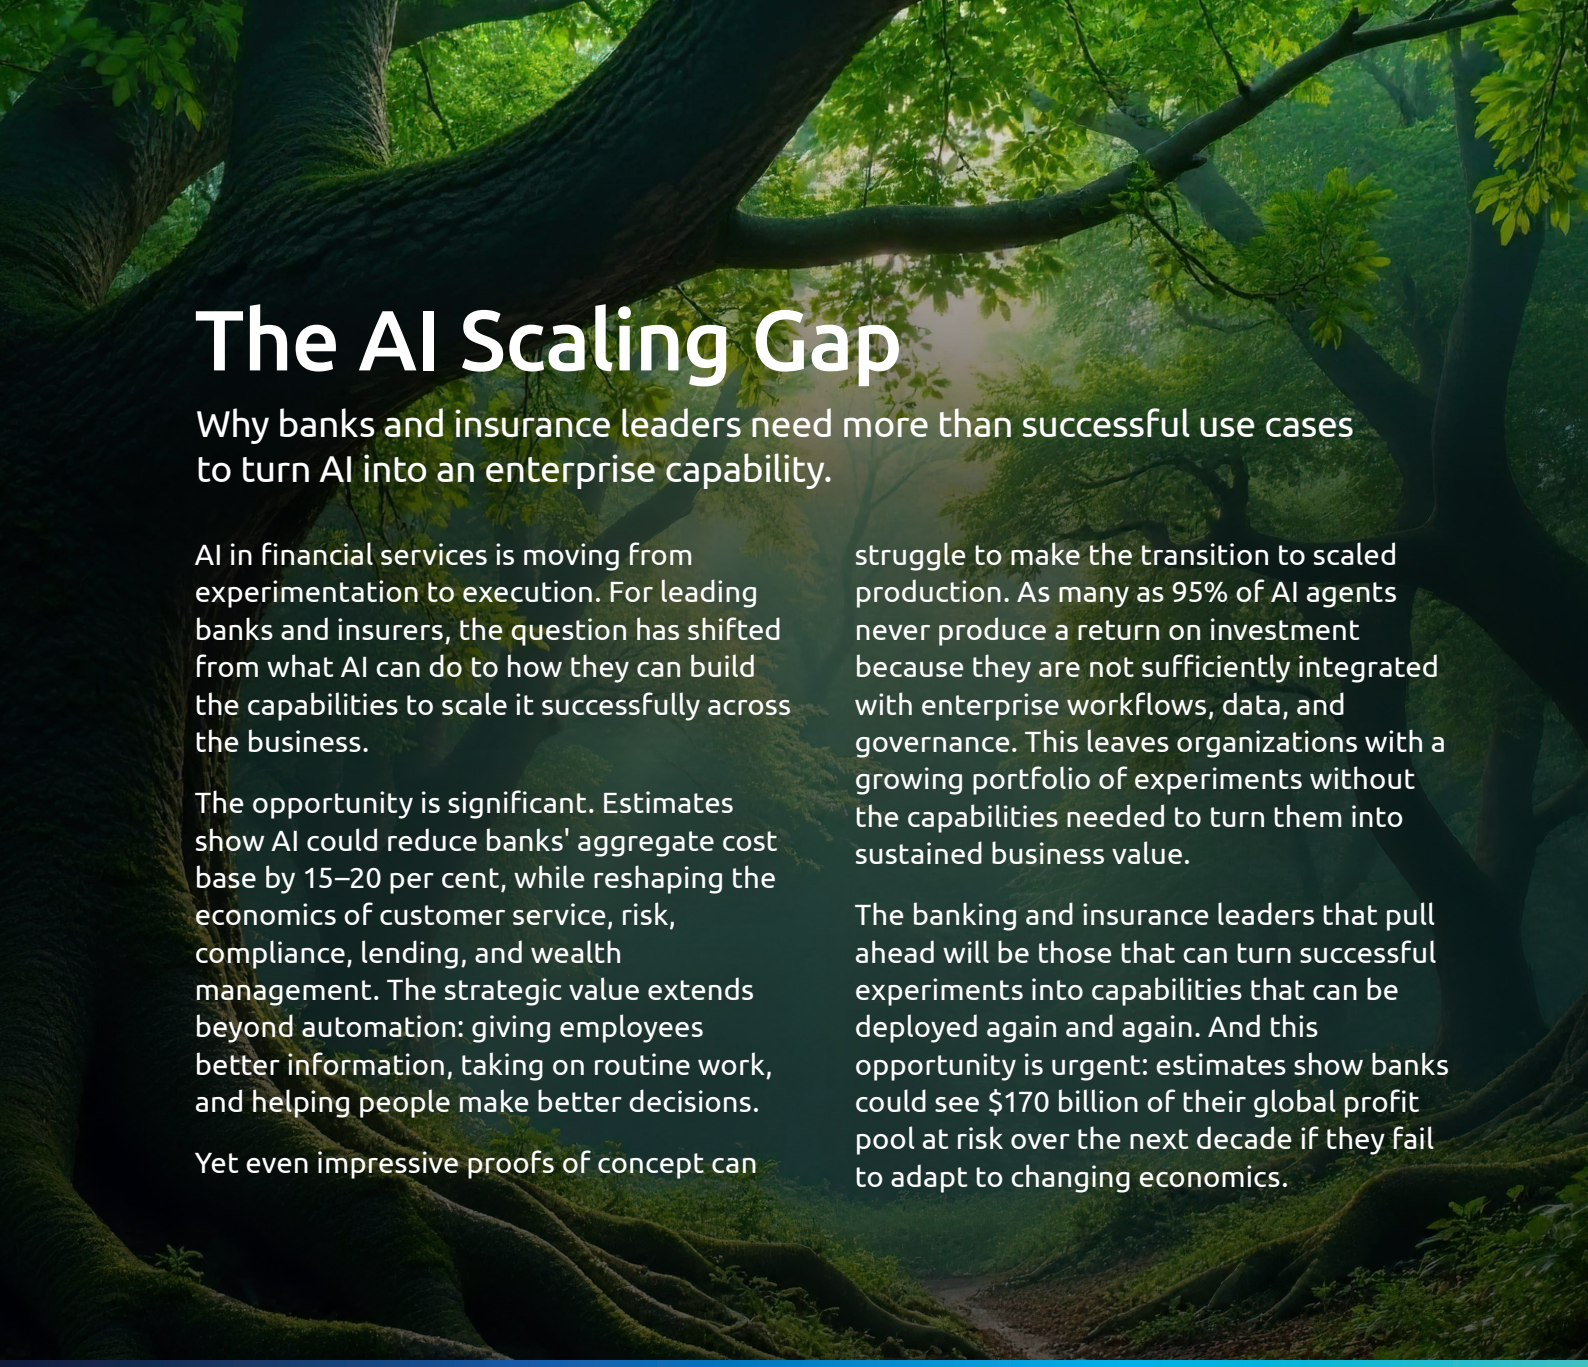
Bridging the AI knowledge gap

The critical role of business-ready data, reusable architecture, and operating models



Capgemini



The AI Scaling Gap

Why banks and insurance leaders need more than successful use cases to turn AI into an enterprise capability.

AI in financial services is moving from experimentation to execution. For leading banks and insurers, the question has shifted from what AI can do to how they can build the capabilities to scale it successfully across the business.

The opportunity is significant. Estimates show AI could reduce banks' aggregate cost base by 15–20 per cent, while reshaping the economics of customer service, risk, compliance, lending, and wealth management. The strategic value extends beyond automation: giving employees better information, taking on routine work, and helping people make better decisions.

Yet even impressive proofs of concept can struggle to make the transition to scaled production. As many as 95% of AI agents never produce a return on investment because they are not sufficiently integrated with enterprise workflows, data, and governance. This leaves organizations with a growing portfolio of experiments without the capabilities needed to turn them into sustained business value.

The banking and insurance leaders that pull ahead will be those that can turn successful experiments into capabilities that can be deployed again and again. And this opportunity is urgent: estimates show banks could see $170 billion of their global profit pool at risk over the next decade if they fail to adapt to changing economics.

15-20%

potential reduction in banks' aggregate cost base

95%

of AI agents never produce a return on investment

$170B

global profit pool at risk over the next decade

From use cases to enterprise capability

Scaling AI across a bank or insurer requires more than extending individual applications. A strategic approach starts with the development of an enterprise-wide foundation that makes each subsequent use case easier to build, govern, deploy, and operate.

