



Redefining possible: Banking and agentic AI

kyndryl[®]

Contents

PAGE

05	Geographic trends
09	From promise to practice: Challenges to the adoption of agentic AI
14	Opportunities: How agentic AI builds resilience, efficiency and personalization
19	Recommendations
24	How Kyndryl helps



Banks offer savings accounts to five-year-olds and mortgages for their parents. They finance trade between companies that have never met, manage a corporate treasurer’s liquidity across time zones, and move money from a saver in Hamburg to a supplier in Santiago. Banks do all of this under the highest expectations for security, stability and resilience.

But customers want more. They want it all to be easy, tailored precisely to their individual preferences, and done in real time. A retail customer wants to see cash, digital assets and investments on a single screen, with the right product suggested at the right moment. A corporate treasurer wants a real-time consolidated view of positions across entities and currencies. A regulator wants the data behind any decision, on demand.

Answering customers’ demands—whether those customers are internal or external, retail, corporate, or regulatory—means removing friction from banks’ systems and making a seamless experience a source of competitive advantage. All this must be done without compromising the security and resilience that are the hallmarks of today’s banking systems.

When banks remove friction, transactions settle in real time. Retail and corporate customers can easily conduct business across borders, in fiat currencies and digital assets. Customers get the services they need exactly when they need them. Banks notice their networks slowing or stuttering, and address issues before a breach or other failure occurs. Compliance becomes a competitive advantage.


None of this will be accomplished simply by deploying a new user interface or even by moving fully into the cloud. It requires a modern, secure, flexible infrastructure, operational agility, and, in many cases, a rethinking of work. It will need to use open banking and real-time payments. But agentic AI redefines what’s possible. Agents can sense, reason, and act across the organization in real time, shifting from rigid legacy technology, manual workflows, organizational silos, and enterprise applications as systems of work, to agile autonomous operations, open ecosystems, agentic workflows with human oversight, and enterprise applications as systems of record.

To get there, banks will use agents for more than just targeted productivity gains — they’ll need to scale agents throughout their organizations. For agents to scale, they need to be placed in a governed operational environment that supports explainability, human oversight, and operational accountability. Increasingly, that governance is operationalized through policy as code, translating regulatory, business and operational requirements into executable controls that govern how AI agents act in production. It also means addressing fragmented architecture, governance, data silos and data lineage, security, and operational trust.

Few organizations are ready for this. Some 91% of leaders say their infrastructure is best-in-class, yet they show concern that best-in-class isn’t good enough: only 41% say that their infrastructure is ready to handle future risks.¹

Banks are moving ahead, and quickly: 65% of financial services institutions overall say they’re actively using AI.² Incumbents are no longer trailing: fintech companies are slightly less likely to be using AI than capital markets firms.

AI adoption across financial services



All this is happening under increased pressure, because customers, regulators and bad actors are pushing forward as well. Newer banks have an advantage: neobanks (such as Revolut) are purpose-built for digital assets and real-time settlements. In much of Asia, digital assets are commonplace, and cash is optional; the European Central Bank is working to support both tokenized assets and a digital Euro.³

But there are other factors at work. Geopolitical concerns are prompting a desire for data sovereignty. Bad actors are using AI as a tool for industrialized fraud; frontier AI models are identifying vulnerabilities faster than humans can fix them, and regulators globally are asking banks to respond. Security improvements made today must also ready a bank for an age of post-quantum cryptography. And bank leaders know that those who fail to create competitive advantage out of this maelstrom will find themselves acquisition targets for those who have.

In this strategic guide, we'll lay out a framework for banks that want to continue steering their own destinies and develop the infrastructure, data architecture, agents, and cultures that will support the next evolution of banking.



A vintage-style map with a compass and a divider tool. The map is light blue with various geographical features and text in Cyrillic. A silver compass with a black face and green and red markings is in the foreground. A silver divider tool with green handles is in the background. A red diagonal line runs from the top right to the bottom left, separating the image from the text.

Geographic trends

While many trends affect banks globally, other challenges—and banks' abilities to meet them—vary substantially by region. Here are some of the dominant trends and watershed events influencing banks in selected regions of the world.

