Redesign for enterprise AI

Focus on critical value streams to drive growth

IBM

Contents

Foreword 1

Introduction
What's holding AI back? 3

Part one
Breaking through bottlenecks 6

Part two
The real size of the prize 9

Part three
Renovate the enterprise for AI 13

Action guide
Get more value from your AI investments 19

Foreword

AI doesn’t eliminate bottlenecks. It exposes them.

Chris Wright

Chief Technology Officer and Senior Vice President, Global Engineering

Red Hat

When we first started talking about generative AI in the enterprise, much of the focus was on quick productivity wins. How we could make individual tasks faster? But over the last year, the conversation has completely shifted.

Optimizing a single task only gets you so far. If you automate one step in a process without looking at the bigger picture, you haven’t actually solved the problem—you’ve just pushed the bottleneck to the next step downstream. The real value of AI is completely rethinking end-to-end workflows, from how we handle handoffs to how we manage dependencies and make decisions.

Our current workflows were built around historical constraints and legacy organizational boundaries. Now we have technology that fundamentally challenges and changes those boundaries. And it’s moving incredibly fast. Models improve, tools evolve, and approaches that felt like science fiction a year ago are now delivering repeatable, measurable results.

Waiting for the technology to stabilize before you jump in is a losing strategy. If you stay on the sidelines, you miss out on building the organizational muscle memory required to use them. The iteration is where the real learning comes from. The companies succeeding today are the ones learning how to evaluate, deploy, and iterate on AI in real time as the tech evolves. But you can build that capability only when your teams are doing the work.

This study comes exactly when engineering and business leaders are grappling with this shift. Most organizations have what I call an “infinite backlog” of ideas, features, and customer requests that they never get to because they simply run out of time or capacity. AI gives us a chance to finally dig into that backlog. It’s not just about doing the same work faster; it’s about taking on work we could not attempt before.

The following pages offer a practical look at this journey. They show why real value happens when you fix the entire workflow, why your underlying infrastructure matters, and why your culture has to evolve alongside your code. The technology will keep changing next month and next year, but the responsibility for redesigning how our teams actually get the work done is entirely up to us.

Key takeaways

AI investments are struggling—but removing friction and focusing efforts can improve returns.

Most AI initiatives are still failing.

At the end of 2025, only 37% of AI initiatives had delivered the business outcomes expected by senior leadership.

Complexity is a drag.

In 2025, 21% of AI ROI was lost due to friction within the IT and business environment.

Process reinvention is the missing link.

Executives cite the failure to redesign business processes as the top obstacle to high-ROI AI implementation.

Improving AI outcomes requires investing in IT upgrades.

Yet 58% of organizations are under pressure to cut non-AI enterprise tech spending, pitting critical investments against each other.

Enterprise AI is still managed piecemeal.

Only 17% of executives say the process for approving and funding AI initiatives is consistent across the board, and just 25% say business value and ROI are measured consistently.

Introduction

What's holding AI back?

AI investment is accelerating. But where are the results?

Despite increased spending, most AI initiatives are still failing. Only 37% had delivered the business value expected by senior leaders by the end of 2025, according to a new survey conducted by the IBM Institute for Business Value (IBV) in partnership with Oxford Economics. According to our research, the AI portfolio of an average USD17 billion enterprise generated USD115 million in net value—an ROI of 51% in 2025 (see "Research methodology" on page 24).

Looking to 2028, executives have set their sights much higher. On average, they're aiming to deliver USD462 million in net value and boost ROI to 105%. However, the 1,000 technology CxOs, AI transformation officers, and line of business sponsors we surveyed expect to deliver on these targets while taking just 29% of AI initiatives into production at scale (see Figure 1). That's where the math goes sideways.

The fact is, agentic AI can't deliver better outcomes unless initiatives are deployed at scale. But fragmented ownership, disconnected systems and data, and unclear decision rights are limiting AI's reach. The organizations that realize the greatest returns from agentic AI will be those that redesign end-to-end value streams to reduce friction and allow agents to carry work from one stage of a process to the next.

Not all friction is technical. Some of it is human. Even when AI systems are available, employees may not understand how they create value, where they fit into existing workflows, or why specific tools were adopted. Adoption friction can prevent organizations from capturing returns even when the technology itself performs as intended.

AI agents also break traditional span-of-control models. Leaders will need to build hybrid team structures that clearly outline how human and digital labor should work together to spur effective adoption at scale.

> "Treating IT spend as investment rather than a cost can create uncomfortable conversations. But it's the only way to demand and achieve real returns."
>
> **Gavin Wood**
> CIO, APA Group

FIGURE 1

AI outcomes are limited by friction and a lack of focus

2025 outcomes, obstacles, and 2028 projections reflect executive value assessments of their AI portfolios. 2028 potential outcomes are based on IBM IBV analysis.

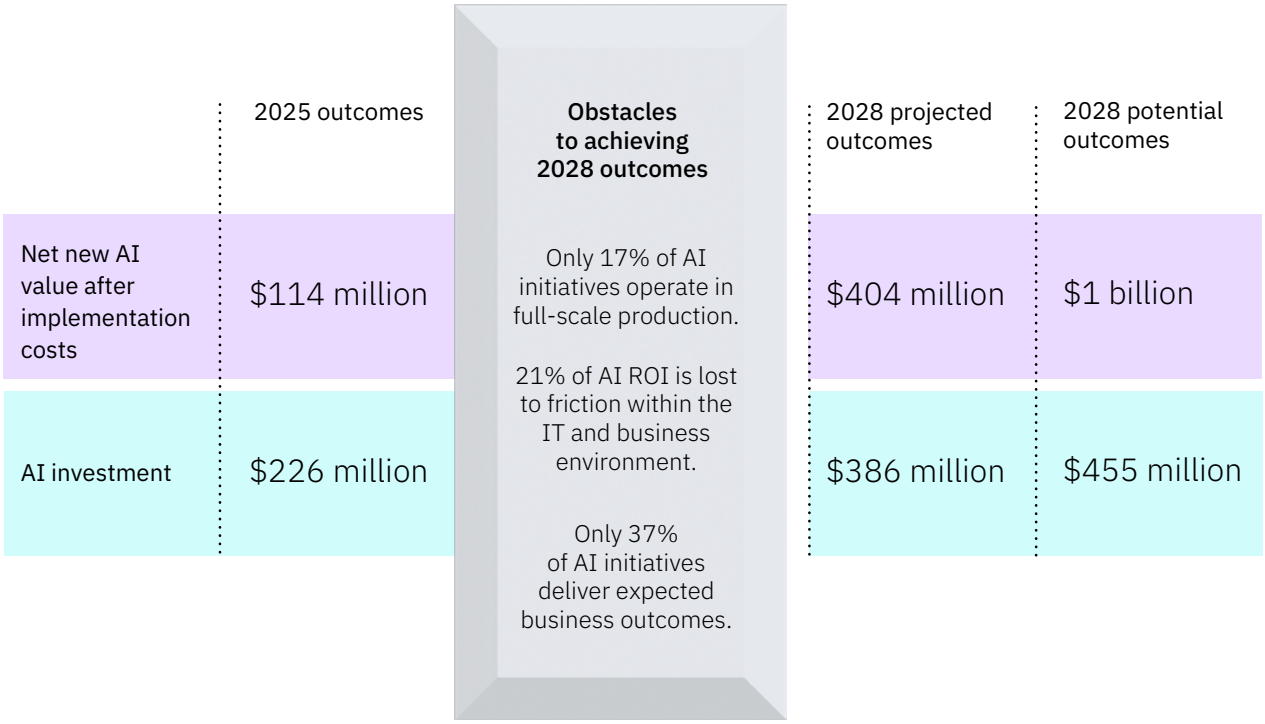

> "With AI-driven automation, the link between spend and return isn't deterministic—it's inferred. That makes the business case harder to explain, and it requires more experimentation and proof to justify investment."

