

Enterprise generative AI

State of the market



Harnessing the power of generative AI

Executives know they need generative AI—but there’s a gap between awareness and the ability to deliver value at scale.

The generative AI debate is often reduced to extremes. As enthusiasts and skeptics struggle to control the conversation, the middle ground often gets lost. And that’s a big problem for business leaders, who need to base their approach to AI on cold, hard facts and strategic relevance—not hubris and hyperbole.

To offer a more sober assessment, the IBM Institute for Business Value (IBM IBV), in partnership with Oxford Economics, surveyed almost 400 executives across the US, Australia, Germany, India, Singapore, and the UK in May 2023. At the same time, we also surveyed 200 CEOs in the US. We asked business leaders about their plans for adopting generative AI, benefits they expect to see, and barriers holding them back.

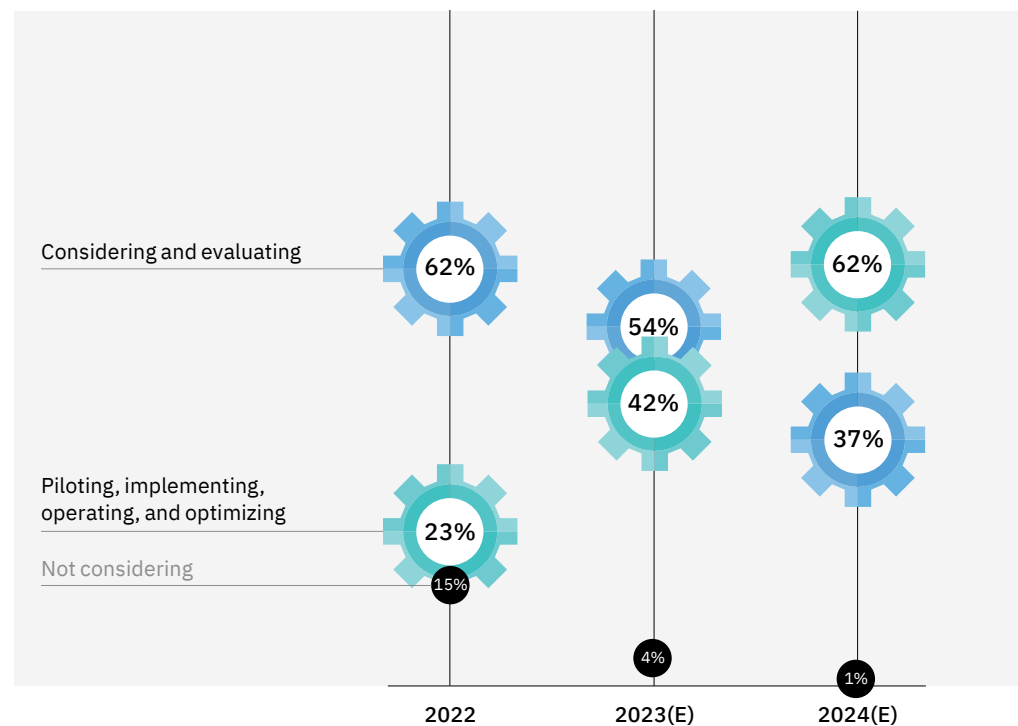
We found that, surrounded by the continued hype, executives are indeed embracing generative AI—but with caution.

The average ROI of generative AI projects is edging up, with execs telling us they expect it to exceed 10% by 2025, continuing a rising trend of surpassing their cost of capital with AI.¹ As a result, enterprises are planning to boost generative AI adoption over the next two years. While only 23% of executives say their organization piloted, implemented, operated, or optimized generative AI in 2022, that figure is expected to rise to 62% by 2024 (see Figure 1).

FIGURE 1

Generative AI comes of age

Execs expect adoption to nearly triple between 2022 and 2024



This rapid uptick is partially fueled by familiarity: Leaders today are better acquainted with generative AI than they were with traditional AI in 2016, amid the initial AI hype cycle.² As a result, executives have a much more focused view of where to deploy generative AI.

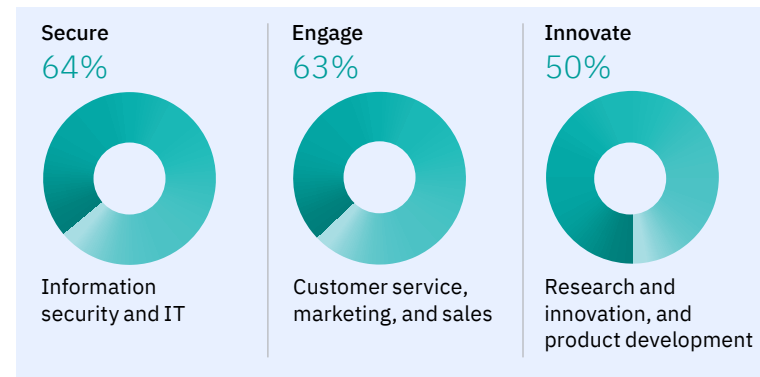
In 2016, business leaders were unfocused, prioritizing adoption of AI across as many as 13 different business functions. But strategy is as much about what you choose to do as what you choose *not* to do. Today, business leaders are much clearer on what they see as driving the most value. Our data shows that, at least initially, leaders are focused on three key priorities (see Figure 2).

