Capgemini

From pilots to *real impact:*

Business outcomes through scaling AI

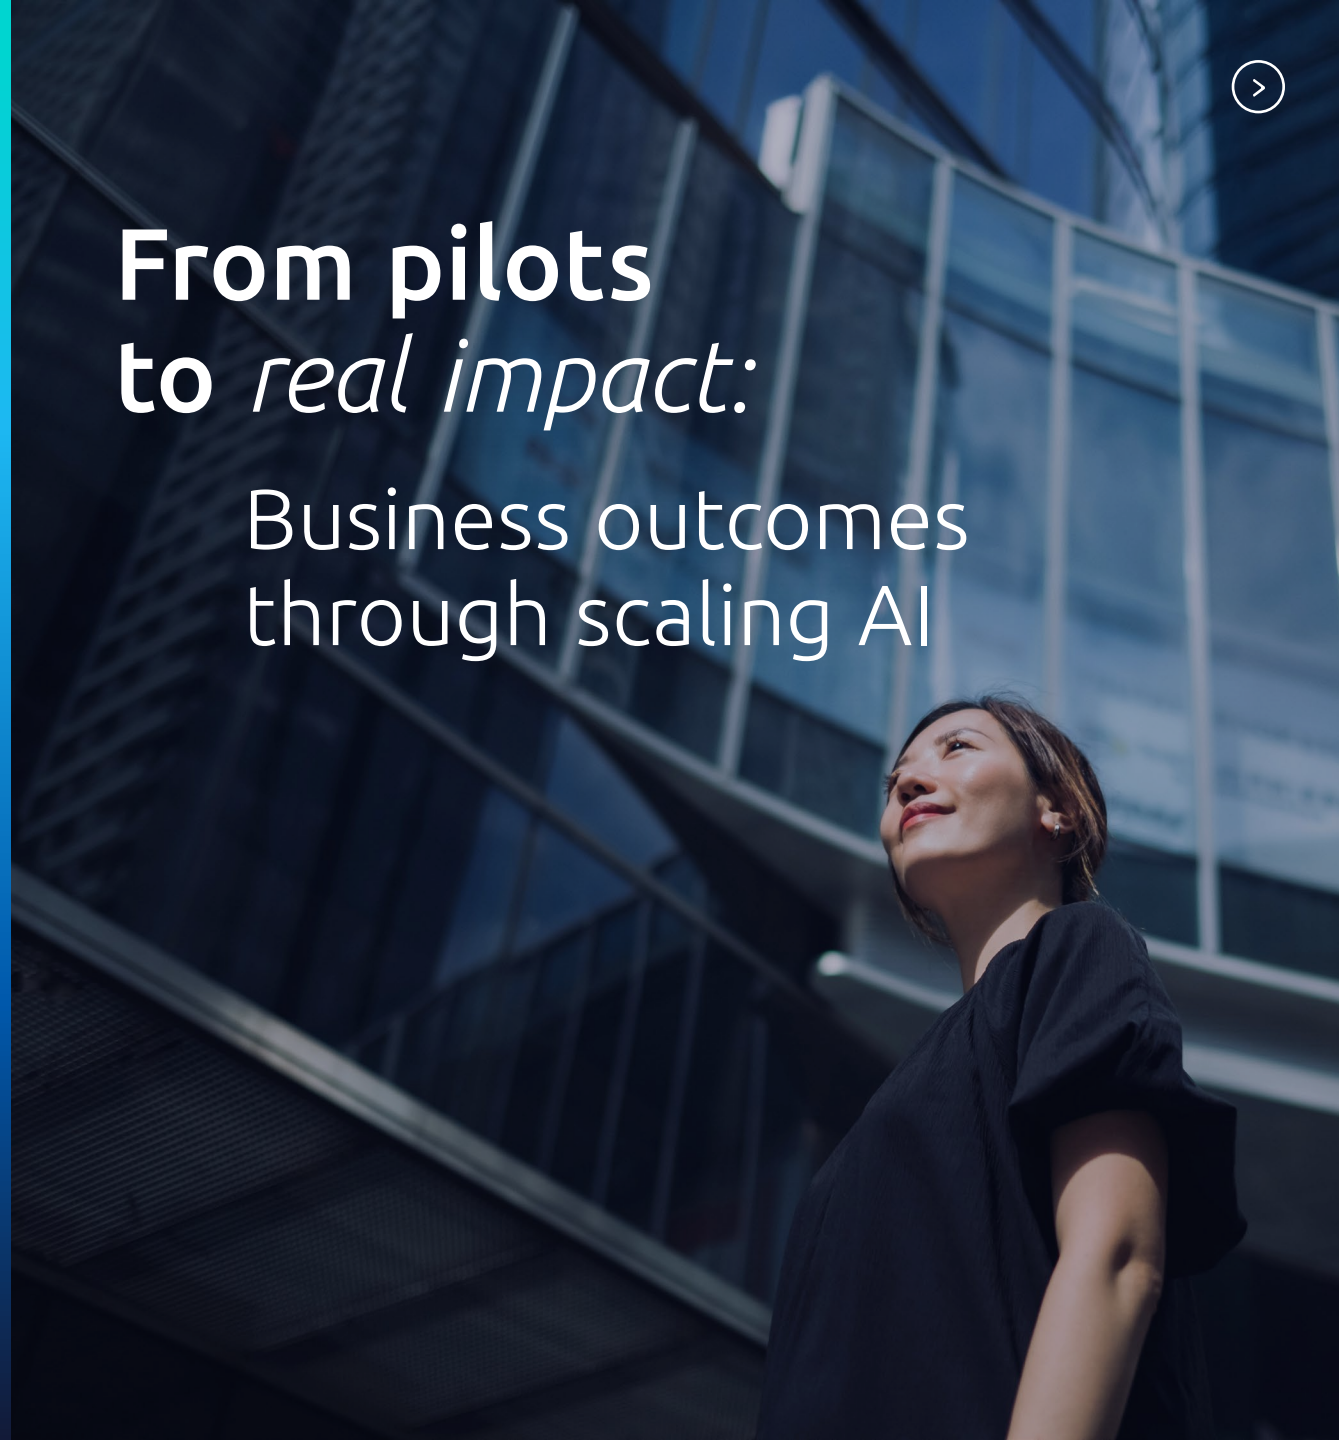

Table of Contents

Enterprises have proved they can build AI pilots – but scaling them into enterprise-grade, agentic systems remains the real challenge. Here are three steps to do exactly that.

Most enterprises are running multiple AI pilots. **Very few can scale one.**

Even with impressive proofs of concept, **95%* of AI agents never reach production** because they're not wired into enterprise workflows, data, or governance. Instead of transformation, organizations end up with a graveyard of experiments.

Scaling agentic AI isn't a bigger pilot – it's an entirely different challenge. It means embedding AI into end-to-end processes, designing the right balance between autonomy and human oversight, and operating agents securely across complex, multi-cloud environments. Without the right foundations, even the most promising initiatives stall.

But the 5%** who are succeeding aren't just building better pilots – they're engineering AI into the enterprise itself. Understanding what they're doing differently is the key to making agentic AI scalable for everyone.

Where scaling agentic AI fails quietly

Agentic AI projects most often fail when they encounter the complexity of existing enterprise systems.

The contributors to failure include:

- ROI remains theoretical because agents aren't reused
- Complexity explodes immediately after the demo phase
- Integration is the #1 bottleneck
- Operating agents is the hidden cost center

