


Agentic AI's strategic ascent

*Shifting operations from incremental
gains to net-new impact*

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



Key takeaways


By 2027, twice as many executives expect AI agents will make autonomous decisions in processes and workflows, compared to today.

- More than three-quarters of executives say most of their AI investment has focused on improving existing processes rather than developing net-new capabilities.
- Yet, 78% of C-Suite executives say achieving maximum benefit from agentic AI requires a new operating model.
- Companies that excel in three key AI adoption areas are 32 times more likely to achieve top-tier business performance than those with minimal implementation.



Contents

Rethinking how work works with agentic AI	1
AI transformation isn't a forecast, it's a deadline	4
Are leaders focused on the wrong AI challenge?	7
The trust deficit: AI systems need to show their work	9
Operating model redesign: The competitive separator	10
Action guide	14



Rethinking how work works with agentic AI

As companies continue to explore how to perform better with artificial intelligence, a cohort of executives are asking a question that goes far beyond simple deployment: “What if we rebuild our entire operating model around AI?” A growing number already are—and they’re creating business outcomes that were previously impossible.

These leaders aren’t just implementing AI tools—they’re pioneering entirely new operating models around autonomous decision-making. The radical premise: that agentic AI systems will make an increasing number of decisions, while humans make the decisions that matter most, preserving human agency and expertise where it is most crucial.

The performance gap is widening rapidly. These transformation-driven organizations are creating net-new business capabilities, fundamentally changing what’s possible at enterprise scale.

By

2027

twice as many executives expect AI agents will make autonomous decisions in processes and workflows compared to today.

By

2030

fully autonomous robotic systems with embodied AI will be operational realities across industries.

The companies winning in the autonomous AI future won’t be those that simply deploy the most sophisticated AI agents. They’ll be the organizations that abandon existing operating models for entirely new ones designed around autonomous decision-making capabilities.

In the emerging AI transition, speed matters, but not all efforts are equal. Some organizations are reaching truly strategic peaks while others chase tactical AI.

800

C-suite executives

20

Countries

19

Industries


To better understand the real-world impact of autonomous AI, the IBM Institute for Business Value (IBM IBV) surveyed 800 C-suite executives in 20 countries across 19 industries.

This report reveals key elements of an autonomous AI operating model—from management mindset to workforce evolution to the role of trust and transparency. Alongside, we highlight changing team structures, and examples of specific operating model transformations in progress.

We finish with an action guide that synthesizes what to do now to build a synergistic operating model that can optimize agentic AI's competitive benefits.

FIGURE 1

Measurable impact: AI is expected to drive double-digit improvements across the enterprise



Perspective

Agentic AI's split screen

In this recent survey of C-suite operational executives, we uncovered a fascinating divergence in how companies are approaching agentic AI. We asked two very telling questions, which revealed a sharp divide in strategic focus.

First, we asked what percentage of their agentic AI initiatives were focused on improving existing processes and how successful they were.

Second, we asked what percentage of their agentic AI initiatives were designed to create entirely new workflow capabilities and how successful they were.

The responses divided our respondents into two distinct groups.

01


The process-focused

This group is technically getting it done; they're laser-focused on optimizing existing workflows and are achieving measurable business impact. In short, they are proficient at efficiency, but they have yet to crack the code on true transformation.

02

The transformation-driven

This more strategic cohort is advancing agentic AI with a dual mandate. They are not only improving existing workflows, but they are also successfully creating and building net-new capabilities. Their vision extends beyond optimization; they are actively reimagining their operating model. The data shows it's working. They are achieving real-world value of AI and automation across their business operations.



AI transformation isn't a forecast, it's a deadline

Today, 24% of executives say that AI agents take independent action in their organization. By 2027, 67% expect that to be the case. Similarly, twice as many executives expect autonomous decision-making from agentic AI in processes and workflows by 2027 (57%) compared to today (28%).

Even in highly regulated areas, transformation-driven organizations are moving fast—executives expect agentic AI to automate an average of 29% of risk and compliance operations by 2027. And many anticipate that AI agents will automate the innovation process itself, with 19% of executives saying that is happening today and 48% of executives expecting this level of innovation-related automation in 2027.

These benchmarks represent a fundamental change in the DNA of business. It's no wonder that

78% of our respondents agree that achieving maximum benefit from agentic AI requires a new operating model.