North America

Innovation driven by competition and regulatory changes

- Encouraged by federal legislation, banks are working aggressively to accommodate stablecoin and other digital assets.
- Bank consolidation is accelerating, prompted by a more accommodating regulatory stance and a reemergence of shareholder activism.
- The Nasdaq is moving to 23/5 trading and the NYSE Arca to 22/5, putting pressure on activities such as clearing, margining, and liquidity management.
- FedNow, running alongside Real Time Payments (RTP), pushes banks to upgrade their infrastructure to both send and receive real-time payments.
- Clearing, margin, and collateral management systems need to be upgraded to meet new SEC requirements for the central clearing of US Treasury cash and repo transactions.
- Wealth managers need access to advanced tools to help clients more easily understand and manage newly diverse portfolios (private equity, hedge funds, and public securities).

E.U. and U.K.

Resilience, security and AI top CEO agendas

- Operational resilience is a board-level conversation, thanks to the Digital Operational Resilience Act (DORA), concerns about digital sovereignty, and several high-profile cyber attacks.
- The EU AI Act treats AI systems for credit scoring as “high risk,” with strict requirements for data governance, risk management, transparency, and human oversight.
- Instant transfers of money within the Eurozone are now mandatory, thanks to the single euro payments area (SEPA) regulation, although cross-border payments can still take days.
- The U.K., E.U. and Switzerland will shorten settlement cycles to one day (T+1) as of October 2027.
- U.K. tax changes are driving demand for wealth management services and tax-efficient investing strategies.
- Digital currency is mostly limited to pilot programs; international payments still lag.
- CEOs and top leadership are enthusiastic about agentic AI, but business leaders are struggling to match strategy with reality. Explainability is a hurdle, as banks need to explain agents’ actions both internally and to regulators.
- Younger customers do not conform to typical savings and spending patterns, frequently switch banking providers, and often would prefer to bank without personal interactions.

Japan

Security concerns overshadow innovation

- Pushed by regulators, banks are devoting significant resources to documenting and mitigating security flaws in their systems uncovered by frontier generative AI models, dampening their ability and enthusiasm for innovation.
- Most Japanese banks are inclined to use the managed services versions of LLMs, and are now actively trying to optimize to reduce token cost. They are also taking a new look at the costs of their cloud environments.
- After nearly a decade of experimentation, banks are working with newfound urgency to move stablecoin and other digital assets into production.
- The role of fintechs has shifted from disruptors to big-bank collaborators seeking to be acquired by the banks they serve.
- Many small banks and an aging and shrinking population are driving bank consolidation.

India

A push for financial inclusion spurs competition and innovation

- In an effort to jumpstart financial inclusion, the Indian government has licensed Small Finance Banks to encourage savings practices in rural and semi-urban communities.
- The compliance load is growing, as regulators ask for more frequent reports.
- India's Unified Payments Interface replaces cash, credit, and debit cards; its central bank operator recently released its own credit card.
- Increased transaction volumes are putting stress on banking infrastructure.
- Largely free of legacy mainframe estates, newer Indian banks can direct their technology spend on innovation, targeting growth in mobile banking. They benefit from the country's growing talent pool in engineering and in AI, as well as a digital public infrastructure.

Latin America

Innovation in payments

- Open Finance Brazil is the most advanced open-banking regime globally, and the Brazilian Central Bank has run a pilot (Drex) to test a platform for wholesale central bank digital currency and tokenized assets.
- Frameworks erected by central banks have facilitated the widespread adoption of real-time payments within countries, such as Mexico's SPEI/ DiMo and Brazil's PIX — arguably the world's most successful instant-payments initiative. But innovation in cross-border payments is lagging.
- AI sovereignty is becoming a concern, as is the ability to decouple from hyperscaler clouds if necessary.
- Concern about token costs is growing. Countries and banks in the region do business in volatile local currencies, but tokens are priced in U.S. dollars. This makes many AI use cases uneconomic and demands a sophisticated FinOps function.
- While the region's largest banks are heavy users of mainframes, newer banks tend to be cloud-native, often enabling them to offer a more seamless customer experience.

