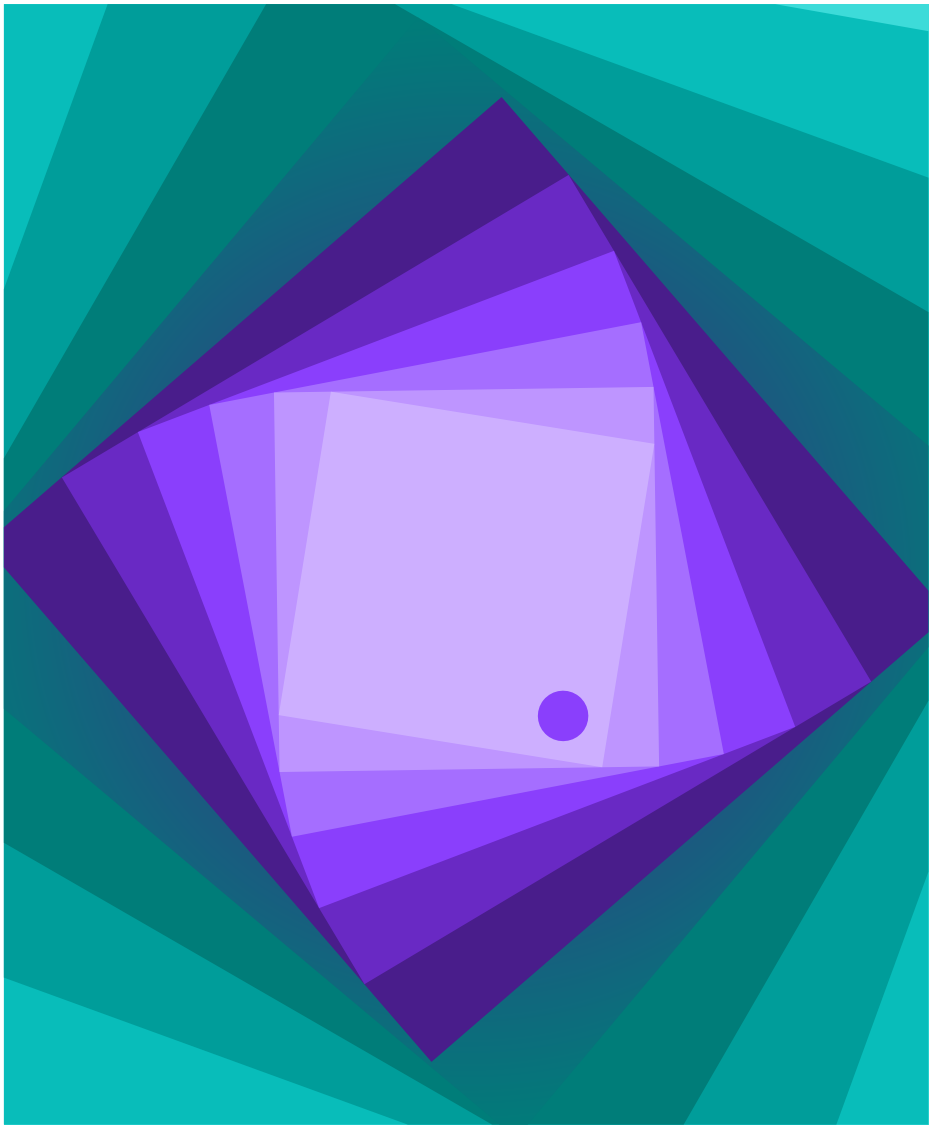

AI in motion

Orchestrating AI at scale for sovereignty and resilience

IBM | مؤسسة دبي للمستقبل DUBAI FUTURE FOUNDATION

Contents

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 using IBM's deep industry, functional, and technical expertise; enterprise-grade technology solutions; and science-based research innovations.

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-intelligence3

Foreword

AI is scaling faster than the systems built to govern it. That gap, between deployment and control, is where value is lost and risk accumulates. This report puts numbers to what many leaders already sense but cannot yet see clearly.

The findings are striking. Nearly seven in ten executives say they don't have full visibility into the AI their teams are using. Only one in four organizations knows where their AI assets are being deployed. And fewer than half the systems in use today are delivering the outcomes expected of them. These are not edge cases. They are the norm.

Dubai has long understood that the ability to shape the future depends on the ability to govern it. That conviction sits behind our investment in institutions, in regulation, and in the infrastructure of foresight. As AI becomes as foundational as any physical infrastructure, the same logic applies. Sovereignty — knowing where your data lives, controlling how your systems operate, retaining the ability to act without external dependence — is not a constraint on ambition. It is what makes ambition sustainable.

What this research shows is that governance, done well, is not a constraint on performance. It is a driver of it. Organizations that govern AI through a coordinated, visible layer — what this report calls orchestration — are scaling thirteen times faster, seeing six times greater productivity impact, and achieving twenty percent higher ROI on their AI investments. The case is no longer theoretical.

The challenge now is adoption. Only twelve percent of organizations have orchestration platforms in place today. The gap between knowing what good looks like and building it remains wide.

The gap will not close on its own. This research is part of a deliberate effort to understand what effective AI governance looks like at scale, and to share those findings with organizations navigating the same questions. We hope it serves as a practical resource for leaders who are ready to move from awareness to action.

H.E. Khalfan Belhoul

CEO
Dubai Future Foundation

Saad Toma

General Manager
IBM Middle East and Africa

Key takeaways

Governance without orchestration doesn't hold. It defines what should happen but can't enforce it.

AI orchestration is now one of the biggest dividers between AI leaders and everyone else.

Organizations using orchestration-led governance are 13 times more likely to be scaling AI. This shift turns AI from scattered experiments into a coordinated capability that compounds value.

AI sprawl is eroding value and control.

A $20 billion enterprise loses roughly $140 million a year to AI irregularities—half due to governance gaps. Without system-level visibility, organizations are building portfolios they can't control.

Compliance alone delivers diminishing returns—coordination delivers impact.

Holistic orchestration is associated with six times greater productivity impact than compliance only approaches.

Most companies are still governing AI with tools built for a different era.

Only 12% have orchestration platforms in place and only 18% maintain a complete AI inventory. The majority are operating complex, autonomous systems without the oversight required to steer them.

Introduction

The quiet rise of AI sprawl inside the enterprise

AI didn't arrive in the enterprise as a coherent system—it took hold in pieces. An assistant in HR. A model in marketing. An agent picking up operational tasks. Each delivered value on its own. Together, they've created something far larger: a sprawling digital workforce making thousands of decisions each day—sometimes aligned, often not.

The issue is no longer just efficiency—it's value. Small failures, misaligned decisions, or model drift don't stay contained; they compound across the enterprise, eroding performance, distorting outcomes, and quietly undermining the value AI was meant to create.

This is where most organizations now find themselves, just as expectations peak. Two-thirds expect to operate AI at scale by 2030, with average organizational investment set to rise 138%. Yet AI sprawl is already eroding value, governance failures can cost as much as the AI itself, and nearly a third of executives say their investments have been wasted.

Sovereignty pressures only raise the stakes. Ninety-three percent of executives now factor sovereignty into strategy—governing AI across national security rules, data residency laws, and independence requirements—even as agentic systems outpace review processes.

There is an alternative: orchestration-led AI governance—coordinating AI through platforms that turn fragmented deployments into a system that compounds value instead of diluting it.

This report, based on an IBM Institute for Business Value study of 1,006 senior executives in 20 geographies and 23 industries, breaks down the new AI governance landscape in five sections. In Part One we define orchestration-led governance and why timing matters. Part Two explores the often-hidden cost of AI complexity versus clarity. In Part Three, we spotlight orchestration best practices and demonstrate the positive business impact. Part Four addresses what partner relationships mean for governance. Part Five offers an Action Guide with specific how-to steps for any enterprise.