Amid global economic uncertainty, 69% of them report a pressing need to develop more agentic AI-enabled predictive and simulation modeling capabilities—for example, for financial forecasting, human resource talent and skilling analysis, dynamic transportation and distribution routing and scheduling, and proactive mitigation strategies with built-in resilience.

Customer-touching operations have emerged as primary arenas for agentic AI implementation. Organizations are prioritizing sales forecasting, dynamic pricing based on inventory scenarios, and intelligent customer order processing—areas where autonomous system intelligence can deliver immediate competitive advantages and measurable returns.

Because they are focused on developing net-new capabilities, transformative organizations are asking different questions: What becomes possible when systems can make decisions autonomously?

The appetite for change is real, but a critical chasm exists between aspiration and execution. The vast majority of AI investment—78%, has been funneled into improving existing processes. This is the difference between optimization and transformation. Process-focused organizations are stuck in a loop of making their current operations better, rather than making transformative leaps a reality. These enterprises must also recognize the shifting landscape AI commoditization is driving. In other words, the window for first-mover advantages in process innovations is narrowing. Therefore, while investing in current operations isn't to be underestimated, the true key to competitive differentiation now lies in bold, transformative leaps forward.

How decision-makers think differently about agentic AI

The organizations achieving genuine transformation are those that view agentic AI not as a tool to do existing work faster but as a catalyst to do entirely new work efficiently. Consider the compliance challenge: in highly regulated industries, agentic AI isn't just flagging risks—it's using teams of specialized agents to continuously scan, interpret, and act on regulatory changes across multiple jurisdictions and industries. Imagine one set of agents monitoring shifting financial rules in Europe, another tracking evolving trade sanctions in Asia, and yet another adjusting healthcare compliance workflows in the U.S. All of them work in concert, automatically updating communications, documentation, and approval processes in real time. The result? Risks can be mitigated before they materialize, compliance teams spend less time on repetitive reviews, and organizations can operate more confidently across geographies without drowning in legal complexity.

Because they are focused on developing net-new capabilities, transformative organizations are asking different questions: What becomes possible when systems can make decisions autonomously? How do we redesign our value creation processes around this capability?

The measurement gap reveals this strategic divide clearly. Only 42% of the process-oriented organizations identified in our analysis have developed new key performance indicators (KPIs) to monitor AI agents' impact on business targets, compared to nearly half of transformational organizations (see Action Guide for examples of new KPIs). These leaders aren't just implementing agentic AI automation—they're measuring different outcomes entirely.

“When you don't have good, objective KPIs, older ones take over. We see that technology is often deployed around creating improvements with the infrastructure you currently have.”

Board advisor, United Kingdom

Perspective

Leaders can experiment, swap out underperforming agents, and stitch together multi-agent networks that collaborate.

Plug in, level up: The agentic AI marketplace

An agentic AI marketplace is a digital marketplace where, instead of apps or software licenses, you're browsing virtual shelves of autonomous problem-solvers. Each "product" is an agentic AI—a specialized algorithm that can take action, make decisions, and adapt to changing conditions without waiting for human approval.

These marketplaces are more than procurement hubs; they're ecosystems of ready-to-deploy intelligence. Need a compliance watchdog that scans global regulations in real time? A dynamic pricing engine that recalibrates with every inventory shift? A supply-chain planner that reroutes shipments before bottlenecks appear? In the agentic model, these aren't R&D dreams—they're plug-and-play capabilities you can activate in days.

Once futuristic, agentic AI marketplaces are already operating today—from Agent.ai and the Lyzr AI Marketplace to enterprise platforms like AWS Marketplace and Google Cloud's Vertex AI Agent Builder. These hubs offer curated, ready-to-deploy agents across finance, healthcare, marketing analytics, and conversational AI, making the "app store for AI" a commercial reality.

For executives aiming beyond incremental process tweaks, marketplaces offer speed and flexibility to compose entirely new capabilities for redesigned agentic AI operating models. Leaders can experiment, swap out underperforming agents, and stitch together multi-agent networks that collaborate—with humans stepping in where judgment and expertise matter most.