Middle East

Global ambition with sovereign support

- Helped by models such as Falcon, from the Technology Innovation Institute, The United Arab Emirates (UAE) is a world leader in sovereign AI. Saudi Arabia is close behind, through large infrastructure projects and state-backed AI businesses such as Humain.
- The UAE and Saudi Arabia are applying their AI ambition to the finance sector, piloting blockchain solutions for trade finance and payments, and launching initiatives such as FinTech Hive (UAE) and Open Banking Program (Saudi Arabia).
- Embedded finance is becoming more common, surfacing contextual banking services at the moment people need them.
- Financial inclusion has improved rapidly, in large part because of the wide availability of digital banking.



From promise to practice: Challenges to the adoption of agentic AI

While banks understand the importance and potential of agentic AI in helping them to become more efficient, resilient, and innovative, they also face difficulties in areas as disparate as regulation, infrastructure, finance, and change management.

Compliance, auditability, and observability

Compliance is the biggest hurdle to the use of agents enterprise-wide. It's clear that agents can do certain tasks more efficiently than people, but it's not clear that they can be relied upon to do so correctly. Policy as code translates business rules, regulatory requirements and operational controls into machine-readable language and promulgates them throughout a software environment, providing important guardrails. It determines how AI agents execute at runtime, determining when actions are permitted, escalated for review, or blocked, while creating auditable evidence of decisions. But most organizations have yet to develop processes that will allow full governance, tracking, validation and observability of their agents, nor are they currently able to audit them.

As a result, compliance departments are not comfortable with agents doing mission-critical work or communicating directly with counterparties, regulators, or customers. To better understand agents and the risks of deploying them, most banks start small: with one agent performing one function, fronted by a chatbot interface. Such internal-facing agents are being rolled out quickly, while external-facing agents lag behind.

Uncertain and escalating token costs

After initial excitement around the use of agents, especially in software development, enterprises are growing increasingly concerned about token costs, which are eliminating the business case for some uses of agents. Agents can consume orders of magnitude more tokens than chatbots, for the very reason that agents can be so useful: they are capable of acting independently, rather than waiting for detailed instructions.

Tokens are not the only cost of agentic AI. Some use cases also incur fees with hyperscalers or data and data storage providers; banks may be incurring increasing MIPs on their mainframes while also seeing higher cloud consumption. In each case, the bank needs a sophisticated FinOps function, or equivalent expertise, to manage these costs.



Fragmented architectures and data

To be effective, agents need real-time access to clean, unified data. Many organizations lack a single source of truth, or they lack the ability to give agents easy but controlled access. Meanwhile, banks are re-evaluating their data strategies as data sovereignty becomes more of a concern. Agentic technology can be a powerful ally in a bank’s journey to a coherent future-state data strategy and the technology that supports it.

Ideally, agents and employees would pull data from a single architecture that encompasses the entire banking group. In practice, fragmented architecture and data make it difficult to identify a single source of truth. Even banks that start over with a new core banking system need to clean and rationalize their data before migrating it.

Just as agents can accelerate mainframe modernization, agents can also be used to speed up the rationalization of an organization’s data architecture. Agents can take on tasks that used to be highly manual, such as creating data catalogs. Once the data has been located (a large task in itself), an agentic layer can pull the data from each system and consolidate it.



Change management

Scaling agentic AI throughout any enterprise requires more than infrastructure and technology. It also often demands a new operating model, incorporating both humans and agents and allowing organizations to evolve their processes rather than simply speeding them up. As those processes change, the rollout of agents will also affect hiring, retention, and training: a department will work differently when an agent performs some of its tasks, and will require a new form of oversight.

Some 87% of leaders say AI will completely transform job roles and responsibilities at their organization in 2026.⁴ Organizations aren’t yet on top of this, and they’re not optimistic: 79% agree that the speed of AI adoption will outpace their organization’s ability to adapt its workforce, governance, and operating model.⁵

While 49% of leaders say skills and talent gaps will be a top challenge in executing their AI strategy, they haven’t laid the groundwork to acquire the necessary skills. Only 34% of leaders say they have an accurate inventory of employee skills, and only 31% say they have a budget and proactive upskilling strategy. A mere 29% of leaders say their organizations are building new AI-related skills such as oversight, prompt design, and output evaluation.⁶

A looming skills gap



Building for flexibility and avoiding vendor lock-in

The next one to three years will see rapid technological advances. While agents and other forms of AI currently offer powerful advantages, tomorrow's models and agents will be even more compelling—and business needs could change equally dramatically in that same time.