Dalton Gouws
Group IT Director and Board Member, VWG UK Ltd.

"If we apply AI inside each silo only to improve individual work efficiency, the benefit disappears at the enterprise level," says Yasuo Kawai, Vice President, Head of Digital Unit, Honda Japan. "To operate AI enterprise-wide, enterprise-wide reform is necessary."

To meet the AI performance projections they've set for 2028, executives must do more than fund the work of pushing last year's pilots to production. They must also reduce the amount of AI ROI lost to friction, sharply focus investments on the business outcomes they want to deliver, and commit to the sustained cycles of reinvestment needed to succeed with AI at enterprise scale.

The good news is, success relies less on chasing the latest AI breakthroughs and more on creating an operating environment that AI can evolve within. Building a system in which AI works effectively can help organizations deliver more value with AI faster—and accelerate past their performance targets.

We see the potential upside for organizations that scale a greater portion of AI initiatives, reduce friction in the IT and business environment, and deliver expected outcomes more consistently to be as much as USD1 billion by 2028, more than double what executives expect (see "Research methodology" on page 25). Organizations would also need to spend more than they have planned, but this investment could increase ROI from 105% to 220% (see Figure 1).

Getting there requires addressing the deeper issues that many organizations have been skirting for years. The business case for process automation and renovation at scale just hasn't been strong enough. Agentic AI changes all that. With the right redesign, It can unlock new sources of revenue and deeper pools of value that can be tapped over the long term. This paper outlines where executives should focus investments and efforts to deliver the most meaningful business results.

In part one, we explore how friction drags down AI performance. Part two explains where organizations have the opportunity to extract the greatest net value with AI. Part three focuses on the technology and operational renovations that organizations should prioritize first. We close with an action guide that outlines steps enterprise technology leaders can take to help improve AI performance in the next 30, 120, and 360 days.

Part one

Breaking through bottlenecks

The biggest barrier to AI impact isn't model quality or data volume—it's friction across systems that weren't designed to talk to each other. In 2025, 21% of AI ROI was lost to friction in the IT and business environment, according to the executives we surveyed.

That's because agents don't sit on top of IT infrastructure. They operate within it. They carry work and decisions between functions that were historically optimized in isolation. When their environment is fragmented, inconsistent, or weighed down by decades of technical debt, even the smartest AI hits a wall.

Executives say the top obstacles to high-scale, high-ROI AI are processes, data, and decision rights that haven't been optimized for AI (see Figure 2). A large-scale renovation is needed to build AI resilience into the IT estate. In addition to technical upgrades, organizations must create new risk, governance, and security protocols, and redesign workflows from start to finish.

More than three in four executives agree: applying AI to isolated parts of a workflow delivers limited value unless the entire process is redesigned. Most enterprises are starting with an IT landscape built from thousands of rational-but-isolated decisions—each one sensible in the moment but not orchestrated as a whole. In fact, recent IBM IBV research found that shadow IT—the software, cloud services, and tools purchased directly by business units without IT oversight—represents 24% of the average IT budget.[1]

"At scale, explainability, accountability, and auditability can't be retrofitted. They must be designed in from the start."

Dena Almansoori
Group Chief Technology and Innovation Officer, ADNOC

The result is an IT environment that spans multiple clouds, platforms, and jurisdictions, with thousands of integration points that no single individual fully understands. Only 10% of executives say their AI development and operations infrastructure is fully standardized across environments on a unified platform.