But the opportunity comes with a caveat: without disciplined curation and governance, marketplaces can flood operations with "black-box" decision-makers that undermine trust. The smartest leaders will treat them as curated innovation networks—demanding transparency in every choice, alignment with strategic KPIs, and readiness to integrate into redesigned operating models.

In the autonomous future, the marketplace won't just be where you buy tools. It will be where you assemble the next generation of business capabilities—the ones you couldn't have imagined five years ago, but will depend on five years from now.

For leaders, this isn't just about what you can buy — it's about defining your role in the ecosystem. Will you participate as a buyer, develop and sell your own agents, or orchestrate an entire marketplace in your industry? The answer will shape how you capture value in the emerging "agentic commerce" economy.

“Critical thinking skills become even more important in the age of AI. We’re seeing it now as people use gen AI and are advancing with agentic AI.”

Senior Vice President, Demand & Digital Transformation,
multinational consumer products corporation

Are leaders focused on the wrong AI challenge?

The C-suite primarily has seen AI through a single lens: a tool for productivity, a way to cut costs and automate jobs for efficiency. The assumption has been that the biggest challenge is the tech itself—the algorithms, the data, the deployment.

But as AI becomes truly autonomous and agentic, that assumption is crumbling. The readiness of the workforce to evolve with AI is emerging as a make-or-break factor for transformation. But readiness isn’t just about skills—it’s also about mindset. Cultural resistance can quietly derail even the most sophisticated AI deployments if people feel threatened, excluded, or unclear about their role in an AI-driven future. The shift to agentic AI introduces new risks of workforce disruption, from fears of job loss to uncertainty about career paths. Leading organizations are tackling this head-on, pairing AI rollouts with intentional change management strategies: transparent communication about how roles will evolve, early involvement of employees in AI-enabled workflow design, and clear pathways for reskilling and upskilling. Rather than treating AI as something done to the workforce, they position it as something built with the workforce—turning potential skeptics into active champions.

While a staggering

47%

of organizations cite inadequate employee skills as a barrier to agentic AI implementation, the most successful companies are making a strategic bet on their people.

An overwhelming

79%


of leaders now believe they need to protect and value the very skill that will help differentiate them as algorithms become commoditized: human critical thinking.

In an agentic AI enterprise, human critical thinking reveals itself in the moments when people challenge, refine, or even override autonomous decisions—and can articulate why. Measuring it means tracking the quality of those interventions and the outcomes they shape—which is why new KPIs are fast becoming critical.

All organizations are facing new talent issues due to AI and many are focused on creating entirely new roles. Think human roles—AI orchestrators and collaborators who actively decide when and how AI agents engage in a process, or autonomous system auditors who trace decision paths and validate compliance before execution. In legacy models, these checks and choices were buried in slow, linear chains of approvals; now they happen dynamically, in parallel with AI, turning humans into strategic directors rather than task executors. This is a profound shift from productivity to value creation.

FIGURE 2**Executives are already changing team structure with new roles**

Top 5 things execs are already doing to staff an operating model designed with agentic AI automation in mind.



“We do not want to attract the best talent; we want to attract the talent that has the best attitude to grow their skills.”

Julián Mora Gómez, Corporate Vice President, Bancolombia

The trust deficit: AI systems need to show their work

How does autonomous AI make its decisions? This logical question is often the key stumbling block to embracing AI at scale. In fact, 45% of executives in our study cite a lack of visibility into agent decision-making processes as a significant implementation barrier. Yet, this isn't just a technical limitation—it's a design choice.

Leading organizations tackle this with business-friendly practices: Machine learning operations (MLOps) to keep AI models healthy and up to date, observability and logging to track and explain every decision, and A/B testing to compare outcomes and see which AI-driven approach moves the needle on key KPIs. All of this toward a desired result—AI that's transparent, measurable, and continuously improving.

Organizations serious about transformation are engineering transparency into their systems from the ground up, ensuring professional oversight remains meaningful even as agentic AI becomes more autonomous across enterprise operations. It's easy to write off the “black box” problem as a technical curiosity; but it's really a trust deficit. Leaders who want to scale agentic AI must address this head-on, building systems where every automated decision, every transaction, and every customer, partner, or employee communication can be understood, audited, and explained.

This workforce evolution requires leadership commitment that extends beyond technology budgets. C-suite executives in leading organizations are actively identifying, recruiting, and training new roles for autonomous automation orchestration. They're not just hiring for today's needs—they're developing capabilities to support an operational model that doesn't yet exist—as they build it.