Once organizations move beyond experimentation, they must make deliberate choices about how agents will connect to enterprise data, logic, and tools, and how they will consistently reflect their unique business context.

Enterprise leaders must decide where AI can act autonomously, where humans should remain in the loop, and how risk is managed across sensitive processes. They also need to ensure agents can operate within the existing IT ecosystem, with clear ways to measure performance, enforce governance, and maintain security.

But many AI initiatives move forward without AI-ready data or a clear operating model. As a result, agents struggle to access usable data, decisions lack traceability, and leaders lose visibility into how and why AI systems are acting.

Together, these challenges explain why so many agentic AI efforts remain stuck at the pilot stage – and why scaling requires a more deliberate and holistic approach.

Three keys to scaling agentic AI

Successful organizations engineer for enterprise complexity from the start.

For example, we're helping enterprises navigate three key steps using an agentic AI delivery toolkit that supports the full AI lifecycle – from traditional AI and machine learning applications to generative and agentic AI. Successful organizations are starting to move from pilot phase to making it real, by providing a structured way to design, build, deploy, and operate AI across complex enterprise environments.

1. Engineering agentic AI into real enterprise workflows

Enterprises that see lasting impact start by decomposing business processes task by task. Once they understand exactly which tasks are most suitable for AI agents, they can design multi-agent systems that align with existing business logic.