Of course, organizations are still tailoring their work in these broad functional areas to their own strategic capabilities and business priorities. For example, IBM’s Chief Analytics Office—a critical part of its broader Transformation & Operations area—is focusing on applying generative AI in IT and application modernization, customer service, and employee engagement, all underpinned by automation.³

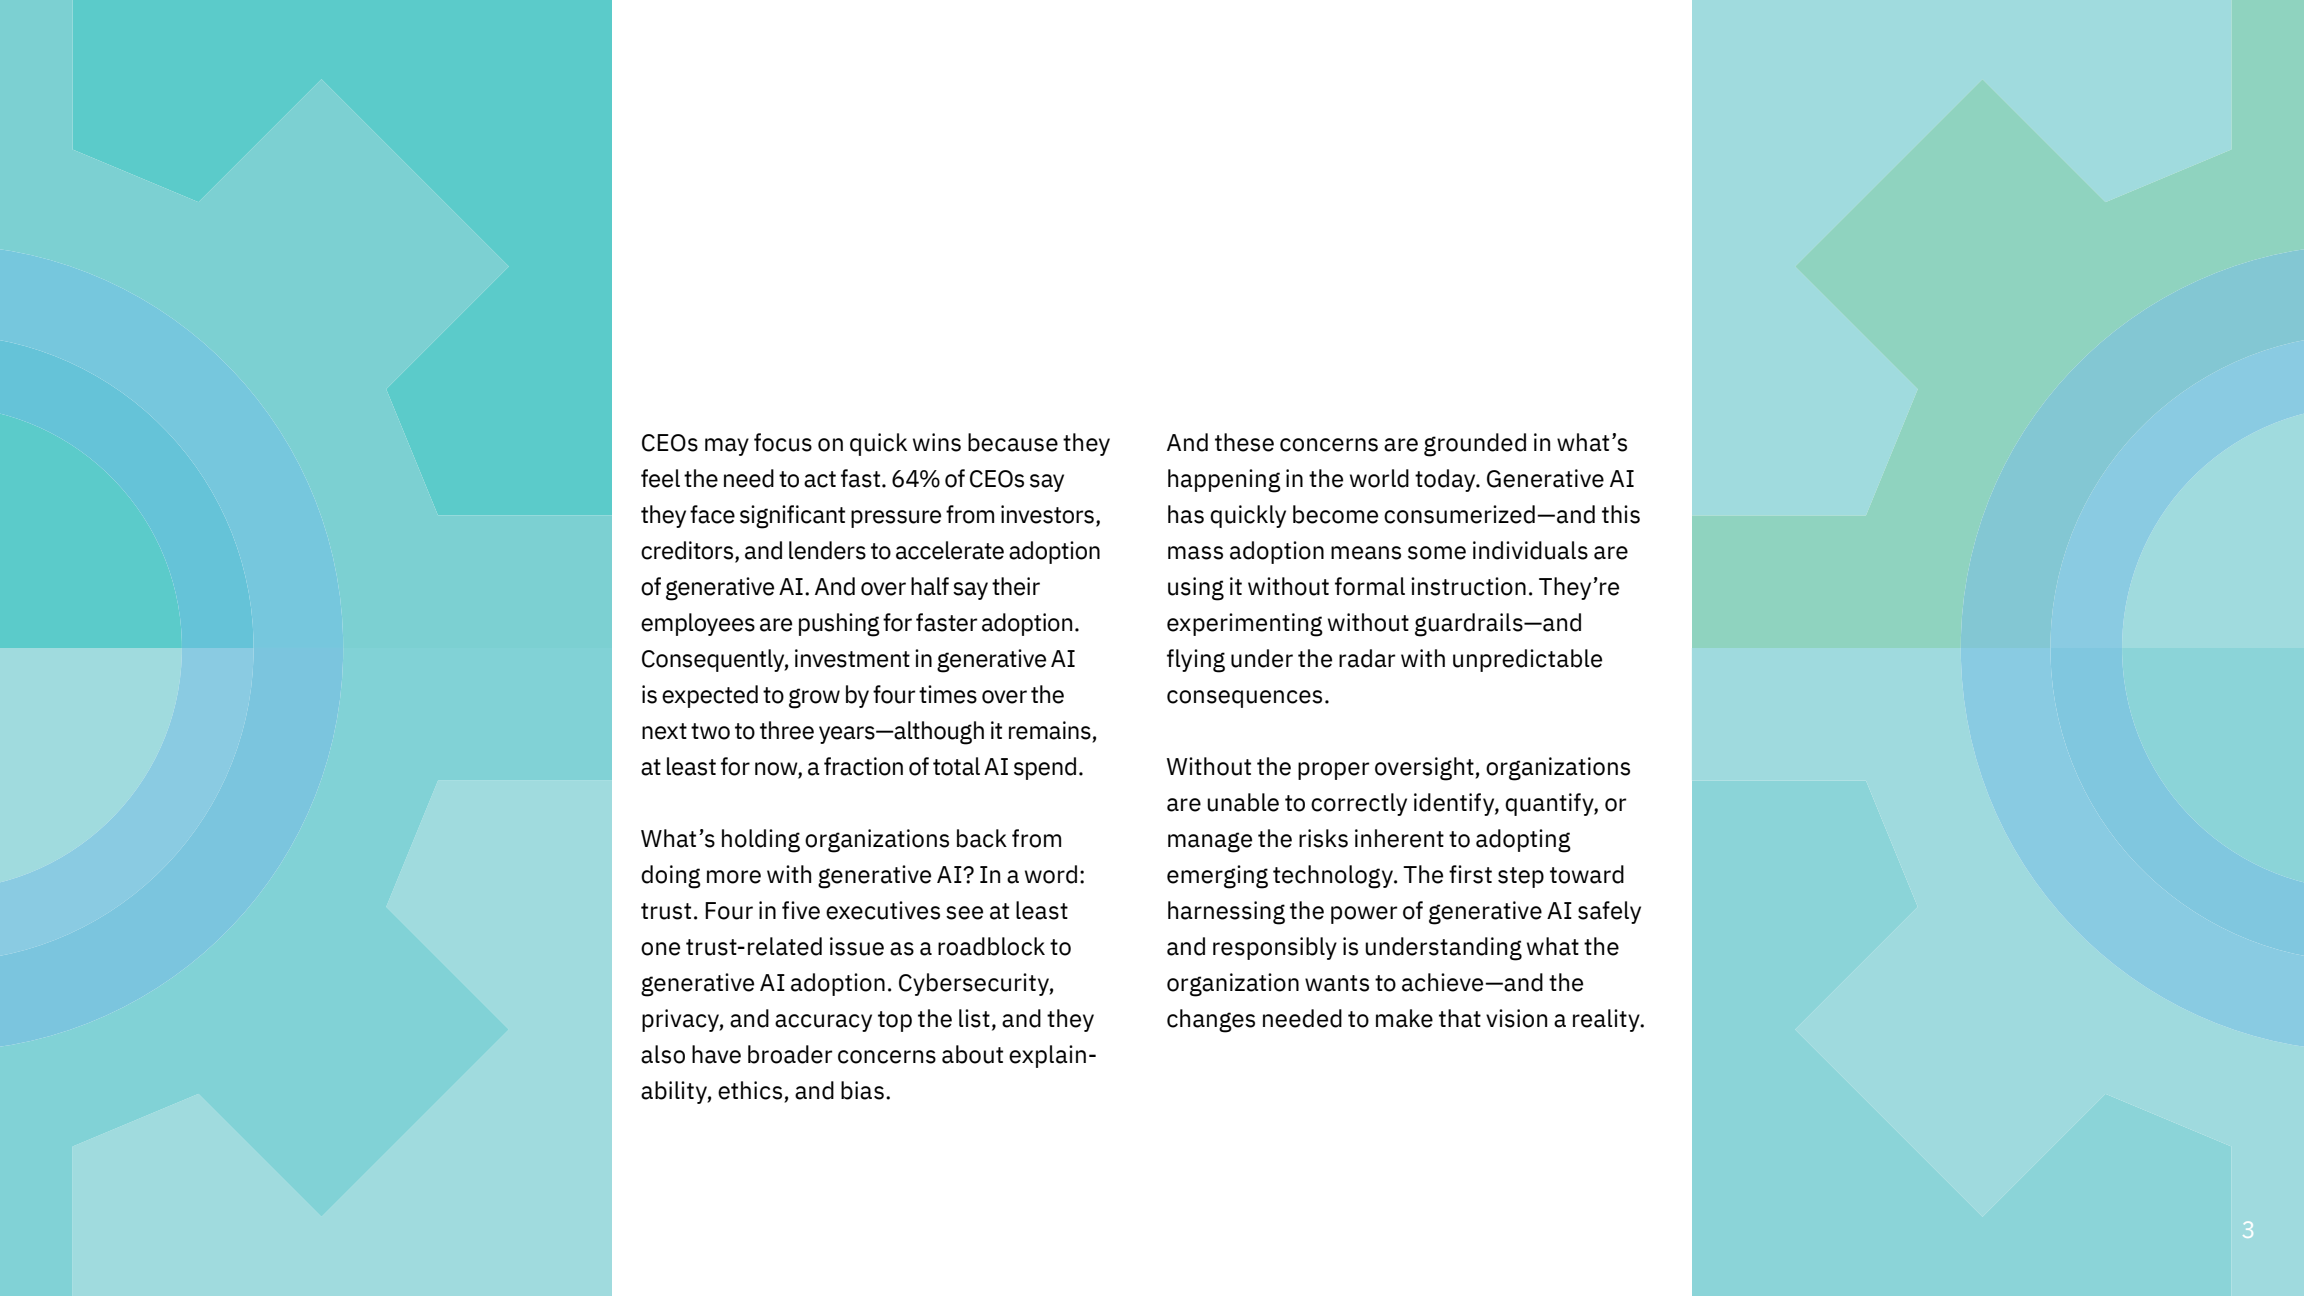
FIGURE 2

A clearer focus

Execs have identified three priorities for generative AI adoption



Awareness has reached the C-suite quickly, but there is a gap between this awareness and the ability to deliver value at scale. The priority areas that executives identified are those with the most mature AI capabilities—not necessarily strategic pain points. This means many organizations are falling short of executing AI initiatives that can deliver strategic value for the entire enterprise.



CEOs may focus on quick wins because they feel the need to act fast. 64% of CEOs say they face significant pressure from investors, creditors, and lenders to accelerate adoption of generative AI. And over half say their employees are pushing for faster adoption. Consequently, investment in generative AI is expected to grow by four times over the next two to three years—although it remains, at least for now, a fraction of total AI spend.

What's holding organizations back from doing more with generative AI? In a word: trust. Four in five executives see at least one trust-related issue as a roadblock to generative AI adoption. Cybersecurity, privacy, and accuracy top the list, and they also have broader concerns about explainability, ethics, and bias.

And these concerns are grounded in what's happening in the world today. Generative AI has quickly become consumerized—and this mass adoption means some individuals are using it without formal instruction. They're experimenting without guardrails—and flying under the radar with unpredictable consequences.