Operating model redesign: The competitive separator

The organizations pulling ahead aren't merely experimenting with agentic AI; they are fundamentally redesigning their operating models around it. The focus is shifting from automating transactional processes to tackling the most difficult, high-leverage problems.

Roughly three in four leaders (76%) indicate that focusing on these complex challenges is more likely to yield a competitive advantage. This insight cuts to the heart of transformational thinking: using agentic AI to solve previously unsolvable problems rather than simply surmounting solvable problems faster.

Data infrastructure becomes the foundation of this new operating model, but the requirements extend far beyond traditional data management. Agentic AI implementation faces challenges from:

**Data privacy
and security**

65%

**Data integration
complexities**

60%

**Data quality
issues**


56%

These aren't just IT problems—they're strategic barriers that require enterprise-wide solutions.

Transformational organizations are realizing measurably greater impact across every meaningful business metric: productivity, efficiency, revenue growth, brand perception, R&D ROI, customer loyalty, and employee sentiment all exceed performance in transformation-focused organizations. This isn't correlation—it's the direct result of operating model choices that maximize AI's transformational potential rather than limiting it to process improvement and optimization.

FIGURE 3

Evolving a synergistic “always on” agentic AI operating model



Differentiators of a synergistic agentic AI operating model

- Data is a continuous feed and refresh. As people and agents execute workflows, data is refreshed by new outcomes—a continuous loop.
- People have new AI impactful roles, new expertise, new careers and growth paths.
- Autonomous agents, learn, adapt and evolve in multi-agent communities.
- People and agents learn together as they perform work in collaboration with a focus on outcomes and impact.
- New KPIs monitor the impact of agentic AI on business targets.
- Performance is constantly monitored in workflows, experiences and outcomes.
- All enterprise workflows are seamlessly interconnected. This model transcends functional boundaries—functions no longer exist, but expertise does.
- Ethics are embedded into AI deployments.
- Infrastructure is secure, interoperable and transparent.
- Cybersecurity is integrated into AI initiatives.
- Specific small language models are implemented into workflows for precision, transparency, and speed to action.
- As new technologies, market trends and disruptions occur, the operating model synergistically provides feedback loops, adjusts algorithms, readies, and innovates itself.