Figure 1

Benefits for organizations that lean into orchestration-led governance

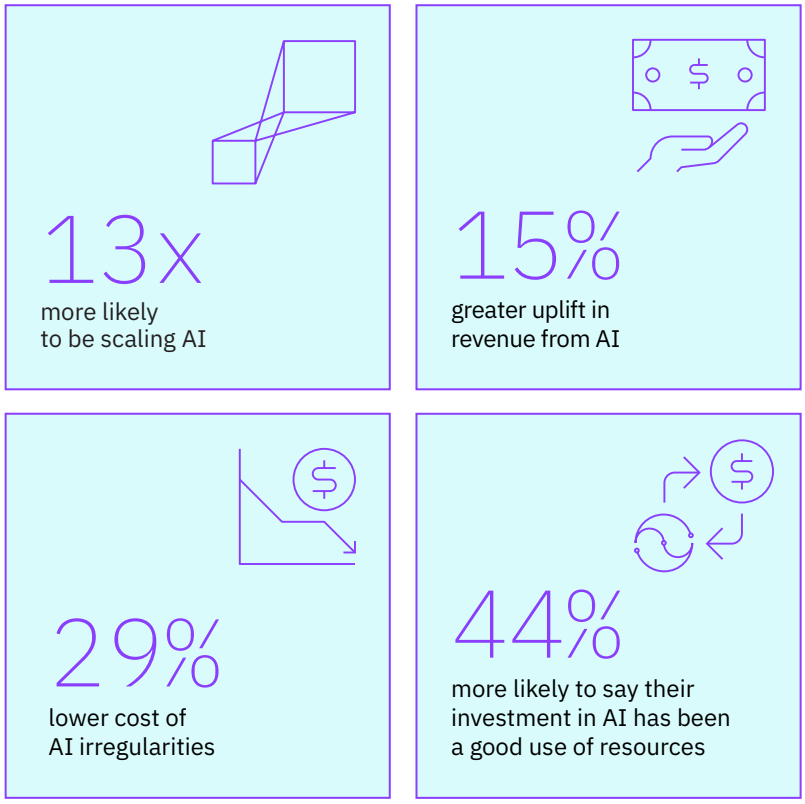

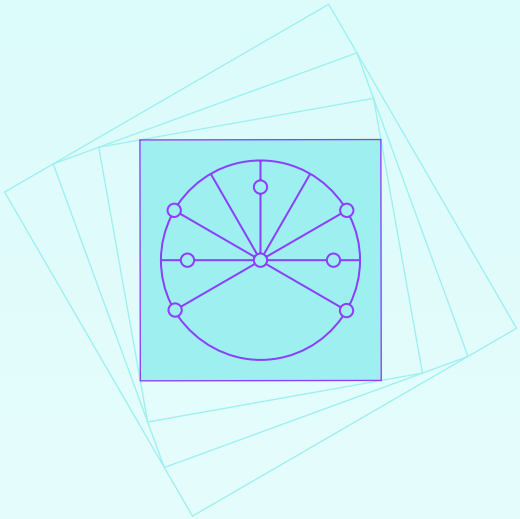

Part 1

What orchestration-led governance is and why it matters right now

For years, companies treated AI governance and orchestration as separate concerns. Governance set the rules—data use, accountability, risk. Orchestration ran the system—how models, data, and workflows operate.

That separation no longer works.

In the age of agentic AI—where systems act autonomously, collaborate, and make real-time decisions—policy alone can't control a fleet of AI agents. Orchestration does.

Governance without orchestration doesn't hold. It defines what should happen but can't enforce it. Orchestration without governance is worse—it scales quickly without guardrails, amplifying risk and weakening control over data, models, and decision rights. Nearly seven in 10 executives now admit they lack full visibility into the AI their teams are using—or where it operates.

As adoption accelerates, the definition of responsible AI is expanding to include security, ethics, financial oversight, and environmental impact—placing new strain on governance frameworks. Without transparency, AI investments risk becoming costly experiments rather than durable sources of value.

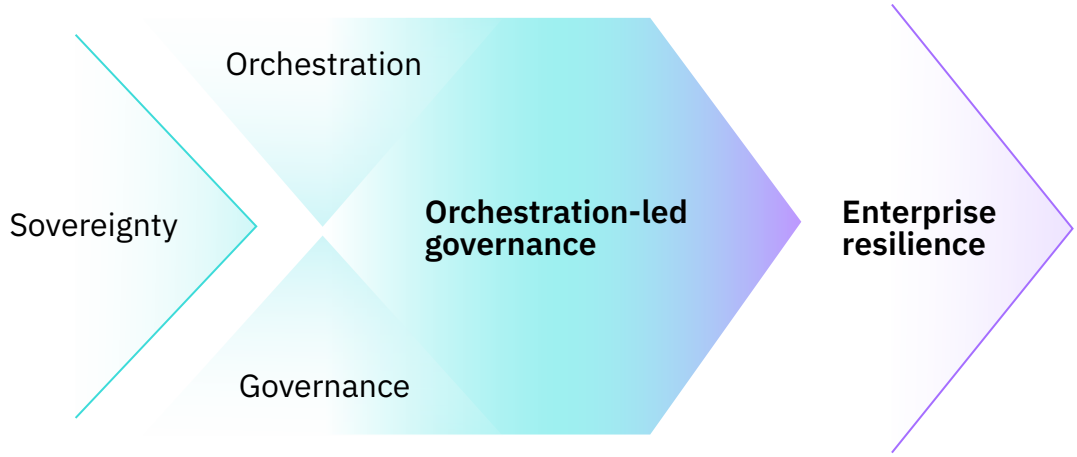

Putting control where AI actually runs

Orchestration-led governance changes where control actually lives. Instead of writing policies and hoping teams interpret them correctly, companies build guardrails into the systems running AI. Rules aren't separate from execution—they shape it in real time.

That shift gives leaders something they've been missing: visibility. Orchestration platforms create a live view across models, agents, data flows, and decision pathways—making it possible to catch drift early, enforce standards, and adjust as systems evolve. The payoff is significant. Organizations using a full orchestration approach see more than six times the productivity impact of those focused on compliance alone (see figure 2). Yet only 12% have these platforms in place today, leaving most enterprises scaling AI without a clear line of control.

As AI becomes more autonomous and embedded across the business, governance can't sit on the sidelines. It has to operate inside the system—within workflows, decisions, and day-to-day execution. Otherwise, it stays theoretical. Orchestration-led governance closes that gap, turning governance from a policy exercise into a working capability.

Figure 2

Extending governance through orchestration yields 6 times higher productivity gains from AI