Without the proper oversight, organizations are unable to correctly identify, quantify, or manage the risks inherent to adopting emerging technology. The first step toward harnessing the power of generative AI safely and responsibly is understanding what the organization wants to achieve—and the changes needed to make that vision a reality.

What is generative AI? What are foundation models?

To make AI investments more cost-effective, companies need flexible, reusable models that can be applied in a variety of ways—including generating new content. Generative AI refers to deep-learning models that can generate high-quality text, images, and other content based on the data they were trained on.⁴ At a high level, generative models encode a simplified representation of their training data and draw from it to create a new work that's similar, but not identical, to the original data.

This technology began emerging in the early 2010s, when variational autoencoders, or VAEs, became the first deep-learning models widely used for generating realistic images and speech. Autoencoders work by encoding unlabeled data into a compressed representation, and then decoding the data back into its original form. Plain autoencoders were used for a variety of purposes, including reconstructing corrupted or blurry

images. Variational autoencoders added the critical ability to not just reconstruct data, but to output variations on the original data.

This ability to generate novel data ignited a rapid-fire succession of new technologies, from generative adversarial networks (GANs) to diffusion models, capable of producing ever more realistic—but fake—images. In this way, VAEs set the stage for today's generative AI. They are built out of blocks of encoders and decoders, an architecture that also underpins today's large language models (LLMs). Encoders compress a data set into a dense representation, arranging similar data points closer together in an abstract space. Decoders sample from this space to create something new while preserving the data set's most important features.