FIGURE 2

Top obstacles to high-scale, high-ROI AI

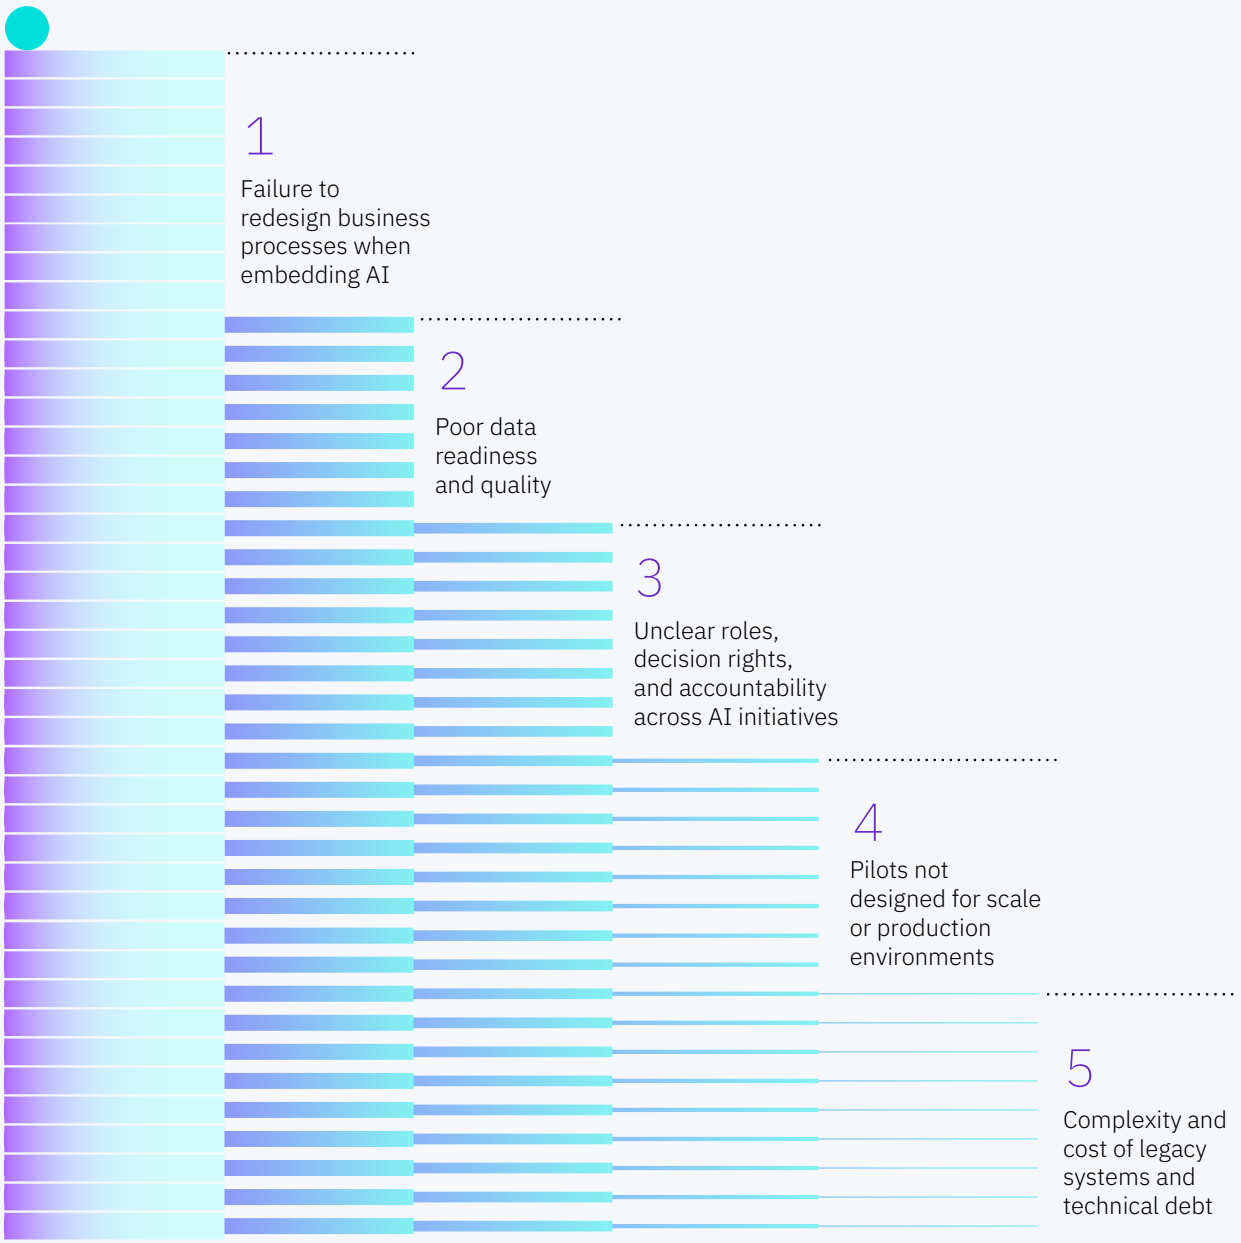

The first step toward solving the problem is acknowledging it. Recent IBM IBV research found that organizations that fully account for the cost of addressing technical debt in their AI businesses cases—modernizing brittle systems, integrating siloed data, and fixing long-ignored vulnerabilities—project ROI that's 29% higher than those that don't.[2]

The next step is measuring where value is delivered—and where it is lost. Our analysis of survey data found that organizations that measure the value of their processes tend to report higher AI returns. When organizations develop stronger ways to quantify process performance, they are more likely to identify where those processes break—especially where work slows down or stalls. Those bottlenecks, once visible and addressed, become the points where AI can accelerate value (see "Research methodology" on page 24).

But not every point of friction is worth fixing. The goal is to remove the friction that keeps agentic AI from reaching its full potential. The processes that work the same way every time, that thrive on structured data and repeatable logic, have already been automated. What remains are the processes and value streams that have resisted standardization. The ones that don't follow a script.

These are the places where humans still make the call—because the problem was never simple enough to solve with traditional automation, such as decision governance. "AI decisions are probabilistic, but our governance frameworks still demand certainty," says Benedikt Schraik, CTO, City of Vienna. "We accept probabilistic human decision-making in government, backed by appeal and correction. AI governance needs the same acknowledgment."

These complex, variable, context-dependent value streams are where agentic AI finds its sweet spot. They're the decisions that shape customer outcomes, drive competitive advantage, and unlock new revenue. Reducing friction in these value streams is where enterprises should focus their AI investments.

"With technology advancing on a weekly basis, speed and adoption require a top-down management team aligned on which bets will truly move the needle."

Michael Voegele
Global CDIO, Philip Morris International

Part two

The real size of the prize

AI spend is rising, but not in a vacuum. Between 2025 and 2028, executives in our survey expect IT spend to balloon from USD1.4 billion to an estimated USD2.2 billion, on average. At the same time, 58% of executives are under pressure to cut non-AI tech costs.

That creates a direct conflict: the same legacy environments that slow AI progress are now competing for funding with the initiatives meant to replace them. Without clear prioritization, organizations risk spreading investment too thin—funding incremental improvements across the enterprise without delivering meaningful value anywhere.

But our survey data show that most organizations are struggling to prioritize AI initiatives enterprise-wide. Only 17% of executives say the process for approving and funding AI initiatives is consistent across the enterprise, and just 25% say business value and ROI are measured consistently.

Defining a clear success threshold for all AI investments—and stopping projects that don't reach it—is necessary to make AI transformation financially sustainable. In fact, recent IBM IBV research found that 70% of executives say they're looking to use the value gained from AI to fund investment and growth across the organization.[3] Early gains matter and must be tracked. But they are only the starting point.

Today, organizations are exploring a broad and diverse set of value propositions with AI. These include business process reinvention, product and service innovation, productivity, headcount reduction, and business model innovation.

There's value in each approach, but organizations should focus reinvestment on the initiatives that are best positioned to deliver meaningful returns. That's the only way to keep the transformation flywheel going. "We've made a conscious choice not to do AI everywhere—only where it actually creates value," says Gavin Wood, CIO, APA Group.

How can organizations deliver the greatest returns with their AI investments? By improving the customer-facing value streams that generate revenue rather than optimizing smaller administrative activities.

Think of it this way: a 1% improvement in selling, general, and administrative (SG&A) costs yields approximately USD24 million for a representative USD17 billion enterprise. A 1% improvement in core delivery processes, where more resources are invested, yields USD102 million. But cutting costs will only take you so far. To deliver revenue growth, AI must be applied to the customer-facing value streams that drive profitability, as well. A 1% revenue uplift would yield USD170 million in top-line growth.[4]

Yet investment often gravitates toward areas where implementation is easier, governance is clearer, and results are faster to demonstrate. That bias is understandable—but limiting. In fact, 85% of executives say many AI initiatives fail to deliver ROI because they aren't designed to improve the outcomes of critical processes.

This is where reinvestment becomes critical. Productivity gains in overhead functions should not simply reduce costs; they should create capacity to tackle the harder problem—transforming the work that drives revenue, margin, and competitive performance. By focusing efforts on revenue-generating value streams, organizations have the potential to deliver 7x more value than if they cut administrative costs alone (see Figure 3).

The greater their early gains, the bigger investments organizations can make in new revenue-generating opportunities. As organizations move from disconnected AI experiments toward smaller portfolios of initiatives tied to measurable business outcomes, concentrating investment in a handful of critical value streams can make a bigger financial impact.

Yes, succeeding with these investments requires greater investment in technology, governance, and people—and that could mean funneling a higher portion of IT spend to AI. But some IT upgrade costs should also be considered central pieces of AI initiatives. Otherwise, ROI figures overestimate AI's true returns.

"We need a clear view of how AI investment grows over time," says José Ovidio Silva, Group Head of Technology and Services, Campari. "I'm not talking about immediately cutting IT but defining a funding model for AI that cannot rely solely on cannibalizing existing IT investments."