It's often critical to have a process decomposer that guides enterprise teams in the identification and decomposition of tasks. For example, Capgemini now has a library of more than 400 prebuilt, domain- and industry-specific agents that can be deployed as is, customized, or combined to support complex processes. Using these orchestration capabilities, organizations can also model workflows, stitch agents together, and coordinate how work moves across systems and teams. These components are modular, so enterprises can pick and choose which tools are most useful.

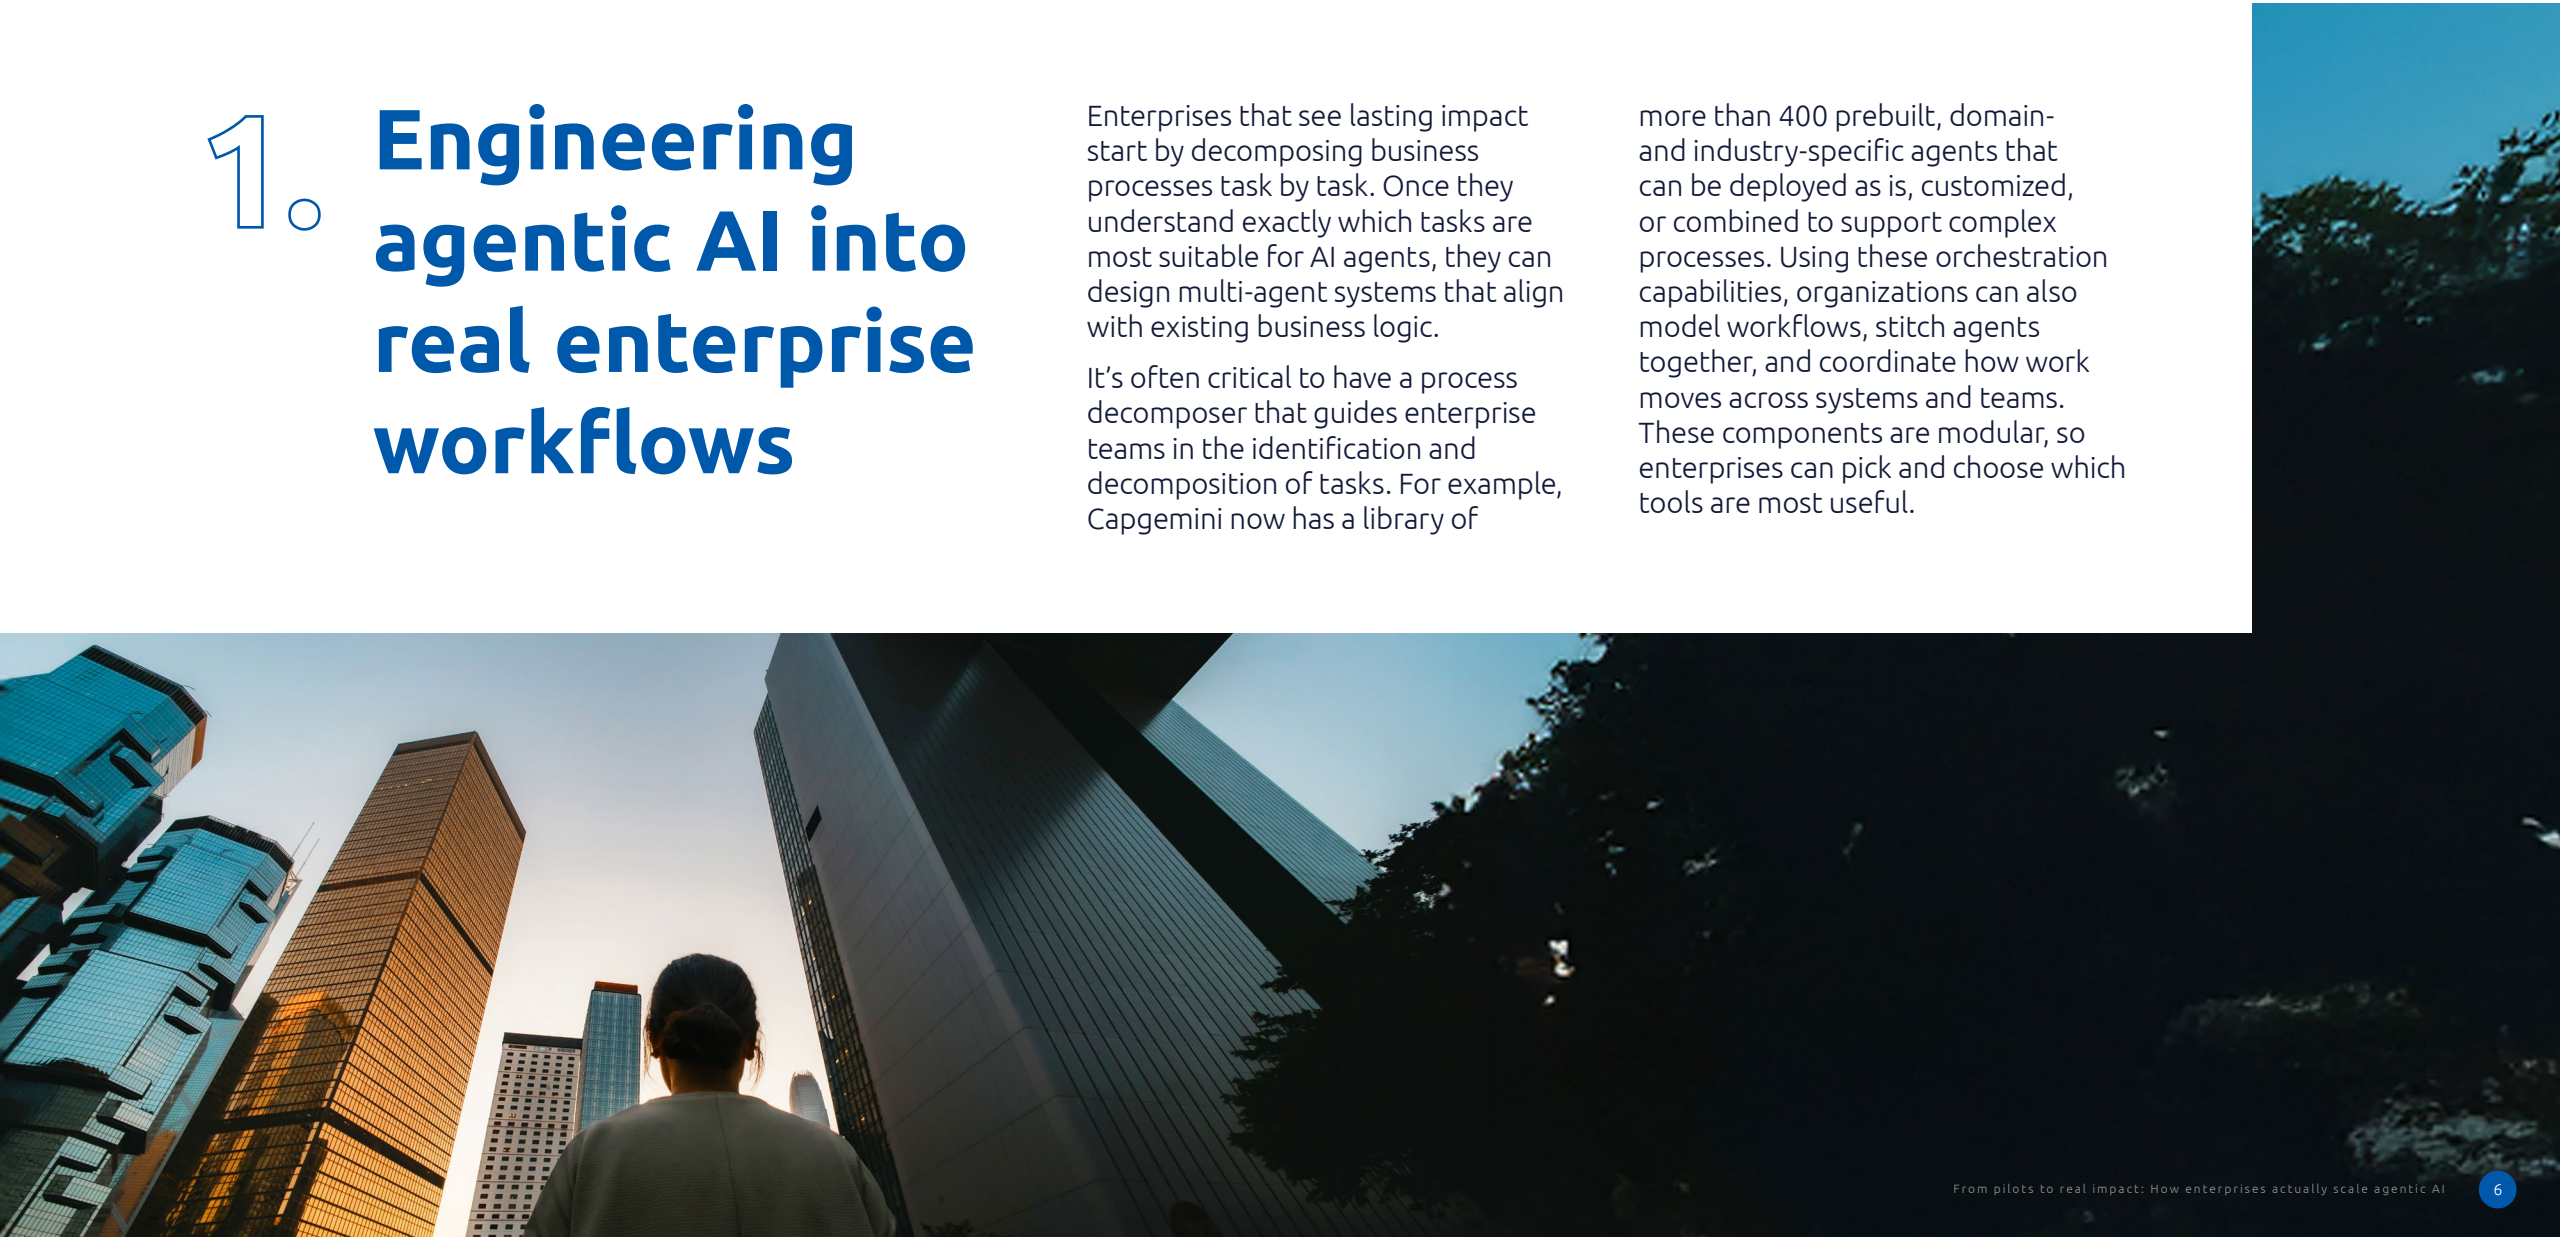

2. Democratizing development while maintaining solid governance

To scale agentic AI, enterprises must dramatically shorten the path from idea to value while expanding access and participation beyond technical teams.

Organizations should be thinking about guided development tools that allow users to design, test, and deploy agents through structured workflows and reusable templates built on open-source frameworks.

These tools leverage native capabilities from hyperscalers where they are strongest, so enterprises can run and manage agentic systems across clouds without being locked into a single provider.

This approach enables faster experimentation and broader adoption. And with clear guardrails in place, organizations can safely empower more teams to build and deploy agents.

