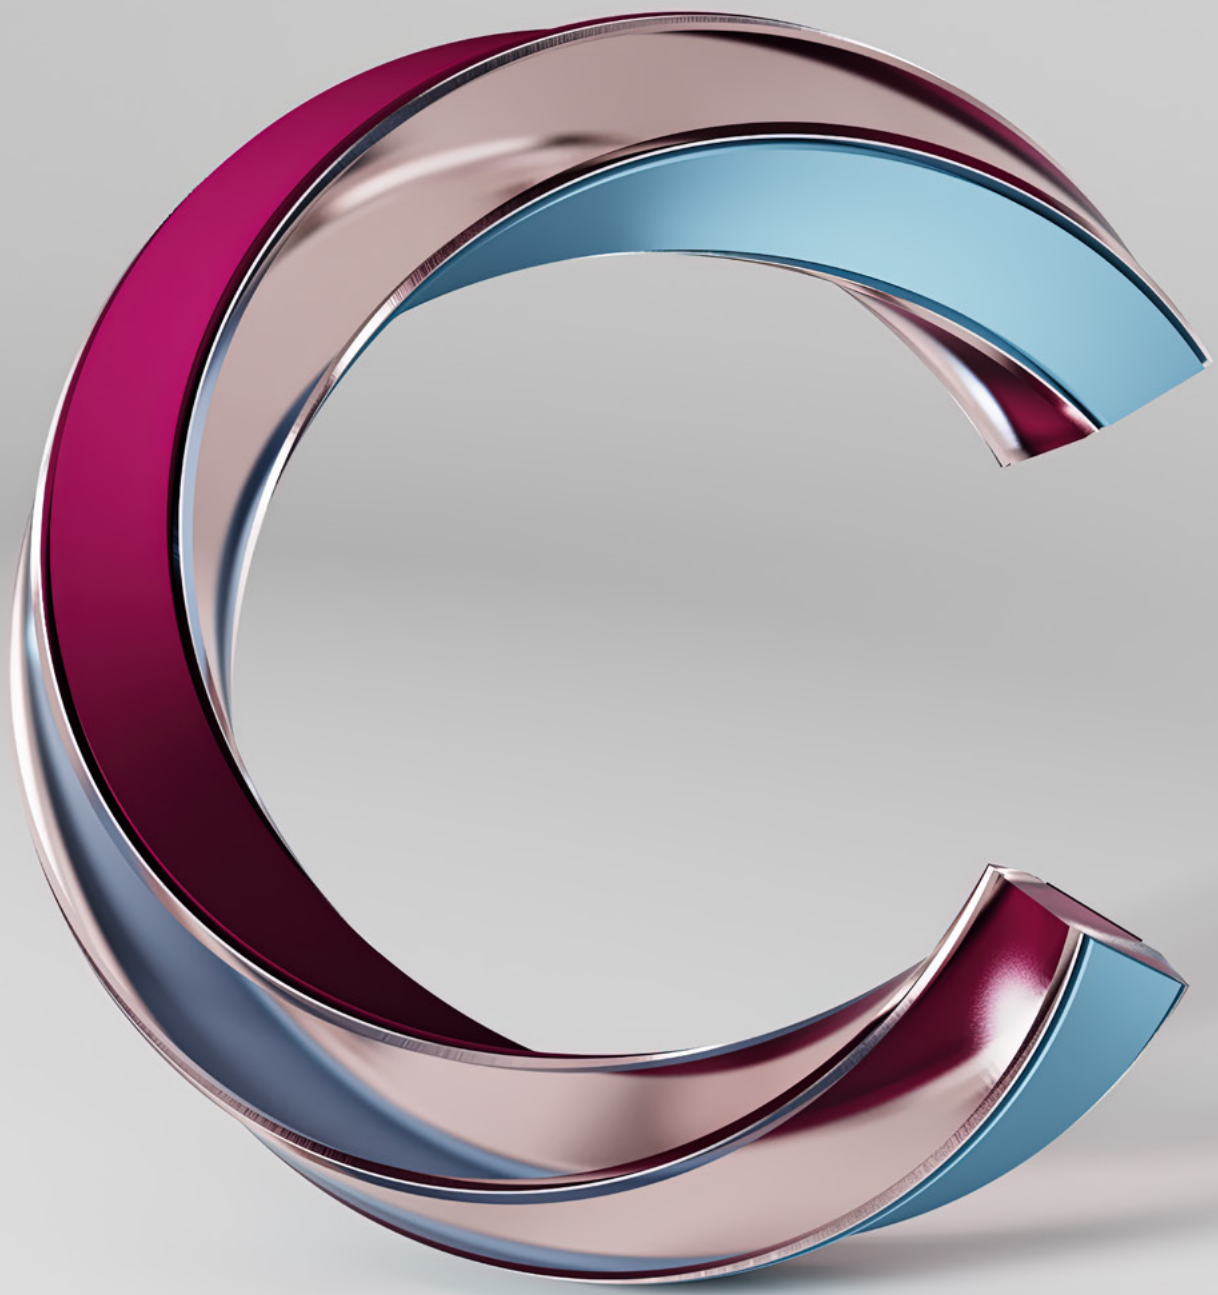
Global C-suite Series
2026 Tech Leader Study



Redefining the tech leader's mandate

Building the IT foundation for agentic AI at scale

IBM

IBM Institute for Business Value

About the study

In Q1 2026, in cooperation with Oxford Economics, the IBM Institute for Business Value surveyed 2,000 CIOs, CTOs, and other C-suite technology leaders across 33 geographies and 19 industries. The study examines how enterprises are managing the financial, operational, and governance challenges that emerge as AI moves from pilots to production. Our analysis combines segmentation and modeling to assess preparedness, control, and efficiency—including cloud portability, AI financial management, and governance maturity—enabling comparative insight into how organizations scale AI with resilience and discipline. For details, see “Research methodology” on page 32.

Contents

Foreword 3

Introduction 4

Part 1
Infrastructure adaptability enables AI at scale 8

Part 2
Governance by design makes AI control possible 16

Part 3
Portfolio discipline redefines AI investment efficiency 24

Foreword

Technology leadership shapes pace and possibility



Matt Lyteson

CIO,
Technology Platform
Transformation, IBM

This is a defining moment to lead technology.

As AI accelerates, CIOs and CTOs are making high-stakes decisions with incomplete information, under constant pressure, as conditions change in real time. These decisions now determine the enterprise's pace and what it can pursue.

The key constraint is adaptability. Most organizations were built for control, standardization, and efficiency. But in an AI-driven landscape, it is now architecture, governance, and investment decisions that establish how fast they can move, what risks they can absorb, and which opportunities remain viable.

This shift pushes the role of the CIO and CTO into critical new territory. While we remain accountable for resilience, security, cost, and trust, we're also shaping the business itself. Technology decisions no longer support business strategy after the fact. They define its boundaries.

The pressure is only intensifying. CEOs want AI scaled faster. Business leaders want autonomy. And expectations for accountability continue to rise. Yet many organizations are still operating with architectures, controls, and funding models designed for human-speed decision making, in systems that now operate at machine speed.

We've had to confront these realities head-on at IBM. Optionality has a cost, manual oversight doesn't scale, and investment logic built for multi-year asset lifecycles breaks down when models evolve in months. These are structural issues, and they're becoming decisive.

This report identifies what technology leaders are facing right now. If you feel a widening gap between the pace of AI change and your organization's ability to respond, know that that gap is strategic, not just operational. And closing it will require redesigning the enterprise itself—its architecture, its controls, and its operating model.

Introduction

IT architecture defines business strategy

Today's biggest constraint on business performance isn't ambition or technology. It's whether IT infrastructure can absorb machine-speed change without sacrificing control, resilience, or economic discipline.

Technology architecture decisions now directly determine what the business can pursue—and how fast it can move.

For CIOs, CTOs, and other C-suite technology leaders, the challenge lies not in deploying new tools but in redesigning an enterprise technology foundation built for a slower, more predictable environment. AI is dismantling any semblance of reliable cadence. And it's exposing the operating models that can't adapt fast enough.

As AI scales, systems make decisions continuously, autonomously, and at volumes no human-centered governance model can realistically supervise. In this report, we use *machine-speed* to characterize that shift in volume, autonomy, and decision frequency—and what it demands from enterprise systems.

