

# The essential guide to agentic AI

*How to compete  
in an autonomous future*



## Contents

### Introduction

Building for an agentic AI edge ..... 2

### Part one

The chassis: Why agentic AI architecture is critical for a multimodel world ..... 6

### Part two

The fuel: How to enable a continuous flow of high-quality data ..... 10

### Part three

The powertrain: How people and agents partner for impact ..... 14

Action guide ..... 19

## How IBM can help

IBM brings deep enterprise operations industry expertise, partnerships, and experience to transform intelligent operational workflows, helping businesses unlock value with AI agents and assets. For more information, visit [ibm.com/consulting/artificial-intelligence](https://ibm.com/consulting/artificial-intelligence)

## How AWS can help

AWS empowers customers to innovate, deploy, and scale agentic AI applications confidently using a comprehensive set of AI services, tools, and resources on a robust foundation of privacy, security, and AI governance. For more information, visit [aws.amazon.com/ai/](https://aws.amazon.com/ai/)



## Key takeaways

Over the next two years, 72% of executives expect agentic AI to enable new technology capabilities that can transform business models and industry structures.

- **Most leaders (76%) know they need an open, secure-by-design architecture, but few have built it.**

Agentic AI demands a secure, open framework that connects agents, models, and systems enterprise-wide—but fewer than one-third of organizations have implemented the interoperability and scalability features needed to make it work.

- **Organizations have strong data foundations, but agentic AI changes the game.**

74% have well-defined frameworks for managing and governing data, while 60% say their practices are future-ready. Agentic AI requires data to flow continuously in real time and data systems that adapt as agents learn and evolve.

- **Executives value human-AI collaboration, but they aren't measuring its impact.**

64% say people and multiagent systems work in tandem to drive strategic objectives, but less than half are measuring the impact of AI agents on employees, customers, and the business. When both people and agents are accountable for shared results, every decision and action is focused on delivering value.



## Building for an agentic AI edge

Organizations are ready to embrace the promise of agentic AI, but scaling its impact demands more than applying AI to old processes or tweaking existing systems. Agentic systems are autonomous by design: they act, learn, and adapt across workflows in real time.

As discussed in previous IBM Institute for Business Value (IBM IBV) research, agentic AI needs a new operating model<sup>1</sup>—engineered from the ground up for real-time strategic decisions and seamless human-AI collaboration across domains (see Figure 1). It’s time to build the high-performing machine of the future—where agents operate autonomously in cooperative multiagent systems across enterprise operations.

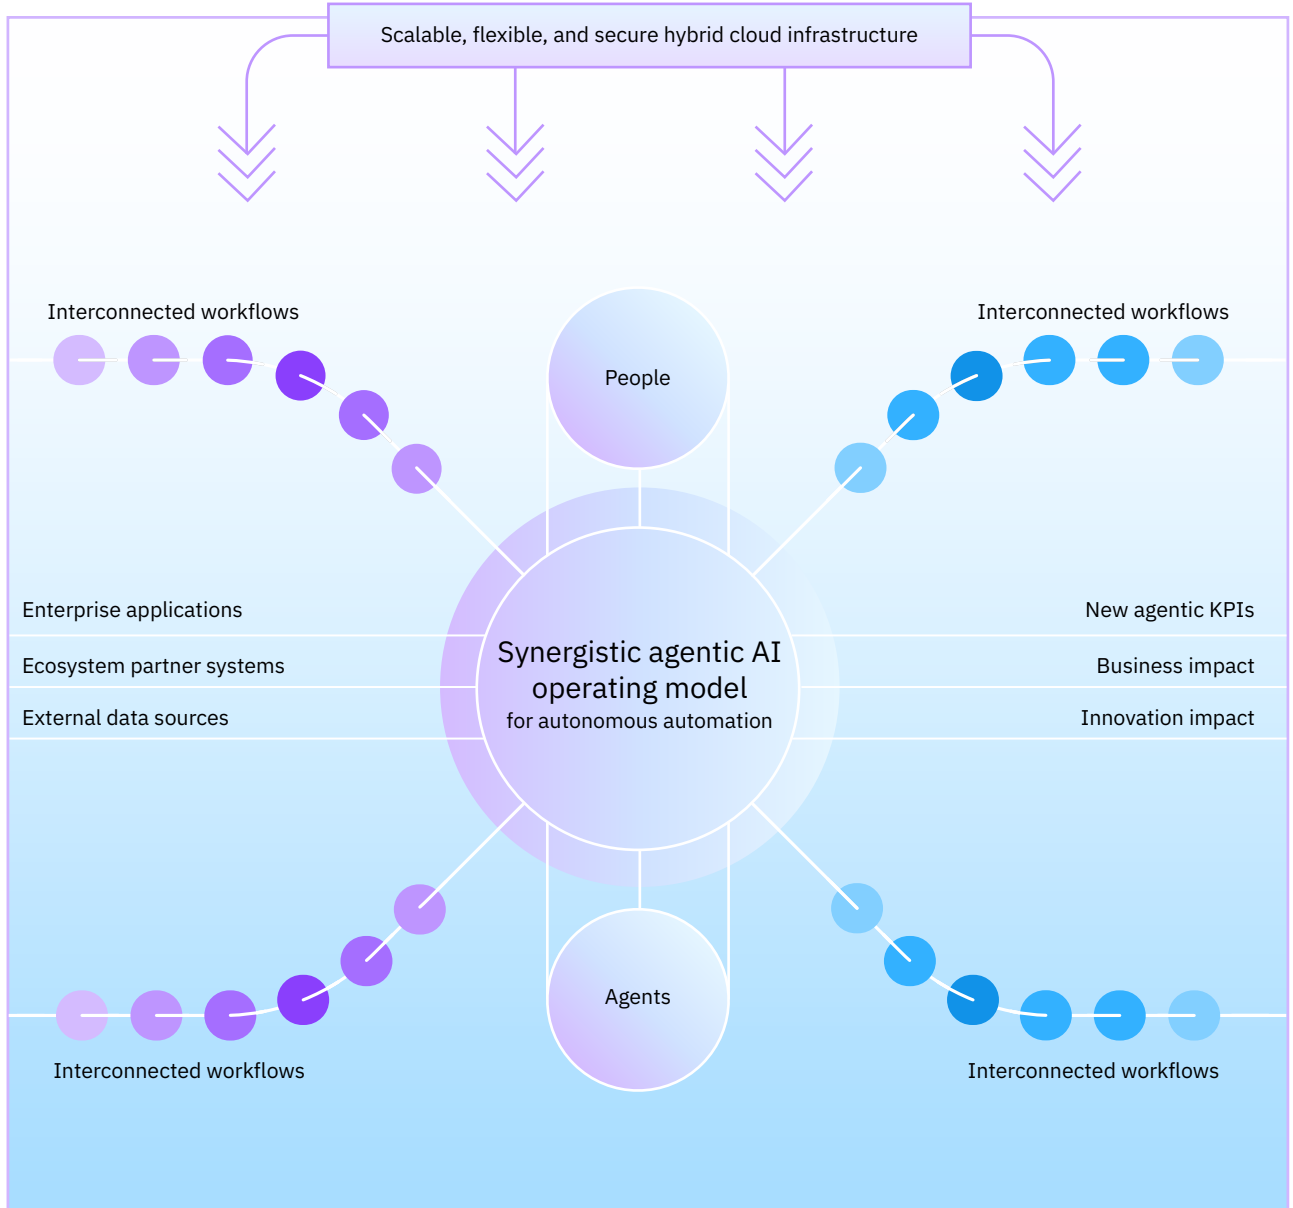
The competitive gap is already forming. The IBM IBV surveyed more than 1,000 global C-suite executives and found that while most organizations are exploring agentic AI, few are ready to scale it. Across every critical element of the operating model—from infrastructure to governance to workforce readiness—fewer than half of enterprises report they are scaling or optimizing today.