Foundation models can help enterprises accelerate and scale generative AI adoption.

Transformers combine the encoder-decoder architecture with a text-processing mechanism. An encoder converts raw text into representations known as embeddings. The decoder takes these embeddings together with previous outputs of the model, and successively predicts each word in a sentence. Through fill-in-the-blank guessing games, the encoder learns how words and sentences relate to each other, without anyone having to label parts of speech.

Transformers can even be pre-trained without a particular task in mind. Once these powerful representations are learned, the models can later be specialized—with much less data—to perform a given task. Transformers have come to be known as foundation models for their versatility.

Foundation models offer an opportunity to accelerate and scale generative AI adoption, as they can, in theory, be applied to many domains.⁵ For example, massive-parameter LLMs can transform how information is generated and shared across an organization. (Parameters are the variables used in training a model, useful in inferring new content.) It just needs to be

adapted for semantic search, classification, prediction, summarization, and translation.

The adoption of foundation models is also supported by a set of emerging AI engineering best practices that have gone mainstream. From model development to prompt engineering, these common practices and approaches streamline collaboration across the enterprise—and the ecosystem.

While foundation models offer real promise and potential, they also come with new challenges. For one, they require significant compute, storage, and network resources, which makes them energy intensive. Training one large natural language processing model has roughly the same carbon footprint as running five cars over their lifetime.⁶

Like other disruptive technology, there are trade-offs that come with adopting generative AI and foundation models.⁷ Success only comes from experimentation and iteration. Especially for enterprises, this journey involves balancing the value generative AI can create against the investment it demands. The future of AI will be defined by those who hit the right mark.

Case study

Lloyds Banking Group improves customer service with large language models

When people need information about their finances, they don't want to be kept waiting. But customers of Lloyds Banking Group (LBG) were not getting the answers they needed when using the search function on its mobile app. Instead, they were deferred to other communication channels to get help, resulting in an increased cost-to-serve and a worse customer experience.⁸

Because search relied on keyword matching, this led to issues when customers made spelling errors in even the most common searches. To address these issues, IBM used generative pre-trained transformer embeddings to vectorize FAQs, identify the appropriate response, and generate a synthetic response back to the user.

An IBM team of data scientists, user experience, and strategy consultants are now:

- Building a proof of concept to prove that large language models radically improve the mobile app search experience through reducing the number of unsuccessful searches, improving virtual assistant integration, and personalizing search responses.
- Outlining an industrialization approach, defining the steps to take the proof of concept to a deployed solution, live in customers' hands. They're considering the risks and opportunities ahead, positioning "search" as the front door for all customer self-service.

IBM and LBG have been defining new and innovative ways to use AI to increase customer engagement since 2017. They have used foundation and generative AI models enterprise-wide to understand evolving customer needs and reduce the manual effort associated with managing, training, and executing AI-driven engagement. This work has allowed the bank to tap a store of data that would have cost over £2 million per year—and require thousands of hours of human effort—to analyze otherwise.⁹

Scaling Generative AI

Three keys to success

There remains a lot of speculation around how generative AI can actually deliver business results. But most executives agree on the outcomes they expect.

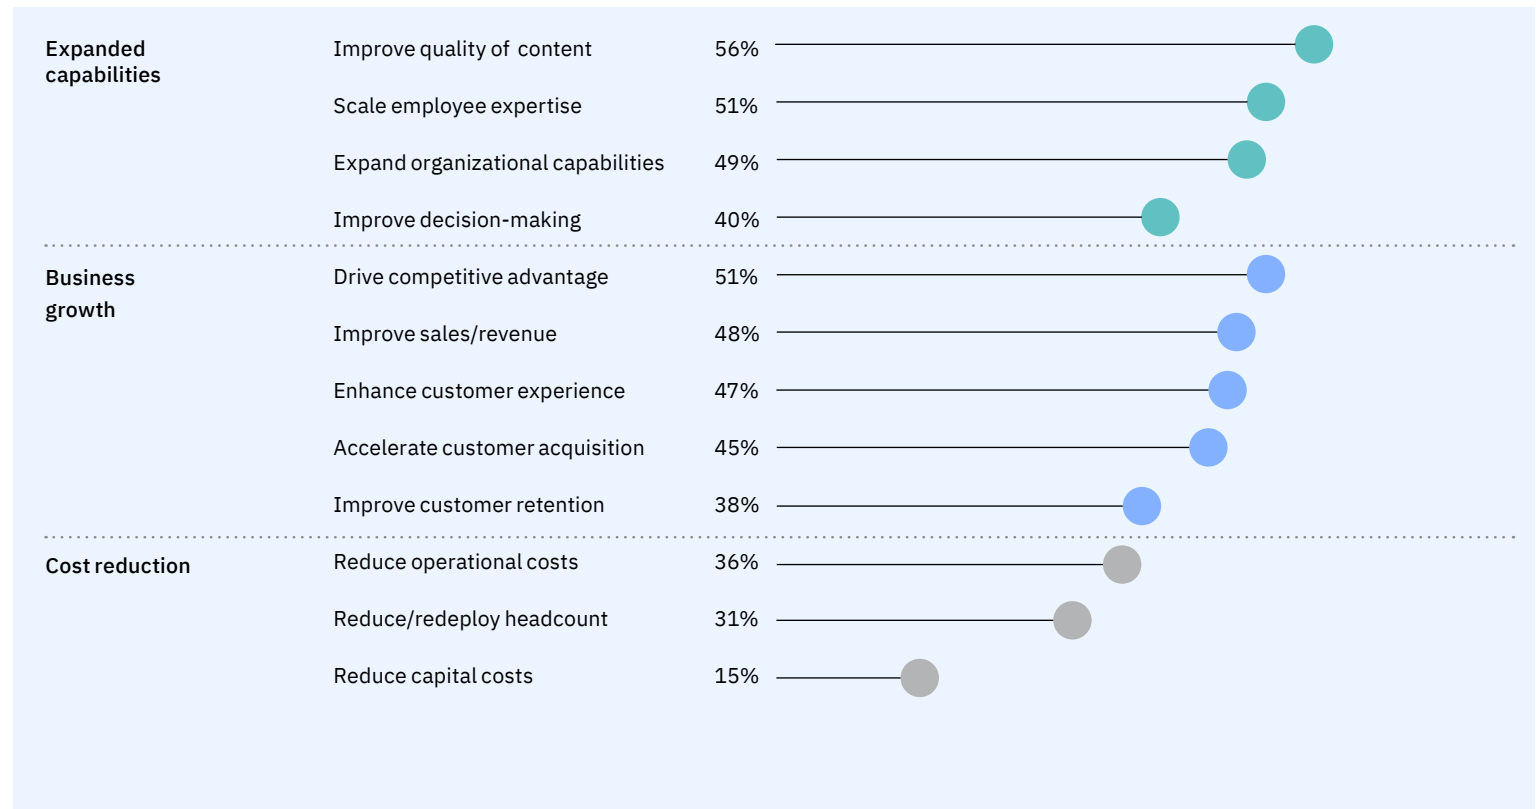
They see a few top opportunities to win by adopting generative AI, including improving quality of content, driving competitive advantage, and scaling employee expertise. Overall, executives are more focused on expanding capabilities and driving growth by enhancing the experience—and less on reducing costs (see Figure 3).