Of course, this isn't the first time technology has reshaped decision making. The internet did it. Then cloud did it. What's different now is the speed—how quickly architectural, governance, and investment choices become structural constraints on business strategy.

"AI isn't just another technology disruption," says Michael Voegele, Global CDIO at Philip Morris International. "Unlike past innovations, AI accelerates its own development, almost like a self-fueling organism that continuously increases the pace of innovation."

As a result, the tech leader's mandate is expanding, often faster than the enterprise around it. CIOs and CTOs are no longer just enabling business strategy. More and more, they're helping define what strategies are even possible. The path will vary by context—regulated sectors, asset-heavy industries, and sovereignty-sensitive regions face different challenges—but the need to redesign IT foundations for machine-speed change is no less urgent.

That expanded mandate now carries real consequences. A recent study from the IBM Institute for Business Value (IBM IBV), *The enterprise in 2030*, predicts competitive pressure will make big bets non-negotiable by the end of this decade.[1] But ambition alone doesn't constitute a strategy. In an AI-driven enterprise—and as quantum capabilities emerge—the differentiator isn't boldness. It's whether the organization is built to place big bets with speed, discipline, and confidence.

"The biggest shift is the pace of change. Long-term roadmaps and monolithic programs no longer work at today's speed."

Rohan Kapoor
CIO, Northern Ireland Electricity Networks, UK

That structural readiness rests on three pillars—capabilities technology leaders must now engineer into their organizations. New research from the IBM IBV, conducted with Oxford Economics, shows what that means in practice:

1 **Infrastructure adaptability enables AI at scale.** For CIOs and CTOs, this means building foundations where switching cloud providers, rotating AI models, or absorbing new capabilities doesn't require replatforming or major re-engineering.

2 **Governance by design makes AI control possible.** When AI agents make thousands of decisions per day, control can no longer hinge on human approval gates. Tech leaders must pre-define and engineer boundaries directly into system architecture, including what agents can access, when they must stop, and how decisions remain auditable.

3 **Portfolio discipline redefines AI investment efficiency.** With shorter AI model lifecycles and asymmetric returns, tech leaders must restructure the IT portfolio to quickly refresh models, exit underperforming bets, and scale what works—without waiting for traditional budget cycles.



"CIOs must be business leaders first. The real value isn't in delivering technology—it's in translating technology into a business model and creating a positive financial impact."

Gavin Wood
CIO, APA Group, Australia

> "Technology leadership in 2026 is about creating the oxygen teams need to keep pace with change."
>
> **Graham Robinson**
> CTO, Smart DCC, UK

When embedded in the context of AI-first operations, these three pillars reinforce one another (see Figure 1). Infrastructure adaptability creates room to move, but only when governance is engineered into system architecture from the start. Portfolio discipline funds both, enabling CIOs and CTOs to refresh AI models, reallocate cloud resources, and scale proven capabilities based on measured evidence. When one is missing, it compromises the others—turning speed into exposure, control into drag, or investment into inertia.

At AI scale, deployment alone is no longer a meaningful measure of progress. What matters is whether organizations define expected outcomes upfront—and continuously measure performance against them. Without that discipline, adaptability, control, and portfolio rigor remain aspirations rather than operating capabilities.

When CIOs and CTOs build all three pillars together, they shift the reinforcing cycle from constraint to advantage. In our analysis, organizations where tech leaders have established infrastructure adaptability, governance by design, and portfolio discipline reported 38% higher expected revenue growth and 7% higher expected operating margin for 2026, and they're already deploying 2.6x more AI agents today.

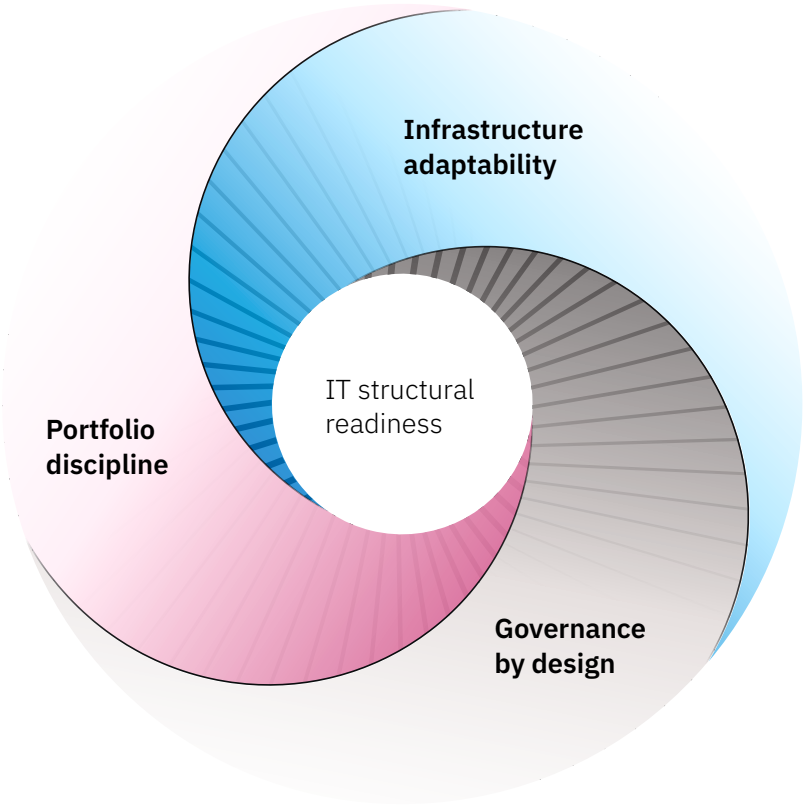
Yet for most, the gap is widening. Many tech leaders are still operating with IT architectures optimized for stability, governance models dependent on manual review, and investment frameworks built for multi-year asset lifecycles. These approaches can't keep pace with AI's speed. In our survey of 2,000 C-suite technology leaders across 19 industries and 33 geographies, only 11% say they're fully prepared for the scale of AI agent deployment expected in the next 12 months—even as 80% report transformation mandates coming directly from the CEO.

This is a structural mismatch, not a crisis of intent. The tech leaders making progress aren't attempting a complete IT overhaul. They're making targeted investments that remove their biggest structural barriers to speed, control, or capital efficiency.

In this report, we highlight how CIOs, CTOs, and other technology leaders can build structural readiness by establishing infrastructure adaptability, governance by design, and portfolio discipline—and where to act first as AI scales across the enterprise.

FIGURE 1

CIOs and CTOs must build all three pillars together to scale agentic AI with confidence



"My role isn't to generate every transformative idea. It's to build the foundation that allows smarter people across the organization to bring those ideas to life."

Chad Jones
CIO, Baylor Scott & White Health, US



Part 1

Infrastructure adaptability enables AI at scale

> "We don't know who's going to win or lose over the next five years. So we're keeping AI models plug-and-play, ready to adapt if the landscape shifts."
>
> **Dalton Gouws**
> Group IT Director and Board Member, VWG UK Ltd, UK

Historically, preparedness has meant stability—alignment to a known operating model optimized for continuity, cost, and predictability. That definition made sense when capabilities evolved slowly enough to allow for planning. At machine speed, it no longer holds. When models, platforms, and tools change faster than planning cycles, stability turns into constraint.

Preparedness today is defined by optionality, or the ability to shift workloads, rotate models, and absorb new capabilities without triggering large-scale disruption. By that measure, many organizations are less prepared than they assume.

What that optionality looks like will vary by context. In regulated sectors and sovereignty-sensitive environments, preparedness may depend as much on data residency and jurisdictional resilience as it does on technical mobility. In asset-heavy industries, adaptability often must be built around long-lived operational systems and uneven infrastructure realities, rather than greenfield flexibility.

That optionality isn't limited to infrastructure. It includes integrating new capabilities without replatforming, replacing models without rewriting workflows, and preventing local optimization from creating enterprise-wide lock-in. Preserving that flexibility requires deliberate investment upfront.

> "Everything we should have already done at the architecture and security levels becomes even more critical now. If those foundations are not in place, we are already at risk."
>
> **Francesc Suero**
> CIO, Audax Renovables, Spain

Most enterprises need optionality they no longer have

As organizations scale from hundreds to thousands of AI agents, immobility stops being a technical inconvenience and starts becoming a strategic constraint (see Figure 2).

Our research finds cloud costs exceeded original projections by 48% on average, and 80% of tech leaders report data transfer costs higher than expected. The root cause is architectural: while 88% of organizations are attempting or planning to move workloads to a different cloud provider, technology leaders say only 25% of those workloads are easily portable.