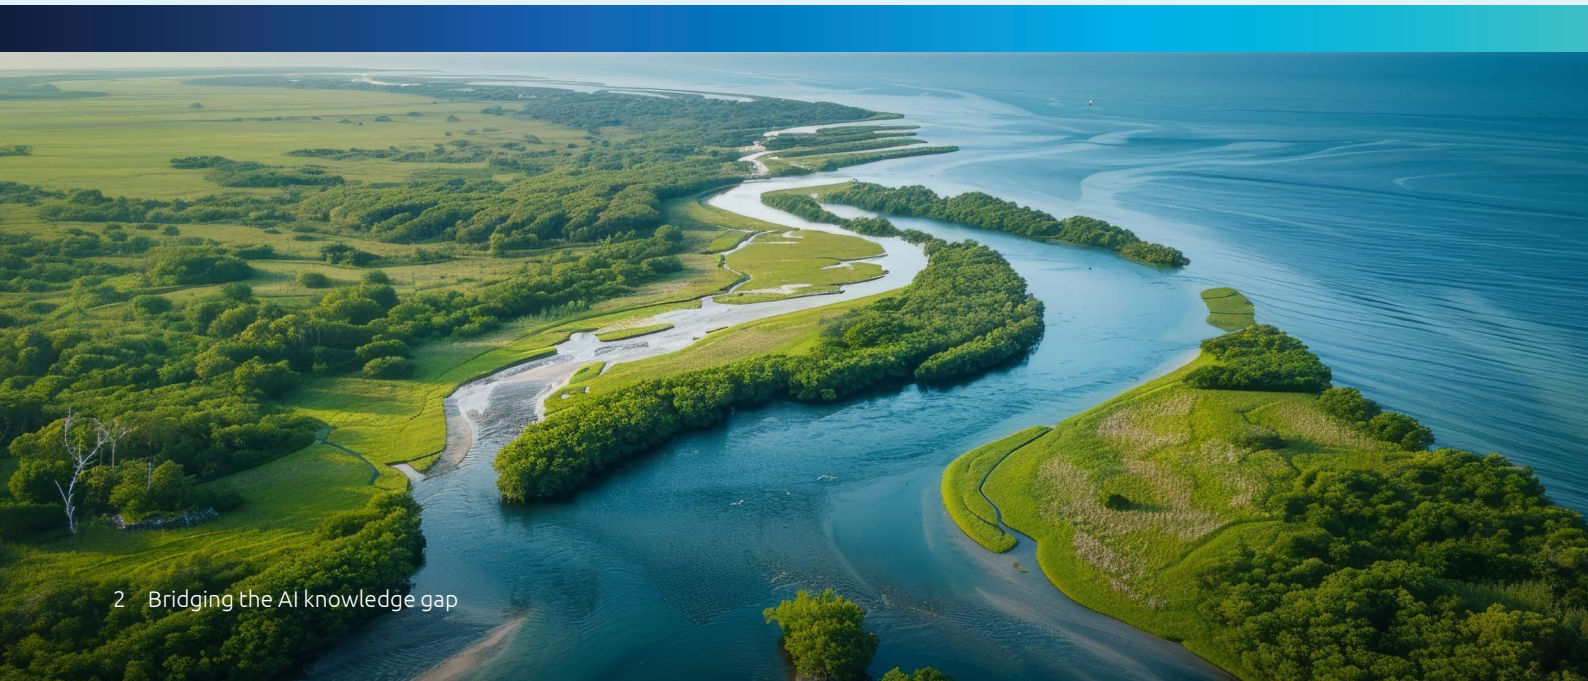
That foundation has several components. Business-Ready Data gives AI the reliable, contextual information it needs to support decisions and processes. A composable architecture provides shared capabilities for integration, security, testing, governance, deployment, and monitoring that teams can assemble and adapt for different use cases. And a repeatable operating model establishes how teams identify opportunities, define ownership, move solutions into production, and manage them over time.

This becomes increasingly important as AI moves from assisting employees to acting more independently across workflows. Agentic systems need clear boundaries around what they can access and do, how their performance is evaluated, when human intervention is required, and who is accountable for their actions. They also require processes for ongoing evaluation, lifecycle management and recovery when systems do not perform as expected.

The goal is to create a governed foundation that allows teams to build for their own needs while drawing on shared capabilities. This gives banks and insurance leaders room to move quickly without creating a collection of disconnected implementations, data, and governance models.

"The value of AI doesn't come from a single breakthrough use case. It comes from creating a foundation that allows the business to learn, adapt, and scale over time. When that foundation is in place, every new use case becomes easier to deliver and more valuable than the last because teams are building on what they've already established, not starting over each time."