While 74% of CEOs believe their teams have the skills to be successful, only 29% of other C-suite executives agree.

FIGURE 3

Beyond cutting costs

Executives want to boost growth and capabilities with generative AI





The enterprise needs tech-savvy business leaders who can drive change, not just operational efficiency.

Yet not everyone views their enterprise as ready for generative AI. While 74% of CEOs believe their teams have the skills to be successful, only 29% of other C-suite executives agree.

CEOs see the potential—and know their organization has the AI experience and skills to reach the next level. But they may not appreciate other operating model considerations: what’s needed to get it done. Another perspective is that some legacy executives may not have the right skills required to compete in a digital world, deferring quickly to their CDO or CIO. A new breed of tech-savvy business leaders—those who can drive change, not just operational efficiency—may be needed.

To close the readiness gap and accelerate generative AI adoption, businesses leaders need to focus on three areas:

[Organization and skills](#) | [Data and platform](#) | [Risk and governance](#)

Strategy is important too—aligning with business direction, prioritizing functional areas, identifying use cases—as is continuing to build broader capabilities.¹⁰ But we focus here on execution. Especially in the current, fluid environment, moving quickly (but thoughtfully) may help determine what role generative AI can play in the near term.

Organization and skills

As many as 80% of executives see workforce roles and skills changing because of generative AI. (We were surprised that number's not higher—and perhaps it will be soon.) Looking forward, executives are prioritizing the skills that both help employees work with generative AI *and* do the work only humans can do. As generative AI becomes more prevalent, 57% of execs expect creativity skills to become much more important. More than half say the same for technology skills and time management and prioritization.

Overall, generative AI is expected to augment far more employees than it will replace—but not all functions face equal impact. A quarter of execs expect to replace employees in marketing and customer service with generative AI, but only 3% of execs in procurement (see Figure 4).


People on the front lines may be impacted the most—but could also benefit the most.¹¹ To help next-gen employees understand their roles and responsibilities in an evolving workplace, leaders should proactively identify the changes that deploying generative AI will require. Enterprises should also reskill and/or upskill employees with in-demand qualifications and launch incentivization plans to jumpstart professional development.

Replace ■
Augment ■

FIGURE 4

Enter the augmented workforce

More execs expect to augment rather than replace employees



Data and platform

Generative AI models are hungry for data—and feeding them responsibly requires collaboration across the organization. However, 60% of organizations are not yet developing a consistent, enterprise-wide approach to generative AI.

Organizations that are farther along on their hybrid cloud journey are in a better position to capitalize on the promise of generative AI, as cloud transformation demands a more holistic approach to data.¹² But organizations that have leveraged cloud primarily to cut costs in isolated areas may have to go back and connect the dots.

Our research shows that roughly three in five executives say hybrid cloud and generative AI are symbiotic in the value they create. The other 40% are still struggling to get their platforms in lockstep.

Unifying data can be a heavy lift—and without a clear goal in mind, the juice may not seem worth the squeeze. However, hybrid cloud and generative AI platforms that are built on a foundation of reliable data can open the door to new, more valuable partnerships. And nearly two in three executives say generative AI can improve and accelerate data sharing with ecosystem partners.



Case study

Bouygues Télécom relies on generative AI to transform customer relationships

Bouygues Télécom has successfully integrated generative AI into its customer relationship processes, focusing specifically on the summarization of conversations within the telco contact center.

This use case directly tackles the persistent challenge the company faced in accurately summarizing 10-minute conversations, aiming to improve client relationships and personalize services. By harnessing the power of large language models (LLMs) and tailoring them to meet the specific requirements and performance needs of the telecom operator, Bouygues Télécom has developed innovative solutions that significantly enhance customer interactions and unlock the inherent value within recorded conversations.

Prior to this accomplishment, Bouygues Télécom directed its efforts toward improving the transcription of conversations conducted in speech-to-text.¹³ Through the evaluation of various LLMs, they successfully generated concise summaries of client conversations, extracting key points with minimal data preparation. Despite the complexities presented in the transcriptions, such as mixed speech from advisors and customers, the trial run produced encouraging initial results.

This can be attributed to the large number of parameters in the LLM's neural networks and the extensive volume of data available for analysis. To optimize the models for Bouygues Télécom's specific needs, the company embarked on an optimization phase that considered factors such as context specification and providing

examples. Going forward, its plan includes retraining the models using company-specific data to facilitate the implementation of these solutions on a larger scale.