That makes it imperative to build a solution that allows flexibility among technologies, partners, platforms and models. Such a solution leans on a modular architecture, making it possible to swap components —accelerators, storage, new tooling and models—without rebuilding infrastructure. A flexible solution ensures proprietary data remains decoupled from specific vendor tools and leverages open-source and open-weight models. Some 83% of banks say open source is important to their AI strategy, with 43% saying it is “very” or “extremely” important.⁷





Customer spotlight:

Managing sudden growth: A leading bank in Asia

A newly designated Small Finance Bank is on a mission to increase financial inclusion. But by 2026, its success—more than doubling the number of customers served in just eight months—had brought operational strain.

Millions more customers meant significantly more inquiries from law enforcement and regulators responsible for investigating suspicious financial transactions and cybercrime. The bank is required to respond to every inquiry but struggled to scale: It soon faced a backlog of communications reaching into the thousands.

To begin to clear this backlog, the bank turned to agentic AI. The bank built an agentic system to read incoming emails, gather relevant information, and compose a draft response for review by a bank employee. Because inquiries often incorporate multiple formats (email, pdf documents, other files) and languages, the bank uses eight large language models to understand and respond to inquiries.

The use of open-source models, coupled with concerns about volume and data privacy, led the bank to host the project with Amazon Web Services—the bank's first venture into the cloud. Today, processing time has been reduced by more than half, with nearly 100% data capture accuracy. The bank has established an agentic center of excellence, developing and rolling out use cases in voice-activated banking, customer onboarding, and small-business underwriting, with more value-add use cases on the way.



Opportunities: How agentic AI builds resilience, efficiency and personalization

Streamlined, efficient compliance

Policy as code can significantly improve efficiency in compliance tasks, allowing agents to automate the massive backlog of compliance work many banks face. At its most basic, policy as code translates rules into code, providing guardrails for agents and the tasks they’re assigned. Once rules are translated into code, promulgated and tested, changes only need to be made once and are automatically passed to downstream business processes, such as auditing and training.

Policy as code can also be used to aggregate supporting data packs for financial decisions. A loan decision, for example, could be accompanied with data showing that it met EU, national, regional, and bank policies, and that the model made the decision based on the applicant’s finances and was not shown personal information about the borrower.

Safer, faster modernization

Agents can be used to leverage existing infrastructure, data and tools to de-risk and accelerate mainframe modernization. Proven agentic workflows designed to handle changing business priorities can help customers modernize iteratively and seamlessly. That might mean putting a light ledger on top of existing infrastructure or moving some applications off legacy platforms and into cloud-native technologies.

Agents may be used to extract business logic from poorly documented legacy code, or to translate that code into a more modern language. Agents can also identify dependencies that make it difficult to uncouple and replatform applications. If new code needs to be written, agents can be used alongside developers throughout the software development lifecycle—from design to testing to deployment and maintenance—to help write that code more cleanly and quickly.

The result is a more flexible, resilient infrastructure, allowing accelerated product and service innovation, faster time to market, and an improved ability to meet the demands of today’s banking customers.

Agents should also be a critical part of organizations’ defense capabilities, especially in their efforts to protect themselves against vulnerabilities and threats identified by and associated with frontier AI models. Agents are an important tool in helping to harden organizations against attacks enabled by new technologies, and to prepare for the coming of quantum computing and its associated risks.

Of those organizations that do modernize, 87% say their efficiency has improved; 69% say they’ve become more innovative.⁸

Those who modernize see improvements in...



A single source of truth: Rationalizing data and data sources

Most banks’ data is fragmented and siloed. That prevents banks from using data fully in service of current customers, and from realizing its broader strategic value in sharper risk pricing, deeper client insight, and new data-driven services, all within the boundaries regulators and customers expect.

Agents can be used to help build a coherent data architecture, or they can sit on top of the current arrangement and pull data from disparate sources.