FIGURE 3

An AI investment portfolio that taps into bigger pools of value can unlock bigger returns

Dollar figures show the impact of a model $17 billion organization achieving a 1% improvement in three different value pools.

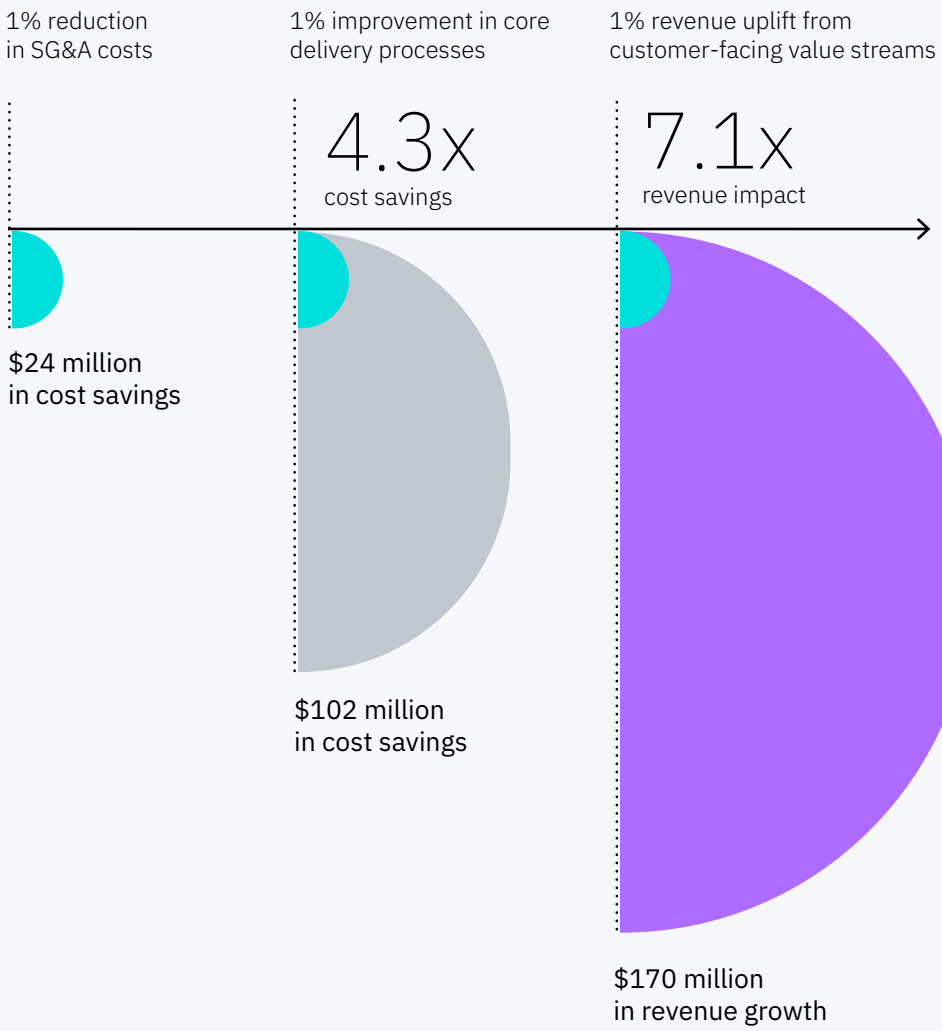

Perspective

Where AI goes to work

Conversations about AI often revolve around four types of activity: tasks, workflows, processes, and value streams. Many people use these terms as if they are interchangeable. They are not.

A **task** is a discrete unit of individual work.

A **workflow** is a connected sequence of tasks within a function.

A **process** spans multiple functions and produces an internal output.

A **value stream** is the complete sequence of actions that converts customer demand into a fulfilled customer outcome. Also known as an **end-to-end workflow**, it crosses organizational borders to deliver on customer needs. It is bounded on one end by an initial customer request and on the other by the customer's realization of value. Everything inside is execution.

The distinction matters because the level at which AI is deployed determines the level at which value is created. A task automation produces a faster task. A workflow automation produces a faster workflow. Neither produces a business outcome unless the value stream into which they are embedded was designed to connect them to customer demand and performance targets. It's only in the value stream that organizations can expect to drive revenue growth.

The same logic applies to team design. Task-level automation may require an individual to use an AI assistant. But value-stream reinvention requires hybrid teams of people and agents playing specific roles.

The higher AI moves in the value stream, the more explicitly organizations must define these roles, including:

- **agent designers** who define the work agents are meant to perform,
- **agent supervisors** who monitor decisions, exceptions, and escalation,
- **agent trainers** who improve agent behavior using feedback from real work, and
- **business owners** who remain accountable for value, risk, and customer outcomes.

Part three

Renovate the enterprise for AI

Most of the work required to drive growth with AI isn't related to building new systems. It is rethinking how technology is managed and how teams work with AI. That begins with clearing away layers of accumulated complexity and removing organizational friction.

The opportunity is significant: revenue-generating value streams offer the greatest potential return. But they also carry the greatest risk. Leaders need to balance both—ensuring that what moves into production is ready, resilient, and governed.

Successfully redesigning a value stream for agentic AI requires changes in four closely related practices: **architecture, data, decision rights, and business and IT operations.** How an organization renovates each area—and how it redesigns the teams, roles, and controls around them—determines not just the value AI can deliver, but how safely and sustainably it can scale.

Architecture

Most organizations today operate hybrid environments in practice, but not by design. Workloads reside in multiple clouds and on-premises systems for valid reasons—cost, performance, regulation. But without a consistent execution layer across these environments, every boundary becomes a barrier.

Operating consistently across environments means establishing a unified model for how workloads run, how integrations are managed, and how changes propagate across the estate. To scale AI responsibly, organizations also need to embed consistent controls across the execution layer. This helps ensure that models, data flows, and integrations are governed, observable, and production-ready before they scale. When engineers, developers, and operators follow consistent processes, organizations reduce friction. They can also scale AI repeatedly rather than reinventing the approach for every new initiative.

Architecture economics matter, as well. Training models, running inference, processing sensitive data, and integrating operational systems may each have different cost profiles and performance requirements. Organizations that can place workloads deliberately—balancing capital investments, operating costs, and business outcomes—will have a significant advantage over those that simply default to a single environment.

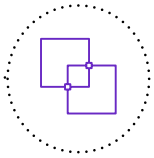

Data

Most high-value process data is distributed, with 74% living outside public cloud environments. This data is also maintained in formats optimized for storage rather than action. Only 41% of executives say data sources are available in real time outside of batch feeds.

Agentic systems require something different: continuous, contextual access to the state of a process. Without that, agents work under incorrect assumptions. This leads to small errors in the best case and completely misguided actions in the worst case.

Creating "warm" data—data that is continuously updated, well-governed, and accessible across the value stream—is an operational shift. It requires streaming event data across systems, establishing clear ownership of critical data sources before agents reach production, and building the engineering discipline to present enterprise data in consistent, structured forms. This data must not only available, but reliable, traceable, and governed end-to-end.

