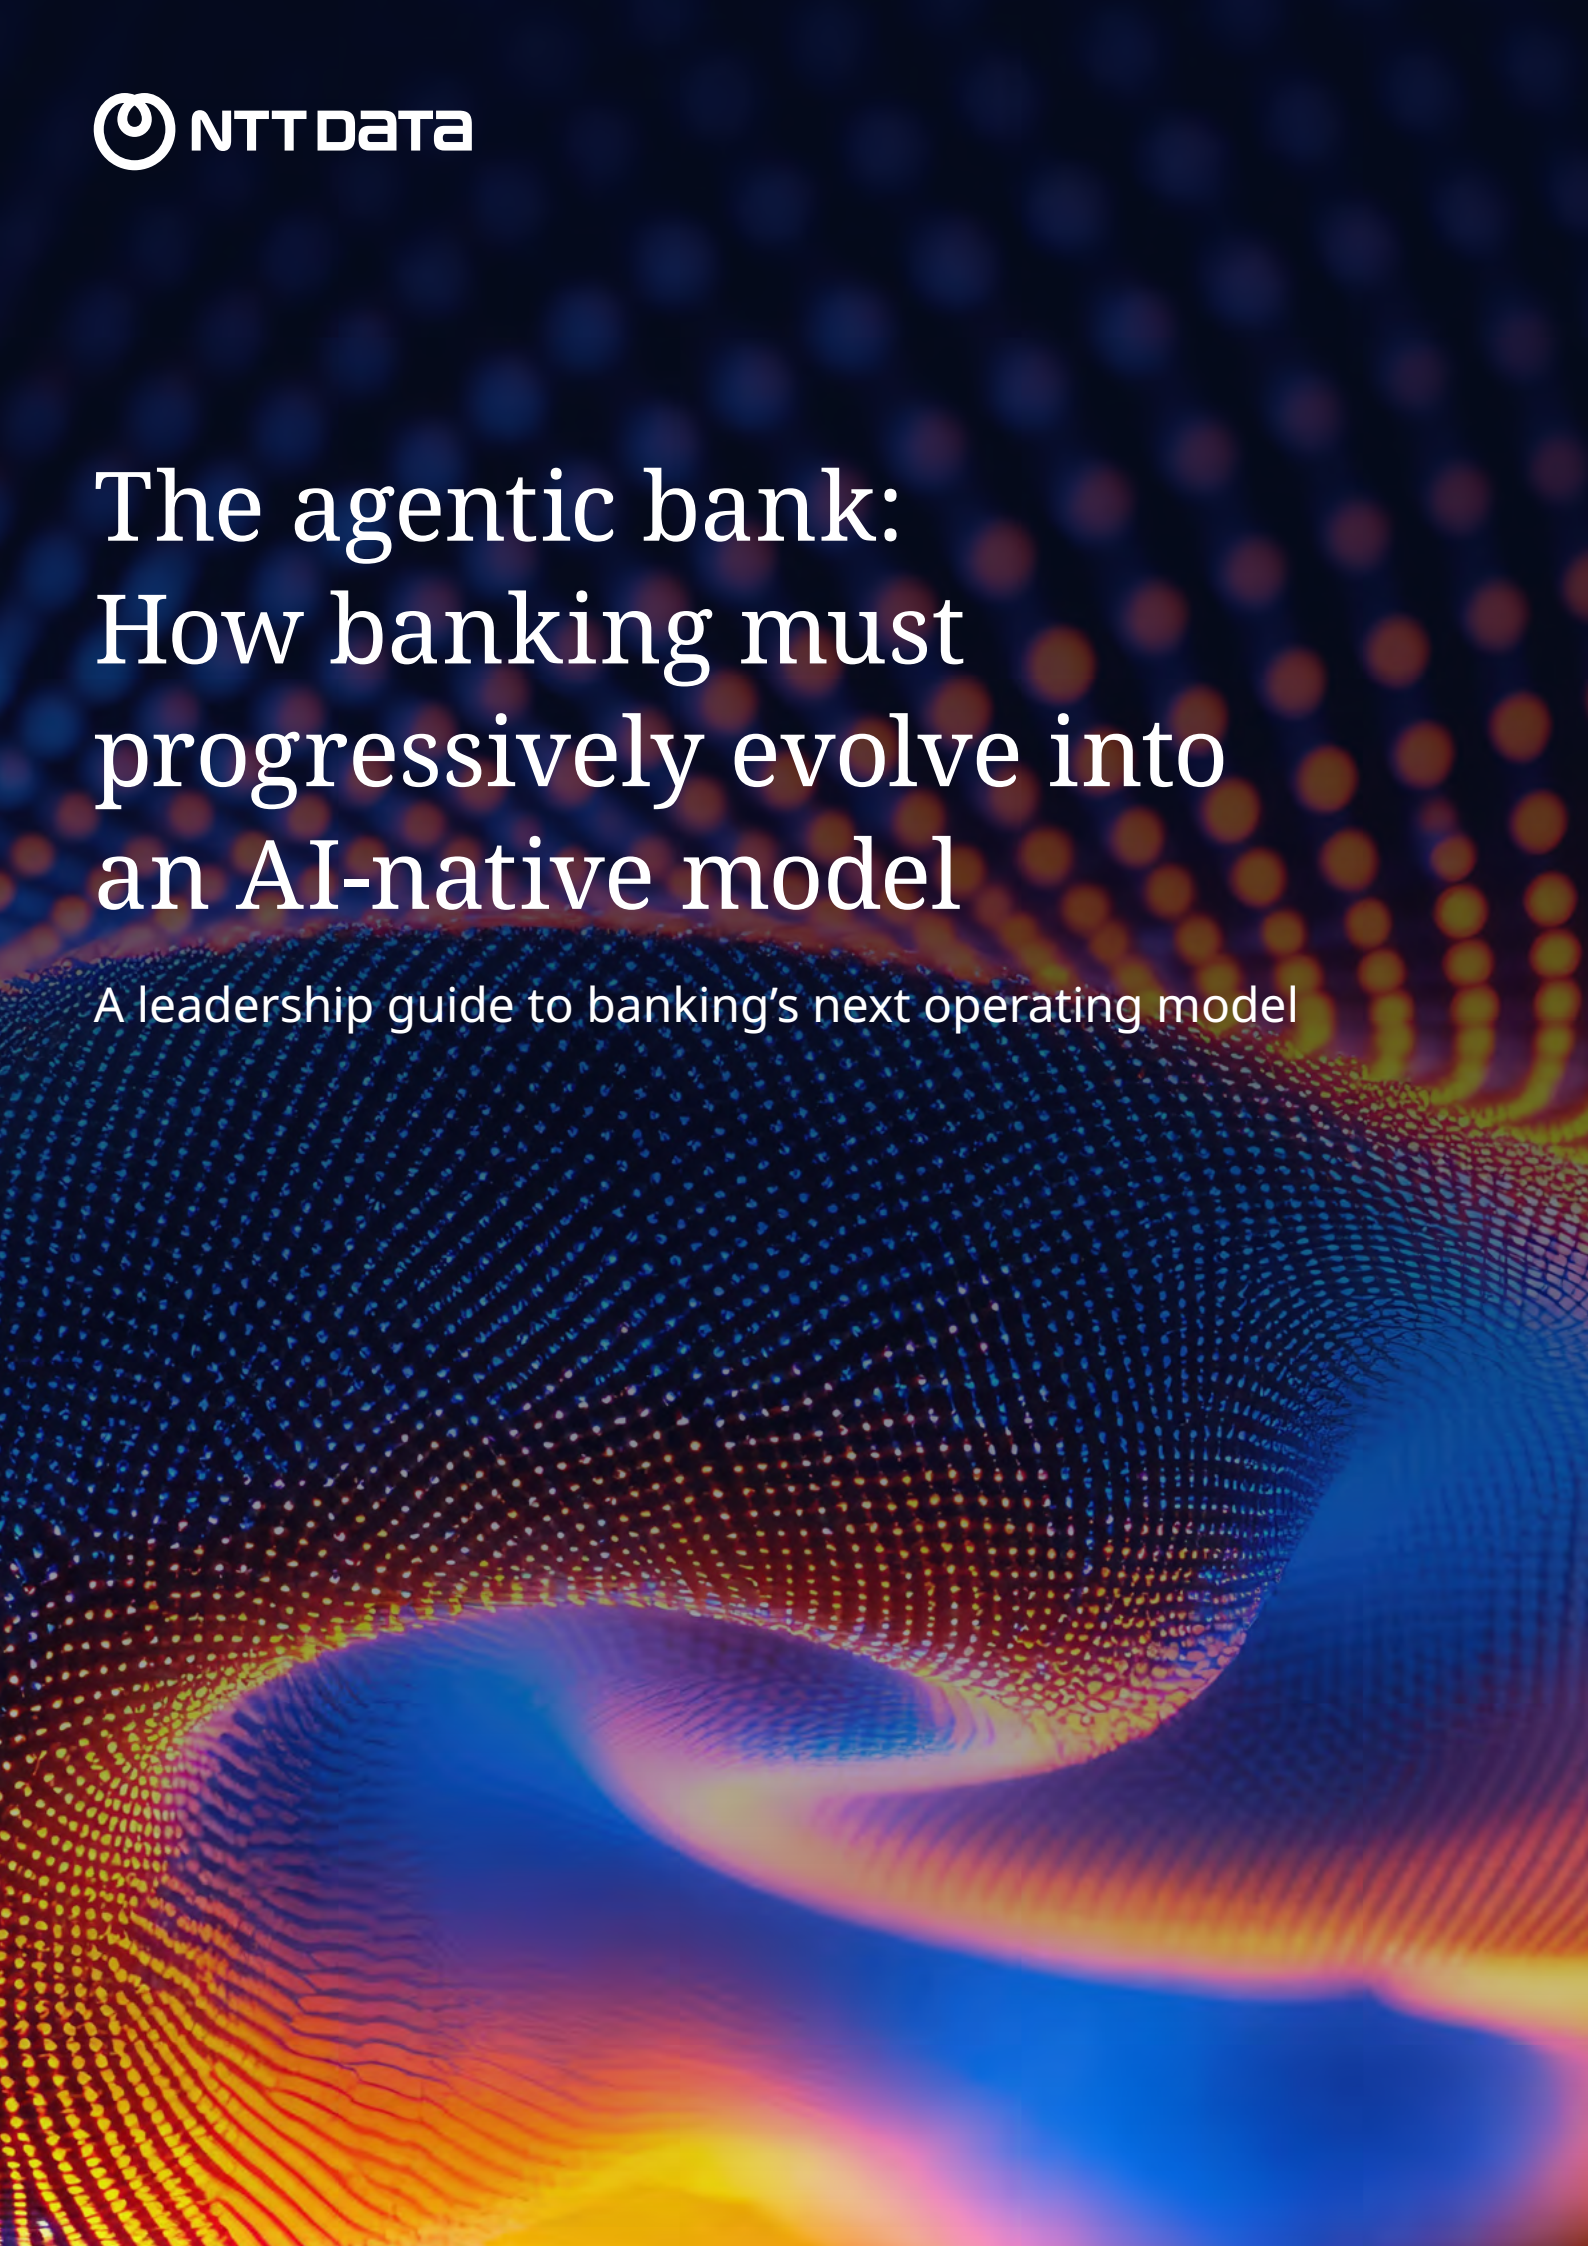
NTT DATA

The agentic bank: How banking must progressively evolve into an AI-native model

A leadership guide to banking's next operating model



Contents

03 Foreword

04 The agentic imperative

05 Dimension 1: Customer relationship and experience

06 Dimension 2: Product and service evolution

07 Dimension 3: Organization and talent

08 Dimension 4: Information and knowledge layer

10 Dimension 5: Governance, identity and agent lifecycle

11 Dimension 6: New lines of defense

12 Dimension 7: Technology

13 Dimension 8: Operations

14 Leading the agentic bank

15 From ambition to action

15 Why NTT DATA

 © 2026 NTT DATA, Inc.