This aligns perfectly with Bouygues Télécom's strategic vision for contact center augmentation through data and AI. Generative AI enables accurate and concise summarization while maintaining coherence and contextual understanding. Furthermore, it effectively addresses the key requirements set by Bouygues Télécom, including customizability, seamless integration into existing systems and workflows, and scalability to handle large volumes of conversations. With generative AI, Bouygues Télécom can fully leverage the valuable insights derived from recorded conversations, thereby enhancing its overall customer relationship management process.

Foundation models for time series data

Time series analysis is the use of statistical methods to reveal patterns, trends, and correlations in sequential data, predicting future values based on previously observed data points. It's not just about understanding the past but shaping the future, which is crucial for staying competitive in the ever-evolving business landscape. This powerful tool enables proactive decision-making and strategic planning, helping organizations foresee changes in market trends, customer behavior, resource needs, and more.

Foundation models are AI-powered frameworks that have been trained on broad, diverse data (see “What is generative AI? What are foundation models?” on page 4). What makes them stand out is their flexibility—the same foundational model can be fine-tuned and adapted for various tasks, making them a one-size-fits-all solution for an array of challenges.

Time series foundation models blend the power of time series analysis and the flexibility of foundation models. They are designed to analyze temporal data, extracting hidden patterns, underlying trends, and temporal dependencies. They serve as an “AI Swiss Army Knife” capable of tackling a wide range of decision-making tasks across various business functions.

These models are pre-trained on comprehensive, diverse time-series data sets. IBM's models, for example, use open-source data from as far back as 1992 across domains such as retail, manufacturing, financial services, government expenditures, and environmental data. Time series data includes stock prices, exchange rates, electricity demand, temperature trends, and more. This rich, diverse data set allows our models to generalize across multiple scenarios.

Time series analysis helps organizations anticipate market trends, customer behavior, and resource needs.

Armed with time series foundation models, businesses can achieve precise forecasting, efficient resource allocation, and informed strategic planning. These models can also be fine-tuned to predict future stock prices, electricity needs, or product demand. Furthermore, enterprises can monetize their data by transforming them into relevant pre-trained models for specific use cases.

With outputs akin to GPT-style queries and generative language explanations, these models can augment various business domains, including marketing, sales, supply chain, IT, product development, customer service, financial reporting, and HR.

Risk and governance

Using generative AI in operations and product development creates new risks and exposures that business leaders need to navigate (see Figure 5). Some are the same issues as in traditional data science, such as poor predictive accuracy, deficiencies in fairness and equity, lack of explainability, model uncertainty, poisoning attacks, evasion attacks, inference attacks, and lack of model transparency. Yet there are entirely new risks with foundation models, such as hallucinations, lack of factuality, lack of data lineage and provenance, toxicity, inability to reason, privacy leakage, and prompt injection attacks.

Topping the list of executive concerns: operational risks increasing legal liabilities, cybersecurity risks introducing new vulnerabilities, and explainability risks arising when data sources aren't identified.

Generative AI risks are introduced through both inputs and outputs. When models are trained and tuned, teams must navigate a variety of data privacy and intellectual property (IP) rules appropriately. They must also ensure the robustness, fairness, and transparency of the data sets being used.

As the evolving IP landscape continues to unfold, we will see whether 2023 is just generative AI's "Netscape moment"—or perhaps also its "Napster moment"—as policies, rules, and lawsuits proliferate about fair use of public data and content ownership. Indeed, these risks may drive enterprises to focus on more proprietary data and AI models.

FIGURE 5

An evolving risk landscape

Generative AI creates new and amplified risks for enterprises to navigate

	Traditional	Amplified	New
<i>Model preparation (Training and tuning phase)</i>	Robustness Data laws Privacy	Fairness Intellectual property Transparency Privacy	Value alignment
<i>Model operationalization (Inference phase)</i>		Robustness	Privacy Intellectual property Robustness
<i>Model outputs</i>	Fairness	Misuse Explainability	Fairness Intellectual property Value alignment Misuse Harmful content generation Privacy Traceability

Source: Selection of risks from "AI Ethics and Foundation Models: Opportunities, Risks, and Mitigation", IBM Corporate AI Ethics Board, 2023

Fairness is also central to machine learning techniques and generative AI outputs. Teams must check responses for bias and harmful content, and make sure the source information is traceable and explainable. Organizations must be able to prove—on demand—that no data has been misused.

For example, Adobe recently rolled out Firefly, a generative AI tool trained on stock and public domain images, that lets designers experiment with colors using simple text prompts.¹⁴ To combat concerns that using Firefly could expose companies to copyright claims, Adobe is assuming some risk by offering IP indemnity to its enterprise customers to demonstrate that Firefly is safe for commercial use.¹⁵

However, to combat what the World Economic Forum highlights as the “intention-action gap,”¹⁶ organizations need to be mindful that they do as much as they say. For example, AI teams remain substantially less diverse than the organizations’ total workforces. One recent IBM IBV study revealed there were 5.4 times fewer women on AI teams than in the overall organization; 4 times fewer LGBTQ+ individuals; and 1.7 times fewer Black, Indigenous, and People of Color (BIPOC).¹⁷

Case study

Mitsui Chemicals discovers new applications for generative AI

Mitsui Chemicals, Inc. (Mitsui) and IBM Japan, Ltd. started working on new application discovery by combining a generative pre-trained transformer (GPT). This initiative aims to increase sales and market share of Mitsui products by advancing digital transformation in the business domain.¹⁸

Since June 2022, Mitsui has been implementing a foundation model for new application discovery in the whole Mitsui Chemicals Group. More than 20 business units are utilizing the foundation model, and more than 100 new applications have been discovered.

However, it still takes a long time to discover new applications. To address this, IBM plans to implement a third-party large language model requested by Mitsui to generate new applications from a huge repository of text data, including everything from patents to business news, to drastically increase the discovery of new applications with agility and accuracy.