The ambition is there: 80% of executives are increasing investments in agentic AI, with spending projected to nearly triple by 2027. But without the right technical frameworks, data strategies, and execution plans, much of that investment will remain trapped in proofs of concept—delivering pockets of value without the enterprise-wide transformation envisioned by leadership (see Perspective, “Executives have high expectations for agentic AI”).

Our analysis identifies a group of more advanced organizations that are translating investment into impact in customer engagement and innovation (see Research methodology on page 23). They combine committed leadership, advanced operating models, and thoughtful human-AI collaboration for purposeful integration of sophisticated agentic systems. They’re more confident in multiagent capabilities, more strategic in model management, and significantly ahead in building scalable, interoperable foundations.

FIGURE 1

**Agentic AI demands a new operating model, engineered from the ground up for real-time, strategic decision-making.**



Developed in partnership with Amazon Web Services (AWS), this report will help organizations at all maturity levels advance their readiness for autonomous agents. We examine three essential components that determine whether agentic AI becomes a performance engine:

1

**First, the chassis**, or the strategic technical framework.

It's the secure, open framework that connects agents, models, and systems enterprise-wide, supporting orchestration, interoperability, and scalability.

2

**Second, the fuel**, or the high-quality, governed data streams that power adaptive decision-making through integration, accessibility, and real-time flow throughout the enterprise.

3

**Third, the powertrain**, or the human-AI collaboration that converts potential into measurable outcomes by aligning actions to KPIs and embedding accountability.

We conclude with an action guide that outlines strategic steps to take now in preparation for agentic AI. Leaders who get these elements right can adapt faster, innovate more deeply, and deliver impact at enterprise scale. The rest will be left optimizing yesterday's operating model.

Without the right technical frameworks, data strategies, and execution plans, most agentic AI investments will only deliver pockets of value, not enterprise-wide transformation.

## Perspective

# Leaders have high expectations for agentic AI over the next two years

Executives are betting on agentic AI as an accelerator for business transformation, not just a tool for incremental operational improvements. Within the next two years, nearly three in four leaders (72%) see it enabling new technology capabilities that can transform business models and industry structures, 70% say it will create entirely new products and services, and 68% expect it to drive personalization at scale for customers.

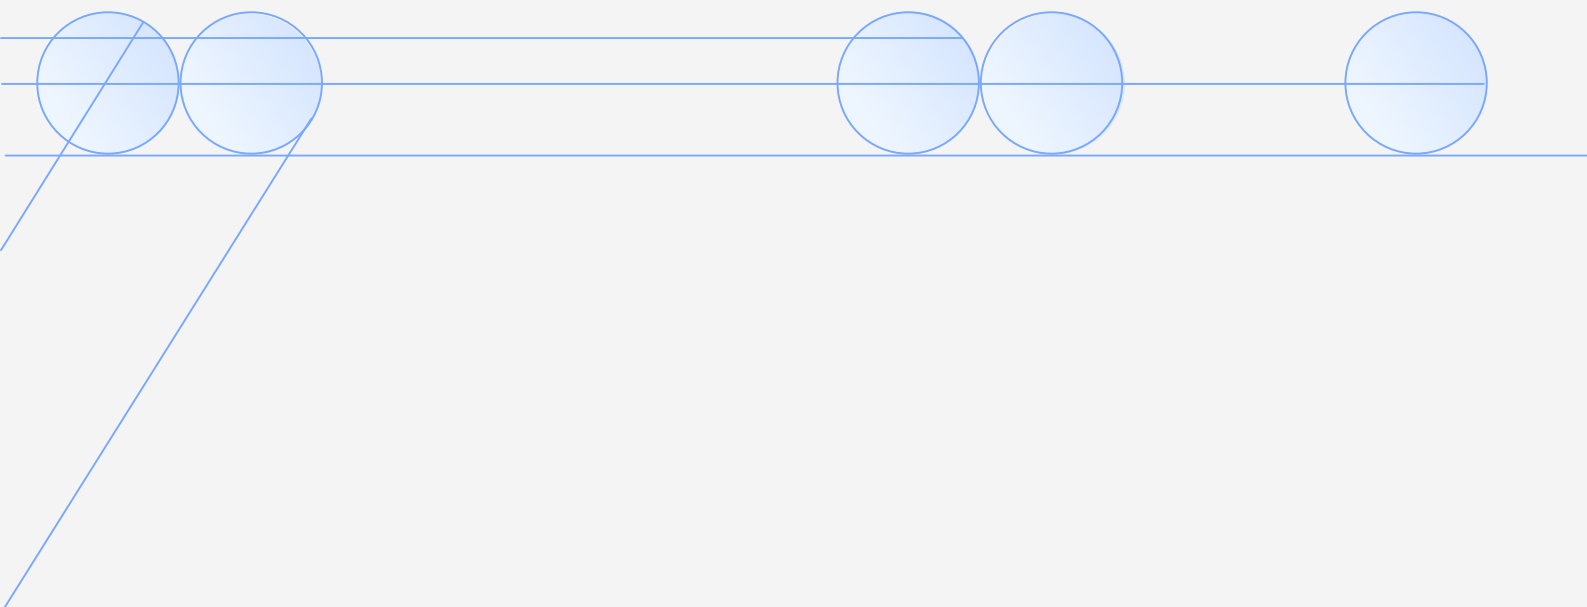
These ambitions rest on a clear vision of how agentic AI will operate. Leaders anticipate multiagent collaboration as a core capability:

