Deloitte.

The Financial Services AI Dossier

A selection of high-impact use cases

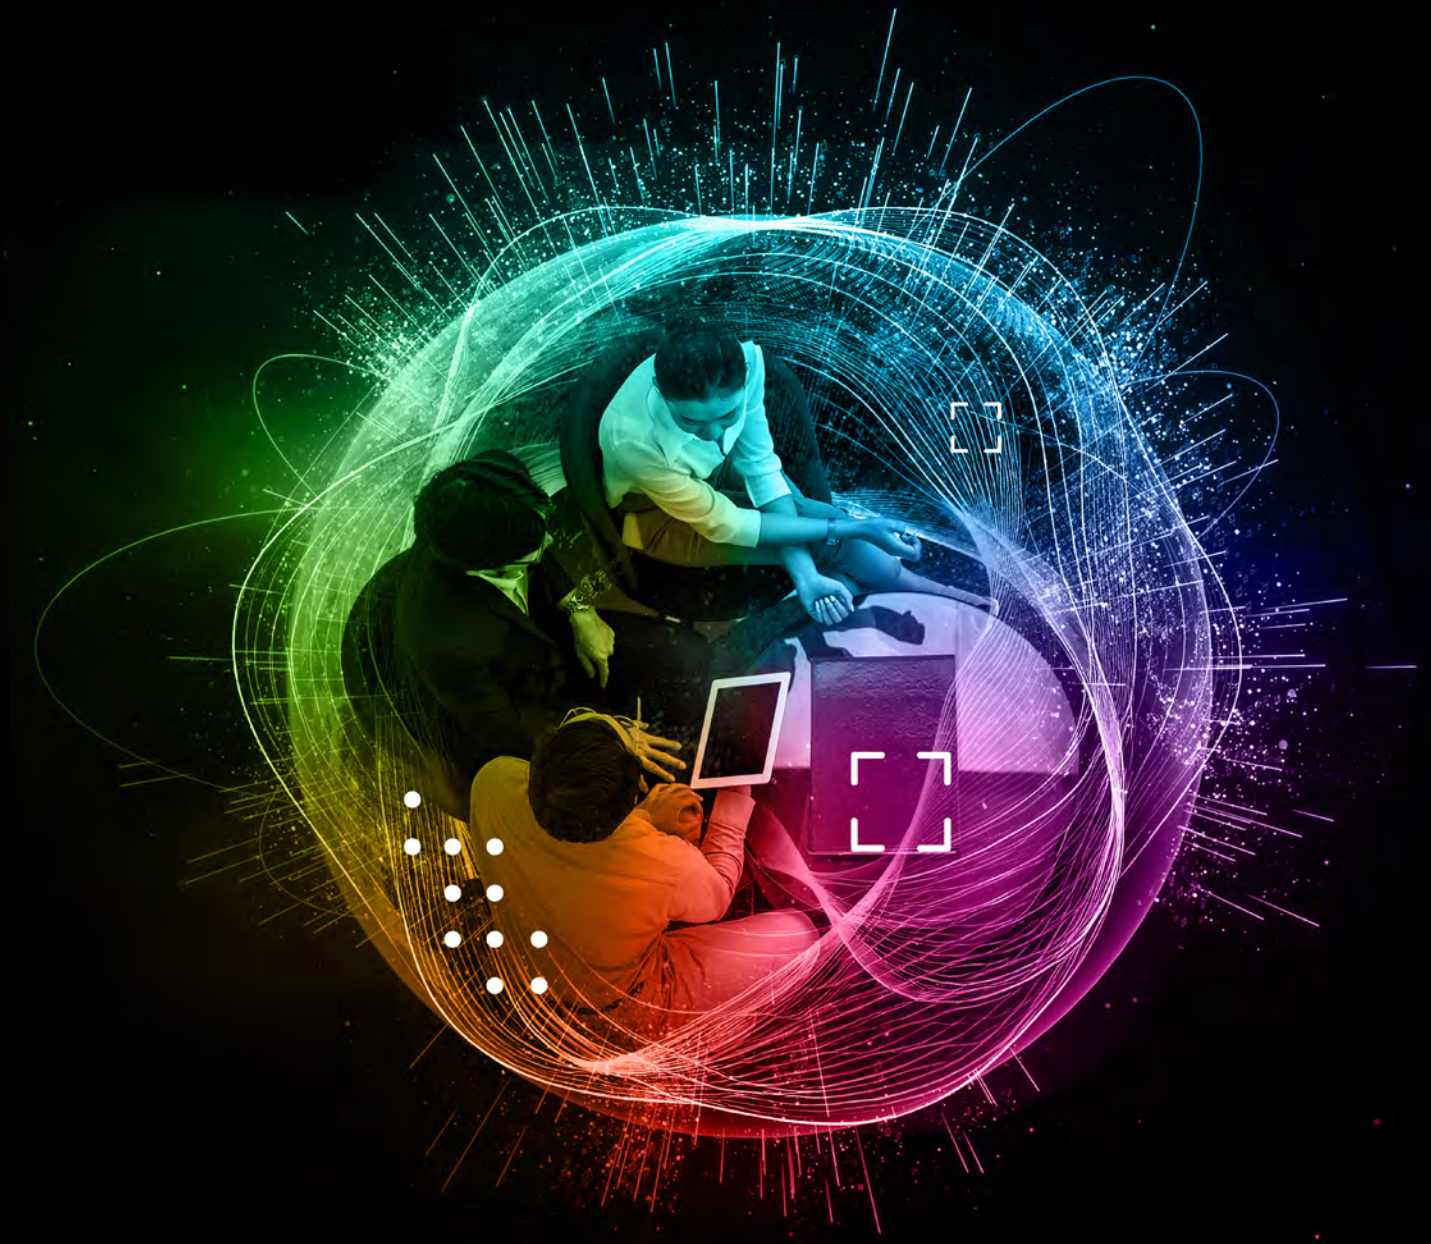

The Deloitte AI Institute™

About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect all the different dimensions of the robust, highly dynamic, and rapidly evolving Artificial Intelligence ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With™."