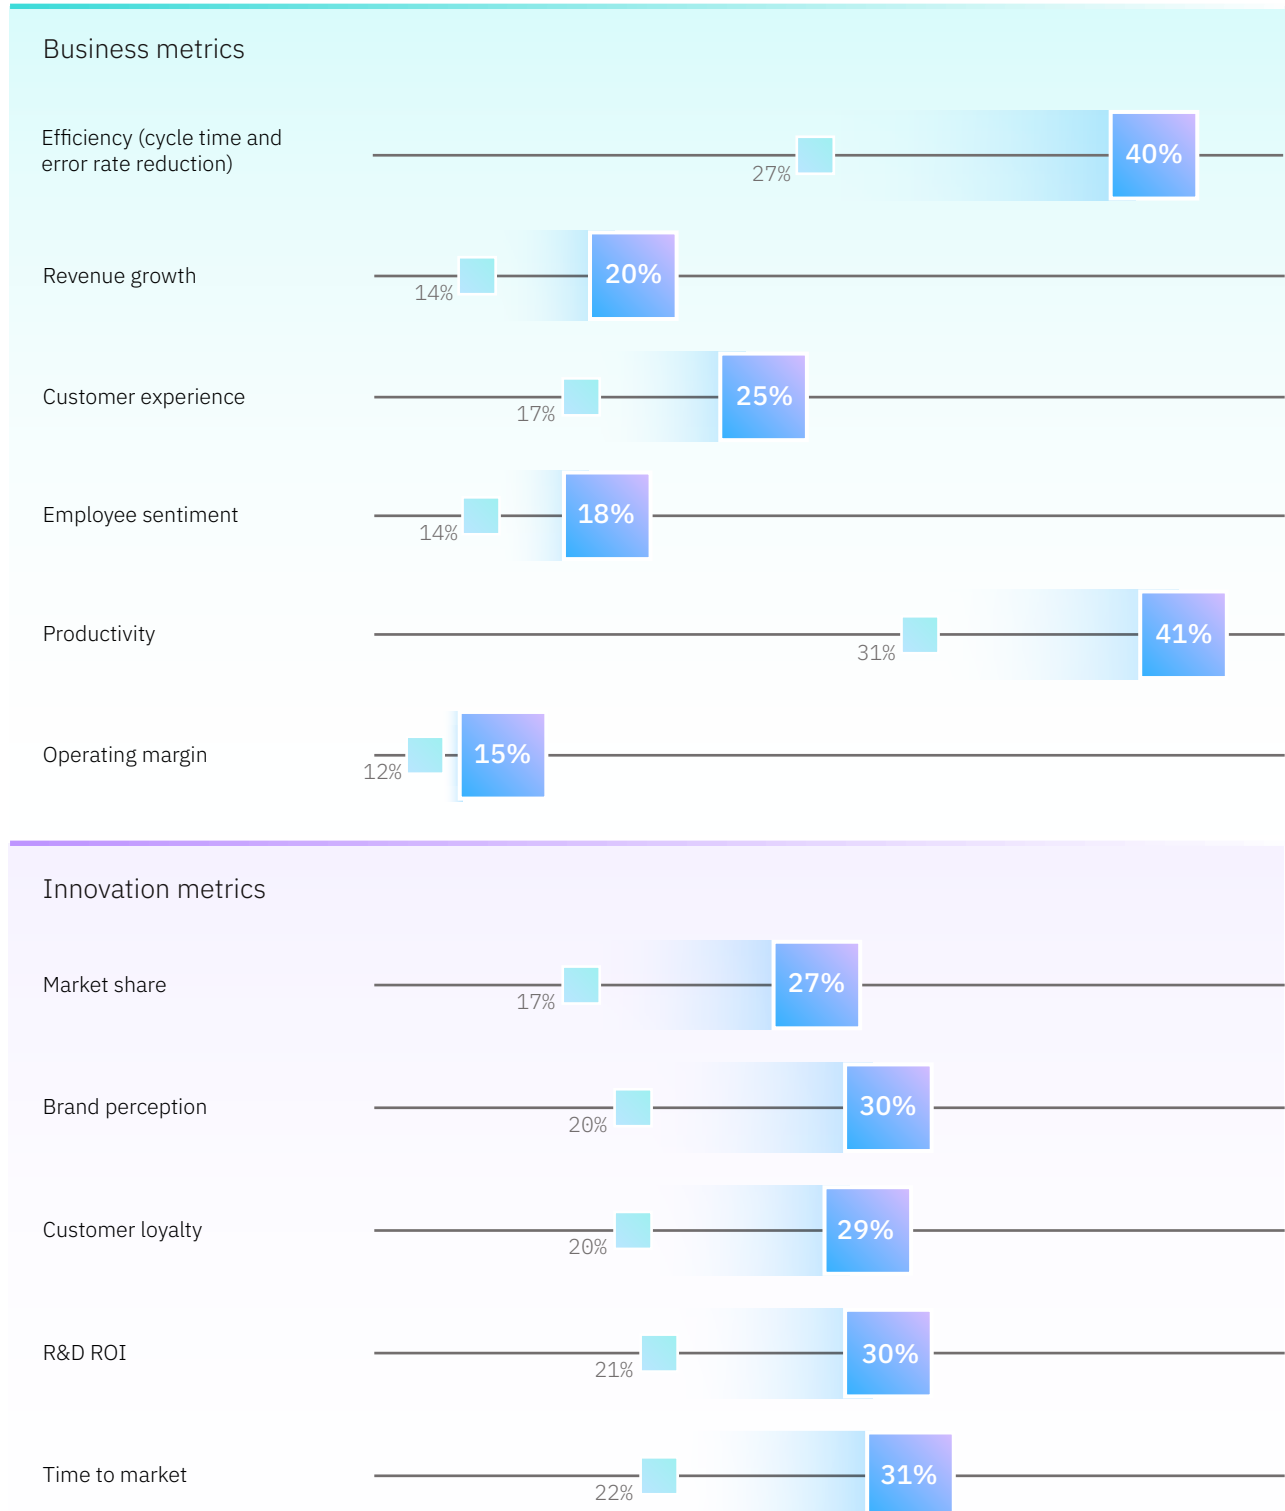
FIGURE 4

Transformers show solid performance across all KPIs

Transformational visionaries are an elite group identified in our advanced analytics regression analysis. They pursue AI initiatives focused on creating net-new workflow capabilities and they succeed in executing those initiatives at a high rate. This group accounts for 17% of the survey sample. They are on plan to achieve significantly greater KPI improvements by 2026 compared to all others.