- 75% say different agents will be able to tackle different parts of complex problems, producing richer solutions.
- 74% expect these systems to deliver comprehensive insights that improve decision-making.
- 73% say multiagent systems will adapt more quickly in dynamic environments, enhancing resilience.

They also foresee agents maturing rapidly to become more autonomous and adaptive:

- 77% expect agents to continuously improve through learning.
- 73% expect them to operate with built-in accountability and transparency.
- 69% expect them to make autonomous decisions in integrated workflows.

In short, leaders are preparing for a near-term future where agentic AI systems learn, collaborate, and adapt at speed—reshaping how work gets done and how value is created.



## Part one

## The chassis: Why agentic AI architecture is critical for a multimodal world

Agentic AI changes the rules for enterprise architecture. Traditional AI often relies on preprogrammed algorithms in point solutions—sometimes powered by a single model—solving specific tasks in isolation. Generative AI and its use of large language models requires more robust infrastructure to handle multimodal (text, audio, video) capabilities and larger computational demands. Autonomous agents demand something fundamentally different: dynamic interaction across multiple models and modalities, real-time information exchange, and contextual task-switching capabilities. Success requires architectures built for orchestration, where agents coordinate autonomously across platforms and systems without constant human intervention.

Most leaders understand this foundation is critical, but construction is still in the early stages. Three in four (76%) agree that an open, secure-by-design architecture is essential for multimodal agentic AI environments, where network effects emerge from collaboration and data sharing between diverse models and agents. More than two-thirds (69%) say this type of architecture promotes transparency, interoperability, and faster innovation. Yet when it comes to implementing interoperability and scalability features, fewer than one-third have done so to a great extent (see Figure 2).

The technical foundation gaps between the most and least mature organizations are substantial. 89% of the advanced group ensures seamless integration of diverse AI models, versus 58% of the least mature group. And 85% have scalable infrastructure to support complex AI workloads, versus 52% of their less advanced peers. These significant differences represent the divide between organizations ready to scale and those still struggling with basic model integration.

FIGURE 2

## Few organizations have the interoperability and scalability architecture features needed for agentic AI.

Percentages represent organizations that have each interoperability and scalability feature to a great extent



### Interoperability features

Open industry standards with open protocols for data delivery between models and external data sources

37%

Ability to incorporate and use third-party AI models or services

30%

Seamless integration and operation of diverse AI models

31%

Flexibility to adapt AI models or architecture without significant disruption

25%

Standardization of APIs to promote better data exchange between different components, models, and systems

23%



### Scalability features

Efficient handling of increasing volumes of data for AI model training and inference

28%

Flexibility to adapt and reconfigure AI system components as requirements change

29%

Scalable infrastructure to support computational demands of complex AI models

28%

Ability to replicate and distribute AI models across multiple machines or cloud instances

23%

Investment in emerging hardware (quantum computing, neuromorphic chips) to advance scalability

12%

In a fast-changing environment where organizations deploy a diverse mix of custom models, third-party foundation models, and open-source alternatives, architecture must deliver these core capabilities:

**Interoperability and orchestration.** Agents must share information, coordinate actions, and hand off tasks across different systems and model types. 62% of leaders understand their architecture requires orchestration layers that manage data flow, model interactions, and dynamic task switching, but 61% say seamless integration and communication between models is a major challenge. A model-agnostic architecture, such as container-based orchestration platforms that package models into standardized, portable units, can support multiple model types and help ensure they operate together without friction.

**Scalability and flexibility.** Elastic scaling is essential for agentic AI's diverse workloads. 47% of executives are turning to cloud platforms for scalability and flexibility, while 44% are conducting infrastructure assessments to close capability gaps. Cloud and hybrid architectures let compute-heavy models run alongside lighter ones, integrate new agents quickly, and connect securely to ecosystem partners—balancing compliance and control with reach and adaptability to accommodate evolving demands.

**Security.** Executives report data privacy and security is the number-one challenge to implementing agentic AI, yet only 38% have embedded security into every stage of the AI lifecycle—from development to deployment. A secure-by-design architecture protects sensitive data, enables compliance management, and creates the foundational layer for autonomous agents to operate safely. It includes integrating encryption, access controls, and threat detection directly into the infrastructure, so security scales automatically as agents are deployed across the enterprise.

**Trust and transparency.** Autonomous decision-making demands visibility and accountability. 73% of executives expect AI agents to enable accountability and transparency in decision-making by 2027. They will need technical architectures that can monitor agent behavior in real time, log decision pathways, and provide audit trails that explain how and why autonomous actions were taken (see Perspective, “Real-time monitoring and auditing for autonomous decision-making”). Without built-in transparency, even the most sophisticated agents become “black boxes” that can erode trust and limit adoption.

In the agentic AI era, your architecture will define how quickly you scale, how effectively you innovate, and how well you compete. Winners make intentional architecture decisions—building the secure, open, and adaptable framework that turns agentic AI pilots into enterprise-wide value.

“Building ethics analysis into agentic AI isn’t just about doing the right thing—it’s about preventing these autonomous systems from making decisions that conflict with human values at machine speed and scale.”

Chief Technology Officer, Electronic products manufacturer

## Perspective

# Real-time monitoring and auditing bolster trust in autonomous decision-making

For trust and transparency in agentic AI, organizations must build monitoring and auditing capabilities into autonomous systems from square one. Effective oversight of autonomous decision-making relies on several key technical capabilities working together:

**Real-time guardrails and policy enforcement.** Automated governance systems can enforce rules and policies across enterprise systems, detecting and remediating non-compliant actions as they occur. These guardrails operate continuously in the background, helping ensure autonomous agents stay within approved parameters without slowing decision-making.