Arindam Choudhury
EVP & Global Head FS Insights & Data



The foundation: Business-Ready Data

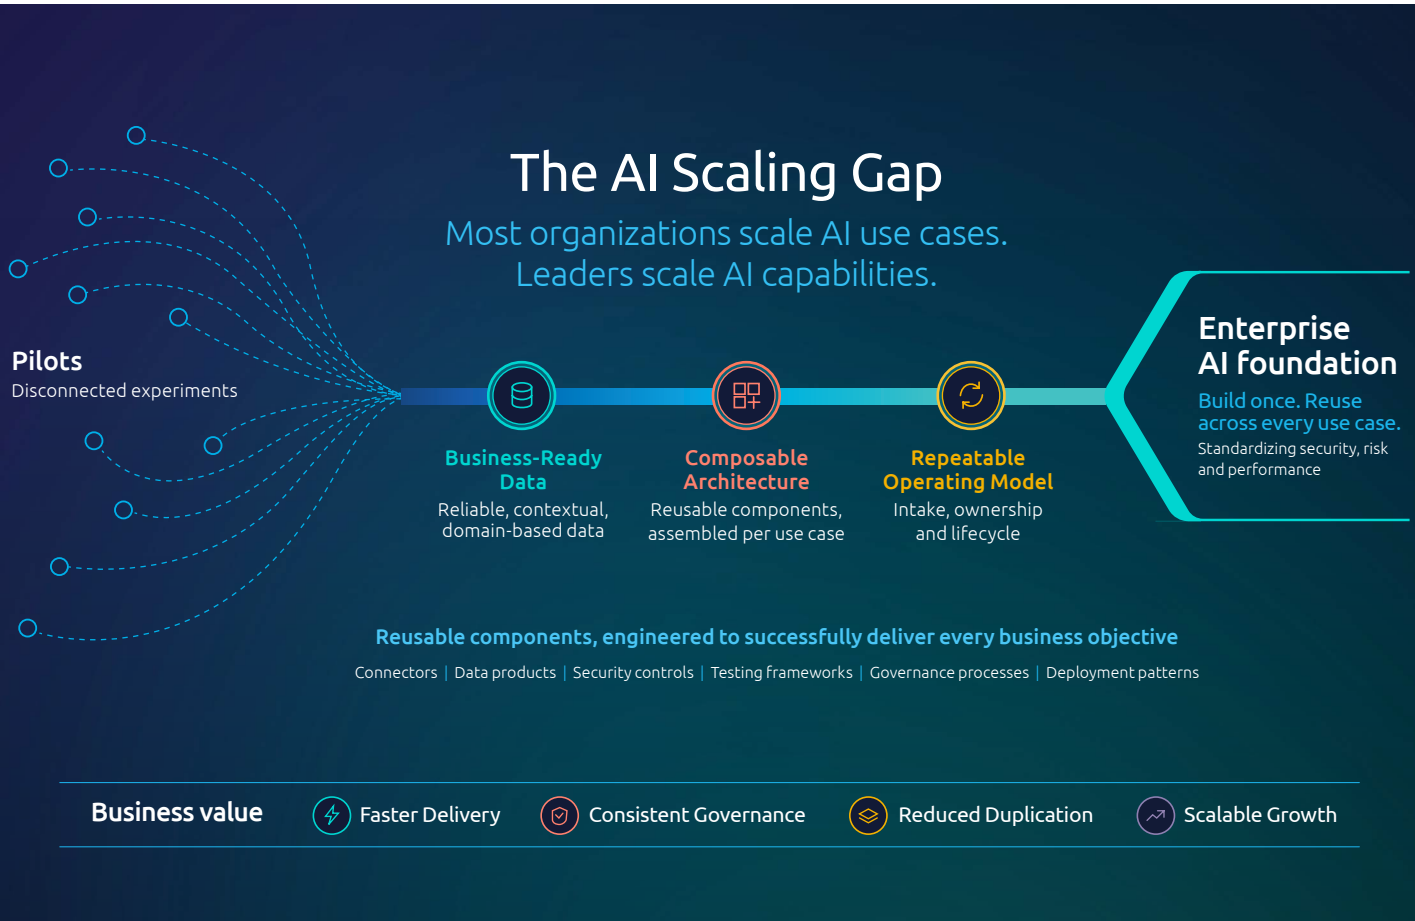
Banks and insurers have no shortage of data. The challenge is making it usable across the business. Customer, product, transaction, risk, and operational data often sit across different domains and systems, with different pipelines, definitions, controls, and levels of quality.