The Deloitte AI Institute aims to promote the dialogue and development of AI, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, the future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you are in: whether you are a board member or a C-Suite leader driving strategy for your organization—or a hands-on data scientist bringing an AI strategy to life—the Deloitte AI Institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet-ups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute

Foreword

Artificial intelligence (AI) continues to advance by leaps and bounds, delivering breathtaking capabilities once thought to be far off in the future. With a remarkable capacity to understand complex inputs and generate valuable outputs—and the rapidly emerging ability to execute real-world actions—AI is opening the door to innovations and new ways of working that were almost unthinkable just a few years ago.

As the AI landscape evolves, so does this compendium. Our latest edition features 86 of the most compelling use cases for AI across six major industries:

- **Consumer**
- **Energy, Resources & Industrials**
- **Financial Services**
- **Government & Public Services**
- **Life Sciences & Health Care**
- **Technology, Media & Telecommunications**

For each of these industries, we explore innovative uses for AI that can address enterprise challenges in new ways, expand and improve capabilities in every business function, and deliver advantages in efficiency, speed, scale, and capacity. To provide further context and clarity, each case specifies the primary business function it supports and whether agentic AI is used. These labels are presented for informational purposes, helping you quickly grasp the intention and scope of each case.

Of course, every powerful tool presents potential risks, and AI is no exception. To help you better understand and manage the risks associated with AI, we use Deloitte's Trustworthy AI™ framework throughout this compendium to illuminate factors that contribute to trust and ethics in AI deployments, and to offer practical steps for strengthening governance and risk mitigation. The specific objective of our Trustworthy AI™ framework is to help organizations create AI systems that are (1) fair and impartial, (2) robust and reliable, (3) transparent and explainable, (4) safe and secure, (5) responsible and accountable, and (6) private.

Given AI's rapidly expanding scope and reach, this compendium offers just a glimpse of what the technology can do. Our goal is to convey what AI is currently capable of, and even more important, to inspire the next wave of AI-driven innovation. As AI technology continues to improve and organizations increasingly embrace it, we anticipate even more impressive and compelling use cases in the future—including those that have yet to be imagined.

We hope the use cases highlighted here will spark new ideas, provide a foundation for successful deployments, and set organizations on a path to harness the maximum value from this powerful new technology.

Nitin Mittal
Global Head of Generative AI
Deloitte Consulting LLP

Beena Ammanath
Global Deloitte AI Institute Leader
Deloitte Consulting LLP

Jim Rowan
US Head of AI
Deloitte Consulting LLP

The Financial Services AI Dossier

As AI capabilities mature and regulatory clarity improves, AI is moving from experimental to essential—powering financial services businesses that are smarter, faster, and more responsible.

In an environment where trust, precision, and speed are paramount, AI is enabling financial services firms to detect risk earlier, serve customers more effectively, and compete with greater agility. Whether it's fraud prevention in real time, hyper-personalized marketing, predictive trading, or automated claim reporting, AI is shifting how institutions create value and manage complexity across banking, capital markets, and insurance.

AI is already having a transformative impact on knowledge work, helping people do their jobs more efficiently and effectively and supporting human decision-making with faster, deeper analysis of complex data patterns. It is also helping behind the scenes to triage and manage cyber threats, develop and debug new computer programs, and even generate synthetic data to train smarter AI models.

Agentic systems go even further, enabling AI to handle tasks such as algorithmic trading while keeping workers in the loop for oversight.

AI's powerful capabilities can help financial institutions address rising pressure to modernize legacy systems, reduce costs, and respond to evolving customer expectations. However, they also bring heightened scrutiny around fairness, transparency, and accountability. Firms that want to harness the full power of AI will therefore need to invest not only in improved model performance but also in foundational elements such as governance, data quality, and organizational readiness.

AI's powerful capabilities come at an opportune time to help financial institutions address rising pressure to modernize legacy systems, reduce costs, and respond to evolving customer expectations.

Note: *The tags below each use case indicate its primary business function and whether Agentic AI is used.*

Tags | **Primary business function** | **Agentic AI**

AI-powered risk management and regulatory compliance

Using AI agents to provide 24/7 risk and compliance monitoring

Agentic AI systems can act as an always-on compliance and risk management team, with specialized agents continuously monitoring risks across a financial institution's operations.

ISSUE/OPPORTUNITY

Financial services firms must continuously monitor for risks—ranging from fraudulent transactions to compliance violations—across a vast landscape of operations. At scale, with millions of customers and complex regulatory demands, risk management can be costly, labor intensive, and subject to error. Missed risks can lead to massive penalties, reputational damage, and systemic vulnerabilities. Yet, human teams are limited in their ability to keep pace with the overwhelming volume of transactions, communications, and rule changes that require their oversight.

Traditional monitoring approaches can be siloed and reactive, identifying issues only after damage has been done. Multi-agent AI offers a proactive, comprehensive solution that can scan for risks, interpret new regulations, and test institutional resilience under different scenarios.

Tags: **Compliance & Risk** | **Agentic AI**

HOW AI CAN HELP

Risk detection and monitoring
Agents that specialize in different risk areas (such as credit card fraud, anti-money laundering (AML), or cybersecurity) can automatically detect and monitor risks. They can also work together, sharing alerts across domains. For example, if a cybersecurity agent detects a breached account, it can warn a transaction monitoring agent to increase scrutiny on that account.

Regulatory change monitoring
Other specialized agents can parse new regulations using natural language processing, interpret rules and guidelines, map new requirements against company controls to identify compliance gaps, and then draft reports or suggest control changes for humans to implement.

Stress testing and scenario analysis
A stress testing agent can simulate scenarios such as interest rate jumps or credit crunches, coordinating with agents representing different departments (e.g., loans, trading, deposits) to assess institutional safety.

AI-powered risk management and regulatory compliance

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Financial firms face massive penalties and reputational damage for compliance failures. As such, multi-agent systems should be validated against historical fraud cases, past audits, and synthetic stress scenarios to ensure their outputs are dependable and comprehensive.