To what extent will agentic AI improve the following metrics in your organization by 2026

Transformers All others



Perspective

Who is 32 times more likely to achieve top-tier business performance?

The analysis reveals a clear path to AI excellence that every executive should understand. Transformational visionaries that excel in three key AI adoption areas are 32 times more likely to achieve top-tier business performance than those with minimal implementation.

Organizations that move from minimum to maximum adoption across these three see exponential performance gains—not incremental improvements. This isn't about doing more; it's about doing the right things exceptionally well.

The three critical drivers

1

Integrating cybersecurity into AI initiatives.

When AI systems can make decisions independently, the cybersecurity stakes change entirely. These autonomous agents don't just process data—they act on it, multiplying both their potential impact and their vulnerability to exploitation. A compromised agentic system becomes a weapon in the hands of malicious actors, capable of making countless harmful decisions before humans even detect the breach. Robust cybersecurity isn't just a technical requirement for these systems—it's the foundation that determines whether autonomous AI becomes a competitive advantage or an existential threat.

2

Embedding ethics analysis in AI deployments.

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. Without proper ethical guardrails, AI agents can inadvertently perpetuate biases, violate privacy, or make choices that seem logical to an algorithm but feel fundamentally wrong to the people affected. Ethics analysis creates the framework that keeps autonomous decisions aligned with fairness, accountability, and transparency. Perhaps most importantly, it builds the public trust that these systems need to succeed—showing that their unprecedented decision-making power has been carefully scrutinized for unintended consequences.

3

Implementing workflow-specific small language models.

Workflow-specific small language models are the specialized translators that make agentic AI systems truly effective in real business environments. Rather than relying on generic AI that struggles with industry jargon and unique processes complexities (like real-time integrated planning), these tailored models understand the specific terminology, procedures, and context that define how work actually gets done. They excel at breaking down complex workflows into logical steps, prioritizing tasks, and allocating resources—essentially becoming the operational intelligence that bridges the gap between AI capability and business reality. The result is agentic AI that doesn't just work in theory, but integrates seamlessly into the messy, nuanced world of actual enterprise operations.

Action guide

Technology deployment is important, and one of the first strategic calls is whether to build proprietary agentic systems or adopt vendor solutions. But the real differentiator comes from moving beyond deployment into the organizational and cultural transformation needed to make a new-era operating model work. Get started with a few key actions:

01

Set a clear AI goal for each area of your business and decide how you'll judge investments based on the value they deliver—making sure every AI effort ties directly to core KPIs and long-term priorities.

Appoint executive sponsors accountable for the P&L impact, not the technology. Assign specific P&L owners (e.g., Head of Operations, CFO) to be personally accountable for achieving the targeted business outcomes (e.g., a 15% cost reduction), holding them responsible for AI-enabled results, not just implementation schedules.

- Initiate parallel projects focused on securing high-quality, standardized data access and aggressively leading the required talent shift (upskilling, critical hiring, external expertise) necessary to execute the prioritized roadmap and sustain long-term growth. This will help mitigate enterprise risk.

02

Use AI agents now in enterprise operations areas. Begin with high-value areas such as finance, supply chain, HR, and customer experience.

Move beyond pilots and deploy autonomous capabilities that drive significant, measurable change.

- Deploy financial modeling AI agents to continuously scrutinize data for precise cash flow projections and budget variance forecasts. Roll out simulation models to run constant, real-time risk and compliance surveillance.
- Introduce dynamic pricing and forecasting tools that leverage ongoing inventory visibility to assess market actions and reactions in real time.
- Empower sales and marketing teams with AI-driven tools for lead prioritization and personalized marketing.
- Deploy predictive models in HR for recruitment, talent skilling, and personalized learning. Ensure your workforce is not only prepared for current needs but for the strategic demands of tomorrow.
- Aim for customized interactions across all transactions and communications to enhance experiences for customers, employees, and business partners. Explore how AI can personalize and streamline every touchpoint.