Risk teams model in isolation. Product teams instrument separately. Compliance runs its own parallel universe of reporting. The result is a gap between the data the business has access to and what it can consistently do with that knowledge.

This kind of fragmentation may be manageable when experienced employees are making decisions. People can fill gaps, recognize anomalies, and apply context that isn't captured in the underlying data. But AI doesn't have the same ability to compensate. It acts on the information it is given, which means poor or incomplete data can be amplified just as efficiently as good data.

This is where Business-Ready Data – data that is high quality, contextual, and optimized for how the business actually operates – becomes critical. Rather than treating data as an input to an individual AI use case, Business-Ready Data provides the core foundation all AI systems need to consistently support decisions, processes, and outcomes across the enterprise.

When data products are designed around financial-services domains and can be used across multiple use cases, banks and insurers don't have to rebuild the same data connections and business context every time a new AI initiative emerges. The same principle applies to integration patterns, security controls, testing, and monitoring. Establishing these shared capabilities once creates a foundation that teams can build on again and again, reducing duplicated efforts and creating greater consistency as AI expands across the enterprise.



Building for repeatability

Repeatability comes from designing AI capabilities to be assembled, adapted, and governed across multiple use cases – not rebuilt from scratch each time.

A composable architecture makes this possible. Common connectors, data products, security controls, testing frameworks, governance processes, and deployment patterns can be combined in different ways depending on the business problem. Teams can adapt established components to fit new use cases while maintaining consistent standards for security, risk, and performance.

In one credit card dispute management initiative, pre-built AI agents, data-product specifications, and an industry-specific data layer were combined to support the dispute lifecycle. Rather than developing each capability from scratch, the solution drew on components that could be adapted for additional processes.

This changes the economics of AI delivery. Instead of every project creating its own infrastructure, controls, and support model, each implementation can build on capabilities already in place. Over time, banks and insurance leaders can develop a library of proven patterns and components that reduces duplicated effort and makes the next use case faster to deliver.



A reusable technology foundation needs a reusable way of working

Technology alone doesn't create repeatability. Scaling AI across banks and insurers also requires a consistent way to decide which use cases to pursue, as well as how to define ownership, assemble the right capabilities, move solutions into production, and manage them over time.

A repeatable operating model establishes practical patterns for use-case intake and qualification, deployment and staffing, roles and responsibilities, adoption, monitoring, and lifecycle management. It also creates a clearer path from experimentation to production, with defined decision points for assessing value, risk, readiness, and ongoing performance.

Business teams can retain the autonomy to identify and develop use cases, while common processes and guardrails provide the consistency needed at enterprise scale. This is particularly important in financial services, where AI systems may influence decisions with regulatory, financial, or customer consequences.

As AI becomes more autonomous, accountability cannot be left to the technology. Banks and insurers need to establish who owns an AI-enabled decision, what level of human oversight is required, and what happens when a system falls outside expected parameters.

A federated model provides a practical balance: teams can move quickly within their domains while shared capabilities provide the standards, controls, and support underneath. The result is greater autonomy without sacrificing enterprise-level governance.



"The Foundry tells banks what to build. The four dimensions of Relevance, Operating Model, Agility and Retrospective determine whether the institution that emerges from that build is genuinely smarter or simply better equipped. Getting the infrastructure right and the management discipline wrong is still failure. The question this paper rightly poses - whether change is designed or imposed, applies equally to the layer above the architecture."



Syed Salman Rabbani
VP, Head of AI,
Leading Financial Services Client

1 Relevance

Ensuring any AI initiative is grounded in genuine business need, feasibility, and value justification — aligning use cases with strategy and calibrating ROI before infrastructure investment begins.

In the Foundry context: Determines which of the five banking domains to prioritise first. Without this discipline, banks pick the most exciting use case rather than the most compounding one — and burn budget before the flywheel starts.

2 Operating model

The people, processes, technology, and governance structures needed to build and sustain AI solutions — with governance embedded as a design constraint from POC, not introduced as a sign-off condition at deployment.

In the Foundry context: Provides the organisational layer that runs the Data Fabric, governs the Model Factory under SR 11-7, and ensures the Insight Delivery Layer reaches the right person — not just the right system.