Safe and secure

Pulling data from various databases to feed AI agents requires significant IT work and raises potential security concerns. Given the sensitive nature of risk management data, agentic systems must be secured against unauthorized access or accidental disclosure of private customer information.

Transparent and explainable

Risk management is overseen by humans who are personally liable and may be wary of trusting "black box" agents. Explainability is key. Agents need to provide clear audit trails (for example, explaining why a transaction was flagged, and including references to rules or past cases). This clarity makes AI outputs easier for human officers and regulators to trust and validate.

Fair and impartial

Biases in data sources or training models can result in unfair or inaccurate risk assessments, raising concerns for customer trust, brand reputation, and regulatory compliance.

POTENTIAL BENEFITS

Higher assurance and lower compliance costs

Proactive and comprehensive monitoring reduces the likelihood of missed issues, helping firms avoid massive penalties while reducing the high labor costs associated with traditional compliance processes.

Reduced risk exposure and faster detection

Always-on agentic systems can sift through transactions, emails, and news at speed and scale, reducing a firm's exposure by identifying risks much sooner.

Improved audit and regulatory relations

AI agents specializing in compliance can speed up audits, doing in hours what might otherwise take human teams weeks to complete.

Ultra-personalized financial advice and wealth management

Using AI agents to deliver highly personalized financial planning

Agentic AI systems can serve as an automated financial advisor, providing tailored wealth management by analyzing an individual's full financial picture and adapting plans to changing needs and market conditions.

ISSUE/OPPORTUNITY

Traditionally, financial advisors segment clients into broad categories—by age, income, or risk tolerance—and provide standardized advice. This limits personalization and overlooks the nuances of an individual's financial situation. At the same time, wealth management firms often reserve high-quality, customized advice for affluent clients because it is expensive and resource-intensive to deliver at scale.

As client expectations shift toward greater personalization and cost-efficiency, firms are seeking ways to deliver more customized service at lower operational cost—without compromising regulatory compliance or fiduciary responsibility.

Tags: **Customer Experience** | **Agentic AI**

HOW AI CAN HELP

Market monitoring
A market scout agent can track real-time developments such as interest rate changes, stock prices, and macroeconomic indicators to identify relevant opportunities and risks.

Personal financial analysis
A client profile analyst agent can aggregate financial data (e.g., bank accounts, credit card spending, loans) and behavioral patterns (e.g., risk-taking versus saving) to build a detailed and dynamic personal profile for each client.

Planning and strategy generation
A planning agent can synthesize market and client inputs to recommend bespoke financial strategies—for example, rebalancing investment portfolios in response to market changes or personal life events such as a home purchase or retirement.

Compliance integration
A compliance agent can help ensure recommendations adhere to fiduciary standards and regulatory guidelines, making the AI outputs safe and trustworthy for both clients and advisors.

Continuous fine-tuning
Unlike static robo-advisors, multi-agent AI systems can provide adaptive guidance, refining recommendations as circumstances evolve, rather than offering static one-time plans. They can identify portfolio risks in real time, flag anomalies for review, and automate activities such as threshold-based rebalancing, factor-based investing, and tax loss harvesting.

Ultra-personalized financial advice and wealth management

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Because errors could cause significant client losses, agents should be validated against historical financial data and tested under varied market and life-event scenarios to ensure they generate sound, resilient advice.
Transparent and explainable	AI-driven decisions—such as rebalancing actions or investment recommendations—should be accompanied by clear rationales and audit trails, enabling accountability and regulatory review.
Fair and impartial	Advisory models should be monitored to ensure they do not steer clients toward biased product selections, specific asset classes, or vendor-preferred solutions. Advice must be applied equitably across client segments.
Private	Because AI agents handle highly sensitive financial data, strong safeguards must be in place to protect information, including anonymization and strict access controls.

POTENTIAL BENEFITS

Personalization at scale

Multi-agent AI makes high-quality, personalized wealth management affordable for mid-market clients, expanding access beyond the traditionally affluent segment.

Better outcomes

Custom investment strategies ultra-personalized by AI to fit an individual's life stages, goals, and preferences can improve customer satisfaction and financial outcomes in a cost-effective way, boosting retention and conversion.

Strategic differentiation

AI-driven wealth management can help a firm establish a premium position in a commoditized market by offering a high-touch yet affordable client experience that is difficult to achieve in other ways.

Greater scalability and efficiency

Automation of portfolio management tasks enables firms to serve more clients at lower cost, freeing human advisors for complex decisions.

AI agents for algorithmic trading and market simulation

Enhancing trading strategies and insights with multi-agent collaboration

Agentic AI systems enable trading firms to use specialized agents to simulate artificial markets and to execute diverse trading strategies, enabling smarter, faster trades and richer insights into market dynamics.

ISSUE/OPPORTUNITY

Financial firms have long been at the forefront of harnessing intelligent technologies such as algorithmic trading and AI to achieve a winning edge. Now, multi-agent AI is taking the game to a whole new level.

Traditional algorithmic trading transformed financial markets. However, behind closed doors, multi-agent systems are already in production and demonstrating better adaptability to market changes—outperforming single-strategy approaches in various timeframes. In fact, it is now common for different algorithmic strategies (i.e., agents) to be running concurrently and even "competing" for capital allocation based on performance.

The latest innovation is to leverage more explicit agent frameworks and inter-agent communication, with multiple and varied trading or simulation agents acting independently yet sharing information to optimize outcomes. Looking ahead, we expect to see greater adoption of multi-agent reinforcement learning for strategy development. We may also see exchanges using agent-based AI to monitor market stability or simulate the impact of rule changes. Without such tools, firms risk being outpaced in highly competitive markets (and regulators may risk missing early warning signs of instability). However, in many cases, the biggest innovations could remain hidden from view for competitive reasons.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Market simulation

In artificial market simulations, agents representing different trader archetypes (e.g., retail investors, institutional traders, market makers) can enable researchers, investment firms, and regulators to observe emergent phenomena and test different strategies and scenarios before they go live.