**Comprehensive activity logging.** Every action taken by autonomous systems—including API calls, data access, and decision triggers—should be recorded in detailed audit logs. This creates a complete trail that allows teams to trace decisions back to their origin, investigate anomalies, and demonstrate compliance during regulatory reviews.

**Continuous monitoring and threat detection.** Advanced monitoring platforms can analyze patterns of autonomous activity in real time, detecting suspicious behavior or deviations from normal operations. When combined with intelligent alerting systems, these tools can trigger automated responses or human intervention before issues escalate.

**Fine-grained access control.** Identity and access management systems enable precise control over what autonomous agents can do, when, and under what conditions. Role-based permissions and dynamic access boundaries help ensure agents operate only within their designated scope while adapting to changing business needs.

**Workflow orchestration with human checkpoints.** For high-stakes decisions, orchestration platforms can integrate approval workflows that require human review before certain actions proceed. These systems balance autonomous efficiency with human oversight, helping ensure critical decisions receive appropriate scrutiny.

**The result:** a layered approach that combines monitoring, logging, access controls, and workflow management to help ensure autonomous decisions are transparent, traceable, and compliant with both internal policies and external regulations—building the trust foundation essential for scaling agentic AI across the enterprise.

## Part two

## The fuel: How to enable a continuous flow of high-quality data

Data is no longer just input. With agentic AI, its role shifts from learning patterns and feeding predictions to becoming the continuous fuel stream that powers autonomous, goal-driven action in dynamic environments. Agents need data flows that update in real time as agents and people execute workflows, creating feedback loops that improve both decision-making and operational outcomes. This constant cycle enables agents to adapt dynamically and deliver context-rich insights that drive sustained business impact.

The stakes of data management are higher with agentic AI. In autonomous systems, data functions as a stream of real-time experience, so its quality, timeliness, and governance are exponentially more critical to performance than for traditional or generative AI. For example:

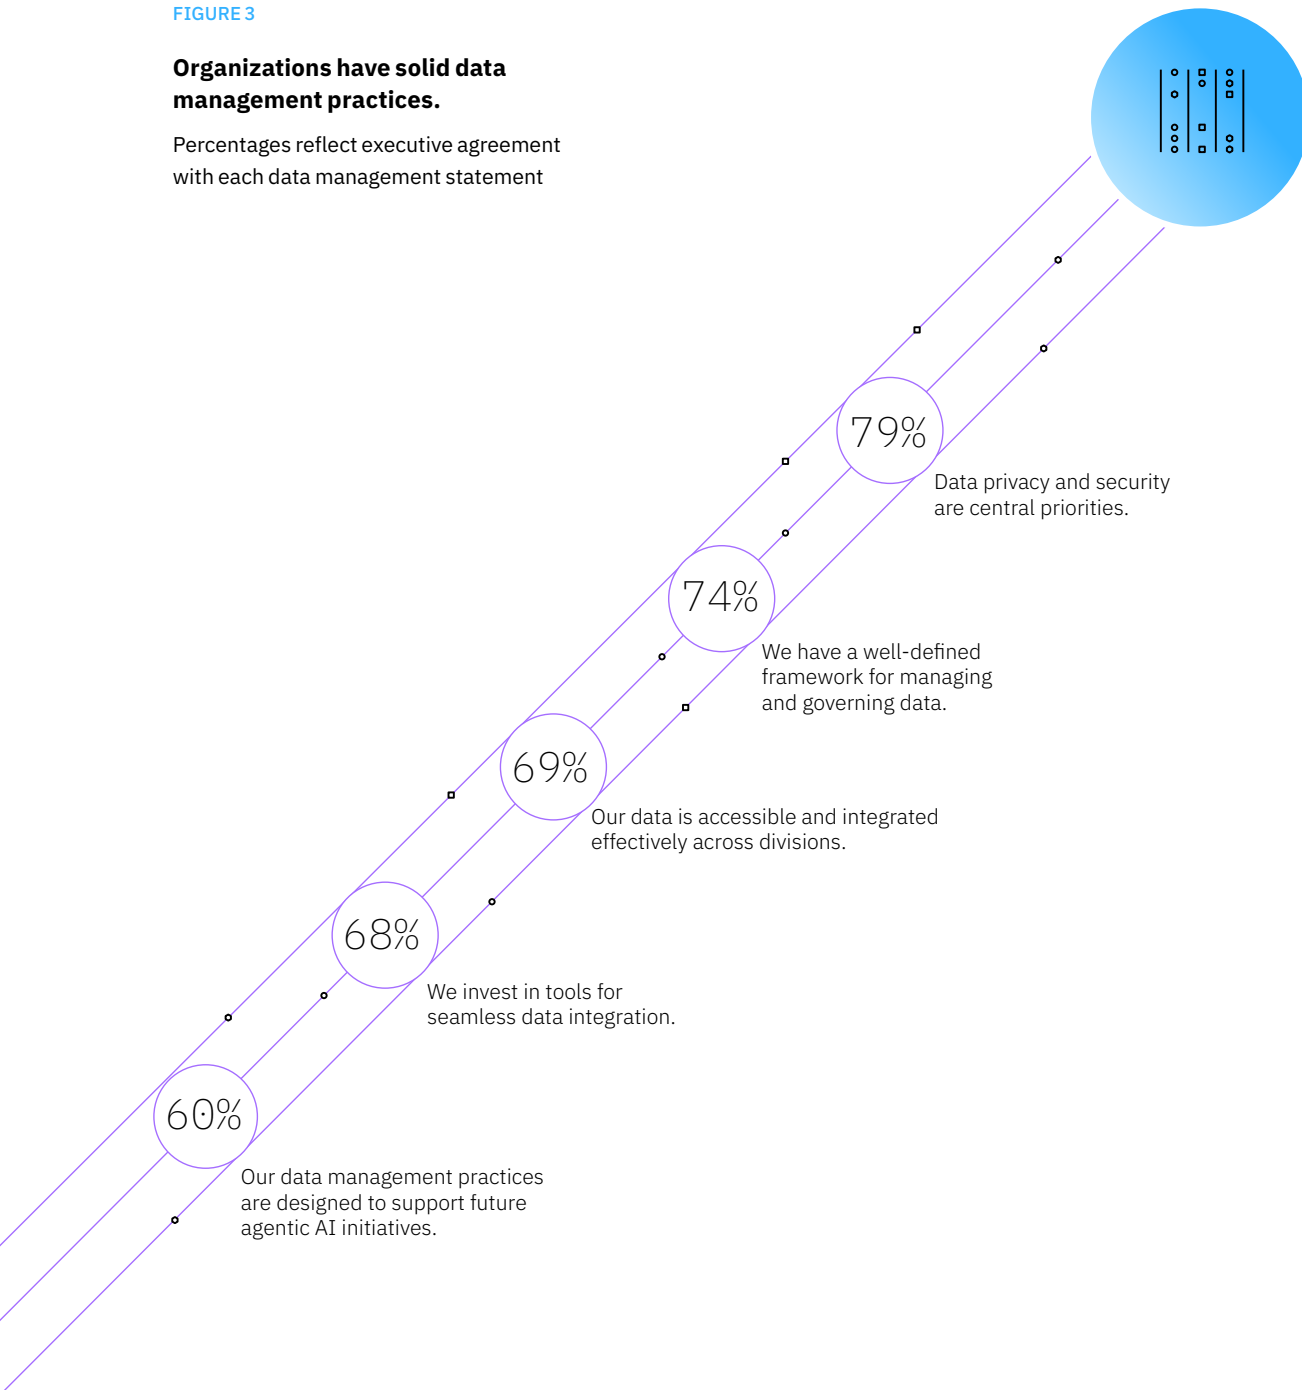
- Inaccurate or incomplete data can mean an agent takes an unpredictable action, such as an inventory management agent automatically over-ordering massive quantities of materials based on faulty demand forecasts.
- Stale data can mean an agent acts on outdated context, such as a trading agent executing a large buy order based on a stock price that is two minutes old, missing a sudden price crash.
- And lack of advanced governance can lead to unauthorized, irrevocable actions taken in the real world, such as an irreversible series of trades that cause significant financial loss.