Foreword

Agentic banking is not simply the addition of new technology to the existing model. It represents a broader transformation of customer relationships, products and services, organizational structures, the way work is executed and the mechanisms of control. The result is a bank where people and agents work together to understand intentions, configure solutions, coordinate capabilities and act in real time. The objective is not unconstrained autonomy, but trusted, governed and accountable automation.

This guide examines eight dimensions of that transformation. For each, we look at what banking does today, what agentic AI changes and what executives must decide. The framework uses NTT DATA's global experience in retail, wholesale, wealth and capital markets. It is based on our view that the organizations building these capabilities today will help define the next decade of financial services.

The opportunity is significant. The decisions being made now will shape what banking becomes next.



 © 2026 NTT DATA, Inc.

The agentic imperative



Banking's holistic transformation

The banking industry has spent the past several years experimenting with AI. The results have been instructive, but isolated pilots alone do not create transformation. The organizations that gain the greatest advantage will not necessarily be those with the most AI use cases, but those that can integrate agentic AI into how the bank operates.

Earlier forms of AI, including generative AI, largely augmented individual tasks. Agentic AI goes further by coordinating activities across workflows and taking actions within defined limits. An AI agent is a purpose-built digital specialist designed to achieve a specific business objective. An AI agency brings specialized agents together to execute an end-to-end business function, sharing context, managing exceptions and operating under common governance.

This changes how banks compete. The ability to onboard a commercial customer faster, detect fraud before a transaction settles or provide a relevant mortgage recommendation at the right moment depends not simply on individual products or technologies, but on how effectively organizations can connect information, decisions and execution.

Evidence is moving beyond experimentation, but maturity remains uneven. The strongest programs increasingly pair near-term use cases with longer-term investments in data, governance, technology and workforce readiness. The strategic question is changing from whether agents can perform useful work to which decisions and actions a bank is prepared to delegate, under what controls and at what economics.

But this transformation is not technological alone. Agentic banking reshapes how banks relate to customers, design products, organize work, manage data, govern AI and control new forms of risk. The following eight dimensions map those changes — and the decisions leaders need to make as agentic AI moves from experimentation toward enterprise adoption.

In the coming years, we will witness the evolution of this technology's adoption as it moves from efficiency improvements to business reinvention. We will pass through a truly exciting phase of transformation in which we are already immersed.

Dimension 1: Customer relationship and experience

From channels to intent

Agentic AI creates an opportunity to redefine the banking relationship. As agents become a new gateway to financial services, third parties can increasingly influence customer intent, recommendations and experiences. This will initially relegate banks to the role of product/service provider or transaction engine. At the same time, agentic AI will allow banks to move beyond channel-based self-service toward more contextual, personalized and proactive relationships.

Focusing on the development of trust as a defining human characteristic represents an intelligent evolutionary step in the customer relationship model. It is in the "moments of truth" that this trust — acting as a natural enabler of decision-making — will take on a distinctly human role.

Today, banking relationships are organized around channels: branch, relationship manager, web, mobile app and contact center. Information and interactions do not always move seamlessly between them, forcing customers to repeat their needs. Personalized advice also depends heavily on access to expert professionals and is often concentrated in higher-value segments.

Under an agentic AI paradigm, shared context can help banks understand intent, maintain continuity and coordinate responses across channels, agents and professionals. The experience can be designed around the customer's needs rather than the channel through which they enter. Personalized guidance can extend further into retail and small-and-midsize enterprise (SME) banking. Relationship managers in corporate banking will gain access to real-time financial and sector intelligence. Human involvement can then focus on trust, expert judgment, negotiation and decisions of greater consequence.

The customer interface becomes a strategic asset. Banks must decide which interactions they want to lead, where they will collaborate with third parties, and how they will differentiate through trust, knowledge, advice and execution. Comparable products and standardized processes face the greatest risk of disintermediation. Complex, high-value relationships remain more defensible when banks combine expertise with strong execution.

The strategic question is clear:
Will agentic AI deepen the bank's relationship with the customer, or allow another platform to own it?

 © 2026 NTT DATA, Inc.

Dimension 2: Product and service evolution

From static products to modular, agentic propositions

Agentic AI transforms financial products from static, predefined offerings into more flexible propositions that can adapt their composition, service level and pricing to each customer's context, risk and objectives. Products must also be designed to be discovered, compared, recommended and contracted by AI agents. It requires banks to rethink distribution across customer agents, proprietary platforms and third-party ecosystems. This increasingly requires machine-readable product rules, eligibility, pricing and service constraints so agents can discover and configure propositions reliably.

Today, offerings are largely built around predefined products and customer segments, with limited flexibility to tailor individual proposals. Banks typically address the financial component of a customer need, leaving customers to coordinate other services and steps themselves. We must also ask ourselves whether market segments will remain an operational concept in the near future, or if their boundaries will become blurred by a more personalized approach to each customer's reality.

Agentic AI allows banks to assess more information and alternatives when tailoring product composition, conditions and pricing. All within the boundaries of suitability, profitability and fair treatment. The proposition can expand beyond individual financial products to combine financial and non-financial services around broader goals such as purchasing a home, starting a business or managing treasury[AS2.1]. Products must also be discoverable, comparable and configurable by agents acting on behalf of customers or through external platforms.