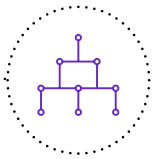

Decision rights

Agentic AI requires organizations to think carefully about decision-making. When an AI system acts autonomously, who owns the outcome of an agent's decision? What thresholds define acceptable autonomy? How is performance measured and corrected?

Few executive teams can consistently answer these questions. Only one-third of organizations formally verify ownership and maintenance responsibility for data sources before pushing an AI agent into production. Decision rights are not only governance mechanisms. They are signals that shape behavior. Employees must understand when AI recommendations should be followed, challenged, or escalated.

Without clear expectations, organizations risk creating uncertainty that limits adoption and slows decision-making. They need to codify decision rights as rigorously as they do financial controls—defining clear accountability, setting boundaries for autonomy, and establishing continuous oversight once agents are in production.

Decision rights also need to account for systemic risk. In an agentic environment, one decision rarely stays isolated. Agent actions can compound across a value stream. Organizations should scale autonomy progressively. This means moving from supervised recommendations to bounded execution to greater autonomy only when decision quality, exception rates, auditability, and business impact prove the agent is ready.

Business and IT operations

Like any complex transformation, renovating an IT estate carries significant overhead. Every improvement comes with setup costs: aligning stakeholders, mapping dependencies, establishing baselines. A lot of work has to happen before there's any chance of success.

That's why the operating model matters as much as the technical plan. The instinct in many organizations is to scale by starting more AI initiatives. But this isn't feasible when projects never end. "Legacy operating models designed decades ago won't hold up in an AI-native world," says Conor Mlacak, CIO, Staples Canada. "The work itself must be fundamentally rethought."

With agentic AI, continuity of effort is key. You want today's work to build on yesterday's work. You want teams that maintain context, deepen expertise, and chip away at the problem over time rather than ramp up, deliver a milestone, and move on.

Continuity matters for another reason, too. People learn how to use AI through hands-on experience. Organizations often treat AI adoption as a training challenge, assuming employees can learn new ways of working through courses, workshops, or periodic change programs. But practical expertise develops through daily interaction. Employees learn when to trust AI recommendations, when to challenge them, and how to identify new opportunities for automation by working alongside AI in real business processes.

This makes operating model design critical. Teams assembled for short-term projects may deliver an implementation, but they rarely develop the organizational muscle memory required to transform how work gets done. In the conventional approach, strategists add expertise to operations but carry no accountability for whether the project produces long-term value. The business outcomes remain someone else's problem.

"Technology alone isn't resilience. In production environments, resilience means anticipation and prevention, not just recovery."

Dena Almansoori
Group Chief Technology and Innovation Officer, ADNOC

A better model is one built around sustained accountability and integrated delivery. This requires a cross-functional team with both depth and breadth of expertise, embedded in the work, accountable for results, and structured to operate continuously rather than episodically (see Figure 4). This team doesn't just advise. It delivers for the long term.

FIGURE 4

Outcomes-driven delivery drives better AI outcomes

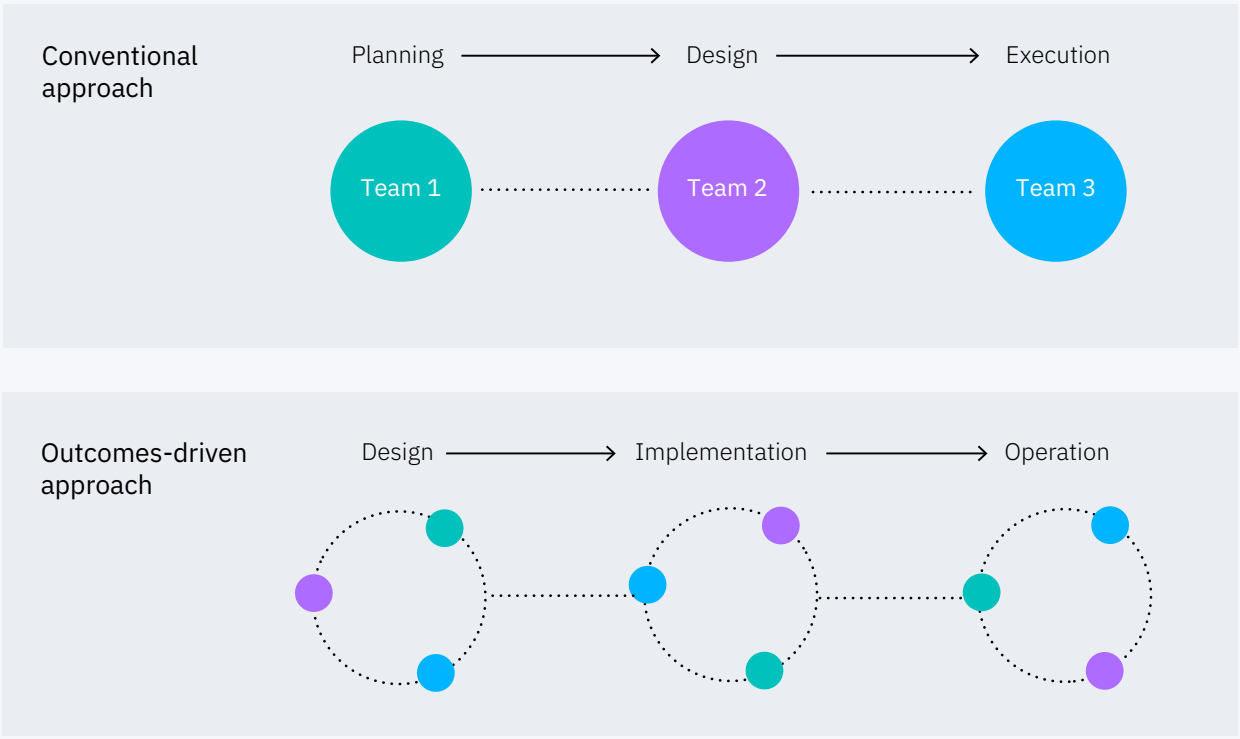

Case Study

Riyadh Air designs a new generation of connected customer experiences[5]

For all the promise of AI-native operations, few industries face a more urgent need to evolve than aviation, where even the smallest inefficiencies carry an immediate and cumulative cost.

Yet many airlines still run on architectures long overdue for reinvention. Legacy systems remain a primary barrier to modernization, constraining innovation precisely when speed and adaptability matter most.

And too often, promising journeys falter under that operational friction. Outdated systems stall, queues swell, and information fragments. Travelers feel every pain point, and inevitably, their journey starts to lose altitude long before takeoff.

Riyadh Air is determined to change that. Riyadh Air designed its AI-native operating model from the ground up, free from legacy systems and built for the demands of the digital era. Each design decision challenged long-standing assumptions that have shaped the airline industry for a century.

Reimagining an industry at this scale requires perspective, pattern recognition, and the ability to translate ambition into coordinated execution. Its operational foundation is a platform that anchors how intelligence operates across the enterprise. Models, workflows, and governance converge, enabling the business to scale as it evolves while keeping insight closely connected to outcomes.