Specialized trading agents

A firm can simultaneously deploy live agents that apply different strategies in the market—such as short-term arbitrage, medium-term trend following, or options hedging—sharing signals to avoid conflicts and improving portfolio resilience.

Coordinated action

Agents can operate semi-autonomously under the supervision of a coordinator agent or human risk manager, ensuring portfolio alignment and preventing overexposure. In some cases, different AI agents might even collude or negotiate with each other to create a competitive advantage.

Reinforcement learning and communication

Advanced systems can use multi-agent reinforcement learning, where agents learn from each other through trial and error in simulated environments. Some of these systems might leverage LLMs to enable agents to communicate and explain their reasoning, making the simulation results easier to interpret.

AI agents for algorithmic trading and market simulation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Because unpredictable agent interactions can create market instability, systems should be stress-tested against diverse scenarios, including flash-crash conditions, to ensure agents behave reliably under pressure.
Fair and impartial	Agents should be monitored to prevent behaviors that could unfairly favor certain counterparties or drift into market manipulation. Simulations should be designed to reflect diverse trader types, avoiding skewed outcomes.
Private	Trading data and strategies are highly sensitive and commercially valuable. Systems should ensure strict data security and access controls to protect intellectual property and prevent leakage of proprietary trading logic.
Safe and secure	Given the adversarial nature of markets, safeguards should be established to prevent agents from being deceived or manipulated by other firms' AI. Strong cybersecurity and resilience measures are essential to maintain safe operations.

POTENTIAL BENEFITS

Diversified trading performance

Multi-agent systems can function like an "AI investment committee," combining different strategies to improve returns across market regimes.

Faster reaction times

Agents can operate in milliseconds, dynamically hedging exposures and adjusting positions to manage risk in volatile conditions. This increased speed can help a firm boost its trading volumes, validate analyses in real time, and potentially drive greater profitability while mitigating risks.

Systemic insight

Agent-based simulations could provide regulators and institutions with cutting-edge tools to identify systemic risks or opportunities, enabling better preparation for disruptive events.

AI agents for credit underwriting

Adaptive, data-driven underwriting through multi-agent collaboration

Agentic AI systems can transform credit underwriting through specialized agents that analyze applicant data, monitor market context, assess risk, and maintain compliance—creating highly personalized lending decisions.

ISSUE/OPPORTUNITY

Traditional credit underwriting often relies on generic segmentation (e.g., age, income, or credit history) and static risk models. This approach can be slow, inflexible, and exclusionary, frequently locking out those with thin credit files or non-traditional income sources. At the same time, lenders struggle with outdated workflows that fail to adapt dynamically to changes in markets or individual circumstances.

As lenders work to expand their businesses responsibly, they need more agile underwriting solutions capable of assessing varied data and real-time conditions and delivering accurate decisions across a broad customer base.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Multisource data aggregation

A data gathering agent can compile structured and unstructured information (e.g., bank statements, tax filings, e-commerce history, location signals) to build an accurate and comprehensive borrower profile.

Dynamic scoring and simulations

Another agent can apply adaptive scoring models and run repayment scenario simulations, guiding approval, denial, or escalation recommendations based on risk capacity and context.

Regulatory alignment

A compliance-focused agent can confirm all assessments meet suitability requirements, fiduciary standards, and traceability benchmarks. This solution layer supports audit readiness and decision traceability.

Workflow orchestration

Agents working in concert—each handling discrete steps in the credit workflow, from data gathering to scoring and reporting—can provide continuous fine-tuning of underwriting decisions as new data or market conditions emerge.

AI agents for credit underwriting

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Credit decisions can have life-changing implications. Agents must be validated continuously against historical performance and varied economic conditions. Simulated stress tests should regularly challenge outcomes to ensure resilience across credit cycles

Transparent and explainable

Regulators and borrowers require clarity on underwriting decisions. Agents should log their reasoning (e.g., which income streams or risk patterns influenced a score), presenting audit-ready explanations that demystify the model's logic.

POTENTIAL BENEFITS

Faster, smarter approvals at lower cost

Multiple data sources, specialized agents, and adaptive models enable faster, more precise underwriting without increasing manual workloads.

Financial inclusion

By incorporating non-traditional signals and actively monitoring for bias, AI agents can extend credit responsibly to underserved populations.

Operational scalability and quality

Multi-agent systems can boost lender efficiency and client satisfaction by enabling scalable, high-volume credit processing with consistent quality and regulatory compliance.

Intraday liquidity optimization

Reconciling breaks in real time to free trapped cash and lower buffers

Agentic AI systems can automatically reconcile breaks and timing mismatches as they occur, so banks can unlock trapped cash faster and operate with a smaller intraday cash buffer.

ISSUE/OPPORTUNITY

Across payment systems, securities depositories, and internal subledgers, breaks and timing mismatches trap cash until exceptions are resolved. The shift to faster settlement cycles (e.g., T+1 in U.S. markets) compresses post-trade windows and raises the cost of delays or late funding. At the same time, supervisory agencies expect better monitoring of intraday flows and positions, even as instant payment operating hours expand the volume and timing of movements that must be continuously managed. Taken together, these trends make same-day reconciliation and proactive liquidity steering a necessity, not a nice-to-have.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Autonomous, event-level matching
Multiple AI agents continuously match payment confirmations, Nostro statements, ledger entries, and settlement messages (ISO 20022 rich data) to resolve breaks as they arise—requesting clarifications, attaching evidence, and escalating only true exceptions. This releases cash sooner.

Intraday forecasting and liquidity routing
A forecasting agent aggregates real-time statuses (e.g., cross-border tracking, queued RTGS payments) and predicts cash peaks and troughs. A treasury agent then recommends actions—sweeps, intraday credit, collateral moves, or CLS pay-ins—to minimize buffers while meeting payment obligations.

T+1 and instant-payments readiness
Exception agents accelerate break resolution and funding decisions under compressed timelines, using enriched ISO 20022 data and end-to-end payment visibility to cut manual follow-ups and reduce same-day funding surprises.