This creates three distinct distribution environments. With customer agents, the bank provides capabilities such as products, financing, data or execution. On proprietary platforms, it manages the end-to-end experience. In third-party ecosystems, it contributes financial capabilities to a broader solution. Across all three, banks face the same strategic choice: Lead the experience or provide the capabilities that enable others to lead.

The opportunity is already taking shape. In retail and SME banking, digital CFO capabilities such as cash-flow forecasting, funding recommendations and working capital optimization are becoming competitive differentiators. In wealth and asset management, AI can support personalized portfolio analysis, reporting and advisor insights. In capital markets, it can support research, trading, surveillance and customer-specific market intelligence.

The shift is ultimately from distributing individual products toward resolving broader customer needs — opening opportunities beyond the traditional product model.



 © 2026 NTT DATA, Inc.

Dimension 3: Organization and talent

From project-based structures to end-to-end ownership

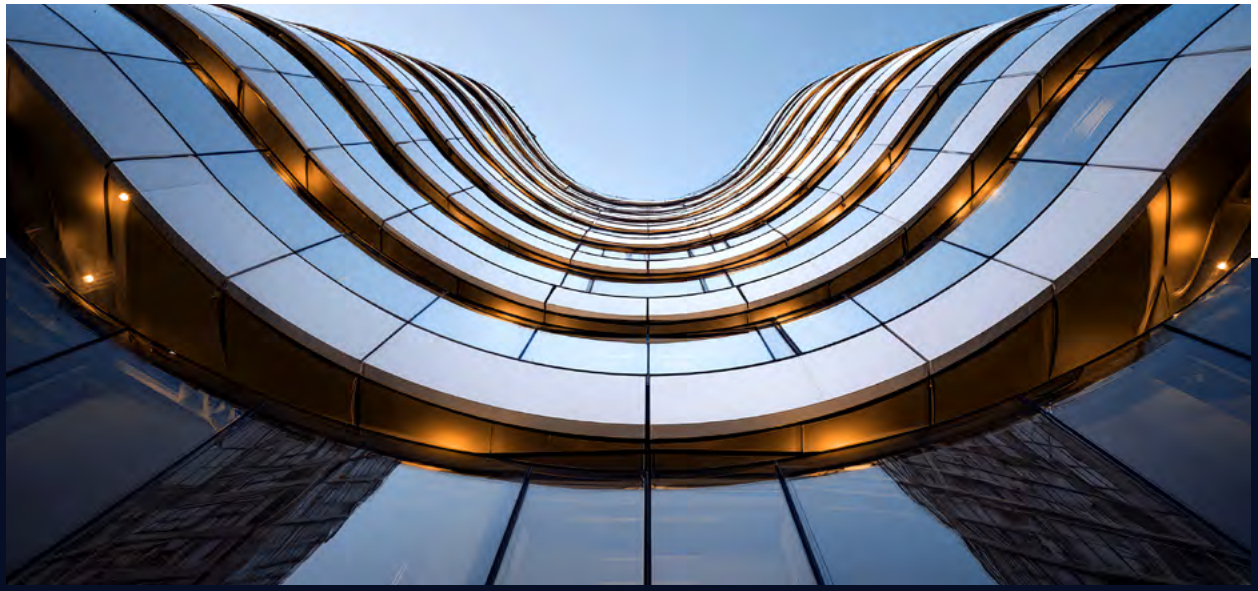
Banks are evolving from functional, hierarchical organizations built around technology projects to models organized around products, platforms and end-to-end processes. These models deliver clearer accountability for business outcomes. AI is no longer solely a technology capability; it is a bank-wide competency requiring new workforce skills and ways for people and agents to work together.

In most companies, the impact of change is significantly greater from an organizational perspective than from a technological one. To simplify the equation, we could say that technology is the "easy part to buy," whereas talent — people who know how to use it effectively, trust it and redesign their way of working — and the organization — governance, processes, incentives and leadership — are the difficult parts to build.

Today, much transformation is organized through temporary projects that deliver a solution and then disband. This leaves ongoing improvement spread across functions. AI expertise is often concentrated within technology or specialist centers.

Agentic AI requires more continuous ownership. Permanent, multidisciplinary teams need responsibility for operating, improving and governing solutions throughout their lifecycle — not simply implementing them. This requires shared platforms, standards and governance across the organization. Teams are responsible for adapting and overseeing agents within their domains.