3 Agility

Iterative, learning-oriented execution — starting small and strategically, building reusable foundations before domain features, and failing fast on experiments before committing at scale.

In the Foundry context: Activates the composability thesis. A foundational MVPO pass builds the shared data contracts and model patterns that make every subsequent use case cheaper. Skip this and the Foundry's cost curve never bends.

4 Retrospective

A structured pause after each deployment cycle — reviewing investment against return, capturing what worked and what did not, scanning emerging models and regulatory shifts, and feeding those findings into the next cycle.

In the Foundry context: Operationalises the north star. Without a formal retrospective loop, the Foundry compounds infrastructure. With it, the institution compounds judgment — which is the actual prize.

Putting the foundation to work

A recent contact-centre transformation illustrates what this can look like in practice. Rather than treating AI as a standalone application, the approach brought together capabilities including agent assist, automated quality assurance, conversational IVR, digital agents, predictive engagement, and AI workforce management within a broader operating model.

Rather than deploying these as isolated AI features, the approach connected capabilities across different parts of the customer-service journey. It also brought together technology and workflow changes rather than treating AI as a technology implementation alone. Shared data, governance, security, and monitoring provided the infrastructure underneath, while the operating model supported adoption, performance management, and continuous improvement.

For banks and insurers, this kind of integration is particularly important. AI-enabled processes need to operate within existing regulatory requirements, customer-protection standards, risk controls, and accountability structures. The technology has to work alongside the people and processes responsible for decisions, oversight, and remediation – particularly where AI influences customer interactions, financial decisions, or other regulated activities.

And the benefits are tangible. The transformation is targeting more than 40 per cent cost reduction, more than 60 per cent self-service through digital agents, and 10–20 per cent revenue growth.



This is where Capgemini's combination of financial-services expertise, technology capabilities, and operating-model experience comes into play. Our work spans data and AI engineering, workflow transformation, operating-model design, and employee adoption, with experience across areas including KYC, AML, risk and compliance, Customer 360, and wealth management. This breadth of experience allows us to address AI in the context of existing business processes, regulatory requirements, and the operating environments in which it needs to work.

We leverage AI Foundry as the execution framework to bring those capabilities together, helping banks and insurers move from individual AI initiatives toward a more repeatable approach to developing, deploying, and operating AI. This can be combined with Business-Ready Data and other Capgemini accelerators and industry assets to build on established patterns rather than starting from scratch with each new initiative.