Intraday liquidity optimization

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Because reconciliation errors can trigger fails and liquidity penalties, agents should be tested on historical breaks, simulated queues, and stress days (month-end, index rebalancing). Agents should be designed to fall back to deterministic matching rules and human review for ambiguous cases.
Transparent and explainable	Since treasury, operations, and auditors must understand every break closure, agents should be designed to produce line-item audit trails (source message IDs, fields compared, confidence scores) and plain-language rationales for funding recommendations and payment releases.
Fair and impartial	Because intraday funding decisions affect business lines and clients differently, agents should be designed to apply consistent prioritization policies (e.g., client SLAs, regulatory cutoffs) and be monitored to avoid systematically disadvantaging lower-volume corridors or smaller counterparties.
Safe and secure	Given the high value of payment instructions, agents should be designed with strong cybersecurity controls (segregated runtimes, signed message handling, anomaly detection) and resilience testing to prevent spoofed confirmations or malicious instruction replay.

POTENTIAL BENEFITS

Less trapped cash and lower buffers

Faster break closure and proactive sequencing reduce precautionary intraday buffers and daylight overdraft usage, improving liquidity efficiency and lowering carry costs.

Fewer fails and fees

Real-time matching and RTGS-aware scheduling reduce late payments and fail charges while improving counterparty confidence and SLA performance.

Audit-ready compliance

Automated evidence trails support supervisory expectations for intraday monitoring and reporting, strengthening an organization's control posture with less manual effort.

Transformation with speed and confidence

Code assistant for digital transformation

AI can enable banks to increase digitization at a faster pace through code assistants.

ISSUE/OPPORTUNITY

Many financial institutions are pursuing cloud and data transformations, which are essential steps in preparing the organization for using AI tools of many kinds. In some cases, legacy hardware is retired as data is curated and shifted to the cloud, freeing up humans for more valuable work while bringing down the costs associated with on-premise infrastructure. Yet, these kinds of transformations are significant undertakings that can bring long lead times and high costs. There is also a risk of failure and error.

Tags | **Information Technology**

HOW AI CAN HELP

Supercharge your human capital

AI can be used as a component of cloud and data transformations to empower developers working across the enterprise on applications, data engineering, machine learning, and front-end development.

A helping hand in code development

As organizations explore new digital and cloud capabilities, development teams can accelerate and simplify their work by using AI as a force multiplier when writing, debugging, and documenting code, as well as translating ideas to code.

A shorter path to software

Part of success in transformation hinges on how quickly new enabling software can be deployed. There are opportunities to use AI in software development to shorten the lifecycle and more quickly reach a stable and deployable version, such as by helping rapidly write APIs, ETL, data pipelines, or even front-end code.

Transformation with speed and confidence

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Partial automation of programming-related tasks requires the system to be reliably available and accurate. If availability cannot be guaranteed to an acceptable extent, the benefits of automation must be weighed against the risk of erroneous or buggy code.

Responsible and accountable

The training data for foundation models may create legal risks related to intellectual property or copyright infringement. If the training data contains copyrighted material, the organization deploying the model needs to evaluate whether the presence of intellectual property in the training set could lead to legal challenges against the enterprise. Also, while the use of AI can accelerate the work of developers, without a human in the loop (e.g., to validate and debug code), critical failures may occur. Shoring up accountability may involve documenting and communicating standards and expectations for employees using AI.

Safe and secure

By using an AI system, proprietary code bases may be exposed to third parties, raising questions around the security of the data and controlled access to it. An inadvertent breach of confidential intellectual property could have significant enterprise impacts.

POTENTIAL BENEFITS

Lower transformation costs

By shortening the software development lifecycle, the organization can reduce overall costs for digital and cloud transformation.

Lower the bar to digital entry

Using AI opens the door for financial services organizations of all sizes, capabilities, and technology maturities to digitize and move to the cloud in a way that was previously out of reach for many.

Business intelligence at your fingertips

Enterprise-wide data search and access

Make business intelligence via enterprise data search accessible to all through natural language interfaces.

ISSUE/OPPORTUNITY

For many banking and insurance companies, data is stored in a multitude of locations, from local hardware to cloud storage solutions. This makes it difficult to effectively query different databases and retrieve relevant information quickly and efficiently. Complicating the matter, multiple mergers and acquisitions over time may have compounded the diversity of data locations and databases, which hinders data mining for insights. Financial services organizations are information-intensive enterprises, and without the capacity to easily query all data, the result is poorer or incomplete insights that can increase enterprise risk while threatening customer satisfaction.

Tags: **Cross-functional**

HOW AI CAN HELP

An enabling interface
AI capabilities can be built on top of an existing solution to facilitate the communication of queries from the user to the search layer. This serves as the interface between search layers and databases, allowing users to easily mine all enterprise data, as well as generate structured analytics reports.

Speed to insight
With the ability to query and analyze disparate data sources using AI as an interface, the enterprise can move past traditional business intelligence techniques and dramatically reduce the time required to generate insights while increasing the workforce accessibility to business intelligence.

Business intelligence at your fingertips

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

When it comes to governance and control, granting more data access to a wider segment of the workforce can create a more complex challenge of restricting who in the organization is permitted to access sensitive business data.

Robust and reliable

Given the known challenges with model reliability and the potential for hallucination, banking and financial services organizations face the risk of inaccurate or false AI-derived insights influencing decision-making and leading to negative ramifications even at the market level.

Private

When dealing with sensitive and proprietary information, the organization must contend with securing the data, removing or obscuring it in training and testing sets, and evaluating the model to determine whether it could "leak" protected information, either due to faulty function or a targeted attack.

POTENTIAL BENEFITS

Lower technical hurdles

AI as an interface between search and data enables business users to query databases and obtain tailored results without in-depth programming experience. This gives more of the workforce access to business intelligence without additional burdens on IT and data science teams.

A new level of data-driven decisions

Real-time access to all of an enterprise's data can help organizations become even more insight driven, which supports improved growth prospects through access to the right insight at the right time.