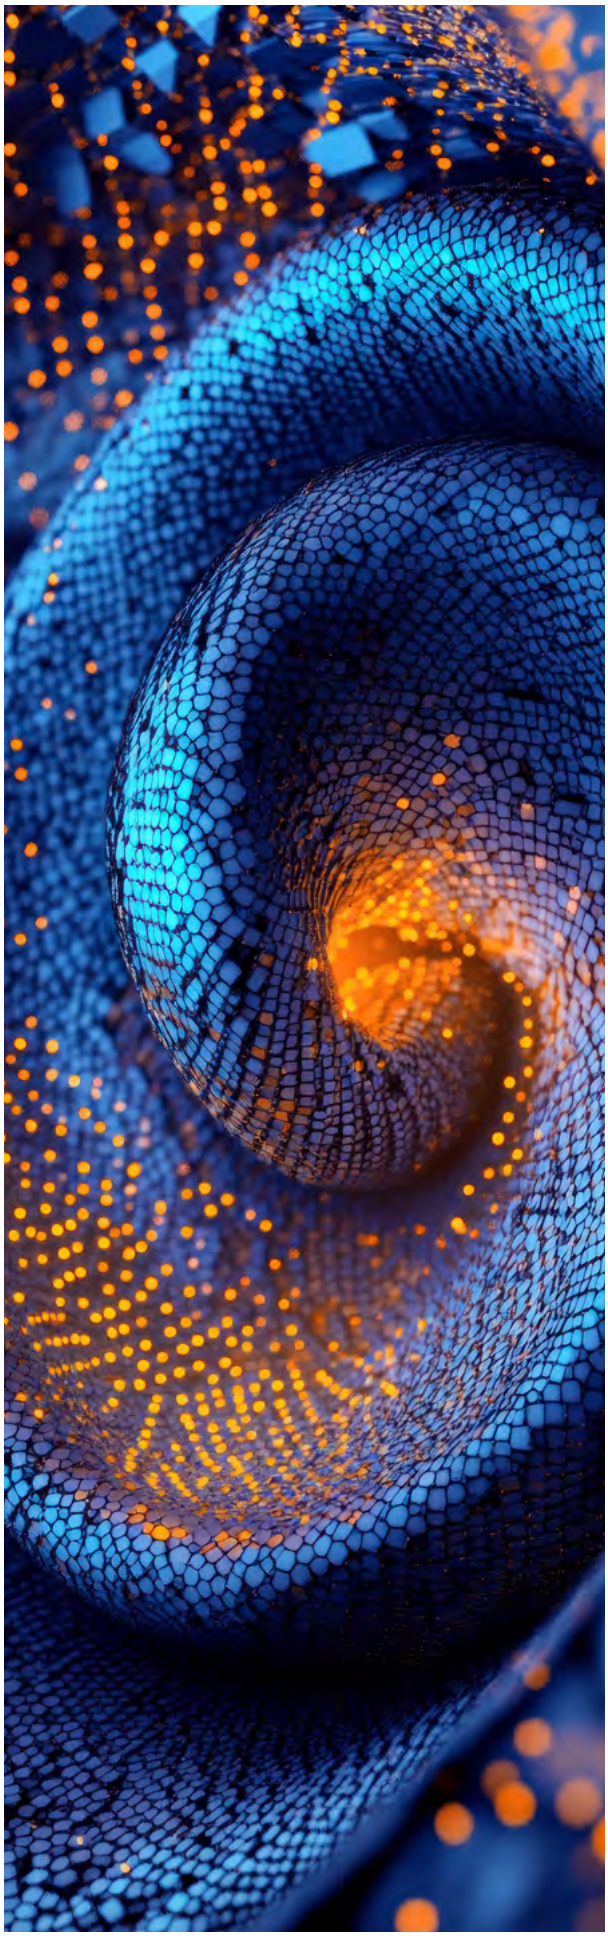
Skills will vary across the bank. Employees need a foundational understanding of AI. Leaders and process owners need to manage AI-enabled operations. Technical specialists remain critical across engineering, data, security, validation and governance. Business experts need to supervise and improve agents. As routine execution declines, operational roles will increasingly emphasize exception management, professional judgment, advisory and oversight.



Organizations that succeed will treat workforce transformation as a strategic priority, not a consequence of technology deployment. Reskilling, internal mobility and the preservation of organizational knowledge will be essential to sustaining the agentic bank.

 © 2026 NTT DATA, Inc.

Dimension 4: Information and knowledge layer



From data to executable context

As AI models become more widely available, competitive advantage will increasingly depend on what each bank brings to them. Strategy, policies, business knowledge, customer intelligence and employee expertise must be converted into reliable, traceable and actionable context for agents.

Historically, transactional data, documents and information about customers, products and risks have often been managed separately. Analytical use cases could also tolerate some incomplete or outdated information because outputs could be reviewed after the fact.

Agentic AI raises the stakes. Two banks using similar models can produce different outcomes based on the context available to their agents. Banks need to connect transactional and analytical data with contracts, communications, procedures, regulations and other sources of organizational knowledge. This will help agents operate according to the organization's priorities, policies and limits.

Ontologies and a semantic layer become essential

This is where ontologies and a semantic layer become essential, not optional. An agent cannot reason correctly over raw, disconnected data any more than an engine can run on crude oil straight from the ground; it needs to be refined into a common language first. An ontology defines what a "customer," a "product," a "risk exposure" or an "obligation" means across the organization, how these entities relate to one another and which business rules govern them. A semantic layer then makes that meaning consistently accessible. It translates scattered systems, formats and terminologies into a shared model that both humans and agents can interpret in the same way. Without this layer, agents are left to infer meaning on their own, silently making assumptions that may be inconsistent, outdated or simply wrong. With it, agents can navigate the bank's knowledge the way an experienced employee would: understanding not simply where data lives, but what it represents and how it should be used.

Data, in this sense, is the fuel of AI and of agentic systems. But fuel that is unrefined, contaminated or poorly stored does more harm than good, no matter how powerful the engine is.

The most advanced models, running on ungoverned, ambiguous or fragmented data, will still make poor decisions or, worse, confidently wrong ones. Ontologies and semantic layers act as the refining infrastructure that turns raw data into something agents can safely and consistently consume: context that is available, as well as meaningful, structured and trustworthy. Investing in this layer is, in effect, investing in fuel quality. And it is what ultimately determines how far, how fast and how safely a bank's agentic AI can go.

As agents take on greater responsibility, the quality, timeliness, source and accessibility of that information become critical. Banks must be able to demonstrate what information an agent used, who had access to it and which version was current. This also applies to agent memory: what is retained, for how long, and under which residency or sovereignty constraints. Stronger traceability and control allow greater autonomy; weaker controls require greater human oversight.

Priority areas include customer and entity intelligence, product catalogs, contracts and transaction history, risk and pricing data, communications and regulatory documentation. Bringing these resources together — organized through a coherent ontology and exposed via a semantic layer — provides a consistent information foundation for AI across the bank.