Building for what comes next

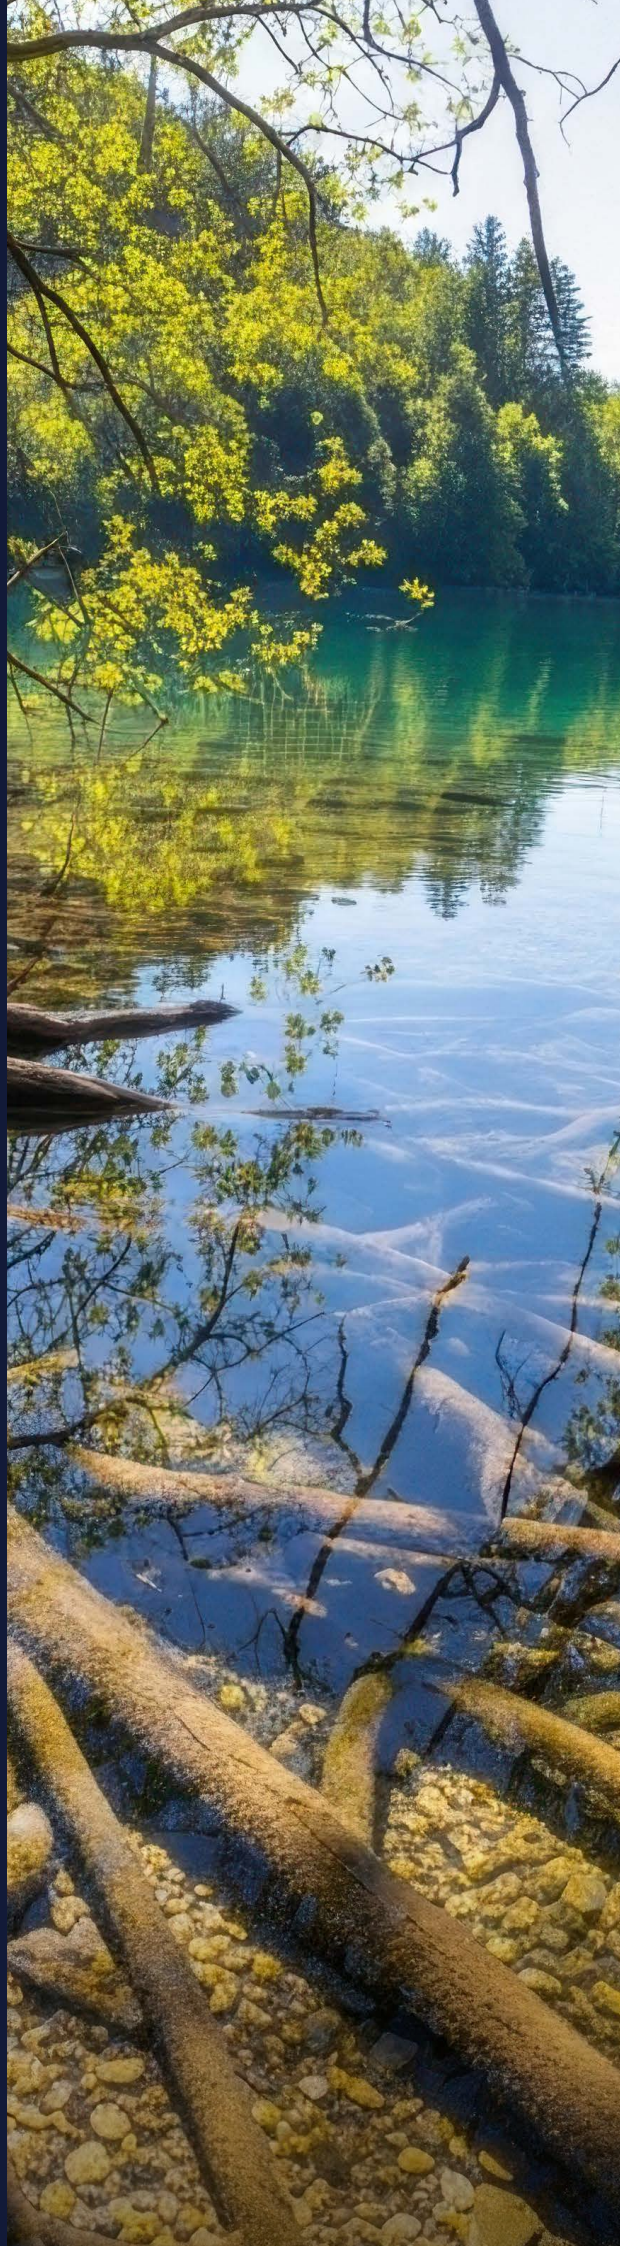
The opportunity now is to make today's AI investments work harder across both existing and emerging use cases, and to create capabilities that can support what comes next.

Banks and insurance leaders can start by assessing whether their data, architecture, and operating models are ready to support AI beyond individual pilots. From there, the priority is to identify which capabilities should be built once, where teams need flexibility to build for themselves, and where enterprise-level governance and controls need to remain consistent.

Capgemini can help put that foundation in place so that banks and insurance leaders successfully build the capabilities and operating models needed to move from isolated AI initiatives to full enterprise AI capabilities.

> Ready to assess whether your AI foundation is built to scale? Connect with Capgemini to identify the capabilities, operating model and reusable patterns needed to accelerate today's AI initiatives while enabling your teams to create and deploy new use cases with confidence in the future.

For more information
contact fsid@capgemini.com



About Capgemini

Capgemini is the business transformation partner for enterprises in the age of AI. We help organizations imagine and build an intelligent, sustainable future, combining AI, technology and human ingenuity to transform how they operate, innovate and grow. With unique end-to-end capabilities spanning strategy, technology, engineering and intelligent operations, we bring together deep industry expertise and market-leading capabilities in AI, cloud and data to turn ambition into measurable business outcomes at scale. Supported by a robust ecosystem of partners and nearly 60 years of expertise, Capgemini is a responsible and diverse global organization of over 410,000 team members in more than 50 countries. The Group reported 2025 revenues of €22.5 billion.

Make it real | www.capgemini.com

© Copyright 2026 Capgemini. All rights reserved.

Capgemini