With this foundation in place, multiple agents are able to manage routing and complete complex tasks together, translating insight into outcomes at operational speed. Contextual data now move continuously across more than 75 systems through more than 1,800 integrations, bringing every domain into a unified flow of information. The result is an integration fabric that allows Riyadh Air to sense, decide, and respond as one system, elevating both guest experience and operational performance.

AI now informs decisions as they happen, turning insight into context-aware action. A real-time engine creates, prices, and delivers offers dynamically as the journey unfolds. Travelers are no longer anchored to a single moment of purchase. AI assistants handle routine work with oversight built in, freeing engineers, operations, and services teams to focus on moments that demand judgment and care.

Offers adapt in the moment. Service adjusts to each traveler. Operations respond as conditions change. With intelligence guiding decisions and automation amplifying impact, Riyadh Air maintains a clean core supported by modular architecture—one that absorbs new capabilities and stays ready for what comes next.

"We had a clear choice—be the last airline built on legacy technology or the first built on the platforms that will define the next decade of aviation."

Adam Boukadida
Chief Financial Officer, Riyadh Air

Action guide

Get more value from your AI investments

Next 30 days:
Establish visibility and focus

See the system before you try to optimize it. Identify where value stalls, where decisions accumulate, and where AI can have the greatest impact.

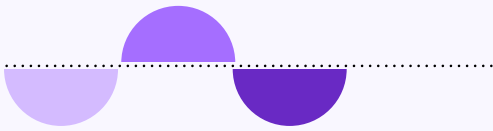

Architecture

- Identify priority value streams with high ROI that would benefit from remodeling for AI. Then map one priority end-to-end value stream to identify critical systems, integrations, and boundaries.
- Inventory where workloads run today (on-prem, private cloud, public cloud, edge) and document gaps in execution consistency.
- Identify the processes experiencing the most friction and redesign those first.

Data

- Locate high-value process data for the selected value stream—including data that sits outside public cloud.
- Assess how data is currently accessed (batch versus real-time); quantify where context reconstruction is required.
- Assign provisional ownership for key data sources—even if governance is incomplete.

Decision rights

- Identify all decisions an AI agent would need to make within the selected value stream.
- Clarify who is accountable today for those decisions—and where ownership is ambiguous.
- Define trust metrics for the selected value stream, including decision accuracy, exception rates, explainability, escalation frequency, user confidence, and business impact.

Operating model

- Assemble a cross-functional core team with representation from IT, data engineering, business operations, and risk/compliance. Ensure they have dedicated time, not just advisory roles.
- Identify and empower a single executive sponsor accountable for the end-to-end success of the selected value stream—not just AI deployment.
- Define success metrics tied to business outcomes. Clarify what “winning” looks like for this value stream in measurable terms (cost, speed, quality, revenue).

Next 120 days:
Redesign a single value stream at scale

Redesign the flow of work, not just the technology, so agents can operate across the value stream rather than within isolated tasks.

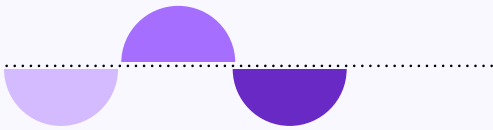

Architecture

- Implement a consistent execution layer across environments for the selected value stream.
- Reduce fragmentation by eliminating redundant or bespoke integrations in critical workflows.
- Establish continuous integration and deployment pathways for AI-enabled processes.

Data

- Connect upstream and downstream systems into a continuous data flow, eliminating batch handoffs that break context.
- Create shared access layers or APIs that abstract underlying complexity, so AI systems don't need to reconstruct meaning.
- Embed data engineers alongside process owners, so changes in the business are reflected in data structures in real time.

Decision rights

- Formalize an enterprise AI governance framework with clear escalation paths, accountability matrices, and performance review cycles tied to business outcomes.
- Require that new or updated agents meet predefined decision performance and risk criteria before moving into production, similar to financial or safety controls.
- Establish a progressive autonomy model that defines when agents move from recommendation to supervised execution to bounded autonomy.

Operating model

- Formalize the delivery team as a standing unit with sustained funding, clear roles, and accountability for both technical delivery and business performance.
- Create a new supervision model that clarifies which agents report into which value streams, who monitors their performance, and how managers remain accountable without becoming bottlenecks themselves.
- Implement a continuous feedback loop with end users and process owners: use their input to refine agent behavior, adjust guardrails, and prioritize next improvements.

Next 360 days:
Expand efforts for sustained enterprise impact

Turn a successful transformation into a repeatable capability by scaling the operating model, not just the AI.

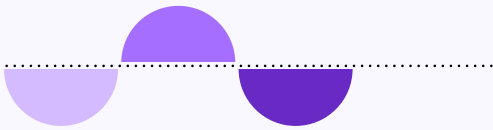

Architecture

- Expand the execution model to three to five high-value streams that collectively represent at least 60% of targeted business outcomes.
- Eliminate architectural debt systematically. Retire legacy systems that cannot support agentic operations or migrate them to modernized platforms with API-first design.
- Build self-healing infrastructure capabilities: automate detection and remediation of integration failures, performance bottlenecks, and configuration drift across environments.

Data

- Achieve enterprise-grade “warm data” across all priority value streams: real-time, governed, and contextually enriched data accessible to agents without reconstruction.
- Implement a federated data mesh or equivalent model that decentralizes ownership while maintaining enterprise-wide discoverability, lineage, and quality standards.
- Establish continuous monitoring of data freshness, completeness, and usage—with automated alerts when quality degrades.

Decision rights

- Measure and optimize agent decision quality at scale: track accuracy, bias, drift, and exception rates—and tie them to process KPIs and financial impact.
- Graduate high-performing agents from supervised to semi-autonomous or fully autonomous operation based on decision quality, risk performance, auditability, user adoption, and demonstrated business impact.
- Embed audit trails, explainability standards, and regulatory compliance in every production agent.

Operating model

- Transition from pilot teams to permanent cross-functional delivery units accountable for end-to-end value stream performance.
- Establish a continuous improvement engine: embed retrospectives, process mining, and value stream mapping into quarterly operating rhythms to identify the next layer of friction to remove.
- Upskill internal teams in agent design, data engineering, and hybrid architecture. Develop role-specific training that includes real, relevant examples of how to use AI in each job role.

About the authors

Varun Bijlani

Global Managing Partner,
Solutions and Delivery

IBM Consulting

linkedin.com/in/varunbijlani/

Varun Bijlani is rebuilding how consulting services get delivered, breaking the model that binds service delivery to hourly labor costs. He's working with clients to restructure delivery around persistent, forward-deployed teams that employ reusable, domain-specific agents instead of building assets from scratch on every engagement. In the new model, forward-deployed teams are accountable for outcomes delivered, not activities performed.

Javier Olaizola Casín

Global Managing Partner,
Hybrid Cloud and Data

IBM Consulting

linkedin.com/in/javier-olaizola-casin/

Javier Olaizola Casín leads industry-specific data initiatives built around a single, recurring theme: AI agents need real-time context to act reliably across the business. He is currently working with clients across sectors to close that gap—building continuously updated, well-governed data that carry context forward. His focus is the ownership, lineage, and hybrid architecture that keep that context trustworthy as AI scales.