Organizations have built strong data foundations—privacy, governance, accessibility, and integration are priorities for most (see Figure 3). The challenge now is evolving those foundations for agentic AI: making data faster, more connected, and ready for agents that learn and adapt in real time. While 60% say their practices are future-ready, they need to turn readiness into a competitive advantage by embedding agentic-specific requirements into their data strategy now.

FIGURE 3

### Organizations have solid data management practices.

Percentages reflect executive agreement with each data management statement



Organizations with advanced maturity have embedded these principles more deeply, such as seamless data integration (73% versus 64% of the least mature). More tellingly, 84% say effective AI governance is essential for agentic AI success (versus 65%).

Enabling the flow of data for agentic AI depends on three strategic priorities:

1

**Data governance that enables agility.** In the agentic AI era, governance must enable speed. That means embedding rules and controls directly into data workflows so the right agents can access the right data instantly, without manual bottlenecks. Automated data governance can enforce quality rules without delays, provide role-specific access and permissions to meet each agent’s needs, and maintain automated lineage and cataloging so agents can evaluate and understand data reliability.

2

**Integration across silos.** Structured and unstructured data trapped in silos loses value. Seamless integration requires two key capabilities: automated data preparation that standardizes formats and ensures quality as information flows between systems, and unified access layers that let agents combine transactional records, operational metrics, partner data, documents, and multimedia content into holistic views that support better decisions and personalized experiences. Modern data integration tools such as a data fabric or a data lakehouse—a combination of a data lake and data warehouse<sup>2</sup>—help dissolve barriers with automated data preparation and unified access layers. Yet in separate IBM IBV research, only 19% of COOs report their organizations have fully developed the components of an enterprise-wide data architecture and scaled data integration across functional areas.<sup>3</sup>

3

**Accessibility for continuous learning.** Once data is unified, it must support constant agent improvement. 77% of organizations expect their AI agents to continuously improve performance over the next two years through learning, but this requires speed, security, and feedback. True accessibility—and adaptability—means real-time data availability, embedded governance in access controls, and closed-loop feedback systems where agent outcomes feed directly back into training datasets. These capabilities enable agents to adapt to changing conditions and deliver ever-better results.

A high-performing fuel line is built for velocity, reliability, and adaptability. Organizations that master continuous, well-governed data flows will give their agents the power to act with precision, learn faster, and deliver outcomes that compound over time. Without a robust fuel line, agentic AI will sputter.

“How can we speed up decision-making? We’re in the fast-moving consumer goods space where an average product cycle from arrival in our distribution center to shipping is 8–9 hours. It would be a huge step ahead if we could use agentic AI to do logistics optimization.”

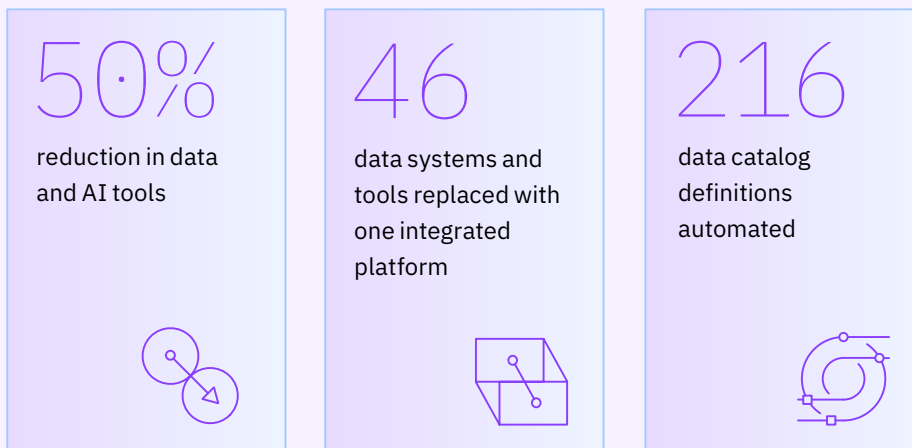
Chief Supply Chain Officer, Consumer goods company

## Case study

# Data fuels AI takeoff at Lockheed Martin<sup>4</sup>

Lockheed Martin—a global leader in aerospace and defense technology—transformed its business with AI, starting with its valuable but disconnected data. The company had previously struggled with data silos due to several data lakes and 46 disparate systems for data management, data analytics, and business intelligence.