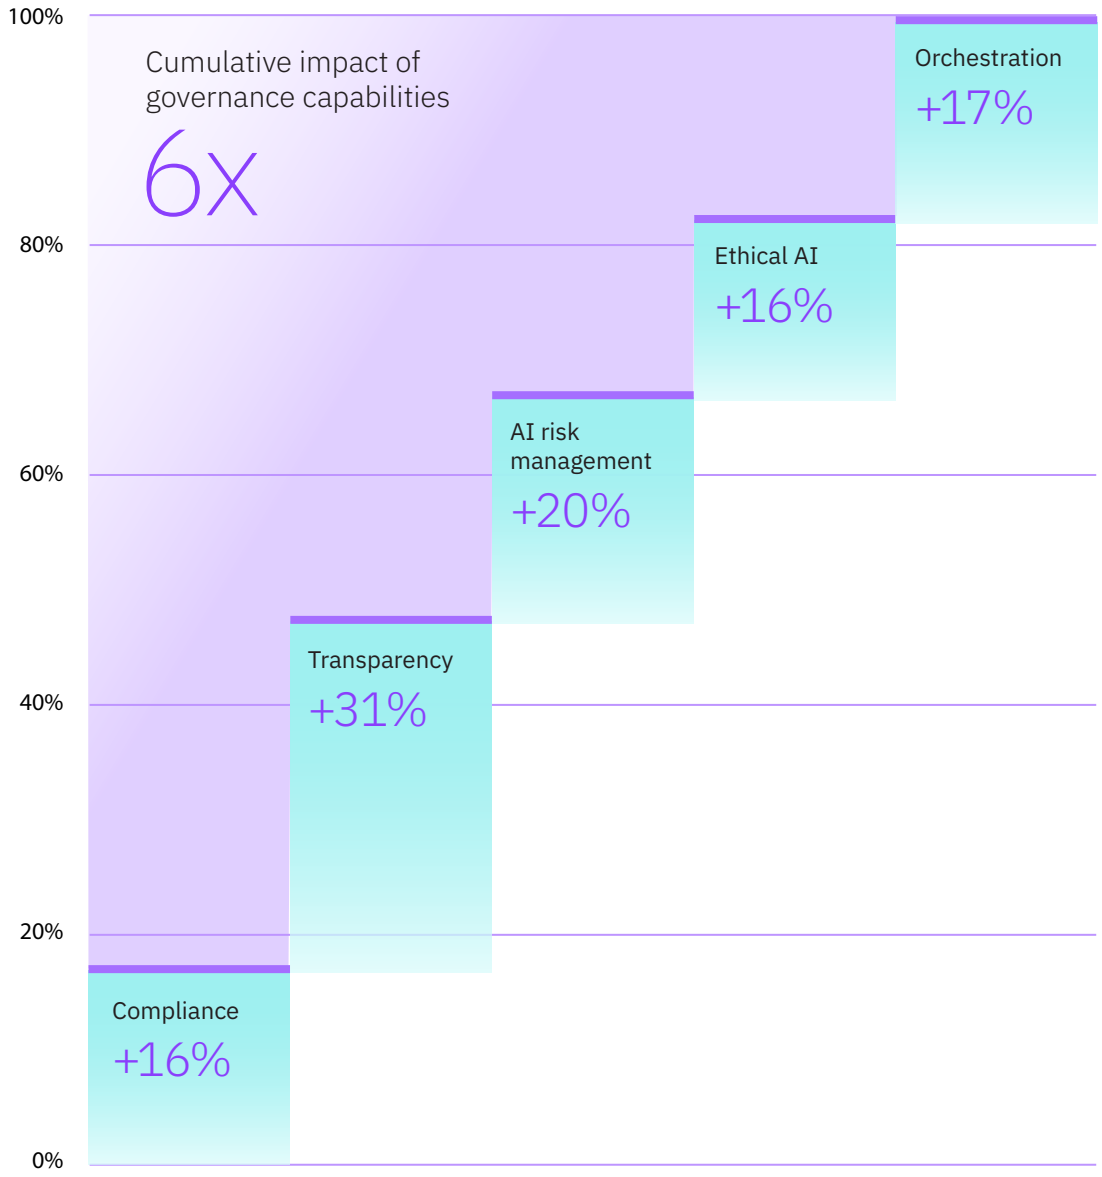

Part 2

The cost of complexity

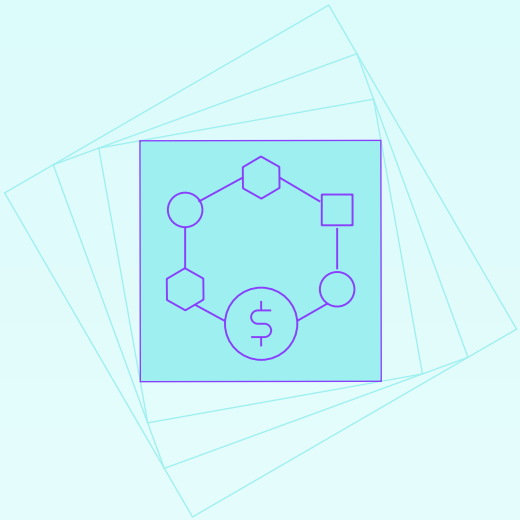

Executives are starting to confront an uncomfortable truth: their AI landscapes are expanding faster than their ability to control them. Complexity has become an adversary in its own right, with half of executives naming it as one of their biggest obstacles to effective AI governance.

That lack of visibility carries a steep price. In a $20-billion enterprise, the hidden fallout from AI irregularities—misfires, bias, duplication, uncoordinated deployments—adds up to roughly $140 million a year. Even more striking: half of that loss is directly attributable to weak governance. Not random errors. Not unavoidable glitches. Preventable oversight gaps that over time don't just erode value but expose how fragile the AI estate has become.

And the challenge is only accelerating. By 2030, organizations expect to manage nearly four times as many AI assets as they do today—yet fewer than half of the systems they currently have deliver expected business outcomes. It's a widening gap between what enterprises are building and what those assets actually contribute.

By 2030, organizations expect to manage nearly four times as many AI assets as they do today—yet fewer than half of the systems they currently have deliver expected business outcomes.

The cost of governing AI projects vs. AI portfolios

This growing chasm is already changing the organizational calculus: recent IBM IBV research shows that the CEO is now taking a direct role in ensuring effective AI governance in 63% of organizations.[1] Even so, responsibility now stretches across the C-suite: the Chief AI Officer advancing transformation, the CIO maintaining the underlying architecture, the CISO safeguarding systems, the CRO managing exposure. But when ownership is distributed without clearly defined decision rights, accountability thins out.

And even this wider executive involvement is insufficient. The practical limits of governance are revealed in everyday use—where systems behave unpredictably, where edge cases accumulate, where intent and outcome begin to separate. The people closest to these moments are rarely in the executive suite. Unless their observations are drawn back into how AI is governed, organizations risk building structures that appear sound at the top but fail to hold where it matters most.

The visibility problem runs deeper than most leaders realize. Responsible AI demands line of sight not just into compliance, but into cost, ROI, risk, workforce impact, and environmental footprint. Yet only 27% of organizations know where their AI is in use, and just 18% maintain a current inventory. This isn't oversight—it's blindness at scale. Fragmentation doesn't come from bad decisions; it builds one reasonable choice at a time until the enterprise can no longer see what it has created.

Figure 3

High-sprawl organizations pay for it in performance vs. those with lower sprawl

This is what happens when companies govern AI projects instead of portfolios—and the cost is measurable. Organizations with high sprawl and low visibility see 22% lower revenue impact, 19% lower operating profit, and 17% lower productivity gains, along with weaker cost reductions. Their success scaling agentic AI is also 11% lower—evidence that fragmentation increasingly constrains value as autonomy grows.

Real impact now depends on something governance frameworks weren't built to do: coordinate a fast-growing portfolio of AI assets so they operate as a system. That means managing more than compliance—navigating sovereignty constraints, geopolitical exposure, supply chain risk, and cross-border data rules—while ensuring machine-speed decisions still reflect enterprise priorities. Seven in 10 executives say inadequate governance is already limiting their ability to scale AI. What they're sensing isn't just a governance gap—it's a resilience failure. When systems can't be coordinated under pressure, scale becomes a liability.

An enterprise can be fully compliant and still fully fragmented—every rule followed, no system coherence.

Compliance is the floor. Too often, it's mistaken for the ceiling.

"Governance is critical—but only if it is implemented at the beginning. If you add governance at the end, it becomes a bottleneck."

Majid Sultan AlMheiri, Chief AI Officer and Director of IT, Dubai Health Authority

Perspective

Sovereignty makes it all urgent

Sovereignty is forcing a rethink of what governance must include. Seventy percent of executives say they must constantly track where data resides, and 71% expect it to matter even more by 2026. Where data sits now determines what companies can build—and where they can sell it.

This isn't a checklist problem. Sovereignty shapes the operating model. AI systems must work within national boundaries while delivering global value, which means governance has to be built into how they run. Done right, sovereign AI doesn't limit innovation—it sets the conditions for scaling it safely.

Agentic AI raises the stakes. Agents don't pause for review; they act and cross jurisdictions at machine speed. The 16% success rate for agentic AI isn't a technology issue—it's a governance gap. Oversight has to move with the system, embedded in architecture and enforced in real time. That's what resilience looks like: the ability to keep operating through disruption, not stopping every time conditions change.

> "The real challenge is balancing speed and sovereignty. You cannot be so controlled that you fall behind, and you cannot move fast into sensitive areas without proper testing."
>
> **Majid Sultan AlMheiri**, Chief AI Officer and Director of IT, Dubai Health Authority

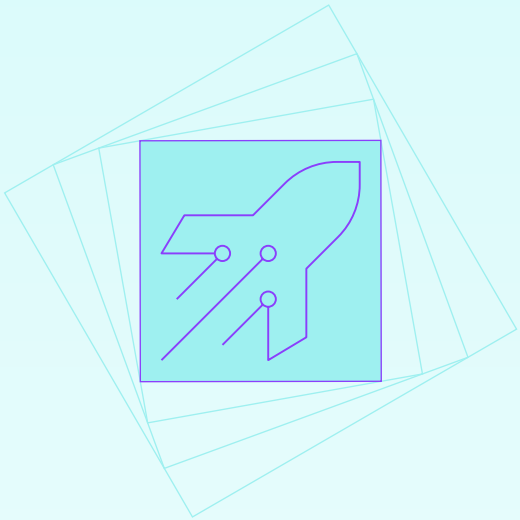

Part 3

Taking control and boosting results with orchestration

The companies pulling ahead aren't winning because they bought more models or hired more data scientists. They're winning because they've solved the coordination problem, turning to governance platforms that don't just publish rules from the sidelines, but sit at the center of the action, directing how models, assistants, and agents interact.