Caroline Roche
Americas Hybrid Cloud and Data Leader
IBM Consulting
linkedin.com/in/caroline-roche/

Caroline Roche leads AI transformation agendas with clients, focused on deploying technology but also thinking about how to re-engineer processes, build incentives, and redesign organizations to support getting value out of AI. Her focus is on deploying multi-AI, hybrid/multi-cloud technologies, as well as making the organizational step-change to operationalize the value of the technology.

Richard Warrick
Global Research Lead
IBM Institute for Business Value
linkedin.com/in/richardwarrick/

Richard Warrick applies decades of global technology transformation consulting experience to research on how large enterprises employ IT to improve performance. His recent research has been focused on the challenges of converting AI investments into earnings-call-ready business value: targeting deep pools of value, managing portfolios of complementary initiatives, and reducing the resistance to change presented by IT complexity and technical debt.

Research methodology

The IBM IBV conducted a global survey to examine how organizations structure, measure, and operationalize their AI initiatives, and how these practices relate to financial outcomes. The study draws on responses from 1,000 senior executives responsible for AI, technology, and business transformation initiatives across 16 geographies and 18 industries.

Respondents represent a diverse set of organizations spanning multiple operating environments and levels of AI maturity.

Analytical approach

To identify the drivers of AI-related return on investment (ROI), we applied a combination of multivariate regression modeling and mediation analysis.

This approach enables:

- Identification of independent predictors of ROI
- Estimation of the mechanisms through which organizational capabilities translate into financial outcomes

Rather than assuming a single linear maturity construct, we analyzed how distinct operational capabilities interact to drive performance.

Regression modeling and validation

We estimated a series of multivariate regression models, using AI ROI as the dependent variable. We incorporated process-related capabilities, measures of operational visibility, and relevant control variables.

We assessed model robustness using standard goodness-of-fit and validation checks, including evaluation of explanatory power (e.g., variance explained), statistical significance of overall model fit, stability and significance of individual coefficients, and assessment of correlation structures across predictors.

These checks ensure that results reflect stable and interpretable relationships rather than model-specific artifacts.

Mediation analysis

We tested both direct and indirect effects because some capabilities are associated with higher ROI but do not consistently drive it on their own. Instead, they appear to create value by enabling more effective execution of AI.

In this context, improvements in operational execution mean having clear visibility into how work is done—such as identifying bottlenecks and understanding how systems connect—so AI can be applied to the right areas and deliver better results. To assess whether capabilities influence ROI indirectly through improvements in operational execution, we conducted mediation analysis using a standard product-of-coefficients framework.

The analysis tests whether capabilities are associated with better process visibility—such as identifying bottlenecks and understanding system dependencies—and whether these improvements are linked to higher ROI.

Indirect effects were estimated by combining two steps: how capabilities improve visibility into processes and how that improved visibility is linked to higher ROI.

We validated statistical robustness using bootstrap resampling, ensuring that indirect effects are stable and not driven by sampling variation.

Results indicate that a significant portion of the impact of process-related capabilities on ROI operates indirectly through improved visibility into process constraints and system dependencies.

Opportunity estimation

To contextualize the financial implications of AI-enabled performance improvements, we created an illustrative value pool-based estimation of potential enterprise impact. This calculation uses survey-reported ROI expectations as a reference point and links them to typical enterprise cost structures sourced from external benchmarks.

We applied external benchmarks of enterprise cost structures to translate percentage improvements into absolute value, using a representative enterprise revenue baseline of USD17 billion as a reference point to ensure comparability across scenarios.

Building on this framework, we assessed how modest, concurrent percentage improvements across multiple value pools can generate substantial aggregate value. Rather than relying on large efficiency gains within a single cost category, we illustrate a blended scenario in which incremental gains across SG&A, COGS, and revenue collectively contribute to enterprise value creation.

This reflects the asymmetric nature of enterprise value pools, where improvements in larger domains (such as COGS and revenue) yield disproportionately higher financial returns relative to smaller cost categories.

The resulting USD1 billion opportunity assumes that key obstacles have been overcome: that 75% of AI initiatives are operating at enterprise scale, that 75% are delivering intended business outcomes, and that only 10% of ROI is lost to IT friction. This opportunity should therefore be interpreted as a scenario-based estimate demonstrating how portfolio-level AI investments—when aligned to high-impact value pools—can translate into material enterprise financial outcomes when they drive plausible incremental performance improvements.