Organizations that manage data and knowledge as business assets — with clear ownership, quality standards and a well-governed semantic layer — will be better positioned to scale agentic AI responsibly. Those that treat them primarily as technology concerns, or that skip the work of defining shared meaning across the organization, will be constrained by the information on which their agents depend. In the end, no amount of model sophistication compensates for poor fuel.



Dimension 5: Governance, identity and agent lifecycle

From documented policies to verifiable, executable limits

Every agent must have an identity, purpose, responsible owner and a clearly defined mandate. Its credentials and delegated authority must be explicit, auditable and revocable. That mandate determines the information it can access, decisions it can make, tools it can use and limits within which it can act. Governance must apply throughout the agent's lifecycle, with autonomy and human oversight proportionate to risk.

Traditional software executes predefined functions, with governance often focused on approval before a solution enters production. Agentic AI requires a different approach because agents can interpret situations, choose between alternatives and act with varying degrees of autonomy.

Each agent requires a defined mandate, risk classification and accountable business owner. Its risk level determines permissions, controls, human validation and escalation requirements. Governance must extend from design and development through deployment, operation, modification and retirement.

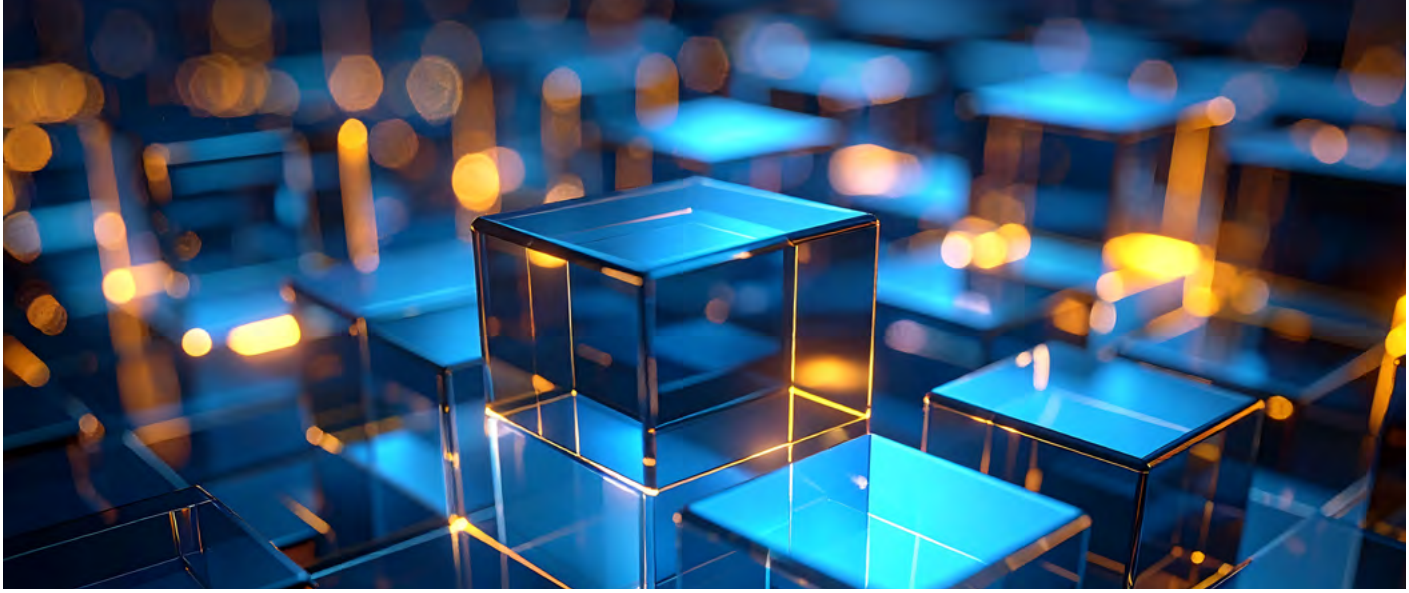
Banks also need clear accountability for who creates, approves, supervises and can intervene in each agent. Common policies and standards can apply across the organization while responsibility for individual agents remains within the relevant business domain.

Regulatory obligations

Existing regulatory obligations remain. In credit, anti-money laundering (AML), fraud, sanctions, suitability, pricing, trading and other areas where AI influences decisions affecting customers or markets, banks remain responsible for the outcome regardless of whether a person or an agent makes the decision. The EU AI Act and the Digital Operational Resilience Act (DORA) add further requirements for managing AI and technology risks.

Other jurisdictions are approaching this differently. The US has no single AI law. Existing rules like the Equal Credit Opportunity Act (ECOA) and fair lending laws apply to AI decisions. However, federal regulators have so far excluded generative and agentic AI from their main model-risk guidance. This leaves a gap that individual states, such as Colorado, are starting to fill with their own requirements. Banks operating internationally should expect this regulatory patchwork to keep evolving.

Governance is not a constraint on the agentic bank. It is what enables greater autonomy safely. Clear accountability, traceability and appropriate human oversight create the conditions for agentic AI to operate in regulated environments.



Dimension 6: New lines of defense

From periodic oversight to continuous control

Risk, compliance, cybersecurity and audit processes must evolve alongside agentic AI. As agents gain greater autonomy, these functions need to define, validate and independently oversee the boundaries within which agents operate.

Today, many policies are documented and verified through upfront reviews or after-the-fact controls. When agents make decisions and take actions in real time, however, controls increasingly need to operate in real time, too.

Permissions, thresholds and restrictions must be built into how agents operate. Control functions must independently validate that those limits reflect the bank's policies and work as intended. This allows banks to prevent inappropriate actions, escalate sensitive decisions and demonstrate compliance.

Agents also introduce new security risks as they interact with models, data, tools, internal systems and third-party services. These include manipulated instructions or context, unauthorized use of tools, agent impersonation, information leakage and errors that can spread across connected systems.