Organizations that adopt orchestration-led governance—built on a strong AI operating model—are on a different trajectory. They reduce losses from AI irregularities by 29% and achieve 20% higher ROI. They also deliver stronger gains in productivity and revenue. In addition, they are more than twice as likely to have full visibility into AI assets, 169% more likely to maintain transparent documentation, and 132% more likely to protect data through anonymization, impact assessments, and strict access controls.

The advantage compounds. These aren't isolated wins—they reflect systems that perform more reliably over time. The same mechanisms that reduce risk also build trust: twice as many organizations report higher employee confidence, and 43% more see increased trust with customers, regulators, and partners.

Orchestration functions like air traffic control for AI—routing capabilities to where they're needed, enforcing sovereignty constraints in real time, and embedding visibility and security into the infrastructure. Governance doesn't sit apart from execution; it travels with the system. Leaders aren't managing disconnected tools—they're directing a portfolio that grows more valuable as it scales.

Figure 4

Organizations with an orchestration-led governance approach achieve higher effectiveness of overall AI governance

Governance that travels changes how decisions are made

Investments can be reallocated dynamically—scaling what works, retiring what doesn't. Organizations can also tap external models and partners without surrendering control or locking into a single vendor.

The result is resilience. With greater visibility and coordination, issues surface earlier and responses expand. Workloads shift, exposure narrows, and operations continue without disruption.

This is the difference: not just compliance, but control—and resilience—that holds under pressure.

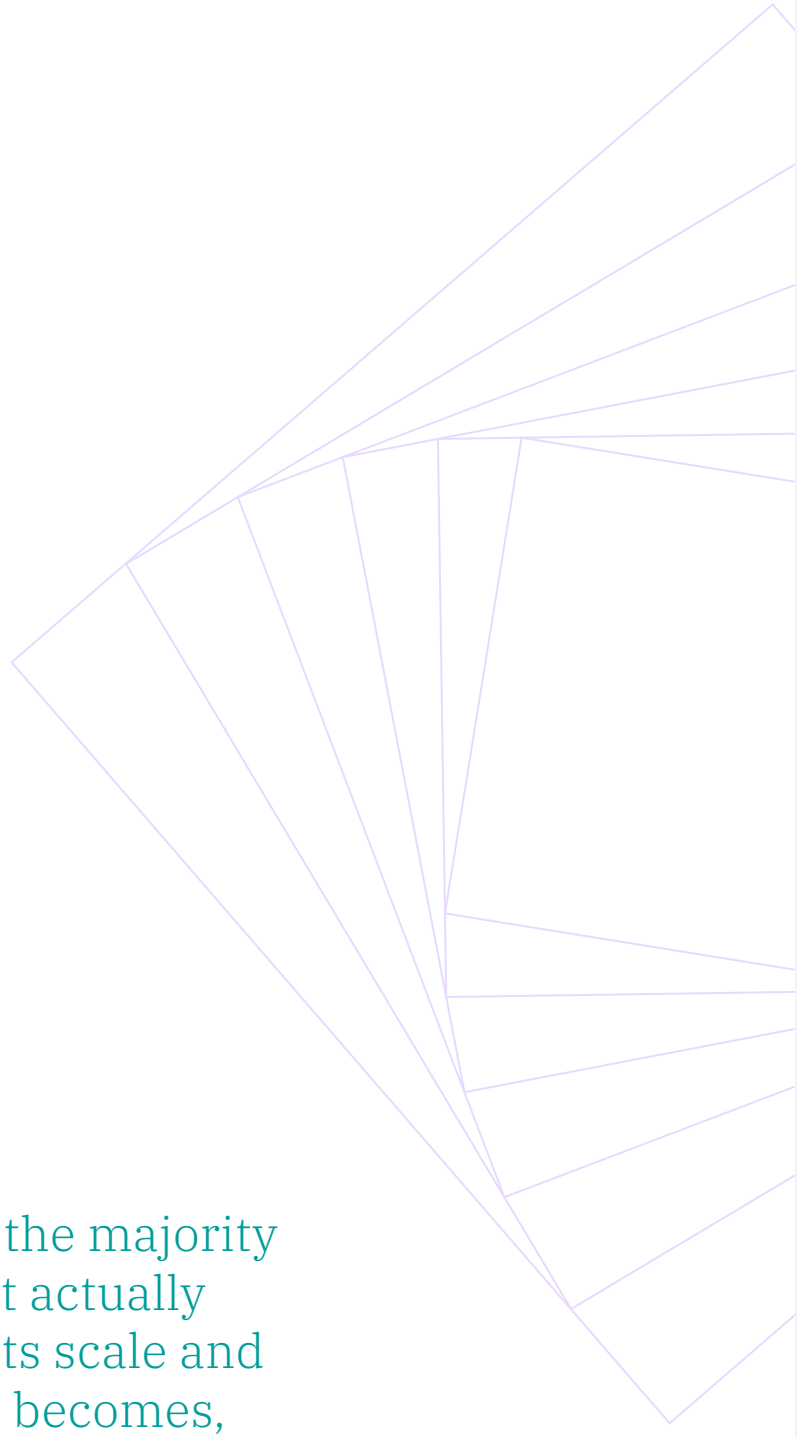

"If you fast forward to 2030, the majority of governance work may not actually be done by humans due to its scale and complexity. So the question becomes, how do you build automated governance tools that humans can oversee—and how long do you have to figure that out?"

Kristie Chon Flynn, Data Protection Officer, Google

Perspective

Orchestration starts with open, hybrid, secure architecture

Orchestration only works if the architecture can carry it. That means open, modular, hybrid-by-design systems that can absorb new models and agents without breaking—with security and privacy built into how those systems run. As agentic systems cross boundaries and expand the attack surface, security has to be enforced in the architecture itself, not layered on afterward. Sixty-two percent of executives now say an open, secure-by-design architecture is essential. Without transparency and interoperability, AI systems can't share context, explain decisions, or collaborate at scale.

Open architecture gives agents flexibility—with discipline. They can select the best model for a task and switch as conditions change. Governance moves inside the system, shaping how assets interact, synchronize, and evolve. That's what turns a collection of tools into a system that performs.

Control hinges on identity. Each agent requires a nonhuman identity, with clear ownership, delegated authority, and least-privilege access enforced at runtime. Full observability—and the ability to revoke access instantly—is critical. Orchestration becomes continuous control over what each agent can do, when, and on whose behalf—anchored in enterprise IAM.

This foundation prepares organizations for what's next: multi-model, multi-agent ecosystems where capabilities shift constantly. Governance sets the rules; orchestration enforces them in real time.

The result is alignment—an AI estate that evolves with the business, absorbs complexity, and keeps autonomous systems accountable as they scale.

Full observability—and the ability to revoke access instantly—is critical. Orchestration becomes continuous control over what each agent can do, when, and on whose behalf.