Despite significant cloud investment, many organizations can't adapt as easily as they thought. The same choices that make workloads harder to move also slow the pace of adopting better models, updating workflows, or bringing in new capabilities. Preparedness is about more than infrastructure. It shapes how quickly the business can respond.

What looks like rigidity now is often rooted in rational choices made earlier. CIOs and CTOs in our study point to cost optimization (60%) and innovation enablement (59%) as the primary drivers of their initial cloud migration decisions. Those choices created value while also introducing dependency.

> "Sustaining legacy systems as complexity grows is like sitting on a ticking time bomb. We may handle a crisis today, but I can't say the same about tomorrow. Without simplification, complexity itself becomes a resilience risk."
>
> **Tawatchai Cheevanon**
> Chief Product & Business Solutions Officer, Krungthai Bank PCL, Thailand

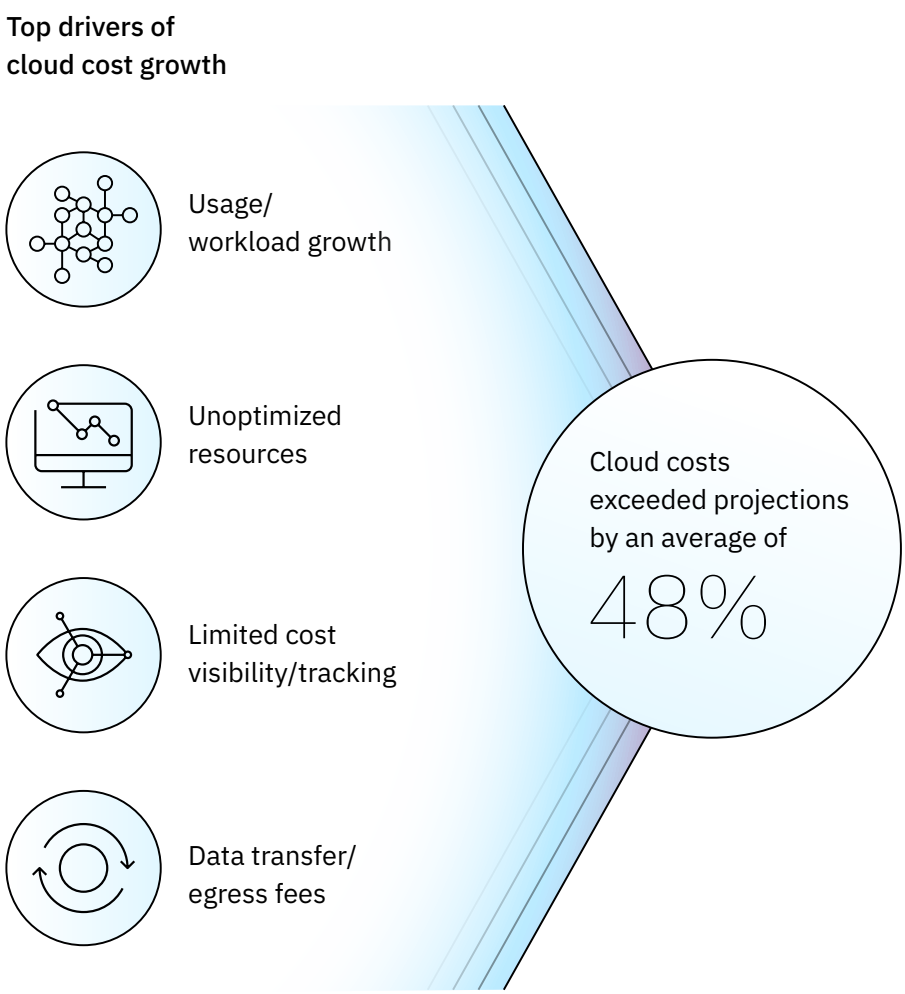
Yesterday's optimization limits today's flexibility

The consequences of limited optionality become clear when enterprises try to move. According to our study, barriers to portability take familiar forms—most often data transfer costs or egress fees (69%), reliance on provider-specific services (61%), and technical complexity (52%). Project issues may surface the problem, but architecture determines how costly it becomes.

The question is no longer whether workloads can eventually be moved. It's whether the enterprise can shift fast enough—when cost structures change, when performance expectations shift, or when new models outperform old ones. When the answer is no, the cost shows up in slower response to market shifts, weaker vendor leverage, delayed adoption of better models, and higher friction when capital needs to move.