Traceability is critical. Banks must be able to reconstruct which agents, models, data, instructions and tools contributed to a decision or action. They must also have mechanisms to stop, restrict or escalate actions when problems arise. Audit and control functions must assess both individual agents and the risks that interactions between agents, systems and external services create. Assurance must include continuous evaluations, red-teaming and runtime monitoring of tool misuse, manipulated context, cascading failures and cross-agent interactions.

Organizations that build effective controls and independent oversight into their agentic environments will be better positioned to increase autonomy safely. Those relying primarily on documented policies and after-the-fact reviews will struggle to maintain effective oversight as agents become more capable and interconnected. At scale, correlated agent behavior may also create risks that are not visible when systems are assessed one by one.

Dimension 7: Technology

From isolated technological ecosystems to a shared agentic platform

Technology evolves from building and maintaining individual applications toward providing the shared platforms, architecture and engineering capabilities needed to create, deploy and manage agentic AI securely and at scale.

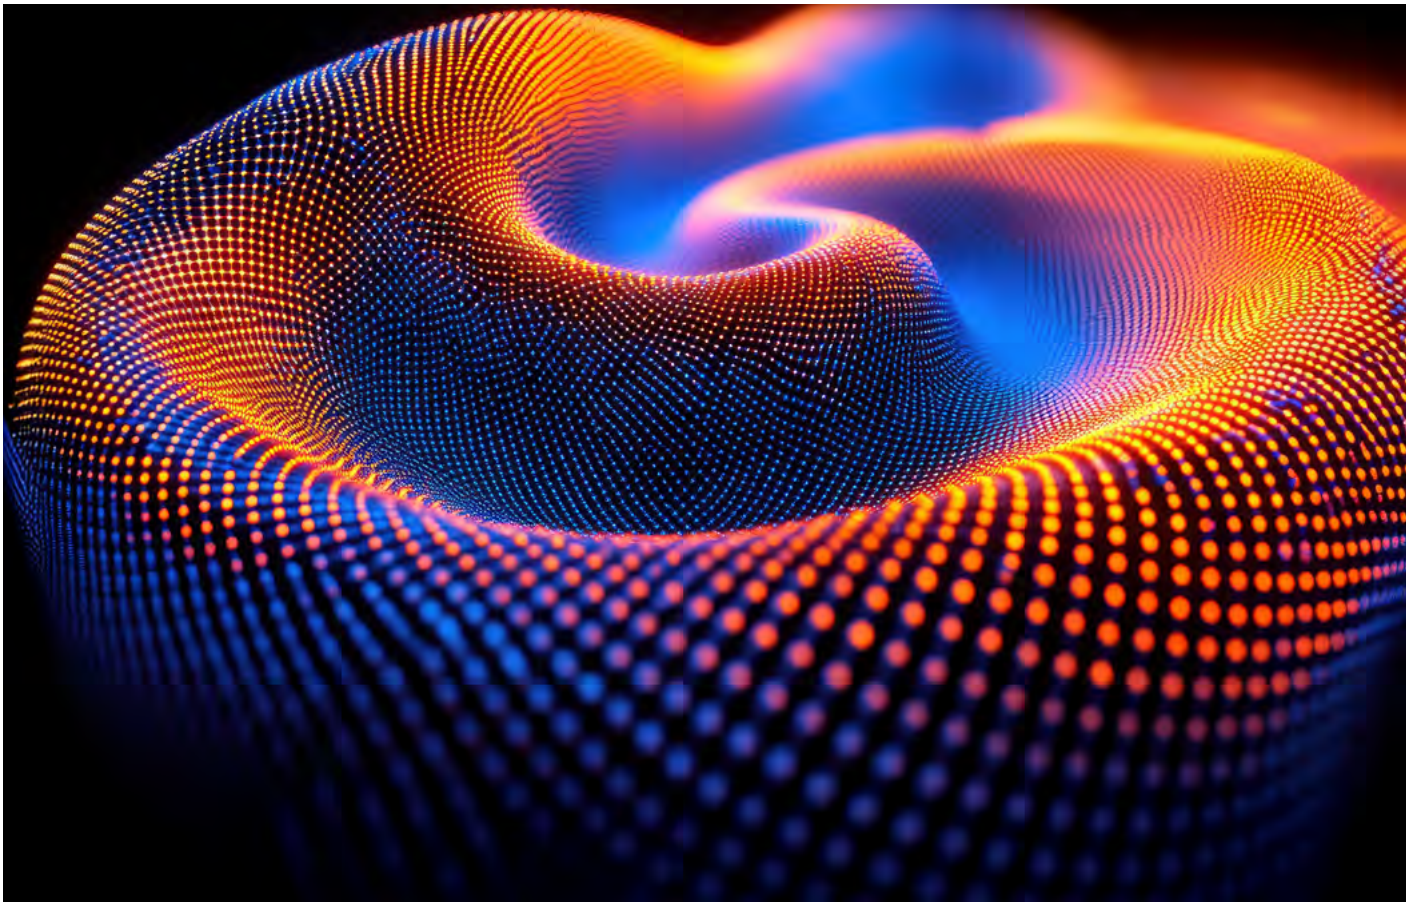
Today, applications and automations often have their own approaches to models, data, tools, security and oversight. This creates duplication and higher maintenance costs. Agentic AI requires a more integrated approach.

A shared platform connects models, agents, data, knowledge, tools and reusable capabilities. It coordinates work between people and agents, provides the information agents need to operate, and connects securely with core banking systems, legacy infrastructure and external platforms. Interoperability matters. Critical workflows should not be hard-wired to a single model, framework or provider. Agent-to-agent and agent-to-tool interactions should use secure, governed interfaces.