With IBM's help, Lockheed Martin replaced those systems with a single, connected, and accessible environment. The solution resulted in:



Once a unified, scalable data foundation was in place, the company was able to access and leverage its highest-quality data to fuel its AI Factory, a secure AI ecosystem where 10,000 engineers can build, iterate on, and deploy large-scale AI solutions—fast. Agentic frameworks and advanced virtual agents helped the company better manage complex workflows, simplify interactions between people and technology, and enhance efficiency through process automation and optimization.

## Part three

## The powertrain: How people and agents partner for impact

Execution in the agentic AI era demands more than new roles and skills. It requires redefining what success looks like for the business, people, and agents. Nearly half of executives (45%) say how they set business goals and objectives will need to change because of agentic AI. Why? Because autonomous agents don't just accelerate existing processes—they also open the door to new opportunities that weren't possible before. Over the next two years, 72% of leaders expect agentic AI to enable new technology capabilities that will transform business models and industry structures.

But goals alone won't deliver impact. The real powertrain is human-AI collaboration, a new approach that blends autonomy with accountability to convert potential into performance. Unlike traditional AI, which often positions the workforce as supervisors of static systems, agentic AI thrives when people and agents collaborate: people provide judgment, creativity, and oversight, ensuring decisions align with strategy and values, while agents surface insights, execute actions at speed and scale, and adapt in real time. The partnership is adaptive, shifting gears based on conditions, priorities, and performance data (see Perspective, "Agentic AI supports application migration and modernization").

Executives value human-AI collaboration. 70% agree people and agents learn from one another, and 64% agree people and multiagent systems work in tandem to drive strategic objectives.

Yet readiness to execute at scale remains limited. Only 42% of organizations are scaling or optimizing workforce requirements for AI-related roles, and just 37% are doing the same for performance metrics and impact evaluation. Less than half have taken steps to measure agentic AI's impact on their organization (see Figure 4). Without these foundations, collaboration may be high on intent but low on quantifiable impact.

FIGURE 4

**By 2027, less than half of organizations expect to have taken steps to measure the impact of agentic AI.**

Percentages represent those executives saying they will have taken each step to maximize agentic AI's impact by 2027

We have KPIs to monitor the impact of AI agents on business targets.



We assign performance metrics to individual AI agents.



We have developed KPIs to monitor the impact of AI agents on employees.



We have developed KPIs to monitor the impact of AI agents on customers.



We assign performance metrics to teams of AI agents.



“My focus is on creating an agile organizational structure. Roles are changing, not just for teams—stakeholders too need to realize they will change. Everyone’s looking to the workforce to embrace AI, but when they do, that changes the foundation of the enterprise as well.”

Chief Operating Officer, Medical devices and disease diagnostics

High-performing human-AI collaboration depends on these mandates:

### **Collaboration aligned to outcomes**

Human-AI collaboration must be tied directly to KPIs and business goals. When both people and agents are accountable for shared results, every decision and action is focused on delivering value. Almost half (46%) of organizations have developed KPIs to monitor the impact of AI agents on business performance.

### **Embedded autonomous performance measurement**

Real-time monitoring of workflows, experiences, and outcomes helps ensure that impact is visible and decisions can be adjusted quickly. This requires systems that track agent performance alongside people's contributions. 45% of organizations have assigned performance metrics to individual AI agents.

### **Trust and ethical alignment**

57% of organizations expect their personnel to work within their AI governance structure to manage ethical challenges. But their frameworks for managing ethical AI practices are the least mature aspect of their operating models: only 30% are scaling or optimizing them. And lack of a formal AI governance framework is their number-two barrier to implementing agentic AI. Trust requires clear role definitions in which people set ethical boundaries and monitor agent behavior, while agents operate within those parameters and flag ethical dilemmas for human review. Organizations need to prioritize embedding governance controls that prevent agents from exceeding their authority and training their talent to recognize and respond to ethical cases.

### **Workforce readiness for orchestration**

Preparing the workforce to guide, refine, and learn from autonomous agents is essential. Executives understand this, with 71% saying they are constantly fine-tuning the relationship between people and agents. 70% say people are always in the loop on high-complexity, high-impact decisions. And as agents continually evolve and take on new responsibilities, change management must be ongoing, not a one-time transition. Roles, processes, and governance must adapt in lockstep with agent capabilities, ensuring accountability, oversight, and collaboration remain aligned with business objectives.


The most mature organizations are more deliberate in managing the human-AI relationship. For example, 77% ensure they protect and value employees' role in critical thinking (versus 66%). This intentional approach strengthens adoption and trust, but even the advanced group has been slow to define KPIs that measure the impact of agents, or the impact of agentic AI on their employees, customers, and business partners. Organizations that invest in outcome-aligned collaboration and real-time performance measurement—along with a skilled pit crew—can turn agentic AI from a promising technology into a competitive advantage that compounds over time.

## Perspective

# Agentic AI in application migration and modernization<sup>5</sup>

Agentic AI can accelerate complex application migration and modernization initiatives by orchestrating specialized autonomous agents alongside humans in the loop. In these workflows, agents handle discovery, analysis, configuration, and testing tasks at speed and scale, while humans provide oversight, governance, and input where judgment or expertise are required.

This hybrid approach reduces manual effort, mitigates risks, and delivers modernized systems that are more performant, scalable, and maintainable.



### Example: Database replatform to cloud

One common modernization journey is moving a database from on premises to a cloud-based platform, such as replatforming PostgreSQL to a managed cloud service.

#### Discovery and planning

Humans define objectives and provide input such as credentials; agents capture current state details, assess the target environment, and identify incompatibilities or dependencies.

#### Solution design and preparation

Agents generate remediation plans, configure the target environment, set up certificates, firewalls, and roles, with human support where needed.

By distributing work between agents and humans, organizations can execute migrations faster, with greater accuracy, and with built-in governance.

#### Migration execution

Agents set up migration tools and processes, perform schema conversions, and manage change readiness; humans approve key steps such as data migration and reconciliation.