Fixing the missing data issue

Synthetic data generation

Generate synthetic data for model training, anomaly detection, and identifying cyber and deception attacks.

ISSUE/OPPORTUNITY

Missing data is a significant challenge for financial services organizations. Datasets may be incomplete, data transfers restricted, and potential anomalies underrepresented in the data. Using synthetic data can help overcome these challenges. In cloud transformation, data transfers may be delayed due to risks and regulations around data governance, and use of synthetic data can enable a smoother and more efficient transformation. Meanwhile, machine learning anomaly detection systems (such as those for identifying fraud, waste, and abuse) must be trained on data from previous events. Yet, the dearth of data from these rare events can make anomalies harder to assess.

Tags | **Information Technology**

HOW AI CAN HELP

Improve model training
Generative AI can be used to quickly create synthetic data to supplement machine learning model training data, which is then used to aid and accelerate digital and cloud transformations. In this way, generative AI complements the enterprise's wider AI initiatives, fueling (rather than replacing) other AI deployments.

Amplify anomaly event detection
The rarity of anomaly events can make it difficult to train machine learning systems to detect instances of fraud, waste, and abuse. Creating synthetic data with generative AI gives ML systems a larger suite of examples that lead to a greater capacity to find patterns and anomalies in the data.

Harden the organization's cyber posture
Just as synthetic fraud data can be used to train models to identify fraudulent activity, synthetic adversarial data can be used to train models to detect and mitigate cybersecurity risks and deceptive behaviors targeting virtual assistants.

Fixing the missing data issue

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

A significant risk when generating synthetic data is that historic biases can creep into the generated data, perpetuating those biases (e.g., in the case of certain communities or socio-economic groups being underrepresented in the data because those groups have had lower levels of banking activity in the past). This bias is not necessarily intentional but is still harmful.

Robust and reliable

Synthetic data created with generative AI can be limited in its scope and scale, and it should not be presumed to be accurate or perfectly reflective of real-world data. An over-reliance on synthetic data may inject problems with data reliability, which can hamper the validity and usefulness of outputs and model training.

POTENTIAL BENEFITS

Faster path to the cloud

Generative AI-created synthetic data can accelerate digital and cloud transformations by making the transition smoother and more efficient.

Tackling fraud

Use synthetic data to train machine learning systems on rare or unknown events, such as novel types of fraud.

Security confidence

Adversarial synthetic data contributes to the enterprise-wide imperative to fortify digital assets against cyber threats.

Getting to know your customer

Research-based report generation

Generate and summarize reports on new customers to inform employee decisions for customer onboarding.

ISSUE/OPPORTUNITY

New customers are the lifeblood of a growing financial services enterprise, but onboarding customers can be a highly manual and time-consuming process. Know Your Customer (KYC) standards and rules require institutions to develop meta reports on customers based on economic outlook, equity research, adverse media, and new prospect due diligence. These are high human involvement tasks, with desk-based meta-research consuming valuable time and resources.

Tags: **Customer Service**

HOW AI CAN HELP

Condensing results for easier consumption

AI can be used to summarize and filter results from existing search engines to inform meta reports, as well as to summarize information for the customer relationship manager.

Research and analysis to inform reports

Generative AI, alongside other machine learning models, can be used to conduct preliminary data searches and meta-analysis, potentially accelerating the KYC process.

Getting to know your customer

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

When using AI to perform research and analysis, there is a risk that it could skip or misconstrue highly relevant information, which could skew the conclusions in the meta-analysis and hamper sound decision-making. If a new customer is erroneously assessed to be higher risk and a relationship manager declines the opportunity to engage the customer, the consequences are missed revenue and diminished customer engagement.

Private

When dealing with a customer's financial or personally identifiable information, the enterprise faces legal and regulatory standards for data privacy. When using AI, the organization should take steps to ensure sensitive information does not inadvertently leak through model outputs, as well as govern who has access to the model, the underlying data, and the customer data it references.

POTENTIAL BENEFITS

Timely insights

Faster and more efficient search and analysis can give decision-makers more up-to-date information and insights that enable better, more timely decision-making around customer onboarding.

Cost reduction

By streamlining and simplifying the report generation process, costly labor hours can be redirected to more valuable work.

Efficiency

The end user saves time and effort by more easily accessing and consuming relevant information.

Enhanced AI support for customers

Financial guardian

A personal virtual assistant powered by AI can cater to daily needs of customers.

ISSUE/OPPORTUNITY

Customer service and engagement is vital for financial services organizations, particularly as they transform from a product-focused to a customer-focused business. As a part of that, there is a pressing need for the enterprise to rapidly and accurately answer both common and complex customer queries and do so with the timeliness customers expect. Yet, increased digitization in financial services has reduced access to representatives who can answer customer questions; this is at a time when financial services customers seek a hyper-personalized experience.

Tags: **Customer Service**

HOW AI CAN HELP

An empathetic, personalized interface

Generative AI can be paired with other models to create a customer interface that delivers a hyper-personalized experience, such as by training the generative AI model to provide answers or insights with empathy. It also overcomes some of the challenges around more traditional chatbots that can lack an empathetic tone.

A more capable digital agent

Using generative AI can improve the usefulness and accessibility of a chatbot interface. With the capabilities of large language models (LLMs), a personal digital assistant can summarize contracts and answer nuanced questions, and the customer may enjoy a range of interface options, including text, audio, and imagery.

Enhanced AI support for customers

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

While a generative AI-enabled solution may provide valuable answers and recommendations most of the time, there is a risk that too much confidence may be placed in the validity of the outputs, both by the organization and the customer. Generative AI is not an infallible oracle, and an overreliance on the AI solution may have a detrimental impact on customer actions, which can in turn increase financial risk.

Private

When confidential or personally identifiable information is inputted via the digital interface by the customer, the financial institution is obligated to follow the laws and rules that dictate how that sensitive information can be transmitted, stored, and accessed. Failing to do so could raise legal peril and potentially subject the enterprise to greater cyber risks.