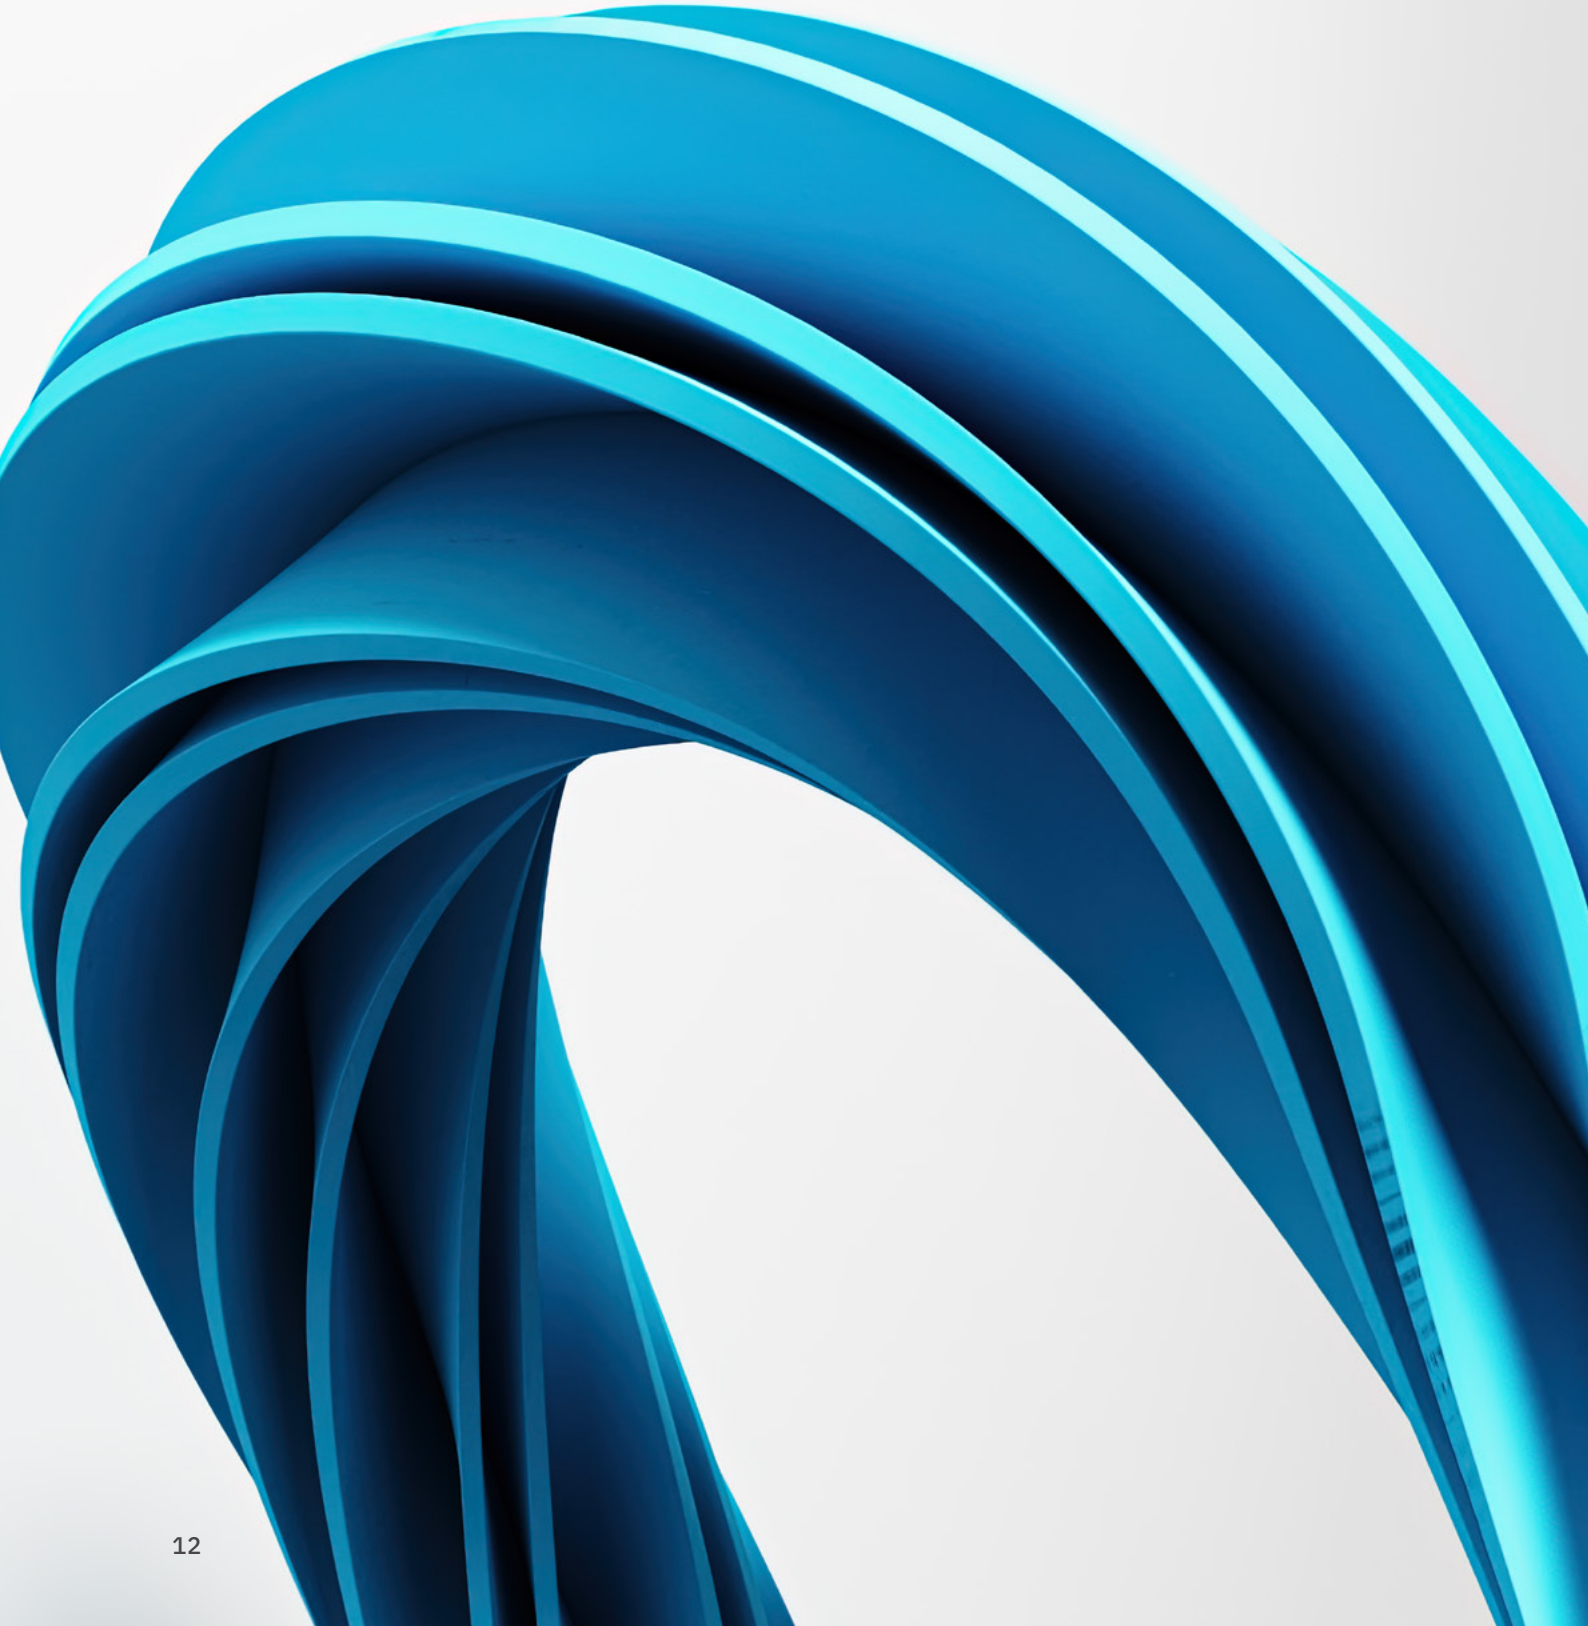
FIGURE 2

Immobility turns cloud cost into a strategic constraint



"Models and hardware should be replaceable. The data and enterprise intelligence built on top of them is what must endure."

Yasuo Kawai
Vice President, Head of Digital Unit, Honda Japan, Japan



Adaptability is the new readiness metric

Organizations that preserve optionality where it matters most approach preparedness differently. They make deliberate choices about where to accept dependency and where reversibility creates strategic advantage—decisions that shape architecture, integration, and investment long before conditions force a change.

Our research reveals organizations that designed for optionality early—by preserving workload portability, limiting hard dependencies, and supporting faster model refresh—reported a 10% higher return on AI investment in 2025. Faster time-to-capability and stronger AI ROI follow when model portfolios refresh continuously and workloads remain portable.

Preparedness is no longer defined by stability alone. It's about the absence of structural barriers to change.

"The most critical architectural capability is integration," says Conor Mlacak, CIO at Staples Canada. "We don't know what's coming next, so the foundation must support constant change."

For CIOs and CTOs, the question isn't whether every workload should be portable. Strategic risk defines the priority—and explicit funding determines whether flexibility exists when conditions shift.

> "We design modular architectures so components can evolve as technology advances, without breaking the overall system. That approach allows us to absorb rapid innovation while supporting products with decades-long lifecycles."
>
> **Boris Alexandre**
> Head of ARP Programme, Airbus, Canada

"Long term, we need a backend architecture that allows us to swap models without rebuilding enterprise systems."

Benedikt Schraik
CTO, City of Vienna, Austria

Quantum as the exception to the rule

Quantum computing presents a fundamentally different preparedness challenge. Unlike AI's plug-and-play apps, quantum requires algorithm and application development, alongside trials to test how an organization could benefit.

Meaningful participation requires long-horizon vision—building skills, partnerships, and operating models well before business value is visible. Once underway, those investments are cumulative and difficult to reverse, especially as business value examples emerge.

Here, preparedness takes a different form. Optionality comes from collaboration with partners that have a wide range of skills and expertise to clearly understand how a quantum modality can successfully execute certain workloads.

It also comes from deciding early on how quantum will impact your future and structuring for that decision in advance. Recent research from the IBM IBV shows 89% of tech leaders investing in quantum computing are motivated to future-proof their computing strategy—underscoring the need for a long-horizon investment posture, given that readiness in this domain cannot be rapidly engineered after the fact.[2]

The lesson is to move fast while recognizing where preparedness must precede deployment—and that delaying structural choices will limit future strategic paths.

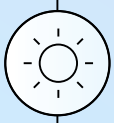
How to build for adaptability, not just optimization

Preparedness requires explicit choices about where having options matters, what level of dependency is acceptable, and how much the organization is willing to invest to preserve flexibility.



Right away:

- Stop treating all workloads as equally important to move.
- Stop approving net-new lock-in by default.
- Stop labeling egress and migration overruns “unexpected.”



Tomorrow morning:

- Set a rule for net-new dependency. New AI deployments shouldn’t create hard-to-reverse lock-in by default. Require justification when provider-specific services are used.
- Fund flexibility deliberately. Portability, resilience, and model rotation are strategic capabilities—not side effects of cost optimization.
- Decide where optionality matters most. Not every workload needs portability. Focus on the platforms, workloads, and data domains where lock-in would limit future business choices.



This month:

- Map your lock-in exposure across infrastructure, data, and AI models—then prioritize 2–3 high-value portability moves that materially improve flexibility without broad disruption.
- Establish cost transparency for egress and data transfer fees so portability decisions are based on actual economics, not assumptions.
- Put a lightweight review mechanism in place for new AI deployments that rely heavily on proprietary platform services.



Part 2

Governance by design makes AI control possible

> "Traditional governance functions design their processes around making it easy for the function to do its job, forcing the business to adapt. AI is flipping that model—pushing governance, IT, and security to think like service providers whose job is to enable the business, not slow it down."