#### Post-migration validation

Agents run sanity and performance tests, with humans reviewing and approving results.

## Conclusion

# Early movers set the pace

Organizations are pouring resources into building agentic AI solutions, in search of greater productivity, operational efficiency, and customer satisfaction—and ultimately, revenues and profits. Our analysis shows that organizations further along the operational maturity curve already report stronger innovation and customer engagement than their less mature peers, and they are more ready for future impact. Three in four expect innovation with agentic AI to enable new capabilities that will transform their business models, versus 64% of the least mature group.

Frontrunners are cementing their ability to scale agentic AI with strategic decisions—knowing when to push, when to conserve, and when to pit. We also see that their leadership plays a decisive role. In the highest-maturity group, 80% say leadership prioritizes AI as a key strategic component (versus 64% of the least mature), 78% report active C-suite involvement in setting AI vision and strategy, and 78% have clear goals and metrics for AI success.

These advanced organizations are positioned to adapt faster, innovate meaningfully, and compete intensely—lap after lap—in the agentic AI era. The rest must shift from a focus on fine-tuning to rebuilding their AI engines with renewed intention and speed.

## Action guide

# Turning vision into measurable progress

Building maturity in your agentic AI operating model doesn't require sweeping change overnight—but it does demand focused, intentional steps that strengthen your architecture, data strategy, and human-AI collaboration. The following actions can help you move forward in the near term.



1

## Build the architecture for agentic AI

Start by assessing your current architecture for scalability, interoperability, and security gaps.

- Leverage existing cloud or hybrid infrastructure for agentic AI by mapping where autonomous agents will run, what data they'll need, and how they'll connect across environments. Extend and integrate existing capabilities—for example, API gateways or container orchestration—with agent-specific communication frameworks and messaging layers to achieve scalability.
- Run a small-scale orchestration pilot in one workflow or domain to test tools that manage task allocation, data exchange, and performance tracking between diverse agents and models. Learn and scale across domains.
- Begin embedding security and compliance into agentic workflows by extending current enterprise security frameworks. Define agent identity/authentication, role-based access to data, and automated logging for audit readiness.

Only 38% of organizations have embedded security into every stage of the AI lifecycle—from development to deployment.

## 2

## Enable continuous data flow to and among agents

Conduct a quick audit of your data governance framework to identify bottlenecks that slow access to high-quality data.

- Pinpoint high-impact data sources most critical to agent decision-making in priority use cases, and integrate them into a unified environment.
- Enable real-time or near-real-time data access for one pilot workflow using streaming or API-based integration to ensure agents act on the latest conditions.
- Set up a closed-loop feedback process so the outcomes of agent actions are automatically captured and used to retrain or fine-tune models, supporting continuous learning.

## 3

## Redefine impact of employees and agents

Define or refresh your KPIs to measure outcomes from human-AI collaboration.

- Redefine KPIs for one or two pilot workflows to measure both workforce and agent contributions. Include metrics for accuracy, decision speed, and adaptability.
- Establish a collaboration framework that defines how people and agents coordinate—who makes which decisions, how hand-offs occur, and how oversight is maintained.
- Progress autonomous performance monitoring using existing analytics or monitoring tools, and then use workflow impact feedback to track agent activity in real time and flag anomalies.

45% of executives say they will need to change how they set business goals and objectives because of agentic AI.

## About the authors

### *Vamsi Yanamadala*

Worldwide Partner Solutions Architecture Leader  
Amazon Web Services  
[linkedin.com/in/vamsi-yanamadala/](https://www.linkedin.com/in/vamsi-yanamadala/)

Vamsi brings more than 15 years of experience leading and supporting the design, development, implementation, and quality assurance of enterprise SAP and ERP systems. He manages system development and support lifecycles for the implementation, testing, and ongoing administration of technical architectures.

### *Amit Chowdhury*

Global AI/ML Solutions Architecture Leader  
Amazon Web Services  
[linkedin.com/in/amit-chowdhury-b2104520/](https://www.linkedin.com/in/amit-chowdhury-b2104520/)

Amit is a seasoned technology leader with 25+ years of extensive industry experience, leading transformative AI/ML initiatives to develop and implement comprehensive technology strategies and drive measurable business outcomes across diverse industry verticals for global enterprise customers, reshaping how enterprises approach AI strategy and adoption. He spearheads solutions architecture in AI/ML/generative AI and data analytics platforms, driving technology transformation and modernization, supporting multimillion-dollar businesses.

### *Francesco Brenna*

Global Vice President and Senior Partner  
Artificial Intelligence  
IBM Consulting  
[linkedin.com/in/francesco-brenna/](https://www.linkedin.com/in/francesco-brenna/)

Francesco is driving the adoption of generative and agentic AI to help organizations worldwide innovate faster, operate smarter, and compete boldly in a rapidly changing digital landscape.

### *Rabeela Janorious*

Global CTO, AI Integration Services  
IBM Distinguished Engineer  
IBM Consulting  
[linkedin.com/in/rabeela-janorious-04bb5417/](https://www.linkedin.com/in/rabeela-janorious-04bb5417/)

Rabeela leads the technical strategy for AI and AI transformations across IBM Consulting® globally, defining and executing technical strategies for AI solutions, empowering clients worldwide in their AI and agentic AI transformation journeys, and identifying opportunities for AI-driven innovation to deliver tangible business value.

### *Karen Butner*

Global Research Leader, AI Automation  
and Digital Operations  
IBM Institute for Business Value  
[linkedin.com/in/karenvbutner](https://www.linkedin.com/in/karenvbutner)

Karen leads data-driven research and thought leadership in AI, automation, and digital operations, equipping business leaders worldwide with actionable insights on industry trends, emerging technologies, and transformative opportunities.