03

Re-engineer your measurement framework.

Traditional KPIs are artifacts of a pre-autonomous age. They measure human activity and efficiency gains in a world where humans are the primary actors. Agentic AI requires a new set of KPIs that monitor the outcomes of automated decision-making.

- Don't just measure productivity; measure new value creation, business growth, and the velocity of innovation.
- Start by asking a fundamental question: "What new value can we create with autonomous systems that we couldn't before?"
- Create metrics like "agent-to-human handoff rates" to understand where agents are failing to solve problems autonomously.
- Implement new KPIs such as "reasoning coherence scores" or "decision accuracy rates" to ensure the integrity and reliability of autonomous actions.

04

Create a reciprocal learning culture.

Your workforce needs to be equipped with a new set of skills focused on teaching, training, monitoring, and providing feedback to AI systems. This is not about basic literacy but about advanced collaboration. For example, an engineer might train a visualization model to spot defects, and in return, the AI system's analysis of millions of data points reveals a previously unknown micro-fracture pattern, teaching the engineer about a new area of material stress that needs to be addressed.

- Identify and recruit new roles for autonomous automation orchestration. These positions require individuals who understand both business strategy and AI capabilities—professionals who can bridge the gap between human judgment and machine learning.
- Train employees in AI collaboration now for compound value later. As people become more skilled at AI collaboration, the AI systems become more effective at supporting human critical decision-making.

05

Develop granular visibility into agent decisions, ensuring it's done with a governance and ethics lens.

Trust in autonomous systems is not a given; it must be built through transparency. Senior leadership, regulators, or employees will not tolerate the opaqueness of a “black box.” You must demand granular visibility into every agent decision.

- Establish comprehensive “decision logging” that captures every autonomous choice, the confidence score, alternative options considered, and the data inputs that influenced the outcome.
- Deploy observability tools that provide real-time dashboards visualizing agent actions, external tool usage, and information flows leading to decisions. These platforms must offer granular visibility without creating overwhelming data volumes that obscure rather than illuminate system performance.

A tale of two enterprises

While many organizations are still just beginning their journey with tactical AI, the companies that will lead the next decade are already playing a different game. They're not focused on simply deploying the most sophisticated technology. Instead, they're using foresight to rebuild their operating model around autonomous systems, preserving human expertise where it is most crucial. It's a total redefinition of the relationship between people, data, decisions, actions and impact.

**The defining question for today's C-suite:
Which enterprise are you running?**

Authors

Francesco Brenna

Global Vice President and Senior Partner,
Artificial Intelligence, IBM
[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.

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.

Manish Goyal

Vice President and Senior Partner, Global AI and Analytics Leader, IBM Consulting
[linkedin.com/in/goyalm](https://www.linkedin.com/in/goyalm)

Manish blends deep product innovation experience—from pioneering Watson healthcare solutions to advancing enterprise machine learning platforms—with strategic leadership to help organizations transform and capture the full value of AI disruption.

Anant Patel

Product Architect—Oracle Generative AI Agents, Oracle

Anant applies his 10+ years of experience in creating and scaling products in AI, machine learning, software, and healthcare to develop pioneering solutions, drive impactful services, and steer strategic market strategies, always guided by a customer-focused mindset and the data.

Federico Torreti

Sr. Director, Product Generative AI, Oracle
[linkedin.com/in/ftorreti](https://www.linkedin.com/in/ftorreti)

Federico is a product executive driving enterprise AI innovation, leading generative AI and machine learning initiatives that enhance workflows and amplify human decision-making.

Shobhit Varshney

Former Head of Data & AI, IBM
[linkedin.com/in/shobhitvarshney](https://www.linkedin.com/in/shobhitvarshney)

Shobhit leads Citi's global AI strategy and adoption, building on over a decade at IBM Consulting where he spearheaded large-scale data and AI transformations across the U.S., Canada, and Latin America.