Figure 5

What executives point to as the markers of an effective AI platform

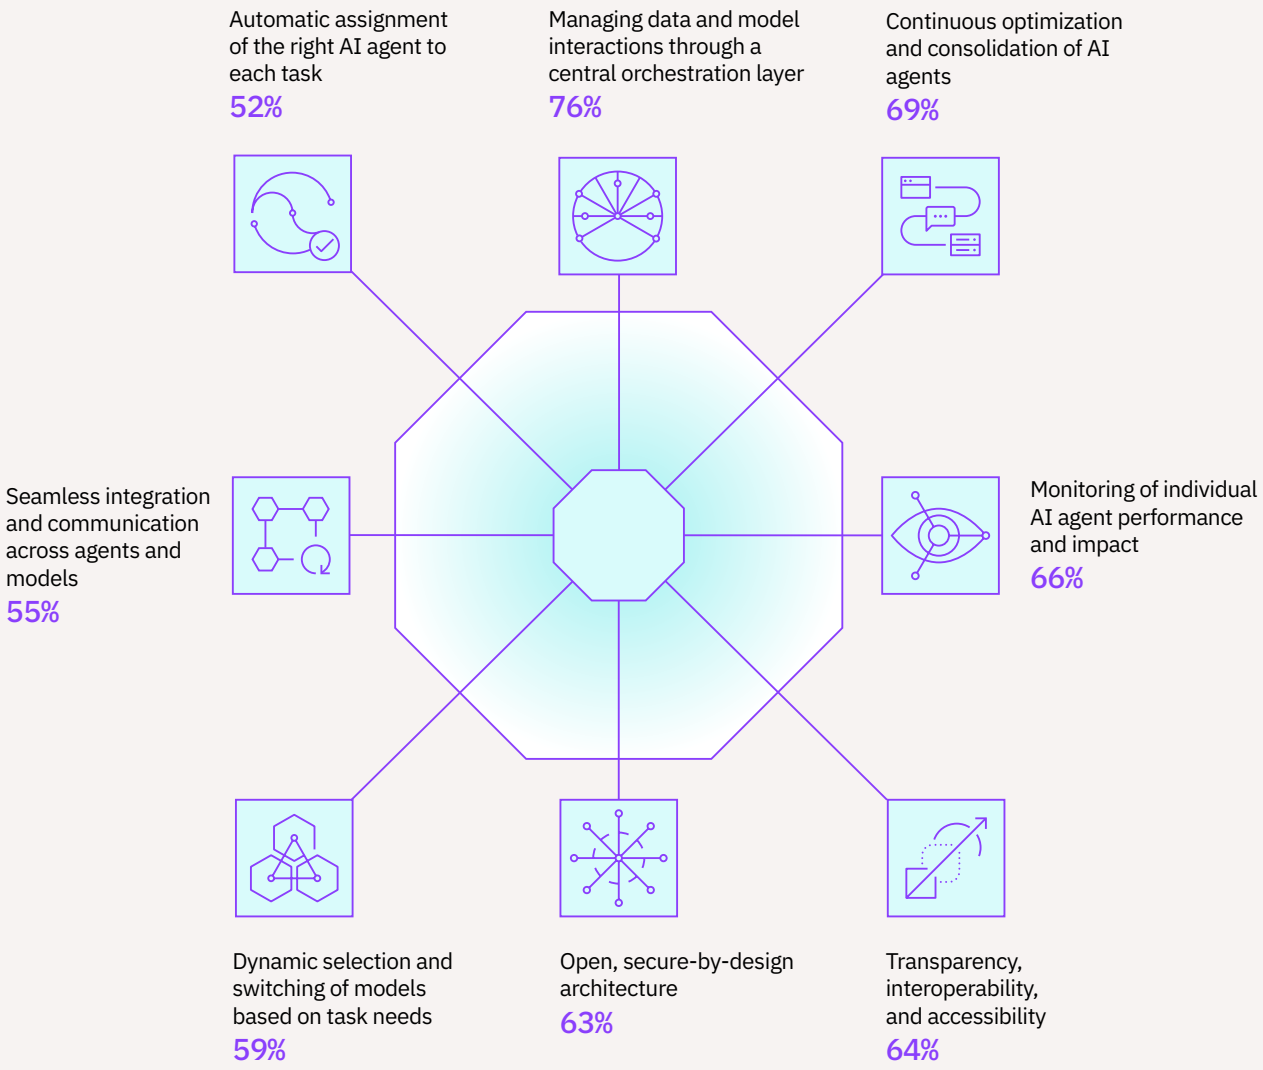

Case study

How AI orchestration turns raw data into real-time sports intelligence[1]

Fight night moves fast. Analysts need compelling statistics, broadcasters need context, and fans expect real-time insights that deepen the drama unfolding in the arena. But behind the scenes, producing those insights often requires combing through decades of fight data—work that traditionally relied on manual research and institutional memory.

UFC, a global mixed martial arts organization that broadcasts to over 950 million households across more than 170 countries, faced exactly that problem. Its editorial and research teams relied on deep institutional knowledge and manual analysis to comb through vast archives of fight statistics and athlete data. The approach worked—but it was slow. As fan expectations grew for real-time commentary and data-driven storytelling across broadcast, digital, and social platforms, the organization needed a faster way to surface meaningful insights.

The answer was an AI-driven insights platform designed to analyze decades of historical fight data and generate contextual narratives about fighters and matchups. Instead of manually digging through databases, editors and researchers can now ask natural-language questions and instantly retrieve analysis grounded in verified performance data.

At the center of the system is AI orchestration—a layer that routes each question to the most appropriate capability, whether structured data queries, generative models, or retrieval systems that pull relevant context from historical archives. Rather than relying on a single AI model, the orchestrated approach coordinates multiple tools working together, enabling faster and more precise responses.

Equally important is governance. Each generated insight includes explainability features and quality scoring, helping editorial teams understand how the information was produced and assess its reliability before sharing it with audiences. In high-visibility environments like live sports broadcasting, that transparency helps maintain trust while allowing teams to move quickly.

The impact has been measurable. Early estimates suggest the platform has reduced the time needed to generate queries by about 40%, freeing editorial teams to focus less on data wrangling and more on storytelling. At the same time, the volume of insights available has nearly tripled, producing hundreds of contextual storylines about fighters, matchups, and event history.

For fans, the change appears as richer broadcasts—real-time analysis, deeper fighter comparisons, and narratives grounded in decades of performance data. For the organization behind it, the initiative demonstrates how AI orchestration combined with strong governance can transform complex data into real-time intelligence, enabling faster decisions, more engaging content, and a new level of storytelling at scale.

3x
estimated increase in insight volume

40%
estimated reduction in query generation time

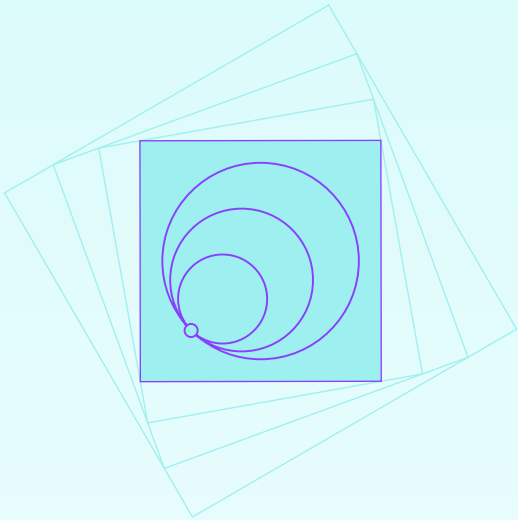

Part 4

Extending governance beyond the enterprise

AI no longer stays neatly inside company walls. It already works across suppliers, partners, platforms, and customers—whether enterprises are prepared for it or not.

As AI systems interconnect across organizational boundaries, resilience becomes a shared requirement. Yet only 48% say their partners are equipped for integrated, future ready AI operations, and just 45% extend governance beyond their own perimeter.

That gap used to be an inconvenience. Now it's a liability. Suppliers, partners, and customers are no longer just exchanging files—they're interacting with enterprise AI systems directly. Data doesn't stop at the firewall. AI agents are beginning to cooperate across entire supply chains and service networks. As these systems intermingle, governance can't remain a single player sport.

But for most companies, it still is. A majority—56%—aren't sharing AI capabilities or assets with their ecosystem at all, effectively capping the potential of multi-agent, cross-enterprise intelligence before it starts.

A majority of enterprises—56%—aren't sharing AI capabilities or assets with their ecosystem at all, effectively capping the potential of multi-agent, cross-enterprise intelligence before it starts.

Orchestration-led governance changes this completely. Orchestration ensures that failures, breaches, or misalignment in one part of the ecosystem don't cascade unchecked across the rest. Once you've orchestrated AI inside the business, you can project that structure outward, opening new frontiers. Guardrails, permissions, identity, and oversight travel with the AI as it interacts with partners. The enterprise doesn't lose visibility the moment data or agents leave home base; the governance layer goes with them.