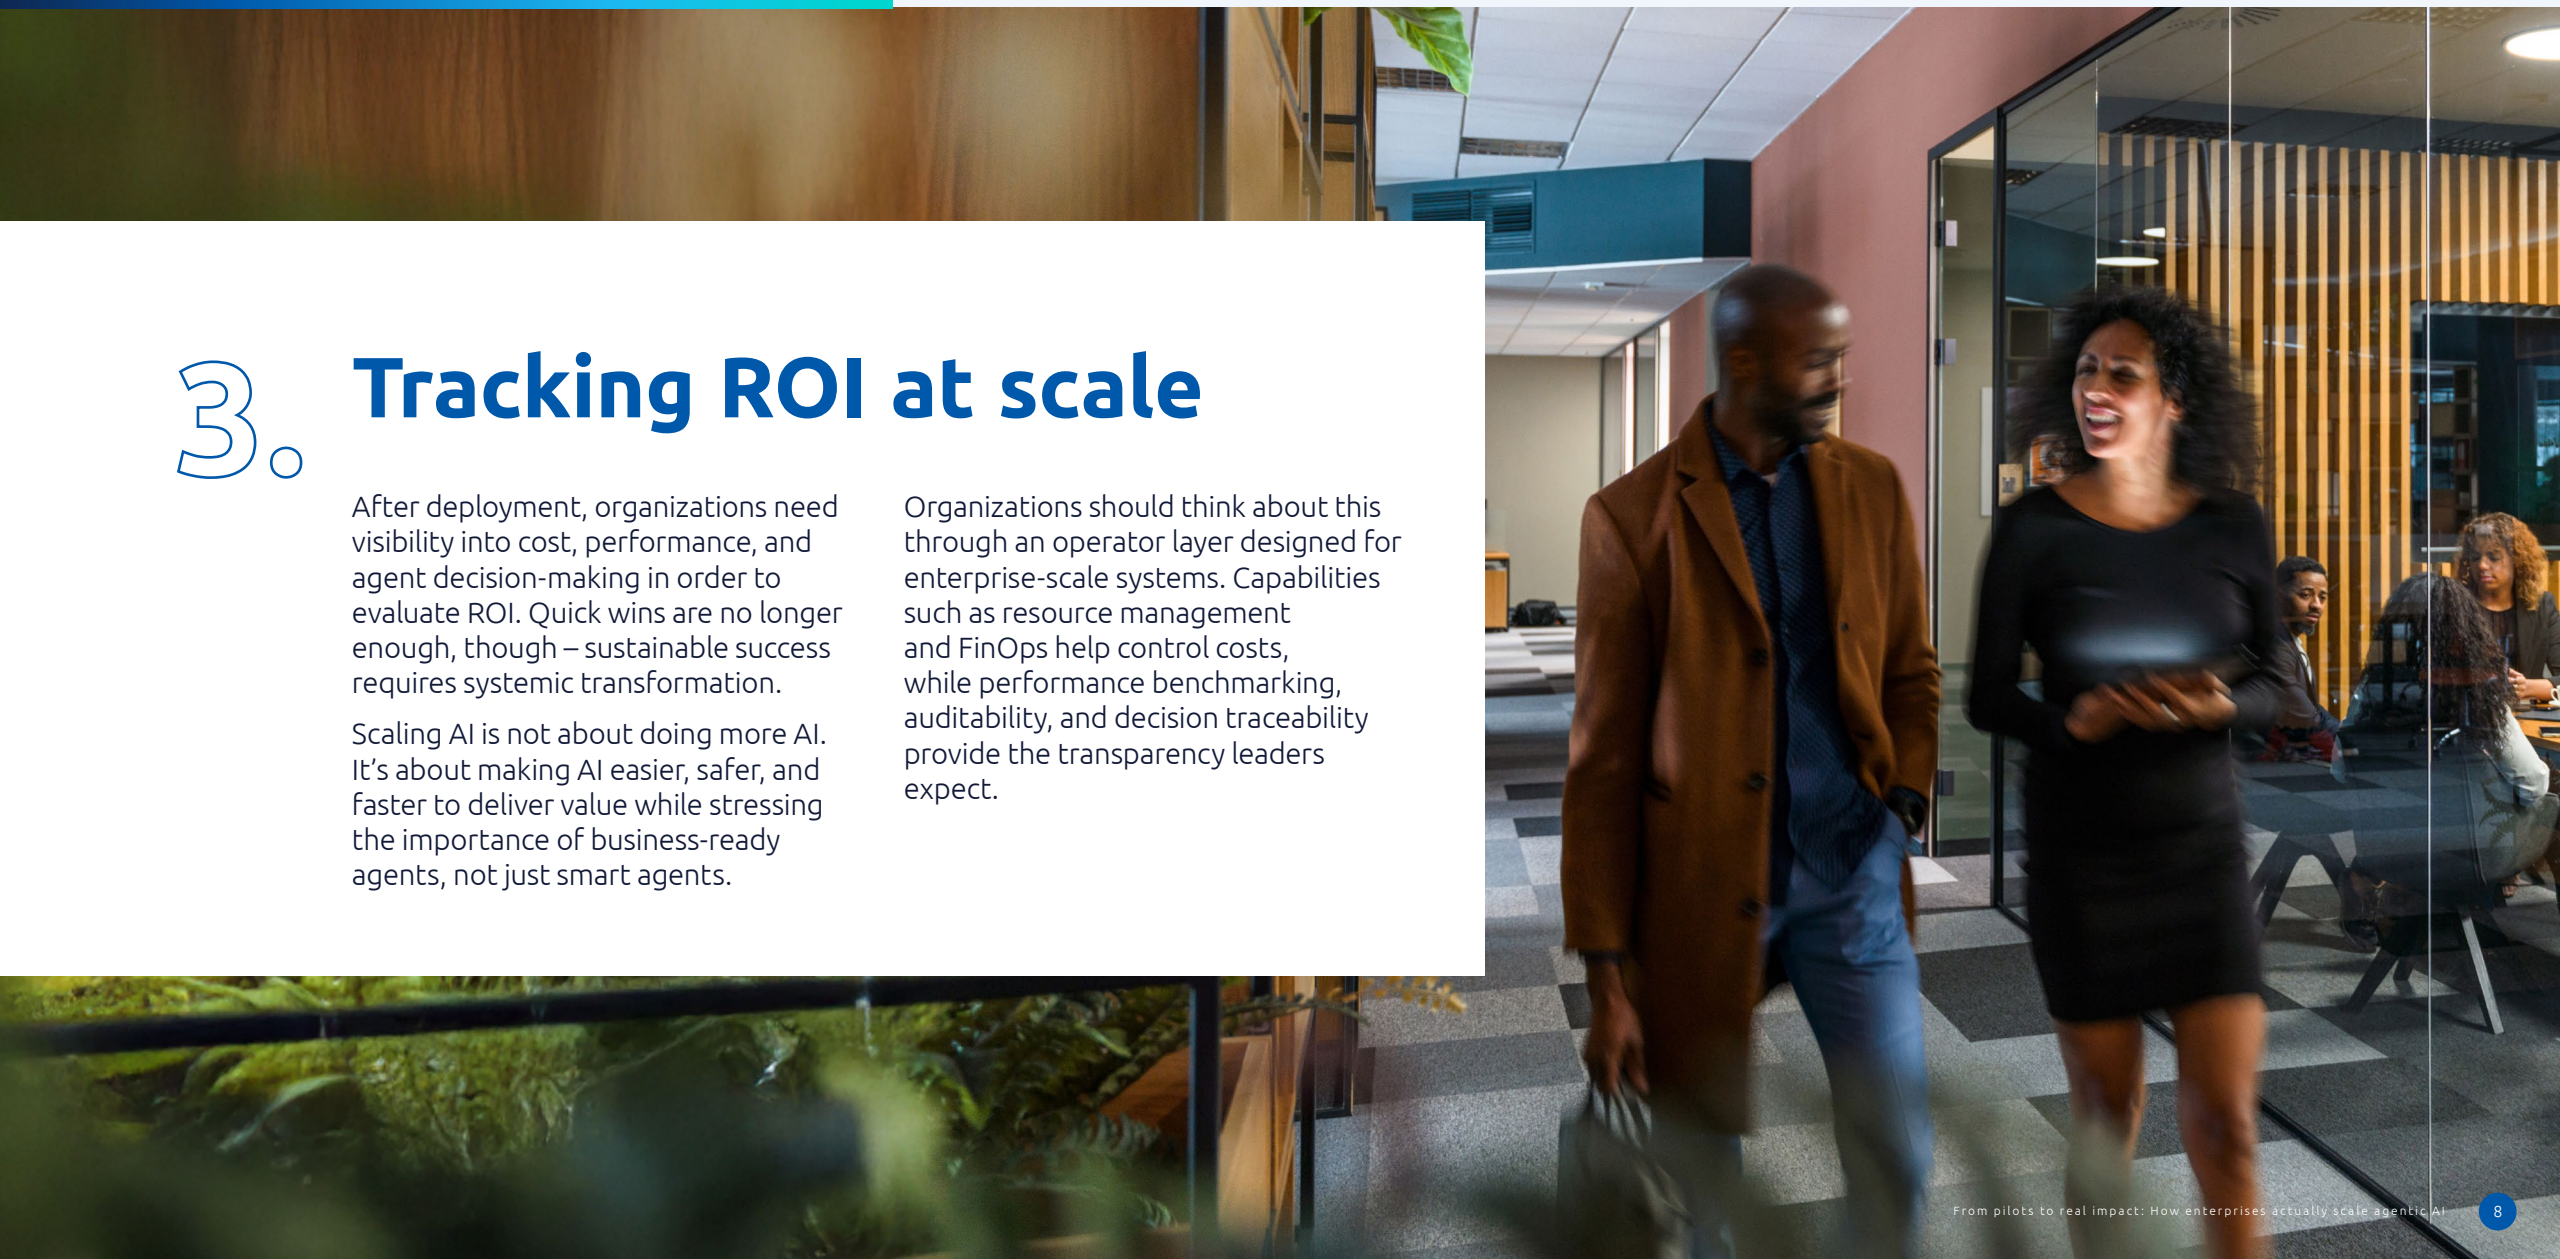

3. Tracking ROI at scale

After deployment, organizations need visibility into cost, performance, and agent decision-making in order to evaluate ROI. Quick wins are no longer enough, though – sustainable success requires systemic transformation.

Scaling AI is not about doing more AI. It's about making AI easier, safer, and faster to deliver value while stressing the importance of business-ready agents, not just smart agents.