The decision to build a new data architecture should be related to a bank’s ambition to gain a competitive advantage from that particular data. If a bank aspires to be the best at something, it’s worth taking the time to build the data architecture and governance to support that goal. If a bank sees a particular function as a commodity service that won’t provide a competitive advantage, it may choose to leave its data in a siloed state and use agents to access it more easily.

A global bank without the scale of the very top tier may choose to concentrate on wealth management, and could focus on developing superior digital capabilities that enable seamless cross-selling. A bank focused on digital inclusion might work behind the scenes to make financial services available at gas stations or supermarkets; still another might develop a super-app experience to bring customers an array of non-banking products and services.

Each of these banks is using technology and data in a very different way and will require different integrations with the wider ecosystem. A bank that understands its ultimate goals will tie its modernization strategy, AI stack, data architecture, and business domains to that north star.

Efficiency and productivity

While agents may not be quite ready for mission-critical tasks or to be exposed directly to customers or regulators, they can bring speed and efficiency to many internal tasks. Relieving employees of repetitive or undifferentiated tasks can allow them to do work that more directly impacts the customer experience.

These internal tasks can include the use of AI throughout the software development lifecycle (not just for coding), retail and client onboarding (including a know-your-customer refresh for corporate clients, one of the most document-intensive processes in banking), trade finance documentation, internal documentation for complex transactions, lead generation, fraud and sanctions screening, and information gathering for loan approvals. Agents can be especially useful for helping teams to understand massive legacy codebases, determine their business logic, and decide how and whether an application should be migrated.

By improving and accelerating internal processes, agentic AI can give banks the agility and capacity to better serve external customers—perhaps via faster response times, improved product development, or more personalized recommendations.

Top banking uses for agents⁹

among banks using agentic AI



Building a single customer view

Banks and retailers have been trying to build truly seamless omnichannel customer experiences for decades. For banks, the task has been made particularly difficult by their significant scale and the fact that a customer's personal information may be stored in multiple places around the world, in disparate systems that don't talk to each other.

Neobanks don't have these fragmented systems and are therefore much better able to offer a consolidated journey. They can be highly personalized, successfully rolling out new products that appeal to different audiences, such as frequent travelers who handle multiple currencies and members of Gen Z with distinctive patterns in money management.

Using agentic AI, banks can conceivably have the best of both worlds, more easily developing new products and accessing data in different parts of the bank to provide customers with information and product suggestions just as they need them. This is especially promising in wealth management, and could be particularly useful in addressing the needs of mid-range clients that are currently very difficult and expensive to serve. In this case, an agent could take on much of the work that surrounds advice—portfolio monitoring, proposal drafting, meeting preparation, suitability documentation—enabling human advisors to serve mass-affluent clients economically, as the advice itself remains under advisor and compliance oversight.





Customer spotlight:

Using agents to accelerate onboarding

Delays in onboarding a new corporate customer are at least as frustrating for banks as they are for organizations that need banking services. But anti-money laundering and know-your-customer regulations produce substantial repetitive manual work, with no obvious way to accelerate it. At one business-to-business bank in Switzerland, onboarding times averaged two to six months.

Seeking to slash that time to two weeks, the bank turned to generative and agentic AI to transform fragmented, manual processes and, where possible, replace them with governed automation. The bank uses generative AI and agents to automate document extraction and reconciliation; screen for sanctions, politically exposed persons, and media; produce risk scoring; prioritize cases; and provide system updates and ongoing client monitoring. Throughout, compliance officers maintain final decision-making authority.

The goal is to allow technology to handle routine cases and to build operations around exceptions that demand a higher level of scrutiny. This would reduce operating costs and review effort, improve data quality and auditability, accelerate onboarding and time-to-revenue, and allow the bank to grow without a proportional increase in compliance headcount.

While the bank continues to expand this initiative, it has already cut customer data-intake time by 85% and improved productivity of compliance officers by 40%.

A photograph of a laptop on a wooden desk with a hand typing on the keyboard. A large, thick red 'X' is drawn over the right side of the image, crossing through the word 'Recommendations'.