This approach is becoming non negotiable as sovereignty rules fragment across regions. Companies that have mastered orchestration can operate globally without losing control of how their AI behaves. Those that haven't are left chasing every new requirement and vendor dependency as a one off—permanently stuck in defensive mode.

"I see governance as a differentiator in 2030, because companies can't choose between doing things quickly or doing them responsibly. It has to be both."

Kristie Chon Flynn, Data Protection Officer, Google

Perspective

Rethinking governance as AI models multiply[2]

At many companies, AI governance shows up late in the process—after models are built, data is sourced, and product teams are ready to ship. The result is predictable: delays, compliance reviews, and last-minute questions about where data came from and whether models can legally or ethically be deployed.

IBM decided to tackle that problem upstream.

Inside IBM's Office of Privacy and Responsible Technology, leaders realized the company's governance capabilities—privacy oversight, data provenance checks, and AI risk management—were strong but fragmented. Different systems handled data intake, regulatory review, and model oversight. As the number of AI initiatives expanded across the company, those parallel processes were starting to slow things down.

The response was a centralized governance framework designed to track the full lifecycle of data and AI models—from sourcing and approval through development and reuse. Instead of governance acting as a checkpoint at the end of development, it became part of the infrastructure supporting it.

The shift produced tangible operational gains. Processing time for data clearance requests dropped 58% for third-party datasets and 62% for internally sourced data. At the same time, more than 2,000 datasets and models were cleared for reuse, giving developers faster access to trusted data and reducing redundant approvals.

As AI systems proliferate across the enterprise, centralized governance is becoming a prerequisite for innovation—allowing organizations to move faster not by bypassing oversight, but by building it directly into how AI gets developed in the first place.

More than 2,000 datasets and models were cleared for reuse, giving developers faster access to trusted data and reducing redundant approvals.

Part 5

A different definition of AI leadership:

An action guide

The future won't belong to enterprises with the most AI, but to those that make it work together—across business units, borders, and systems. Enterprises should stop managing AI as disconnected projects and instead, design for orchestration. That means shifting governance from policy to platform—from periodic oversight to continuous coordination—so AI operates as a coherent portfolio, not isolated experiments.

By 2030, AI will be scaled across most organizations. Spend will rise. Returns will be demanded. Nearly a third of executives already suspect their AI investments have been wasted. That suspicion will either harden—or force a rethink of governance as the capability that makes AI deliver.

"Resilience in AI is no longer just about uptime or disaster recovery. It's about the resilience of the model itself—whether I can switch models, control them, and truly understand their behavior."

Majid Sultan AlMheiri, Chief AI Officer and Director of IT, Dubai Health Authority

To unlock the value of multi-model and multi-agent AI, organizations need an architectural foundation that connects systems, coordinates intelligence, and enforces responsible use. Five priorities matter.

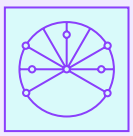

Add the orchestration layer that allows AI to operate as a unified system.

This layer connects platforms, applications, and agents into a coherent environment—replacing fragmented tools with standardized data flows and coordinated actions. By building this into the architecture, not point-to-point integrations, organizations create the conditions for AI to scale and reinforce itself.

Build a governance layer that operates inside the system, not around it.

Embed governance into how data is accessed, models are used, and decisions flow. When built into the architecture, it enforces policy automatically, ensures transparency and auditability, and keeps autonomous systems aligned with responsible behavior.

Maintain strategic control over the enterprise AI stack.

Even with external models and platforms, retain control of the core—data, governance, orchestration logic, and operating rules. Manage data and model locality across jurisdictions, make dependencies explicit, and keep components interchangeable as conditions shift. Control isn't isolation—it's the ability to adapt without losing oversight.

Turn governance principles into enforceable platform requirements.

Principles are too abstract to govern a live AI environment on their own. Translate them into risk-based requirements that can be built into contracts, engineering specs, and vendor criteria. Tie them to decision points across the AI lifecycle and embed them in the orchestration platform—so governance operates as system behavior, not policy language.

Prepare the workforce for the way orchestrated AI changes how work gets done.

As AI begins to coordinate decisions across teams, employees need clarity on where judgment still matters, how to interpret outputs, and how workflows will shift. This isn't just adoption—it requires reworking decision rights, processes, and escalation paths, backed by clear ownership and accountability. It's also a governance input: the people closest to how AI is used—where outputs are accepted, challenged, or overridden—offer the clearest view of system behavior. That insight should be built into how AI is governed.

About the authors

Saeed Al Falasi

Executive Director
Dubai Future Foundation and Director
Dubai Centre for Artificial Intelligence (DCAI)
linkedin.com/in/saeedalfalasi/

The DCAI, an initiative of the DFF, is overseeing the implementation and selection of CAIOs across the Dubai government. With more than 20 years of experience leading digital transformation, technology, and business development, Saeed leads key DFF initiatives such as Area 2071, fostering an ecosystem where innovative companies and government collaborate to cocreate the future.

Shukri Eid

General Manager
Gulf, Levant & Pakistan
IBM Technology
linkedin.com/in/shukriatari/

With 25+ years in the technology industry, Shukri has led regional and global teams across software, hardware, and consulting—helping public and private sector organizations navigate transformation and unlock opportunity through digital innovation. As VP & GM at IBM for the Gulf, Levant, and Pakistan, Shukri leads market transformation efforts and client partnerships that position AI and advanced technologies as forces for growth, resilience, and positive societal impact.

Lula Mohanty

Managing Partner
Middle East and Africa
IBM Consulting
linkedin.com/in/lula-mohanty/

With more than two decades of industry experience spanning strategy consulting, technology implementation, and leadership, Lula has deep knowledge of artificial intelligence, hybrid cloud, and enterprise transformation. She previously played key leadership roles in IBM India and Asia-Pacific, where she led consulting teams responsible for delivering large-scale digital innovation programs to clients in banking,telecom, retail, and the public sector.

Irfan Verjee

Transformation Strategy Leader
Middle East and Africa
IBM Technology
linkedin.com/in/irfanverjee/

Irfan is a technology strategy leader with more than 20 years of experience building global teams and scaling emerging technologies. As IBM's Country Transformation Leader for UAE, Qatar, and Kuwait, he focuses on national priorities—advancing economic growth through AI, data, sustainability, and quantum by integrating technology, research, and ecosystem.

Jacob Dencik

Research Director
IBM Institute for Business Value
jacob.dencik@be.ibm.com
linkedin.com/in/jacob-dencik-126861/

Jacob is responsible for leading research on topics related to AI, innovation, security, economics, and sustainability. He has advised companies around the world on strategy, technology, and operations. He has also advised governments as an expert and economist on competitiveness, sector/cluster analysis, and innovation.

Contributors

We would like to thank the following individuals for their significant contributions to these materials.

IBM Insitute for Business Value:
Rachna Handa, Kristine Rodriguez, Heba Nashaat, Angela Finley

IBM:
Sherry Italia, Phaedra Boinodiris, Helle Valentin, Michael Mojica, Ritika Gunnar, Andy Bentley

Dubai Future Foundation:
Hana Almheiri, Veronica Murguia, Maryan Saeed AlMarri

Research Methodology

To examine how organizations govern AI at scale—and to empirically assess how governance maturity relates to AI adoption, value realization, cost irregularities, and risk—the IBM IBV, in collaboration with Dubai Future Foundation conducted a global, cross sectional executive survey during Q4 2025 and Q1 2026. The survey focused on enterprise AI governance and orchestration. The study was designed to test both linear relationships and potential threshold effects, including whether a critical mass of governance capability accelerates AI value capture.

The survey collected responses from more than 1,006 senior executives responsible for AI strategy, delivery, or governance, including Chief AI Officers, CIOs, CTOs, CDOs, CISOs, Chief Risk Officers, General Counsel, and heads of AI governance. Respondents represent diverse industries, geographies, organizational sizes, and AI maturity levels, enabling robust comparison of governance models across the global enterprise landscape.