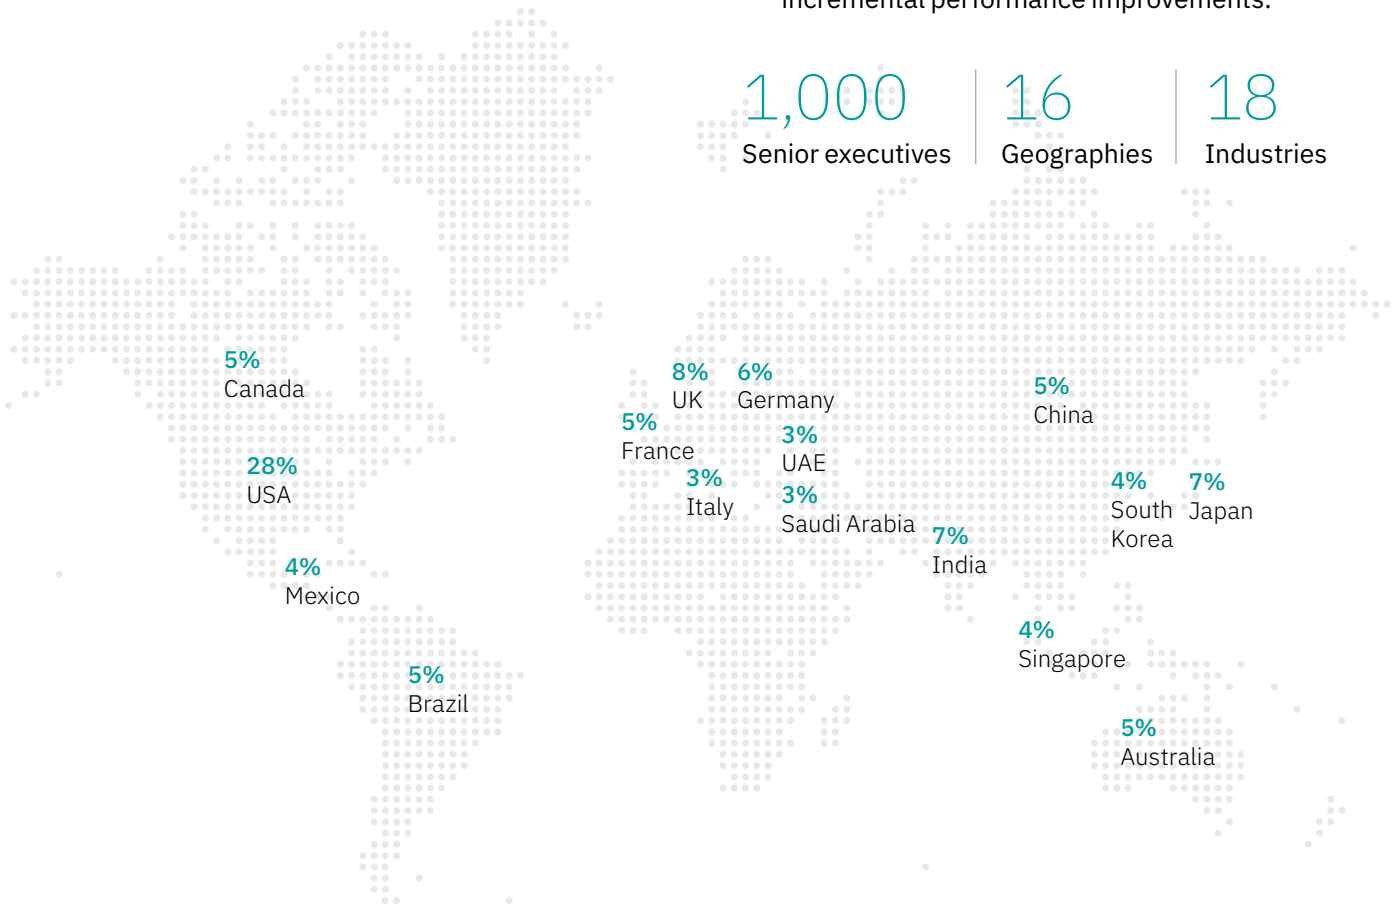

How IBM can help

Transforming business outcomes by using AI to reinvent core business processes requires a diverse ecosystem of assets and expertise. IBM, Red Hat, and IBM's network of strategic partners provide clients with deep industry expertise, a diverse portfolio of consulting services, and enterprise-grade technology solutions. IBM Research keeps clients on the leading edge of technology innovation, including industry-leading quantum computing technology.

For more information on AI services from IBM Consulting, visit ibm.com/consulting/artificial-intelligence

For more information on AI solutions from Red Hat, visit redhat.com/en

For more information on AI solutions from IBM Software, visit ibm.com/watsonx

For more information on AI innovations from IBM Research, visit research.ibm.com/artificial-intelligence

About Research Insights

Research Insights are fact-based strategic insights for business executives on critical public- and private-sector issues. They are based on findings from analysis of our own primary research studies. For more information, contact the IBM Institute for Business Value at iibv@us.ibm.com.

IBM Institute for Business Value

For two decades, the IBM Institute for Business Value has served as the thought leadership think tank for IBM. What inspires us is producing research-backed, technology-informed strategic insights that help leaders make smarter business decisions.

From our unique position at the intersection of business, technology, and society, we survey, interview, and engage with thousands of executives, consumers, and experts each year, synthesizing their perspectives into credible, inspiring, and actionable insights.

To stay connected and informed, sign up to receive IBV's email newsletter at ibm.com/ibv. You can also find us on LinkedIn at ibm.co/ibv-linkedin.

The right partner for a changing world

At IBM, we collaborate with our clients, bringing together business insight, advanced research, and technology to give them a distinct advantage in today's rapidly changing environment.

IBM IBV contributors

Haynes Cooney, Angela Finley, Tegan Jones, Arvind Raj, Thiago Sartori, and Shamasundar S E

Related reports

The enterprise in 2030: Engineered for perpetual innovation.
IBM Institute for Business Value. January 2026. ibm.biz/enterprise-2030

Go further, faster with AI: How governance increases velocity
IBM Institute for Business Value. December 2025.ibm.biz/ai-governance-trends

Agentic AI's strategic ascent: Shifting operations from incremental gains to net-new impact
IBM Institute for Business Value. October 2025. ibm.biz/agentic-ai-ops

Notes and sources

1. Bijlani, Varun, Todd Kathan, Mostafa Zafer, Meenakshi Srinivasan, Minoru Aoki, Jacob Dencik, and Lisa-Giane Fisher. *Cut the cost of complexity: Get more from your technology with intelligent IT automation.* IBM Institute for Business Value. October 31, 2025. https://ibm.biz/intelligent-automation

2. Bijlani, Varun, Javier Olaizola Casín, Suzanne Livingston, Matt Lyteson, and Ajay Patel. *The tech debt reckoning: A practical approach to boosting your AI ROI*. IBM Institute for Business Value. November 7, 2025. https://ibm.biz/debt-ai-roi

3. Baldwin, Andy, Neil Dhar, Ritika Gunnar, Rahul Kalia, James J. Kavanaugh, Salima Lin, and Joanne Wright. *The enterprise in 2030: Engineered for perpetual innovation.* IBM Institute for Business Value. January 16, 2026. https://ibm.biz/enterprise-2030

4. Representative P&L baseline figures were estimated for a typical Fortune 1000 enterprise at the following levels: 12% of revenue for operating income (EBIT), 14% of revenue for SG&A, 60% of revenue for COGS, and 40% of revenue for gross profit. Cross-industry averages were validated using the following sources: "The Hackett Group finds SG&A costs at a five-year high as 62% of US companies struggle to control spending amid slowing revenue growth." The Hackett Group. December 9, 2025. https://www.thehackettgroup.com/sga-costs-five-year-high-us-companies-struggle-to-control-spending-amid-slowing-revenue-growth/. Margins by sector (US). NYU Stern School of Business. January 2026. https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/margin.html

5. *Riyadh Air + IBM: Designing an AI-native airline from the ground up.* IBM case study. Accessed June 22, 2026. https://www.ibm.com/case-studies/riyadh-air

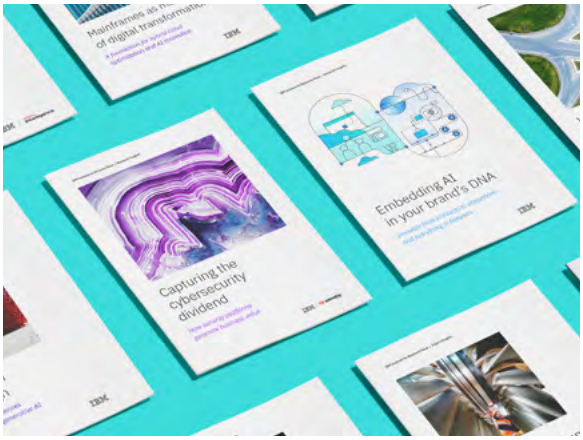

Subscribe to our IdeaWatch newsletter

Just the insights. At your fingertips. Delivered monthly.

Brought to you by the IBM Institute for Business Value, ranked #1 in thought leadership quality by Source Global Research for the third consecutive year.

Research-based thought leadership insights, data, and analysis to help you make smarter business decisions and more informed technology investments.

Subscribe now: ibm.biz/ideawatch

© Copyright IBM Corporation 2026

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | July 2026

IBM, the IBM logo, ibm.com and Watson are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at: ibm.com/legal/copytrade.shtml.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

This report is intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. IBM shall not be responsible for any loss whatsoever sustained by any organization or person who relies on this publication.

16ddab5aa1547458-USEN-01