pioneer innovative payment solutions, and transform enterprise risk management. Learn more at ibm.com/industries/banking-financial-markets

Research methodology

IBM IBV, in partnership with Oxford Economics, surveyed 100 bank executives with responsibilities in risk, compliance, and validation in the US, UK, Australia, India, Singapore, and Germany. Participants were asked about their current use and expectations about AI related to risk management and innovation as applied to the banking industry. Questions were posed in various formats including multiple choice, numerical, and Likert scale. Respondents were equally distributed across the following roles: CRO, CCO, and CVO. The survey was conducted in May 2025. In addition, insights and recommendations in this report draw on case studies and direct extensive work with banking clients around the world.

IBM Institute for Business Value

For two decades, the IBM Institute for Business Value has served as the thought leadership think tank for IBM. What inspires us is producing research-backed, technology-informed strategic insights that help leaders make smarter business decisions.

From our unique position at the intersection of business, technology, and society, we survey, interview, and engage with thousands of executives, consumers, and experts each year, synthesizing their perspectives into credible, inspiring, and actionable insights.

To stay connected and informed, sign up to receive IBM IBV's email newsletter at ibm.com/ibv. You can also follow us on LinkedIn at https://ibm.co/ibv-linkedin.

Notes and sources

1. Marshall, Anthony, Cindy Anderson, Christian Bieck, and Spencer Lin. *The CEO's guide to generative AI: Risk management.* IBM Institute for Business Value. 2024. https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/ceo-generative-ai/ceo-ai-risk-management
2. Ramamurthy, Shanker, John J. Duigenan, Hans Tessellar, Héctor Arias, and Paolo Sironi. *Embedded finance: Creating the everywhere, everyday bank.* IBM Institute for Business Value in partnership with BIAN and Red Hat. September 2023. https://ibm.co/embedded-finance
3. Ibid.
4. Chorev, Michal, Richie Paul, and Joe Royle. *Agentic AI in Financial Services: Opportunities, Risks, and Responsible Implementation.* IBM Consulting. May 2025. https://www.ibm.com/downloads/documents/gb-en/12f5a71117cdc329
5. Ramamurthy, Shanker, John J. Duigenan, Hans Tessellar, Héctor Arias, and Paolo Sironi. *Embedded finance: Creating the everywhere, everyday bank.* IBM Institute for Business Value in partnership with BIAN and Red Hat. September 2023. https://ibm.co/embedded-finance

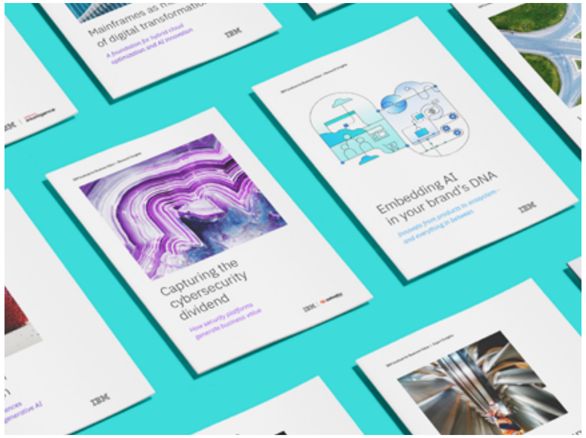

Subscribe to our IdeaWatch newsletter

Just the insights. At your fingertips. Delivered monthly.

Brought to you by the IBM Institute for Business Value, ranked #1 in thought leadership quality by Source Global Research for the second consecutive year.

Research-based thought leadership insights, data, and analysis to help you make smarter business decisions and more informed technology investments.

Subscribe now: ibm.co/ideawatch

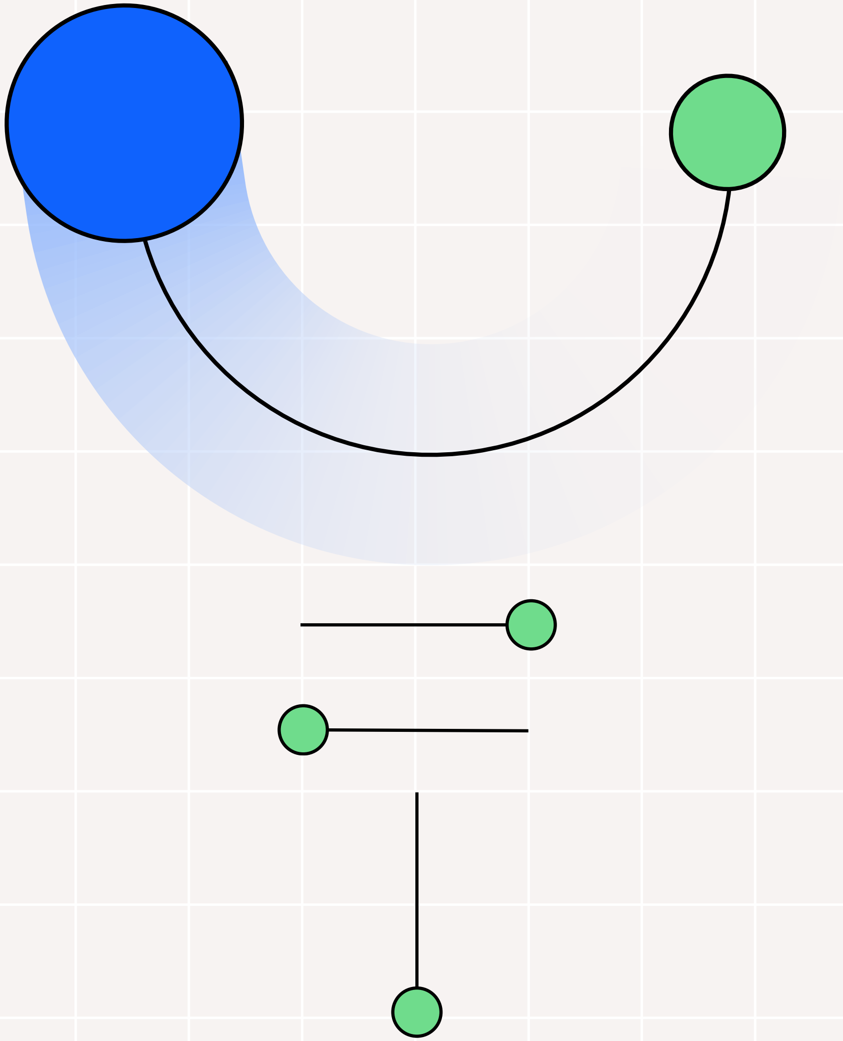

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | June 2025

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

This report is intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. IBM shall not be responsible for any loss whatsoever sustained by any organization or person who relies on this publication.

The data used in this report may be derived from third-party sources and IBM does not independently verify, validate or audit such data. The results from the use of such data are provided on an “as is” basis and IBM makes no representations or warranties, express or implied.

1377429fe5bde65a-USEN-01