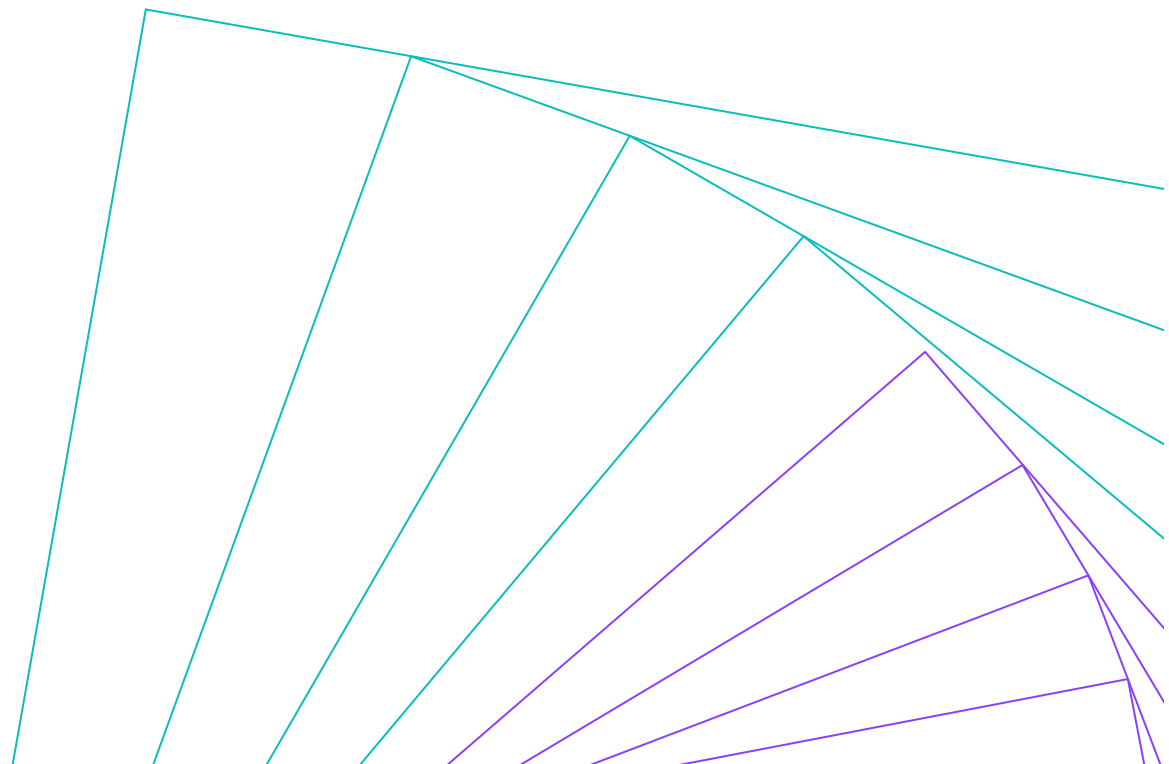

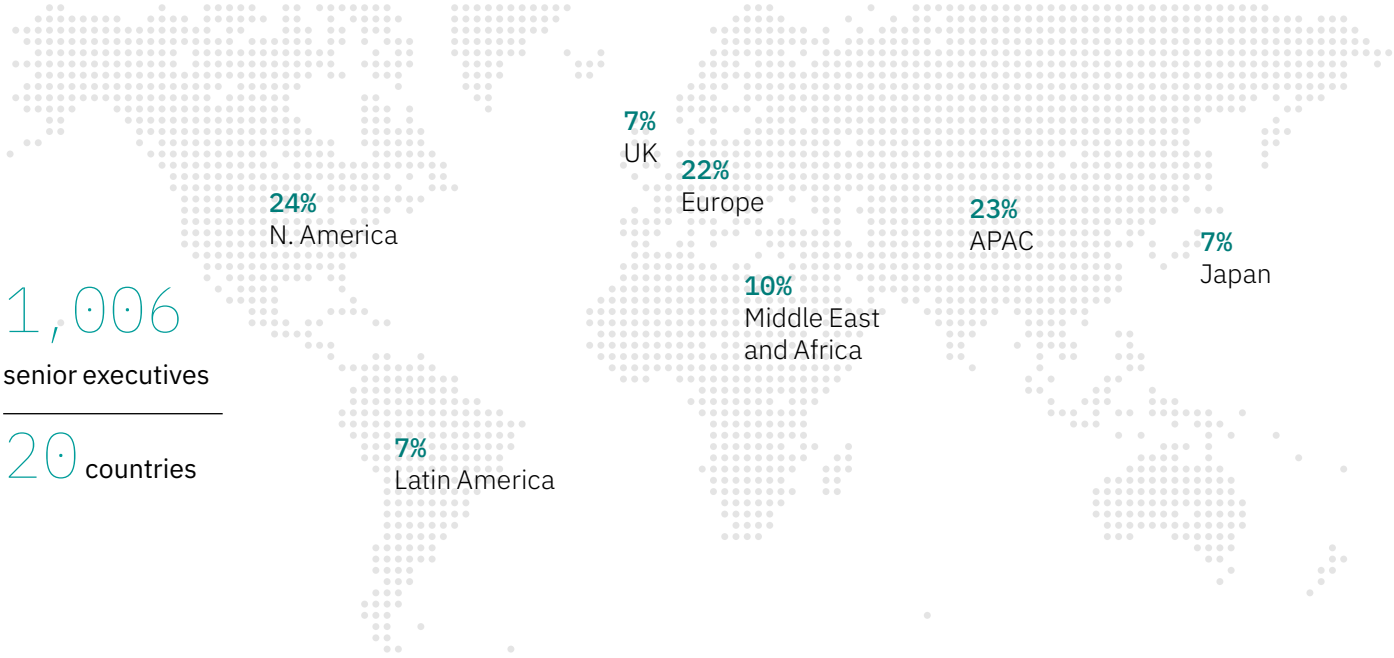

To assess the scale and complexity of enterprise AI portfolios, we analyzed reported counts of AI models, assistants, agents, workflows, governance platforms, and vendors in use in 2025, along with projected growth through 2030. These data quantify the pace of AI asset proliferation and contextualize governance challenges related to visibility, orchestration, and consolidation.

To investigate cost implications of governance maturity, we analyzed self reported AI irregularity costs defined as revenue losses from AI-related compliance, operational, security, or model failures over the prior 12 months. We used organizations without AI orchestration platforms as a comparison group to show how governance maturity changes AI related costs. Incremental and cumulative costs avoided were calculated across successive stages of orchestration adoption to identify nonlinear effects and thresholds associated with materially lower irregularity costs.

Regression analysis was conducted to isolate the incremental contribution of specific governance capabilities to AI attributable productivity. Independent variables included compliance coverage, transparency, proactive risk management, ethical AI practices, and enterprise wide orchestration. Cumulative impact analysis assessed how layered governance capabilities influence outcomes relative to compliance only models.

An AI governance and orchestration index was constructed by standardizing governance maturity measures and orchestration platform adoption metrics and combining them into a composite metric. Organizations were segmented into quartiles from lowest to highest maturity. We calculated average business outcomes attributed to AI for each quartile, and relative differences between the highest and lowest groups were used to estimate the impact of combined maturity on financial performance, productivity, AI ROI, and scaling success.

To evaluate fragmentation and loss of oversight, we developed an AI sprawl index as a composite measure of opaque and uncontrolled proliferation of AI assets. The index captures visibility into what AI assets exist, where they are deployed, and their security implications. We ranked and segmented organizations into quartiles, then examined differences in revenue, operating profit, productivity, cost reduction, and scaling success attributed to AI between high and low sprawl organizations.

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 https://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.