By combining generative AI and foundation models, Mitsui expects to accelerate the product development process by integrating data between different departments, such as lines of business and research and development, and utilizing sales force automation, marketing automation, materials informatics, and robotics.

Action guide

Based on our research and experience with clients, we've identified key actions executives can take in each focal area to speed the scaling of generative AI across the enterprise—assuming a tight coupling with business strategy already has been established.

For more information on how to develop your AI strategy, see the IBM IBV report, “Rethinking your approach to AI: How to ground artificial intelligence in business strategy.”¹⁹

Organization and skills

—**Change the mindset:** Shift from “adding AI” to “starting with AI.” This facilitates a deeper understanding of generative AI, promotes agility to respond when market conditions change, and aligns investments and funding with cross-organizational buy-in at all levels.

—**Set measurable goals:** Define financial and non-financial goals for generative AI adoption and identify specific, quantifiable measures that encapsulate the changes you want your people to embrace.

—**Establish ethical guidelines:** Develop boundaries for the ethical use of AI. Generative AI models are powerful, but they must be used responsibly. This includes respect for privacy, transparency, fairness, and accountability.

—**Promote an experimental culture:** Recognize generative AI is new for everyone. Encourage teams to test, iterate, and improve with generative AI, and track success metrics.

—**Empower middle management:** Middle managers play a crucial role in forming a continuous feedback loop between executives and front-line employees. Train managers to be a bridge by prioritizing generative AI use cases and communicating adoption successes and failures.

—**Clarify role adjustments:** Identify and communicate changes to roles and responsibilities of existing employees transparently in areas where generative AI is being deployed.

—**Revamp the operating model:** To expedite AI adoption, rethink and restructure the operating model to promote cross-functional understanding, streamline the AI deployment process, and ensure that the benefits of generative AI and foundation models are realized across the organization.

- Establish AI integration frameworks for AI to be deployed across your operations seamlessly.
- Establish robust data & AI governance practices that are compliant with regulatory standards and best practices.
- Foster a collaborative environment between different business units, technology teams, data scientists, and decision-makers.
- Develop AI skills academies that prioritize not just technical skills but also collaboration, communication, and empathy. The curriculum should cover proper and improper use of foundation models, thereby fostering responsible AI usage.

For a broader and deeper set of AI-related actions across a holistic set of business capabilities, including organization and skills, please see the IBM IBV report, “Generating ROI with AI: Six capabilities that drive world-class results.”²⁰

Data and platform

- Understand data and platform needs:** Begin by comprehensively assessing the specific data and hybrid platform requirements that foundation models demand in your identified use cases. This involves understanding the types of data you have and the computational requirements for processing and analyzing such data.
- Set selection criteria for platforms:** Establish selection criteria for the platforms that will support your use of generative AI and related foundation models. These may include:
 - Use-case specificity
 - Cost (model development and operational expenses)
 - Availability and accessibility of relevant data
 - Balance between prediction accuracy and computational efficiency
 - Security measures and protocols
 - Degree of customization required
 - Overall system performance
 - Portability across different environments
 - Compliance with legal and regulatory standards

- Identify suitable ecosystem partners:** Once you've established your criteria, identify potential ecosystem partners that can meet your hybrid platform needs effectively and efficiently. These partners should have a proven track record in supporting AI at scale.
- Leverage next-gen data capabilities:** Use state-of-the-art data marketplaces and automation capabilities. This includes platforms that allow for seamless data ingestion, cleansing, and preparation specific to your AI outcomes.
- Scale without disruption:** As you integrate foundation models into your operations, ensure you can scale these models without compromising on business outcomes or causing operational disruptions. This involves robust model management, performance monitoring, and continuous improvement mechanisms.
- Address robust data governance:** Given that foundation models require access to vast, diverse, and potentially sensitive datasets, establish robust data governance practices. This includes data collection, storage, access, processing, and security protocols compliant with regulatory standards and best practices.

For a broader and deeper set of AI-related actions across a holistic set of business capabilities (including data and platforms), please see the IBM IBV report, “Generating ROI with AI: Six capabilities that drive world-class results.”²¹

Risk and governance

- Identify risks collaboratively:** Recognize the specific input/output and usage risks associated with foundation models. This includes model reliability, data privacy, and security, as well as ethical and regulatory compliance issues.
- Deploy mitigation policies:** Develop and implement appropriate policies to mitigate these identified risks. These should be flexible enough to adapt to changes in regulation or industry standards.
- Document AI governance:** Create comprehensive documentation, possibly in the form of fact sheets, detailing all instances of AI use in the organization, along with the current governance structure. This creates transparency and accountability.
- Trace provenance of AI assets:** Implement mechanisms to trace the origin of all AI-generated outputs back to the foundation model, dataset, prompt, and other inputs. This information should be logged in digital asset management and other systems to ensure traceability and accountability.

- Monitor and control risks:** Adopt a comprehensive monitoring and control system that is proportionate to the level of risk and compliance needed throughout the AI system’s lifecycle. This system should continuously track and evaluate model performance, data usage, and compliance.
- Transition to zero trust cybersecurity:** Accelerate the adoption of a zero trust cybersecurity framework across the enterprise and partner network. This approach assumes potential threats can originate from any source and thus ensures all access to systems are verified and monitored.
- Implement AI-enabled security intelligence:** Leverage AI-enabled security tools to enhance your organization’s ability to detect, prevent, and respond to potential security threats. Ensure clear incident escalation policies are documented at every level, all the way up to the board of directors.