Organizations should think about this through an operator layer designed for enterprise-scale systems. Capabilities such as resource management and FinOps help control costs, while performance benchmarking, auditability, and decision traceability provide the transparency leaders expect.

Where real-world enterprises see impact

Enterprises are starting to move from pilots to operating agentic systems at scale. **Here are a few use cases that showcase how it works in global businesses.**

Across all industries enterprises who are doing it right are deploying secure AI foundations and human–agent workflows.

The impacts can be quite powerful: **

- 40–60K service calls deflected
- 60% self-service adoption
- 10–20% revenue lift
- 80% lower application cost
- 50% faster product deployment
- 20% increase in efficiency of supply chain
- 20% reduction in manual loan-processing
- 82K hours saved annually with AI-powered content creation

Building a foundation for enterprise AI

A **biotech company** leveraged agentic AI to build a unified AI foundation spanning R&D and supply chain operations. Rather than deploying disconnected use cases, the focus was on establishing a standardized AI platform that could scale safely across the enterprise.

Within 12** months, the organization had the foundations required to accelerate time to value, manage AI risk with confidence, and prepare the business for more advanced use cases.

Scaling AI across core value chains

A **consumer products company** took the next step: embedding AI across multiple end-to-end value streams including forecast-to-cash, idea-to-profit, source-to-partner, and convert-to-serve.

Agentic AI solutions were designed, built, and integrated across these workflows. The result was coordinated intelligence across functions – connecting data, decisions, and execution in real operational contexts. By engineering AI into how the business actually runs, the company is targeting a 20%** reduction in supply chain costs, while unlocking broader business optimization and growth.

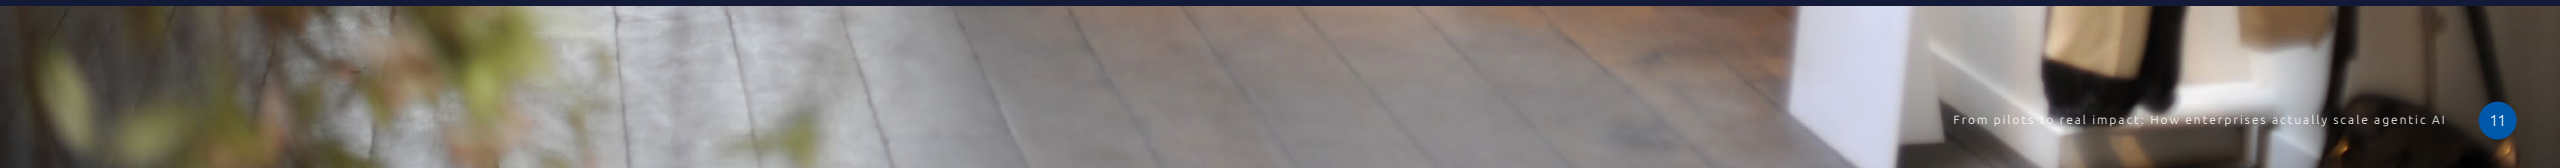

Agentic AI for customer experience transformation

In a regulated environment where trust and control are critical, an **APAC bank** repositioned its contact center into an AI-first customer experience hub.

Agentic AI enabled it to support personalized, omnichannel interactions and self-service capabilities, while maintaining clear governance over when the agents act autonomously and when humans remain in control. The initiative has aggressive targets: more than 40%** cost reduction, more than 60%** self-service adoption, and 10–20%** revenue growth enablement.

Modular Gen AI products at scale

An **oil and gas company** developed modular Gen AI products across finance, HR, and supply chain – focusing on reusability and adaptability.

By standardizing development and managing AI resources through a unified framework, the organization positioned itself to accelerate delivery while reducing cost. Expected benefits include a 50%** reduction in development time and a 57%** reduction in development costs. And the modular approach ensures AI solutions can evolve with the business.

The pattern across industries is consistent and clear: organizations that succeed with AI invest in foundations, engineer AI into real workflows, build trust through governance, and design for scale from day one.

Author

Ashvin Parmar

Global Head of Assets & New Revenue Models, AI & A Global Practice, Insights & Data

Want to turn pilots into real performance?

Connect with us and let's make it happen: *raise@capgemini.com*

** Results represent targets that Capgemini strives to achieve with real client use cases. Results largely depend on client AI maturity, budget and willingness to adopt AI Agent technology in their organizations.

* *MIT Project NANDA, State of AI in Business 2025.*

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real | www.capgemini.com