Charles Newhouse
Vice President & CTO, Leidos UK, UK

Optionality creates speed, and speed creates exposure—unless control is built in.

For years, control in enterprise technology was exercised through policy, approvals, committees, and review cycles. That approach worked when systems moved at human pace. At machine speed, it no longer holds.

As AI agents move from pilots to production, they make decisions continuously and at volumes no escalation path can realistically govern. In that environment, control stops being a permission problem. It becomes a design problem.

Defining AI agent
In this study, "AI agent" refers to an AI-enabled system that can initiate, coordinate, or execute multi-step actions with limited human intervention.

> "It's like flying a plane at 10,000 feet, being told to climb to 12,000, replace both engines mid-flight, and ensure zero turbulence. No one would choose to pilot that plane—but that's exactly what companies are doing today."

Afonso Eça
Executive Board Member, Banco BPI, Portugal

Human-speed governance can't contain machine-speed decisions

Scale is arriving fast. By 2027, enterprises expect to deploy an average of 1,661 AI agents—a 38% increase from today. With each agent making hundreds or even thousands of decisions per day, technology leaders are suddenly managing hundreds of thousands of autonomous decisions daily. Manual governance can't keep up with that math.

Our research shows the strain is already visible. A majority of organizations (77%) report AI adoption is outpacing current governance capabilities. Nearly 60% of tech leaders cite security and compliance concerns as a top barrier to scaling agents, and two-thirds say they're accountable for outcomes in systems they don't fully control. Meanwhile, 70% report teams deploying technology faster than IT can track.

Pressure is rising in parallel. Only 11% of CIOs and CTOs say they feel completely prepared for the scale of AI agent deployment expected over the next year—despite 80% reporting transformation mandates coming straight from the CEO.

At machine speed, policy documents and approval committees can't keep pace. What worked for quarterly software releases breaks down when agents deploy continuously and decisions compound in real time. For tech leaders, governance must shift from reviewing what happened to engineering what's possible—before systems go live.

Manual governance creates a scaling trap

Reliance on human oversight for every deployment decision forces an impossible trade-off. Tech leaders tend to respond in one of two ways—and neither works at scale.

Some prioritize speed. Business units move ahead, and governance follows behind. Local velocity increases, but visibility and containment degrade. More than two-thirds of CIOs and CTOs say business units bypass IT to adopt AI. And with organizations experiencing an average of 54 AI agent incidents last year, 17% of those incidents were high severity, requiring more than four hours to contain (see Figure 3). As for the business impact:

37% resulted in data exposure or security breaches

33% caused cascading system failures

17% triggered compliance issues

13% eroded stakeholder trust

Others prioritize safety. Deployment slows under heavier review and approvals. Immediate exposure declines, but learning stalls, competitive position weakens, and manual oversight becomes unmanageable—especially as enterprises navigate dozens of overlapping AI regulatory frameworks.

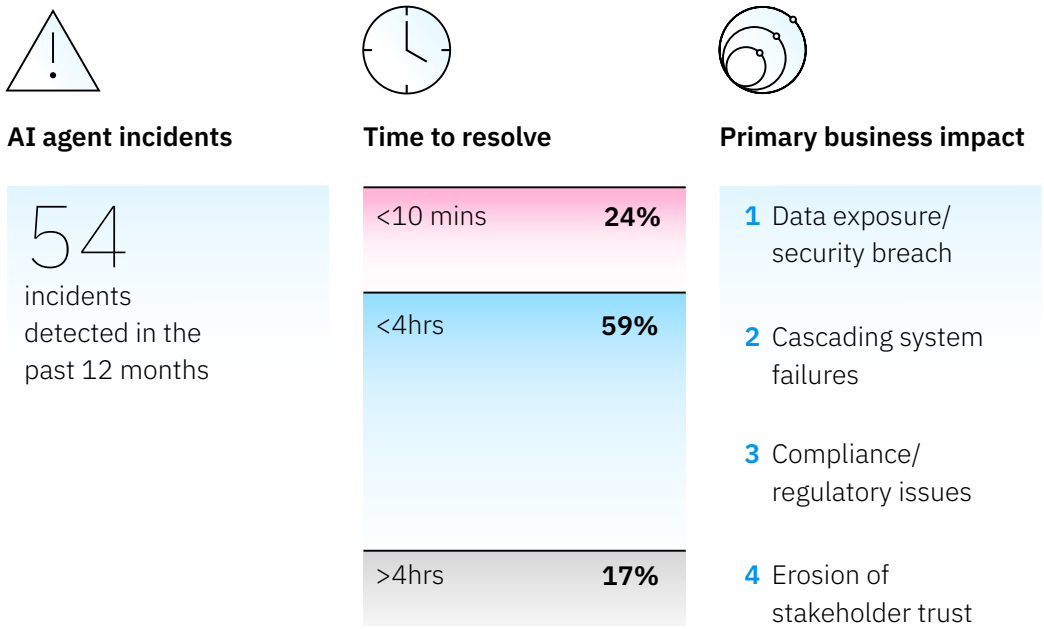
One route trades safety for speed—the other trades speed for safety. Neither path scales, and both accumulate strategic debt.

> "AI has both a light side and a dark side. While most focus on the opportunities, it also introduces new vulnerabilities, and many organizations are more exposed than they realize."

Victoria Medina
Chief Technology & Data Officer, Allianz Spain, Spain

FIGURE 3

Manual governance struggles to contain AI incidents at scale



"You have to let people move at speed, but within clearly defined lanes and guardrails."

Farid Boutaghane
CTO, Philip Morris International, Switzerland

Control is no longer a matter of process

Control is not about centralizing all decisions. It's about engineering enterprise-wide boundaries and visibility that let federated teams move fast inside safe limits. Tech leaders scaling agents safely are redesigning governance accordingly. Instead of relying on downstream review, they embed control into the architecture itself. Guardrails become executable, while decisions remain observable. Systems become bounded and stoppable.

In an AI-first environment, governance is built into how the system operates—at speed, with limits. As Dena Almansoori, Group Chief Technology and Innovation Officer at ADNOC, explains: "Control has shifted from approving inputs to continuously supervising outputs and outcomes—from gates to guardrails."

This shift enables scale (see Figure 4). Our analysis finds organizations we classify as having orchestrated control deploy 16x more agents than those relying on manual governance (see "What orchestrated control looks like in practice"). They do it while spending 4x less of their AI budget and delivering 18% higher operating margins. That kind of designed-in control requires upfront standardization and engineering effort, but it reduces downstream risk and friction as autonomy scales.

Designed-in control doesn't mean centralized command. It's a federated operating model, where platform, risk, architecture, and domain teams each own part of the control system—so autonomy can scale inside shared boundaries.

What orchestrated control looks like in practice

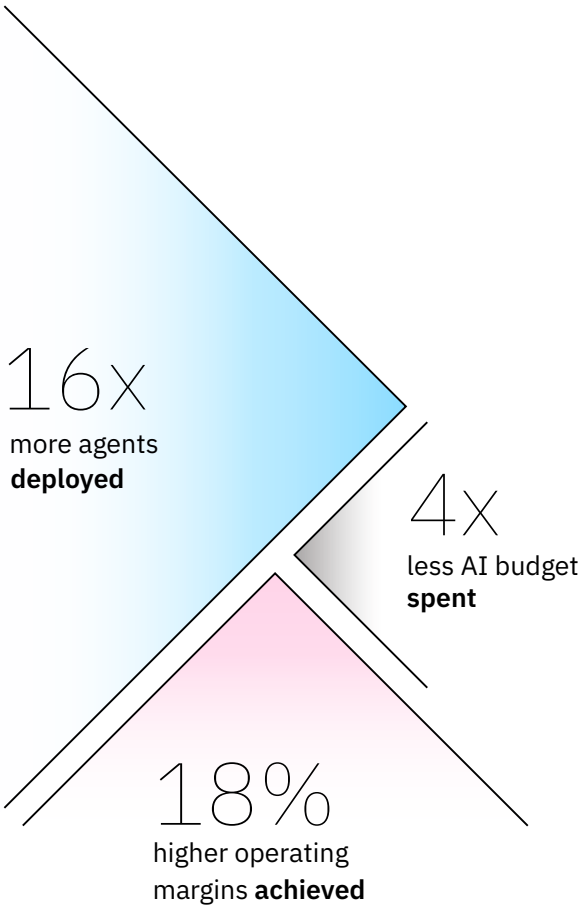
- Platform teams own shared guardrails: telemetry, model registry, identity, logging, rollback, and runtime controls.
- Risk and compliance define policy thresholds, evidence requirements, review triggers, and escalation rules.
- Architecture teams set reference patterns, control points, and interoperability standards.
- Business and domain teams own use-case outcomes, process integration, exception handling, and accountable operation within guardrails.
- Incident response maintains predefined shutdown, containment, and recovery procedures.