AI agents will also play a growing role in technology development, supporting activities from requirements and design to coding, testing, documentation and maintenance. This changes where people add the greatest value. It puts more emphasis on architecture, design decisions and quality control, while experienced engineering expertise remains essential.

Technology economics will change, as well. Agents use computing resources and external tools as they operate, often across multiple steps, making costs more continuous and variable. Banks need visibility into usage, cost and performance to optimize operations without compromising reliability or resilience. Unit economics should be measured per business outcome, not simply per model call.

Without a shared foundation, banks risk creating another generation of isolated pilots that are costly to maintain and difficult to govern. With it, banks can reuse capabilities and controls while exposing stable legacy functions through governed interfaces and modernizing progressively rather than assuming wholesale replacement.



Dimension 8: Operations



From task execution to end-to-end outcome management

Operations is evolving from executing tasks and managing cases toward designing, supervising and improving end-to-end processes. Agents are taking on more execution, and people focus on outcomes, quality, exceptions and continuous improvement.

Today, automation is primarily applied to individual tasks within processes distributed across teams, systems and people. Productivity is often measured by tasks completed, cases processed or time spent on individual activities.

Agentic AI allows banks to coordinate activities across systems, use relevant information and take actions within defined limits while maintaining context throughout a process. Operations can redesign workaround the outcome it needs to achieve, rather than simply accelerating individual steps. This can reduce handoffs and wait times. But it also requires banks to rethink their responsibilities, controls and exception handling.

As agents take on repetitive work, people can focus on oversight, professional judgment and sensitive or unforeseen situations. Exceptions can also provide insight into how processes should be improved. However, they need to be reviewed before being incorporated into rules, knowledge or future agent behavior.

Performance measures must evolve accordingly. Task-level metrics can mask delays, rework, quality issues or excessive human intervention. Banks need to measure outcomes such as resolution time, cost, compliance, customer satisfaction, quality and the level of human support required.

The opportunity extends across complaints resolution, onboarding and know your customer (KYC), credit underwriting, collateral management, loan servicing, payments and disputes, AML alert triage, trade finance, post-trade processing, reconciliations and regulatory reporting. The value comes not from automating individual tasks, but from redesigning how work is managed from beginning to end.

Leading the agentic bank

The executive agenda for 2026 and beyond

The transition to the agentic bank will not happen through isolated AI initiatives. It requires coordinated decisions across strategy, technology, data, governance, operations and talent. Organizations that begin building these foundations now will be better positioned to scale agentic AI responsibly and turn early investments into lasting business value. The business case should be explicit across growth, cost-to-serve, liquidity and risk management, and trust and resilience.

Five decisions define the executive agenda

01

Commit to a unified agentic AI strategy

A portfolio of disconnected pilots is not a strategy. Banks need a clear direction for how agents will be developed, governed and scaled, with investment-building capabilities that can be reused across the organization. The eight dimensions this guide outlines are not independent workstreams. They are interconnected parts of a single transformation that is as much organizational as it is technological.

02

Treat knowledge and data as strategic assets

As AI models become more widely available, differentiation will increasingly come from the context that a bank provides to its agents. This includes its strategy, policies, business knowledge, customer intelligence and employee expertise. Building this foundation should begin alongside agentic AI adoption, not after it.

03

Build governance into the foundation

Every agent needs a defined mandate, an accountable owner, an appropriate risk classification and lifecycle governance. Establishing these controls early allows banks to increase autonomy as capabilities mature while maintaining the accountability and oversight expected in financial services.

04

Prioritize end-to-end outcomes over isolated tasks

The greatest opportunity lies not in automating individual activities, but in redesigning complete processes. Areas such as onboarding and KYC, loan servicing, payments and disputes, AML alert triage, trade finance, reconciliations and regulatory reporting provide opportunities to reduce handoffs and improve outcomes.

05

Build the capabilities to scale

Moving from pilots to enterprise adoption requires more than selecting the right model or technology. Banks need the architecture, integration, interoperability, agent identity, evaluation and observability, governance, operational expertise and workforce capabilities to sustain agentic AI across the organization. Redesigning roles, skills and ways of working is often the harder — and more decisive — part of this shift.

From ambition to action

The journey does not need to begin everywhere at once. It should begin where business value, organizational readiness and the ability to manage risk intersect.

Start by assessing where your organization stands across the eight dimensions of the agentic bank. Next, identify the gaps that could prevent AI from scaling. Then, select a small number of end-to-end domains where agentic AI can address a meaningful business need. Establish the data, technology and governance foundations those domains require. Define where human judgment remains essential. Measure success against business outcomes rather than the number of agents or pilots deployed.

From there, reuse what works and strengthen the shared capabilities behind it. Finally, progressively expand agent autonomy as evidence, controls and organizational confidence mature.

Why NTT DATA

NTT DATA combines banking and financial services expertise with the technology, data, governance, modernization and transformation capabilities needed to help organizations move from agentic AI ambition to enterprise adoption.

The Agentic Bank Blueprint brings these capabilities together across the eight dimensions this guide outlines. It provides a practical framework for assessing readiness, identifying priorities and defining a path forward.

The agentic bank will not be defined by how many agents an organization deploys. It will be defined by how well people, AI, data and technology work together: to improve decisions, operations and customer outcomes; with the governance and accountability financial services demand.

The next step is to determine what that transformation looks like for your organization. NTT DATA can help you assess your current state, identify the right opportunities and build a roadmap for moving from experimentation to responsible scale.

Visit nttdata.com to learn more.

NTT DATA is a $30+ billion business and technology services leader in AI and digital infrastructure. We accelerate client success and positively impact society through responsible innovation. As a Global Top Employer, we have experts in more than 70 countries. NTT DATA is part of NTT Group.

 © 2026 NTT DATA, Inc.


NTT DATA