Transparent and explainable

End users require a clear understanding of how their information will be processed and that they are interacting with a machine. At the same time, the enterprise needs to be able to interpret outputs and understand how and why the generative AI model created a given output.

POTENTIAL BENEFITS

Customer satisfaction

Greater accessibility and more timely answers can lead to a more personalized and satisfying customer experience. This can drive an increase in net promoter score, reflecting increased customer retention and loyalty.

A cycle of efficiency and growth

While customer loyalty and brand reputation fuel business growth, integrating generative AI into virtual assistants can further enhance operational efficiency. AI-enabled systems can serve more customers at scale while the human workforce focuses on resolving the most complex issues or attracting new customers.

Financial inclusivity through hyper-personalization

When customers feel that banking applications are delivering a valuable, personalized experience, it encourages greater customer engagement and interest in service offerings, supporting both the customer's financial wellbeing, as well as that of the enterprise.

Customized marketing for the individual

Hyper-personalized sales and marketing assistant

Regulatory-compliant marketing material generation across different geographies.

ISSUE/OPPORTUNITY

Financial services marketing operations are increasingly coming under regulatory scrutiny for issues such as mis-selling and misinformation. Part of the challenge for multinational organizations is that cultural differences as well as varying levels of customer understanding about products may create regulatory risk for enterprises in a given geography. To overcome this, organizations are investing significant manual labor to maintain a compliant marketing function, which is both time consuming and costly.

Tags **Marketing**

HOW AI CAN HELP

Customized materials for different audiences
AI can be used to create marketing materials that contain the appropriate tone, language, and cultural references, while also supporting consumer understanding of the product to maintain regulatory compliance.

Personalized sales at scale
With AI, financial services organizations have the ability to create marketing materials that are customized to individual customers—and to do so at scale.

Customized marketing for the individual

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

For AI-generated marketing to be valuable, organizations must be able to rely on the validity of the output. AI can be prone to hallucinations, and when tasked with creating marketing that touts one product over another, there is a risk the model will return false statements. This injects potential regulatory violations that could result in fines and other penalties. To shore up reliability with AI, outputs should be validated by a human with the subject matter understanding to do so.

Fair and impartial

Datasets may contain latent bias of which the organization is unaware. This could be due to how the data was acquired, recorded, and curated, and the challenge is compounded when operating in multiple geographies. Organizations should consider the datasets used to train and fuel generative AI systems and whether unknown bias could lead to marketing materials that fall short because they fail to reflect important geographical and cultural differences.

POTENTIAL BENEFITS

Individual-level marketing

Hyper-personalized marketing takes the enterprise to a new stage of customer engagement and enticement that is infeasible through manual effort alone.

Confident compliance

Using AI for marketing development can help ensure the content remains in line with regulatory expectations across many geographies, thereby reducing regulatory risk.

Driving marketing ROI

Personalized marketing can support new sales, strengthen the customer relationship, and reduce the cost of marketing operations while also improving the timeliness of outreach and engagement.

Ensuring the integrity of claims

Automated claims reporting

AI can be used to automatically generate reports based on descriptions or pictures of the relevant subject.

ISSUE/OPPORTUNITY

During claims processing for property and casualty insurance, claims agents must decide whether a destructive event was insured, and if so, the amount and cost of the damage. These processes are complex and can be time consuming, and claims agents have few tools to support their decision-making.

Tags: **Operations**

HOW AI CAN HELP

Virtual damage rendering
AI can be used to help visualize the damage by replicating it virtually. The basis for the replication and visualization can be customer conversations, damage documents, photos, official reports, and other relevant media. In visualizing the data in this way, the claims agent is empowered to make better decisions when assessing the degree and cost of damage.

Automated claims reporting
With AI, claims reports can be generated automatically based on photographic evidence.

Ensuring the integrity of claims

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Damage visualization requires a high degree of accuracy, and erroneous AI outputs could lead to claims being paid incorrectly, potentially leading to overpayment (a detriment to the organization) or underpayment (a detriment to the customer).

Transparent and explainable

If claims agents use AI to automate aspects of claims processing but are unable to articulate to customers how the AI model derived its outputs or contributed to the cost and damage assessment, customers may not accept the outcome of the claims process.

POTENTIAL BENEFITS

Cost reduction

Faster claims processing and more accurate damage assessments can reduce labor costs and claims payments.

Customer satisfaction

By processing claims faster and with less administrative burden, customer satisfaction improves due to quicker, more streamlined adjudication and payment.

Identifying fraud

More expedient claims processing has the advantage of identifying potential fraud more quickly, helping to ensure the integrity of claims and payments.

A virtual bank experience

VR-enabled retail banking centers

Virtual reality customer agents powered by generative AI can change the retail banking experience and interactions.

ISSUE/OPPORTUNITY

Banks seek to provide customers with multiple methods of interacting with their accounts, services, and offerings. As digital services grow, customers may prefer to conduct banking transactions remotely through virtual transactions. However, this approach still requires a human agent, and employees are the most expensive cost component in customer service—particularly when three levels of customer support are necessary. Chatbots can help automate virtual transactions, but existing chat tools are limited to specific, pre-programmed dialogue and options.

Tags: **Customer Experience**

HOW AI CAN HELP

Bring the bank to the customer
With a generative AI-enabled virtual space, customers can use a VR headset to conduct business with the financial institution and interact with a service representative from the comfort of their own home in a way that is convenient for the customer.

Hyper-personalized service
A generative AI agent can provide conversational, tailored responses to questions about customer accounts and financial needs. This helps the enterprise cater to its customers' desires for a personalized experience while also avoiding the costs that come with adding more human customer service workers.

Speed and quality of service
In a virtual space, customer data can be accessed in real time by the conversational agent to provide faster, higher quality service and offerings.

A virtual bank experience

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable

Given a generative AI-enabled chatbot's capacity for conversational outputs, there is a risk that customers may not realize they are interacting with a machine. To promote trust, customers should not only be informed that the chatbot is not human, but they should also understand how their inputs and information are stored, accessed, and used.