"For us, control means safeguarding our strategic assets—especially core platforms and cloud infrastructure, where our dependence is absolute."

José Ovidio Silva
Group Head of Technology and Services, Campari, Italy

FIGURE 4

Designed-in control scales faster—and safer
Organizations that embed control into architecture deploy more agents with less risk



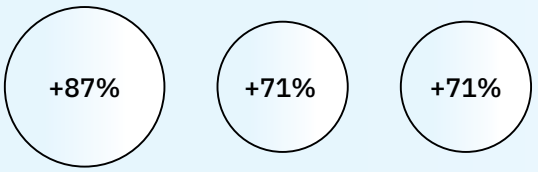
"Architecture is the foundational element of resilience."

Raúl Cals
CIO, Unicaja, Spain

Governance maturity determines how safely AI scales

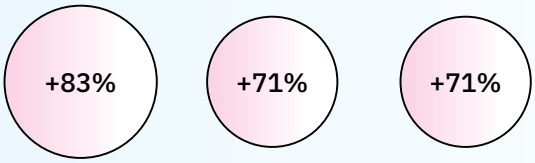
AI agent adoption is surging across regions and industries. But deployment speed doesn't determine success—governance maturity does. Organizations where tech leaders scale agents without strengthening governance see incidents rise in lockstep with adoption. Those with mature governance contain risk as they scale.

Regional leaders in expected AI agent growth, 2026 → 2027



Middle East & Africa | Europe | Asia Pacific

Industry leaders in expected AI agent growth, 2026 → 2027



Energy, Chemicals, & Petroleum | Transportation & Travel | Healthcare & Life Sciences

Across regions, agent growth ranges from 60% to 87%. Yet incident growth varies by nearly 2x. In organizations with weak governance, more agents mean proportionally more incidents. In organizations with strong governance, incident rates stay relatively flat, even as agent deployment accelerates.

Across industries, agent growth ranges from 63% to 83%. But scaling outcomes differ dramatically. Industries with complex, regulated operations—pharma, energy, telecom, retail banking—see governance maturity significantly reduce incidents as agents scale.

The takeaway for tech leaders: Governance investment should match operational risk. In high-stakes environments—regulated sectors, sovereignty-sensitive markets, customer-facing operations—underinvesting in governance turns AI growth into operational exposure.

"You need guardrails that guide AI toward making appropriate decisions. Over time, as confidence grows, maturity follows."

Rafael Salcedo
Head of Technology Platforms, Bankinter, Spain

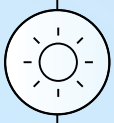
How to embed control, not just supervise it

Real control requires redesigning governance as an engineered system—one that defines what agents can do, when they must halt, and how decisions remain explainable across the AI lifecycle.



Right away:

- Stop treating policy documents as if they're active controls.
- Stop approving autonomous systems without rollback, logging, and clear ownership.
- Stop assuming governance will scale through more reviews and committees alone.



Tomorrow morning:

- Establish minimum production standards. If an agent or model isn't registered, owned, observable, and stoppable, it doesn't deploy. Make this non-negotiable.
- Redesign governance end-to-end for one high-risk domain. Choose claims, customer service, or code generation—then build a model with clear roles, automated monitoring, and defined escalation paths.
- Engineer one manual control into the platform. Move access boundaries, drift detection, escalation thresholds, or kill switches from committee oversight to executable code.



This month:

- Build a unified registry for all production AI agents and models. Track what's deployed, who owns it, what it accesses, and how performance is monitored.
- Automate one governance control across the AI lifecycle. Select model evaluation, deployment approval, or compliance checking—then engineer it into your platform with real-time monitoring and alerts.
- Define and enforce minimum standards for autonomous systems. Establish requirements for observability, auditability, rollback, and ownership—then build mechanisms to ensure compliance before deployment.
- Run an incident simulation to test your governance design. Simulate an agent failure to validate detection speed, containment protocols, and recovery processes.

Case study

Banco do Brasil designs control into AI to scale trust at enterprise speed[3]

As the oldest bank in Latin America and a public financial institution with regulatory duties and a public mission, Banco do Brasil faced a central challenge as AI adoption accelerated across critical banking functions: how to scale AI responsibly, transparently, and with confidence. In a regulated environment where trust and accountability are non-negotiable, existing governance approaches were insufficient as AI adoption expanded across critical functions.

To meet growing challenges related to trust, accountability, and regulatory compliance, the bank committed to a robust AI governance strategy—an evolution of its governance model designed to embed control across the AI lifecycle. Frameworks were established to define clear roles and responsibilities from model design through deployment in critical environments, supported by rigorous evaluation processes for foundational models, including benchmarking and vulnerability assessments.

To support continuous supervision, explainability, and compliance, Banco do Brasil implemented real-time monitoring, proactive alerts, and customized metrics, consolidating a unified system of control and trust. This integrated approach is also being tested in generative AI use cases to help ensure alignment with compliance, performance, and risk management requirements.

Today, the bank operates under a unified AI governance model and has strengthened its ability to manage AI initiatives with precision, transparency, and control—helping ensure models deployed meet rigorous standards for ethics, security, and performance. The result is improved oversight, agility, and institutional trust, demonstrating how designed-in control makes safe scale possible.

Key outcomes include:

- **Automated governance across the AI lifecycle,** reducing manual oversight and accelerating deployment.
- **Real-time monitoring and explainability,** enabling proactive risk and compliance management.
- **Traceability and version control of AI assets,** promoting visibility and accountability.



Part 3

Portfolio discipline redefines AI investment efficiency

"Until now, we focused on automating parts of existing processes. The real challenge now is redesigning processes end to end, based on what AI makes possible—not simply automating what already exists."

Sergio Sánchez Gallego
CTIO, Telefónica España, Spain

Efficiency has long been equated with cost discipline. That definition worked when technology investments were predictable—built to last, depreciated over time, and optimized for stability.

But the IT investment profile has changed. AI introduces a class of assets with short lifecycles and uneven returns, where success is unpredictable and timing is compressed. Value depends less on cost minimization than on speed of learning and reallocation.

In such an environment, efficiency cannot be defined by cost control alone. Tech leaders must manage their AI investments as a portfolio, where capital moves based on evidence, not annual budget cycles.

"As organizations scale, resilience becomes increasingly critical. There is always a trade-off between the cost of resilience and the impact of failure. The larger you are, the more that impact matters."

Atul Bhardwaj
EVP & Chief Digital & Technology Officer, LEGO Group, Denmark

> "Digital transformation will ultimately be a business transformation—enabled by technology but not driven by it alone."
>
> **Christophe Mario**
> CIO, Mutua Madrileña, Spain

Short lifecycles demand a different investment approach

AI models don't age like traditional software. According to our research, average useful life is roughly 14 months—short enough that refresh decisions occur continuously, driven more by iterative improvement than by failure.

This reality undermines the old business-case logic built on three- to five-year asset lifecycles. Most AI models today are retired not because they break, but because organizations learn what works. In our study, 71% of technology leaders cite the availability of better models as the primary motivation for retiring or replacing models, while 60% cite evolving business needs and use cases.

For tech leaders, this means traditional capital planning—where investments are approved once, deployed, and depreciated over years—no longer applies. CIOs and CTOs must now build mechanisms to refresh models, redirect capital from underperforming bets, and double down on what's working—not on annual schedules but in real time, as learning accelerates.

As pilots mature, technology leaders gain clarity into where AI creates the most value. And that clarity must drive capital decisions.

> "Preparing the company for the future means aligning business and technology. Technology is no longer optional—it is a core asset that creates business value."
>
> **Rafael González**
> Corporate IT Director, Logista, Spain

Cost control is necessary, but no longer sufficient

Traditional total-cost-of-ownership (TCO) logic remains essential for operational infrastructure. Proven systems running at scale—including mature AI use cases—still demand optimization, monitoring, and financial rigor.