## Contributors

A B Vijay Kumar  
IBM Fellow, Vice President, Master Inventor, CTO,  
IBM Consulting Delivery

Shiva Bayati  
Global AI Integration Services Offering Lead (AWS),  
AP, IBM Consulting

Tom Utiger  
Global AI Integration Services Offering,  
IBM Consulting

### *From the IBM Institute for Business Value:*

Heba Nashaat  
Data and Content Management Manager

Steve Ballou  
Director, Research Hub

Kathleen Martin  
Senior Managing Consultant

Sara Aboulhosn  
Associate Creative Director

Angela Finley  
Design Lead

Joanna Wilkins  
Editorial Lead

1,000

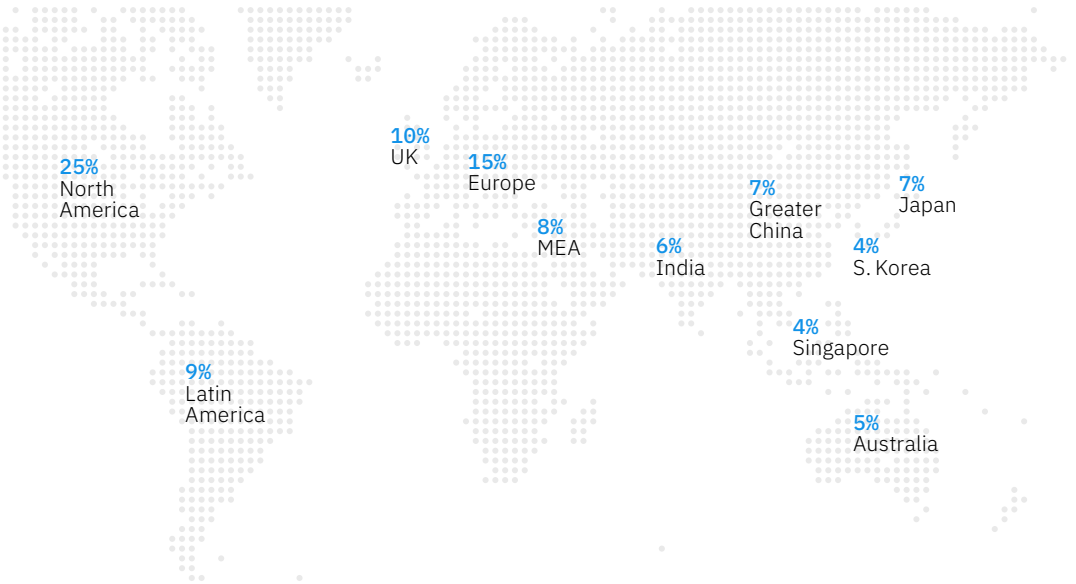
Senior technology and business leaders

11

Countries

17

Industries



## Research methodology

To assess global enterprise readiness for agentic AI, the IBM Institute for Business Value (IBM IBV), in partnership with Oxford Economics, conducted a comprehensive survey of 1,000 senior technology and business leaders. The sampling frame was designed to capture a balanced perspective from large enterprises across 17 major industries, including retail and corporate banking (8%), federal government (8%), and telecommunications (7%). The surveyed organizations were substantial in scale, with a median average of 27,000 full-time employees and \$2.8 billion in annual revenue.

### Analytical approach: The three-dimensional maturity assessment

To evaluate organizational readiness for agentic AI, we constructed a composite three-dimensional maturity index. This holistic assessment is based on a weighted average of three critical domains:

- AI architecture and orchestration (40%), which measures the organization’s technical backbone. Here we assessed open standards, dynamic model orchestration, and scalable, secure infrastructure.
- Data flow and governance (25%), which assesses data management frameworks, quality, accessibility, and future-ready practices.
- Human-AI collaboration and operating model (35%), which captures the evolution of roles, strategic alignment, governance, and performance monitoring between humans and AI systems.

Using the overall weighted score, organizations were segmented into three distinct performance tiers based on the empirical distribution of their results. The bottom 25% of organizations represent those with limited capabilities. The middle 50% consists of organizations actively building and implementing their foundational systems. The top 25% represents organizations that have achieved a sophisticated level of strategic integration and operational maturity in their agentic AI initiatives.

## 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](mailto: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.

## 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.

## About the IBM and AWS partnership for agentic AI

Together, AWS and IBM help connect and deploy the infrastructure, software, and expertise required to scale agentic AI across the enterprise. By combining trusted expertise and essential integrations with a comprehensive suite of solutions around hybrid cloud, AI, data, and automation—now available in 90+ countries in AWS Marketplace—IBM and AWS enable organizations to seize the true opportunity presented with agentic AI.

## Related reports

### ***Agentic AI's strategic ascent: Shifting operations from incremental gains to net-new impact***

IBM Institute for Business Value in partnership with Oracle. October 2025. <https://ibm.biz/agentic-ai-ops>

### ***Orchestrating agentic AI for intelligent business operations***

IBM Institute for Business Value. May 2025. <https://ibm.co/agentic-process-automation>

### ***The intrepid COO: 5 strategies to shatter enterprise inertia***

IBM Institute for Business Value. August 2025. <https://ibm.biz/2025-coo>



## 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.co/ideawatch](https://ibm.co/ideawatch)



## Notes and sources

- 1 Brenna, Francesco, Karen Butner, Manish Goyal, Anant Patel, Federico Torreti, and Shobhit Varshney. *Agentic AI's strategic ascent: Shifting operations from incremental gains to net-new impact*. IBM Institute for Business Value. October 10, 2025. <https://ibm.biz/agent-ai-ops>
- 2 "What is a data lakehouse?" IBM Think website. Accessed October 28, 2025. <https://www.ibm.com/think/topics/data-lakehouse>
- 3 *2025 Global C-suite Series. The intrepid COO: 5 strategies to shatter enterprise inertia*. IBM Institute for Business Value. August 28, 2025. <https://ibm.biz/2025-coo>
- 4 "A new era of AI takes off for Lockheed Martin." IBM case study. Accessed November 5, 2025. <https://www.ibm.com/case-studies/lockheed-martin>
- 5 Ganookar, Vikas, Kannu Malik Seth, Roopali Anand Thapar, and Abhijeet Deshpande. "Reimagining the application migration and modernization value-chain: The agentic way." IBM Think website. Accessed October 28, 2025. <https://www.ibm.com/think/insights/reimagining-application-modernization-migration-agent-ai>

© Copyright IBM Corporation 2025

IBM Corporation  
New Orchard Road  
Armonk, NY 10504

Produced in the United States of America | November 2025

IBM, the IBM logo, [ibm.com](http://ibm.com), and IBM Consulting 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](http://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.

Examples presented are illustrative only. Actual results will vary based on client configurations and conditions and, therefore, generally expected results cannot be provided.