Contributors

Sara Aboulhosn, Associate Creative Director, IBM Institute for Business Value

Jen Basco, Global HR Leader, IBM

Angela Finley Design Lead, IBM Institute for Business Value

Corinne Koppel, Oracle Global Practice Leader, IBM

Christopher Nowak, Research Leader, IBM Institute for Business Value

Emily Paez, Global Partner Marketing, IBM

Kristine Rodriguez, Editor-in-Chief, IBM Institute for Business Value

Nitish Patel, Global OCI Offering Leader, IBM

Arunava Saha, Global Oracle Offering Leader, IBM

Thiago Sartori, Data Scientist, IBM IBV

Kankayan Segupt, Global Oracle CTO, IBM

Angela Ugrotto, Global Marketing, Oracle

Andrew Womack, Creative Director, IBM Institute for Business Value

800

C-suite executives: CFO, CIO, CHRO, CSCO, COO, CPO, Chief customer/Customer service officer

20

Countries

19

Industries

5%
Canada

20%
US

3%
Mexico

5%
Brazil

3%
Sweden

10%
UK

5%
Germany

4%
France

3%
Italy

6%
UAE

5%
Saudi Arabia

5%
India

2%
South Africa

5%
China

6%
Japan

6%
Taiwan

5%
S. Korea

5%
Singapore

4%
Australia

Methodology

Fieldwork

The study was conducted in Q3 2025 with a global survey targeting senior executives who have direct responsibility or significant influence over shaping, managing, and investing in agentic AI initiatives and AI-driven workflow capabilities.

The survey achieved a final sample of $n=810$ respondents across 20 countries and multiple languages, meeting targeted quotas for geography, role, industry, and revenue brackets.

Geographic distribution included: Australia, Brazil, Canada, China (PRC), France, Germany, India, Italy, Japan, Mexico, Saudi Arabia, Singapore, South Africa, South Korea, Spain, Sweden, Taiwan (ROC), United Arab Emirates, United Kingdom, and United States.

Respondent profile: Executives spanned operations, technology, finance, supply chain, procurement, human resources, and customer service roles—including COO, CIO, CFO, CSCO, CPO, CHRO, and CCSO equivalents—and represented 19 industries from aerospace and defense to telecommunications. Organizations ranged from USD\$500 million to over \$20 billion in annual revenue.

Multinomial Dependent Variable—Q25A & Q25B Index

The dependent variable for the analysis was a combined Q25A and Q25B index, representing an ordinal performance outcome.

- Q25A and Q25B captured the organization's current level of performance and adoption maturity in agentic AI initiatives, particularly those creating new workflow capabilities with measurable business impact.
- Responses were binned into four performance tiers: Lowest, Mid-Low, Mid-High, and Highest.

- The model was designed to evaluate upward movement likelihood—identifying factors most strongly associated with transitioning from lower to higher tiers.

Analytical Approach

The research team employed a combination of descriptive statistics and multivariate modeling:

- Descriptive analysis summarized distributions across geography, role, industry, and revenue.
- Cross-tabulations compared differences between performance tiers.
- Multinomial logistic regression was used to estimate the relative odds of being in higher performance tiers compared to the lowest tier (reference category).

Model Validation and Goodness-of-Fit

- Likelihood ratio chi-square tests confirmed the model provided a statistically significant improvement over the null ($p < 0.01$).
- Multiple predictors were statistically significant ($p < 0.05$), indicating they reliably explain variation in performance tier outcomes.
- Pseudo R^2 (Nagelkerke) values indicated a moderate-to-strong effect size, suggesting the model explains a meaningful proportion of the variance in upward movement.
- The model's classification accuracy exceeded baseline proportions, demonstrating practical predictive utility.

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.

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 IBM and Oracle

For nearly four decades, IBM has been partnering with Oracle to help clients navigate their cloud journeys and maximize their business transformations in a complex hybrid and multicloud world. IBM is a leading hybrid cloud, AI, and consulting services provider helping clients in more than 175 countries capitalize on insights from their data, streamline business processes, and gain the competitive edge in their industries. Oracle offers integrated suites of applications plus secure, autonomous infrastructure in the Oracle Cloud. IBM's breakthrough innovations in AI, quantum computing, and business services deliver open and flexible options to our clients.



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





© Copyright IBM Corporation 2025

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | October 2025

IBM, the IBM logo, and [ibm.com](https://www.ibm.com) 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](https://www.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.

The data used in this report may be derived from third-party sources and IBM does not independently verify, validate or audit such data. The results from the use of such data are provided on an "as is" basis and IBM makes no representations or warranties, express or implied.