Recommendations

Build the operational foundations first

It's not productive to bring AI agents into an environment that's not ready for them. Agentic AI is a full enterprise stack, with a set of layers that need to work properly to produce results. The model is one layer; the framework is another. The short- and long-term memory bring context, and tools and infrastructure support performance.

The rest of the organization needs to be ready, too. Traditional governance frameworks are designed for static, predictable environments; policy as code helps your governance adapt in concert with external requirements. Observability and monitoring also need to be in place, along with identity controls, DevSecOps, and trustworthy data sources.

All of this needs to work properly before agents can be successfully scaled throughout an enterprise. The winner may not be the bank that moves the fastest.

Identify the impacts on security and resilience, enterprise-wide

Before beginning a modernization program or agentic rollout, consider the entire bank and think through the impacts on data privacy, transaction security, anti-fraud efforts, and the ability to secure data and applications in a post-quantum world.

Each step in the modernization of the tech stack—from infrastructure up through enterprise platforms, middleware, and applications—brings new capabilities and efficiencies. But every step also has the potential to create vulnerabilities in the form of security issues, latency, and fragmentation. Adding AI further expands the attack surface, introducing model risk and the potential for prompt injection, among other threats. Build resilience into your evolving architecture and test it continuously.



Agentic AI must be a joint project of the business and technology teams

Implementation of agentic AI is a team sport, and it will stretch professionals across the organization, from business and technology to legal, risk, and compliance.

The technology and business teams should collaborate from the very start and be willing to dive into each other's domains. Technology teams will have to see the project through the lens of business outcomes. Business teams will need to understand technology at a level of detail that may be new to them.

Their teamwork will help foster the radical transparency necessary to clearly understand the potential benefits of agents and how they might best be implemented. It will also be critical in managing stakeholder expectations, which can have a distorting effect. A CFO, for example, may seek cost savings from a use case that is designed to increase productivity. Working together, the business and technology teams can clarify the expected outcomes of the use case, make the most persuasive argument for it, and be ready to work towards its success.

Modernization: Decouple before replacing

Legacy technologies weren't meant to support agents, digital assets, or new payment rails. But they've done an excellent job supporting the existing business. Big rip-and-replace programs are unnecessarily risky with uncertain time horizons and follow-on effects.

Yet banks need an agile, secure, modern infrastructure. So focus on APIs, modularization, service abstraction and isolation layers. Together, these allow modernization with resilience, sovereignty, operational control, and governed AI adoption, and help lay the groundwork for an agentic orchestration layer. They also enable isolated changes without affecting the entire system.

Importantly, this approach gives banks options. Banks may still choose to replace their core system. Or they can find a workaround: implementing a cash layer or a digital twin, or building a light ledger to sit on top of the core system. Such a light ledger decouples the primary transaction engine from rigid legacy software. The light ledger processes real-time and high-volume transactions immediately, while leaving complex regulatory functions to batch-based systems.



Actively manage AI enablement

Successful rollouts of agentic AI will change an organization's workflows and the way work is managed and completed. To successfully use agents at enterprise scale requires strong governance, security accountability, and continuous monitoring and auditing.

Empowered organizations are prepared to manage shadow agents and non-enterprise agents, as well as agents that are launched by business units and teams without appropriate organization-wide governance or monitoring. As agents become increasingly autonomous, organizations need visibility into what agents exist, what they can access, what actions they are authorized to take, and who is accountable for them throughout their lifecycle.

An AI enablement team should be established to:

- Establish enterprise-wide discovery, inventory, and governance of AI agents, including clear ownership, approved use cases, risk classification, certification, and lifecycle management.
- Ensure that compliance, regulatory, security, privacy and risk teams are brought in on Day One, with governance requirements embedded into AI design.
- Establish identity and access policies and controls for AI agents and other non-human identities.
- Ensure third-party AI products and services meet enterprise requirements for security, privacy, data protection, regulatory compliance, resilience, transparency, and ongoing risk management.
- Define guardrails for the data, models, tools, APIs, and systems that agents can access and the actions they are permitted to perform.
- Continuously monitor agent behavior and outcomes. Maintain appropriate audit trails and mechanisms for human intervention when required.
- Establish clear accountability for AI-driven actions and decisions, including escalation paths and the ability to suspend, contain, or retire agents.
- Evaluate the probable token and infrastructure costs of agentic AI use cases and establish token budgets, cost-per-task tracking, and business-value measures.
- Make sure that the work of AI projects can be repurposed—building reusable AI capabilities, controls, patterns, and components once and applying them across multiple use cases.
- Prepare the organization to develop the training and awareness of AI risks and the skills required for future AI development.
- Regularly reassess AI agents as models, regulations, threats, business processes, and organizational risk tolerance changes.