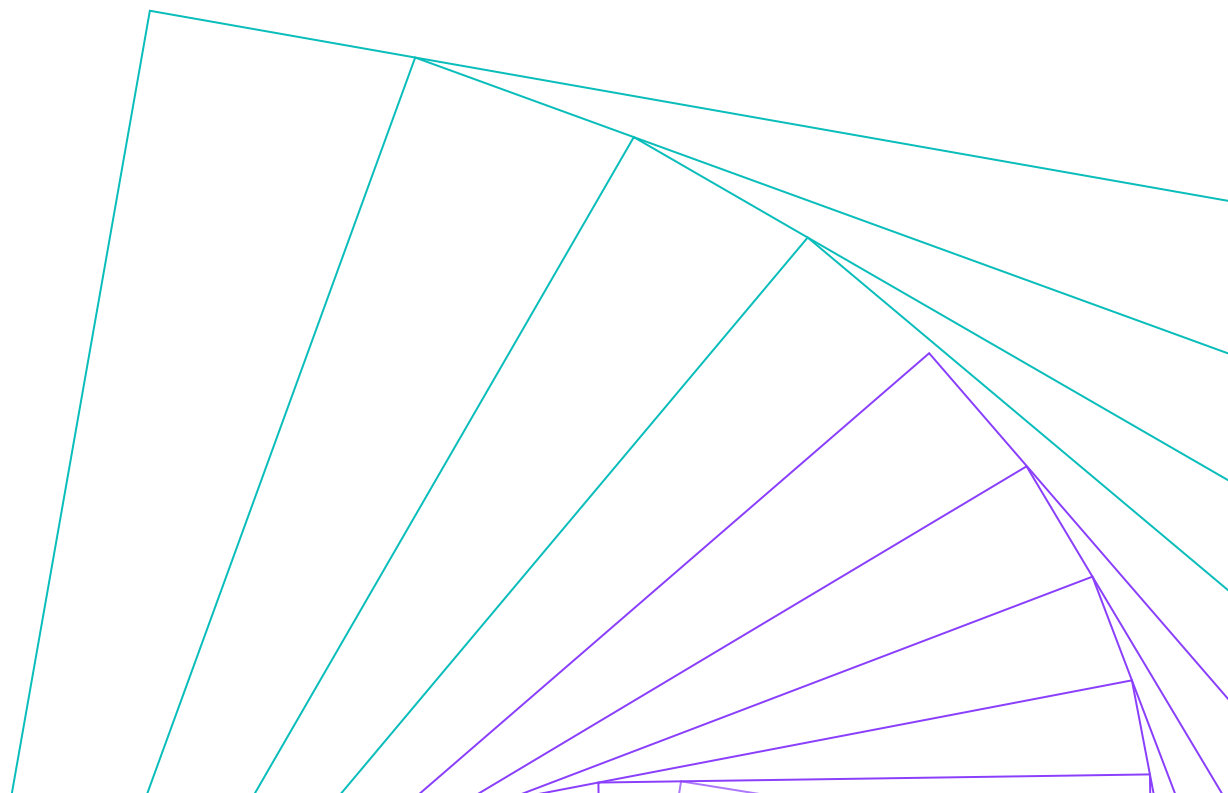

Related Reports

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

Orchestrating agentic AI for intelligent business operations
IBM Institute for Business Value. May 2025. https://ibm.co/agentic-process-automation

Solving the AI ROI puzzle
IBM Institute for Business Value. July 2025. https://ibm.biz/chief-ai-officer

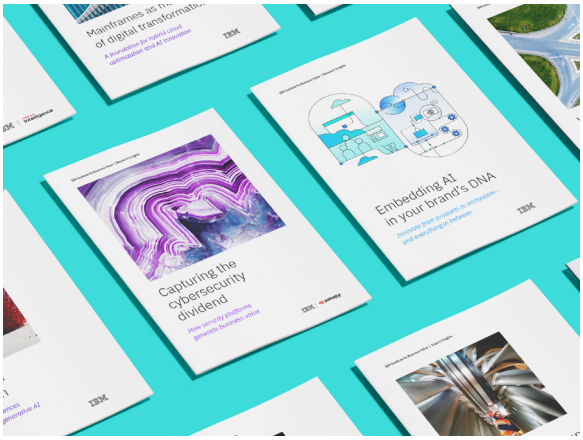

Subscribe to our IdeaWatch newsletter

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 second 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

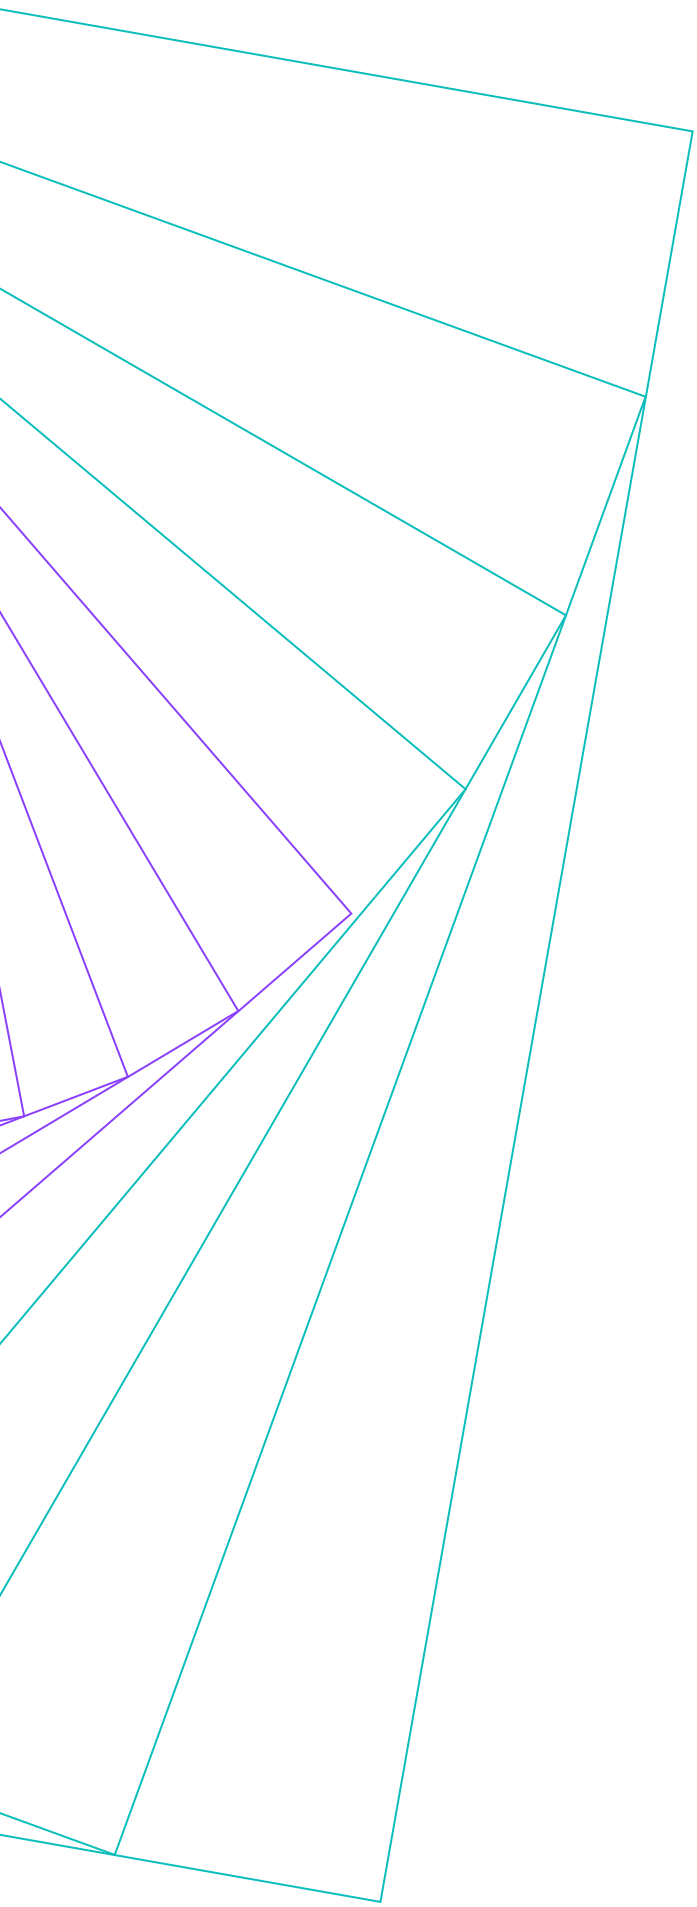

Notes and sources

1 “Boosting fan engagement with AI-powered insights.” IBM. Accessed April 13, 2026. https://www.ibm.com/case-studies/ufc

2 “Accelerating innovation through centralized AI governance.” IBM. Accessed April 13, 2026. https://www.ibm.com/case-studies/ibm-office-of-privacy-and-responsible-technology

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | April 2026

15db45f086d20297-USEN-00