For a broader and deeper set of AI-related actions around risk and governance, please see “AI ethics in action: An enterprise guide to progressing trustworthy AI.”²²

Notes and sources

- 1 Ashoori, Maryam, Brian Goehring, Timothy Humphrey, Mahmoud Naghshineh, and Cathy Rodenbeck Reese. "Generating ROI with AI: Six capabilities that drive world-class results." IBM Institute for Business Value. May 2023. <https://ibm.co/ai-capabilities>
- 2 Abercrombie, Cortney, Rafi Ezry, Brian Goehring, and Anthony Marshall. "Fast start in cognitive innovation: Top performers share how they are moving quickly." IBM Institute for Business Value. January 2017. <https://www.ibm.com/downloads/cas/KA4D6X-NB>; Abercrombie, Cortney, Rafi Ezry, Brian Goehring, Neil Isford, Anthony Marshall, and Hiroyuki Nakayama. "Accelerating enterprise reinvention with cognitive capabilities: How to build a cognitive enterprise." IBM Institute for Business Value. January 2017. <https://www.ibm.com/thought-leadership/institute-business-value/report/acelentreinvent>
- 3 Ashoori, Maryam, Brian Goehring, Timothy Humphrey, Mahmoud Naghshineh, and Cathy Rodenbeck Reese. "Generating ROI with AI: Six capabilities that drive world-class results." IBM Institute for Business Value. May 2023. <https://ibm.co/ai-capabilities>
- 4 "What is generative AI?" IBM Research. April 20, 2023. <https://research.ibm.com/blog/what-is-generative-AI>
- 5 "What are foundation models?" IBM Research. May 9, 2022. <https://research.ibm.com/blog/what-are-foundation-models>
- 6 Ibid.
- 7 Goyal, Manish, Shobhit Varshney, and Eniko Rozsa. "What is generative AI, what are foundation models, and why do they matter?" IBM. March 8, 2023. https://www.ibm.com/blog/what-is-generative-ai-what-are-foundation-models-and-why-do-they-matter?social_post=9128516062&linkId=204694717
- 8 Based on interviews with internal IBM account teams.
- 9 "IBM Consulting with Lloyds Banking Group." MCA. May 11, 2023. <https://www.mca.org.uk/consulting-case-studies/ibm-consulting-with-lloyds-banking-group>
- 10 Ashoori, Maryam, Brian Goehring, Timothy Humphrey, Mahmoud Naghshineh, and Cathy Rodenbeck Reese. "Generating ROI with AI: Six capabilities that drive world-class results." IBM Institute for Business Value. May 2023. <https://ibm.co/ai-capabilities>; Prabhakar, Aparna, Veena Mosur, and David Cox. "Rethinking your approach to AI: How to ground artificial intelligence in business strategy." IBM Institute for Business Value. September 2021. <https://ibm.co/ai-business-strategy>
- 11 Eastwood, Brian. "Workers with less experience gain the most from generative AI." MIT Sloan School of Management. June 26, 2023. <https://mitsloan.mit.edu/ideas-made-to-matter/workers-less-experience-gain-most-generative-ai>
- 12 Granger, John, Shai Joshi, Thais Lima de Marca, Varun Bijlani, and Shue-Jane Thompson. "Mastering hybrid cloud." IBM Institute for Business Value. June 2022. <https://ibm.co/mastering-hybrid-cloud>
- 13 Fléchaux, Reynald. "Bouygues Télécom relies on generative AI to transform its customer relationship." CIO. April 12, 2023. <https://www.cio-online.com/actualites/lire-bouygues-telecom-mise-sur-l-ia-generative-pour-trans-former-sa-relation-client-14869.html>
- 14 "Adobe Firefly Now Powers First Ground-breaking Generative AI Capabilities in Adobe Illustrator." Adobe. June 13, 2023. <https://news.adobe.com/news/news-details/2023/Adobe-Firefly-Now-Powers-First-Ground-breaking-Generative-AI-Capabilities-in-Adobe-Illustrator/>
- 15 Morrison, Ryan. "Adobe offers IP indemnity from generative AI copyright claims." Techmonitor. June 9, 2023. <https://techmonitor.ai/technology/ai-and-automation/adobe-firefly-generative-ai>
- 16 "How businesses can create an ethical culture in the age of tech." World Economic Forum. January 7, 2020. <https://www.weforum.org/agenda/2020/01/how-businesses-can-create-an-ethical-culture-in-the-age-of-tech>
- 17 Goehring, Brian, Francesca Rossi, and Beth Rudden. "AI ethics in action: An enterprise guide to progressing trustworthy AI." IBM Institute for Business Value. May 2022. <https://ibm.co/ai-ethics-action>
- 18 "Combined Generative AI with IBM Watson, Mitsui Chemicals Starts Verifying Higher Agility and Accuracy for New Application Discovery." IBM Newsroom. April 12, 2023. <https://jp.newsroom.ibm.com/2023-04-12-Mitsui-Chemicals-starts-verifying-new-application-discovery-generative-ai-and-ibm-watson>
- 19 Prabhakar, Aparna, Veena Mosur, and David Cox. "Rethinking your approach to AI: How to ground artificial intelligence in business strategy." IBM Institute for Business Value. September 2021. <https://ibm.co/ai-business-strategy>
- 20 Ashoori, Maryam, Brian Goehring, Timothy Humphrey, Mahmoud Naghshineh, and Cathy Rodenbeck Reese. "Generating ROI with AI: Six capabilities that drive world-class results." IBM Institute for Business Value. May 2023. <https://ibm.co/ai-capabilities>
- 21 Ibid.
- 22 Goehring, Brian, Francesca Rossi, and Beth Rudden. "AI ethics in action: An enterprise guide to progressing trustworthy AI." IBM Institute for Business Value. May 2022. <https://ibm.co/ai-ethics-action>



IBM Institute for Business Value

The IBM Institute for Business Value (IBV) delivers trusted, technology-based business insights by combining expertise from industry thinkers, leading academics, and subject matter experts with global research and performance data. The IBV thought leadership portfolio includes research deep dives, benchmarking and performance comparisons, and data visualizations that support business decision making across regions, industries and technologies.

For more information

Follow @IBMIBV on Twitter, and to receive the latest insights by email, visit ibm.com/ibv

© Copyright IBM Corporation 2023

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | July 2023

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

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

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

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

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.