For smaller and mid-size institutions: consider leapfrogging to cloud-native

The right infrastructure depends on scale. For smaller institutions such as credit unions and building societies, AI and agentic AI may present a compelling reason to move from an older, batch-based core to a cloud-native platform. The largest institutions may conclude that only their existing systems can provide the necessary throughput and resilience. For those organizations, decoupling is the preferred path.

Design for regulatory divergence

Banks operating across the E.U., U.K., Middle East and Latin America face diverging AI and payments regimes. Policy as code can help codify, test and apply differing regulatory requirements across jurisdictions, tying regulations such as DORA, the EU AI Act, and national and local rulemaking into a bank's existing framework.





How Kyndryl can help

The Kyndryl Agentic AI Framework is an enterprise-grade approach to designing, governing, and scaling agentic AI systems—AI systems that don't just generate insights but also plan, decide, and act autonomously across complex, mission-critical environments, while remaining secure, auditable, and aligned with business intent.

The framework is composed of key capabilities:

- **The agentic core:** This secure-by-design capability orchestrates, secures, and scales AI agents throughout the organization. It delivers new operational processes, unlocking efficiencies and growth opportunities at speed.
- **Agentic ingestion:** This process extracts and analyzes a customer's entire technology estate, including policies, processes, data schemes, and workflows, uncovering hidden complexities and integration challenges.
- **Agent catalog:** This library of validated agents and proven blueprints helps organizations determine if their current state should be transformed into enterprise-grade agents, modern code, or a combination.
- **Agent builder:** The agent builder leverages industry and domain reference architectures and the agent catalog to build, design, engineer, and deploy agents tailored to each enterprise.
- **The future workforce model:** This architecture defines the roles agents will play within an organization and how they will collaborate with employees.
- **Policy as code:** This capability translates organizational rules, regulatory requirements and operational controls into machine-readable policies that govern how AI agents execute. Policy as code brings governance into the workflow itself, at runtime, permitting, blocking, escalating and logging agentic actions before they negatively affect production systems.






To help customers transform quickly, safely, and iteratively, Kyndryl has introduced agentic modernization services-as-software. With this model, Kyndryl has codified its trusted mission-critical and engineering expertise, together with capabilities from AI ecosystem partners, into pre-defined agentic workflows that help customers unlock business outcomes, de-risk and accelerate their modernization objectives, while maintaining enterprise guardrails and cost visibility throughout their AI adoption journey.

The Kyndryl Agentic AI Framework helps organizations:

- Reimagine workflows for AI native execution
- Scale AI safely across enterprise and hybrid environments
- Embed trust, governance, and compliance into agentic workflow execution and AI-assisted decisions
- Translate AI capability into real operational outcomes

Whatever the current state of your modernization and AI initiatives, Kyndryl has the experience and expertise to help you innovate quickly, adapt to changing business needs, accelerate your time to value, and maximize the impact of your AI investments.

[Schedule a consultation →](#)



Citations

- ¹ The Readiness Tipping Point, Kyndryl, 2025.
- ² State of AI in Financial Services: 2026 Trends, Nvidia, January 2026.
- ³ Stablecoins and the future of onchain finance, Visa, 2026.
- ⁴ The Readiness Tipping Point, Kyndryl, 2025.
- ⁵ Kyndryl People Readiness Report, Kyndryl, 2026.
- ⁶ Kyndryl People Readiness Report, Kyndryl, 2026
- ⁷ State of AI in Financial Services: 2026 Trends, Nvidia, January 2026.
- ⁸ The Readiness Tipping Point, Kyndryl, 2025.
- ⁹ State of AI in Financial Services: 2026 Trends, Nvidia, January 2026.