Fair and impartial

The datasets used to train and inform the chatbot may contain latent biases, such as underrepresented customer groups or semantic deficiencies in some languages but not others. As a result, the model may simply not work as well for some customers, leading to a variety of negative customer impressions and complaints.

Responsible and accountable

With issues around reliability and data quality, if the chatbot outputs erroneous data or recommendations, a human stakeholder needs to be accountable for the outcome. Organizations can promote accountability by keeping a human in the loop and documenting roles and responsibilities.

POTENTIAL BENEFITS

Customer-focused experience

A generative AI-enabled virtual banking experience connects customers with tailored digital representatives that can converse in the customer's preferred language in a timelier manner.

Liberated human capital

With conversational chatbots satisfying most customer inquiries, the human workforce can focus on the most complex or value-driving customer service activities.

Cost avoidance

Virtual agents allow the bank to serve more customers without expanding the human workforce, which helps limit customer service costs.

Focused cyber

AI-driven vulnerability management that focuses on real threats

AI can help filter, analyze, and prioritize security alerts, allowing organizations to focus their resources on true vulnerabilities—which can help reduce their cybersecurity risk.

ISSUE/OPPORTUNITY

Financial institutions and other large enterprises rely on multiple security tools that generate millions of alerts daily—many of which are false positives or non-actionable.

Security teams and developers spend time triaging and remediating vulnerabilities that pose little or no real risk. This inefficiency delays responses to real threats, creates operational fatigue, and increases the likelihood of breaches—threatening business continuity and brand reputation.

Tags **Compliance & Risk**

HOW AI CAN HELP

Evaluating alerts
AI models can evaluate alerts by calculating breachability (likelihood the vulnerability can be tapped in a real-world scenario) and exploitability (likelihood it can be actively taken advantage of to inflict harm).

Setting priorities
An AI-driven system can use enterprise-specific context such as workload risk profiles, application architecture, and network exposure to separate high-risk from low-risk vulnerabilities. AI automatically identifies which threats demand attention and which can be deprioritized, providing a data driven, risk-based prioritization process.

Focused cyber

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	The vulnerability management system should be continuously retrained and tested against live alert data and historical incidents to validate its prioritization logic. A human-in-the-loop approach provides important oversight on critical decisions, especially during early-stage deployment.
Safe and secure	The models should be hosted in secure environments (e.g., on-prem or restricted cloud instances), with strict access controls and audit trails. Systems must be designed to prevent data leakage and be regularly tested for robustness.
Transparent and explainable	AI outputs should include the rationale for prioritization with traceable decision paths that help analysts understand and trust the recommendations.

POTENTIAL BENEFITS

Less noise

Moving from thousands of daily alerts to a handful of real threats enables security and engineering teams to focus their time and effort more effectively.

Improved security and reputation

Decreasing security incidents strengthens a company's standing with customers, investors, and regulators.

Cost savings and improved efficiency

Better resource allocation allows security and engineering teams to spend less time and money protecting the organization from cyberthreats.

Enhanced productivity and morale

Minimizing time spent on false positives reduces friction between the security and engineering teams, improving productivity and morale.

Contacts

Beena Ammanath
Global Deloitte AI Institute Leader, Deloitte AI Institute United States, Lead
Deloitte Consulting LLP
bammanath@deloitte.com

Audrey Ancion
Deloitte AI Institute Canada, Lead
Deloitte Canada
knuttall@deloitte.com

Naser Bakhshi
Deloitte AI Institute Netherlands, Lead
Deloitte Netherlands
nbakhshi@deloitte.nl

Dr. Bjoern Bringmann
Deloitte AI Institute Germany, Lead
Deloitte Germany
bbringmann@deloitte.com

Martin Cabrera
Deloitte AI Institute Chile, Lead
Deloitte
mcabreraa@deloitte.com

Jefferson Denti
Deloitte AI Institute Brazil, Lead
Deloitte Brazil
jdenti@deloitte.com

Tiago Durao
Deloitte AI Institute Portugal, Lead
Deloitte Portugal
tdurao@deloitte.com

Richard Eudes
Deloitte AI Institute France, Lead
Deloitte France
reudes@deloitte.fr

Roman Fan
Deloitte AI Institute China, Lead
Deloitte China
rfan@deloitte.com

Tomas Meca Figueres
Deloitte AI Institute Spain, Lead
Deloitte Spain
tomecafigueras@deloitte.es

Alfredo Maria Garibaldi
Deloitte AI Institute Italy, Lead
Deloitte Italy
agaribaldi@deloitte.it

Nicolas Griedlich
Deloitte AI Institute Luxembourg, Lead
Deloitte Luxembourg
ngriedlich@deloitte.lu

Geert Hallemeesch
Deloitte AI Institute Belgium, Lead
Deloitte Belguim
ghallemeesch@deloitte.com

Jan Hejtmanek
Deloitte AI Institute Central Europe, Lead
Deloitte Central Europe
jhejtmanek@deloitte.com

Prashanth Kaddi
Deloitte AI Institute India, Lead
Deloitte India
pkaddi@deloitte.com

Sultanbek Khunkaev
Deloitte AI Institute Middle East, Lead
Deloitte United Arab Emirates
sukhunkaev@deloitte.com

Tomotake Kozu
Deloitte AI Institute Japan, Lead
Deloitte Japan
tomotake.kozu@tohmatsu.co.jp

Carlos Labanda
Deloitte AI Institute South-Latin America, Lead
Deloitte Colombia
clabanda@deloitte.com

Chris Lewin
Deloitte AI Institute Asia Pacific Lead
Deloitte Singapore
chrislewin@deloitte.com

Wessel Oosthuizen
Deloitte AI Institute Africa, Lead
Deloitte Africa
woosthuizen@deloitte.com

Sulabh Soral
Deloitte AI Institute United Kingdom, Lead
Deloitte United Kingdom
ssoral@deloitte.com

Dr. Elea Wurth
Deloitte AI Institute Australia, Lead
Deloitte Australia
ewurth@deloitte.com.au

www.deloitte.com/us/AIInstitute

Deloitte.