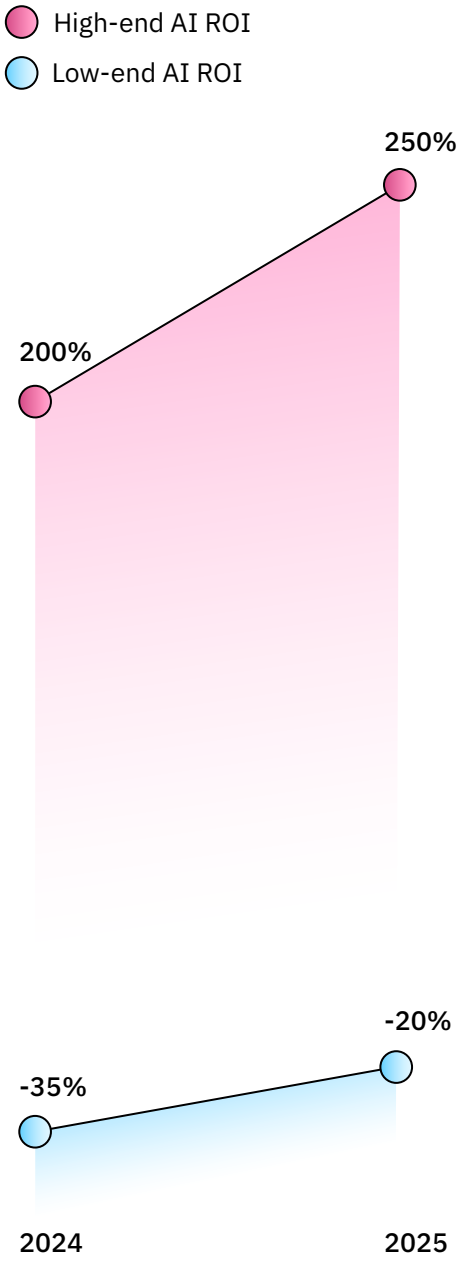
But exploratory AI investments behave differently. Many AI initiatives are experimental by design. Value emerges unevenly. Some bets fail quickly, while others scale rapidly and reshape business performance. Evaluating all of them through a single cost-optimization lens systematically undervalues investments with asymmetric upside.

When CIOs and CTOs manage AI investments as a portfolio rather than optimizing them individually, the enterprise learns faster and can scale proven capabilities without traditional approval delays. And those dynamics show up in performance outcomes.

From 2024 to 2025, organizations at the top end of AI ROI increased returns from roughly 200% to 250%. At the low end, ROI improved only modestly, from -35% to -20%. The gap is widening—not because organizations spend more, but because they allocate and refresh capital differently (see Figure 5).

FIGURE 5

The AI ceiling is rising faster than the floor



Portfolio discipline defines modern efficiency

Cost discipline hasn't disappeared, but the tech leader's role in applying it has fundamentally changed. Operational AI—proven use cases running at scale—still demands tight cost control. Infrastructure supporting agents still requires optimization, monitoring, and vendor discipline.

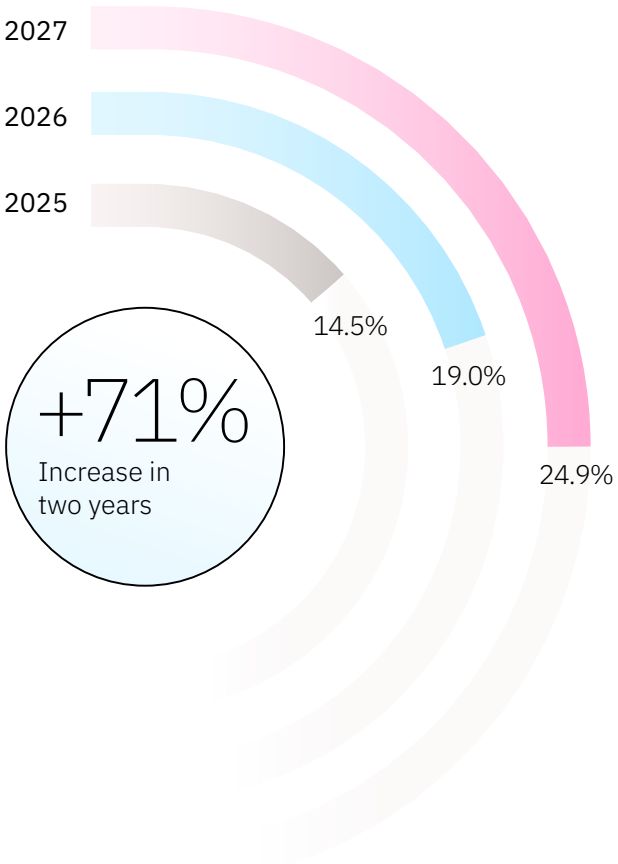
But strategic AI investment is different. Experiments, capability-building bets, and model refresh decisions require portfolio logic: explicit owners, clear success criteria, defined exit points, and the ability to reallocate capital as evidence changes. Higher per-unit cost is accepted in exchange for speed, optionality, and asymmetric upside. Portfolio discipline adds management complexity, but it's the necessary trade-off for reallocating capital at the pace AI economics demand.

Some organizations are already putting mechanisms in place to manage this actively.

"We're building a control plane for AI, so we can identify what's productive, elevate what's valuable, and retire what isn't," says David Treat, CTO at Pearson. "Some experimentation waste is unavoidable—but that's the price of learning fast."

FIGURE 6

AI is becoming a portfolio-scale investment
Figures reflect AI spend as a percentage of IT budget



"We're starting to see how quickly costs can get out of hand. Building agents is powerful, but the key question is this: how much is each agent actually costing us?"

Miguel Santos
CTO, MasOrange, Spain

Pressure to get this right is mounting. AI spend is projected to grow from just under 15% of IT budgets in 2025 to nearly 25% by 2027—a 71% increase in two years' time (see Figure 6). That translates into hundreds of millions of dollars moving into a category of investment that behaves quite differently from traditional IT.

For CIOs and CTOs, managing this shift means building financial visibility and coordinating tightly with Finance to reallocate capital as new evidence emerges. The dual mandate of optimizing infrastructure while actively rebalancing a portfolio is what separates tech leaders that scale AI successfully from those that don't.

At AI scale, efficiency is no longer just about spending less. It's about allocating better.

And yet, most organizations are not set up to manage AI this way. Our research finds 84% of technology leaders have not yet fully operationalized AI financial management, and 85% still lack full visibility into real-time AI spend. Without that visibility, portfolio decisions slow down—and efficiency defaults back to cost cutting, even as AI spend accelerates.

"The real challenge is understanding how technology initiatives impact the end-to-end cost of business processes."

Javier Revuelta
CIO, CAF, Spain

Where financial discipline creates advantage

Organizations where technology leaders have built strong financial discipline outperform their peers—by building early visibility, coordinating more tightly, and allocating capital more deliberately.

+17%

more likely to have cost-tracking in place from the start of cloud migrations.

+16%

more likely to report joint IT-Finance decision making and shared accountability.

1.5x

more likely to report cloud initiatives exceeding expectations.

2.4x

more AI agents deployed—with no higher AI/IT budget.

3x

more likely to say they're fully prepared for the AI surge.

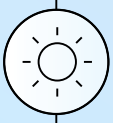
How to manage AI as a portfolio, not a cost center

Portfolio discipline requires separating operational efficiency from strategic optionality—and building the visibility to manage both.



Right away:

- Stop evaluating exploratory AI bets using infrastructure payback logic alone.
- Stop funding pilots without accountable owners and explicit success criteria.
- Stop reviewing AI economics only during budget cycles.



Tomorrow morning:

- Separate AI spend into two categories. Manage operational AI with TCO rigor. Manage strategic AI investments as a portfolio.
- Demand use-case-level visibility. If spend, owner, model, business objective, and expected return aren't visible by use case, portfolio management isn't happening.
- Establish a shared IT-Finance cadence. Capital reallocation needs to be an operating rhythm, not an annual event.



This month:

- Build an initial portfolio view of AI spend by use case, owner, model, and expected value.
- Align with Finance on clear criteria distinguishing operational AI from strategic investment.
- Launch a recurring portfolio review to reassess funding, refresh decisions, and exit criteria as evidence changes.

"Moving the organization through change is the one CIO responsibility that has never changed."

Chris Pesola
CIO, Roush, US



Research methodology

Study design and sample

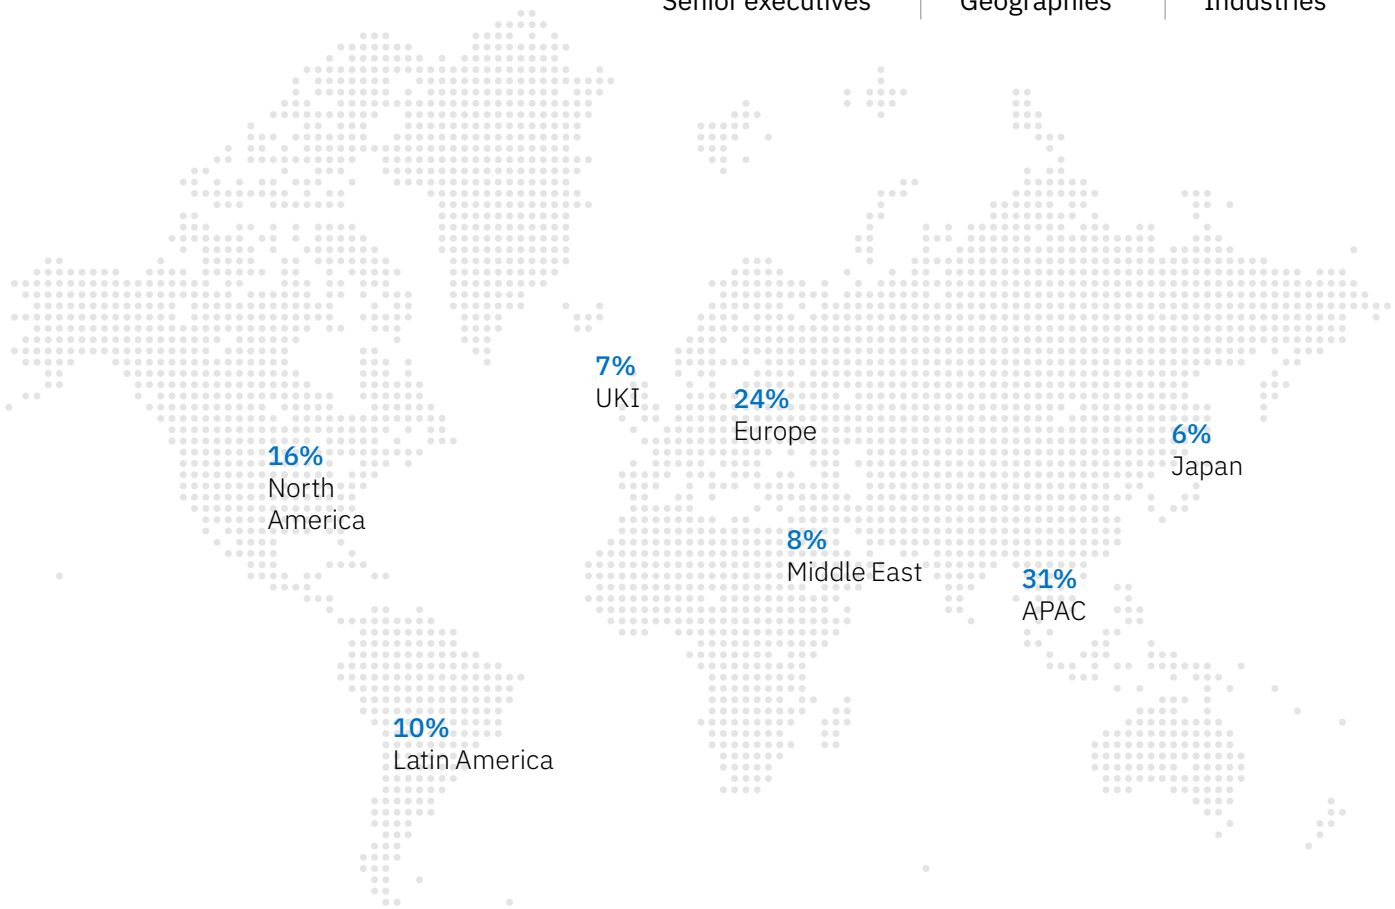
The IBM Institute for Business Value conducted a global survey from January to April 2026 to examine how organizations are managing the financial, operational, and governance challenges associated with scaling AI. The study draws on responses from 2,000 senior executives responsible for IT, technology, or AI-related decision making within their organizations. Respondents represent 33 geographies and 19 industries, reflecting a broad range of operating environments and stages of AI adoption.

The sample is weighted toward large, complex organizations. The average annual revenue —or budget for governmental organizations—is USD14.4 billion. Organizations report an average workforce size of approximately 33,500 employees, with the full range spanning 250 to 800,000 employees. This diversity in scale and industry composition enables robust comparison across organizational contexts while maintaining a strong focus on enterprises actively investing in AI and digital transformation.

2,000+
Senior executives

33
Geographies

19
Industries



**Percentages may not sum to 100% due to rounding.*

Analytical approach

The analysis applies a two-stage framework to identify organizations that have built the structural capabilities required to scale AI effectively. First, we segmented organizations based on preparedness and efficiency. Preparedness reflects whether architectural optionality—particularly cloud workload portability—has been deliberately designed into migration and operating decisions. Efficiency captures the maturity of AI financial management, including real-time spend visibility, cost allocation, ROI tracking, and enterprise-wide portfolio governance for AI investments. Threshold-based classification identified organizations that pair operational flexibility with disciplined AI investment practices.

We then assessed governance maturity through a model-based analysis. Using structural equation modeling, we treated AI agent governance as a latent capability reflected through seven observable design features, including ownership, access controls, activity monitoring, override mechanisms, pre-deployment risk assessment, rollback procedures, and formal governance frameworks. This analysis generated governance maturity scores used to classify organizations along a continuum from reactive to engineered and adaptive control.

The final segment represents the intersection of these two layers. Organizations in this group demonstrate both structural readiness (through preparedness and efficiency) and high governance maturity. This approach isolates enterprises that combine operational adaptability, financial discipline, and designed-in control—enabling comparative analysis of AI scale, resilience, and performance outcomes.

Unless otherwise noted, relationships described in this report reflect observed associations in survey and modeled data and should not be interpreted as proof of direct causality.

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 ibv@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.

Related reports

The 2026 CEO Study.
Rewiring the C-suite: The fast track to 2030.
IBM Institute for Business Value.
May 2026.
ibm.biz/ceo-2026

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

The 2025 Chief Data Officer Study.
The AI multiplier effect: Accelerate growth with decision-ready data.
IBM Institute for Business Value.
November 2025.
ibm.biz/2025-cdo

How IBM can help

IBM has been providing expertise to help organizations win in the marketplace for more than a century. Clients can realize the potential of AI and better navigate a changing world using IBM's deep industry, functional, and technical expertise; enterprise-grade technology solutions; and science-based research innovations.

IBM is also a global leader in quantum computing, ready to aid clients as they embark on their own quantum journeys.

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

For more information about AI solutions from IBM Software, visit ibm.com/watson.

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

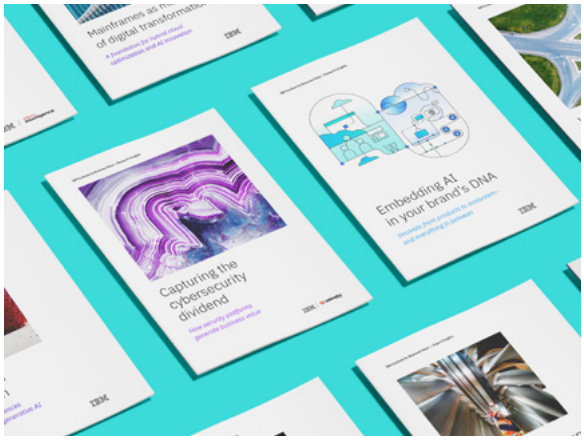
For more information about quantum computing from IBM, visit ibm.com/quantum.

Notes and sources

1 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

2 Higgins, Heather, Petra Floorizone, and Veena Pureswaran. The 2025 Quantum Readiness Index. *Quantum is coming: 5 realities shaping the race to advantage.* IBM Institute for Business Value. December 2025. https://ibm.biz/quantum-readiness

3 "Banco do Brasil leads in innovation with responsible AI and ethical oversight." IBM case study. Accessed May 11, 2026. https://www.ibm.com/case-studies/bb-ey



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 | June 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.

16ddce7b4954875d-USEN-01

