



A new era of generative AI for everyone

The technology underpinning
ChatGPT will transform work
and reinvent business



Table of Contents

03	Welcome to AI's new inflection point
04	How did we get here? Milestones in the journey to generative AI
05	Consume or customize: Generative AI for everyone
08	A look ahead at the fast-paced evolution of technology, regulation and business
12	Embrace the generative AI era: Six adoption essentials
19	The future of AI is accelerating
21	Glossary and References
22	Authors

Introduction

Welcome to AI's new inflection point

ChatGPT has woken up the world to the transformative potential of artificial intelligence (AI), capturing global attention and sparking a wave of creativity rarely seen before. Its ability to mimic human dialogue and decision-making has given us AI's first true inflection point in public adoption. Finally, everyone, everywhere can see the technology's true disruptive potential for themselves.

ChatGPT reached 100 million monthly active users just two months after launch, making it the fastest-growing consumer application in history.¹


A foundation model is a generic term for large models with billions of parameters. With recent advances, companies can now build specialized image- and language-generating models on top of these foundation models. Large language models (LLMs) are both a type of generative AI and a type of foundation model.

The LLMs behind ChatGPT mark a significant turning point and milestone in artificial intelligence. Two things make LLMs game changing. First, they've cracked the code on language complexity. Now, for the first time, machines can learn language, context and intent and be independently generative and creative. Second, after being pre-trained on vast quantities of data (text, images or audio), these models can be adapted or fine-tuned for a wide range of tasks. This allows them to be reused or repurposed in many different ways.

Business leaders recognize the significance of this moment. They can see how LLMs and generative AI will fundamentally transform everything from business, to science, to society itself—unlocking new performance frontiers. The positive impact on human creativity and productivity will be massive. Consider that, across all industries, Accenture found 40% of all working hours can be impacted by LLMs like GPT-4. This is because language tasks account for 62% of the total time employees work, and 65% of that time can be transformed into more productive activity through augmentation and automation (see Figure 3).

How did we get here?

Milestones in the journey to generative AI



Machine learning: Analysis and prediction phase

The first decade of the 2000s marked the rapid advance of various machine learning techniques that could analyze massive amounts of online data to draw conclusions – or “learn” – from the results. Since then, companies have

viewed machine learning as an incredibly powerful field of AI for analyzing data, finding patterns, generating insights, making predictions and automating tasks at a pace and on a scale that was previously impossible.

Deep learning: Vision and speech phase

The 2010s produced advances in AI’s perception capabilities in the field of machine learning called deep learning. Breakthroughs in deep learning enable the computer vision

that search engines and self-driving cars use to classify and detect objects, as well as the voice recognition that allows popular AI speech assistants to respond to users in a natural way.

Generative AI: Enter the language-mastery phase

Building on exponential increases in the size and capabilities of deep learning models, the 2020s will be about language mastery. The GPT-4 language model, developed by OpenAI, marks the beginning of a new

phase in the abilities of language-based AI applications. Models such as this will have far-reaching consequences for business, since language permeates everything an organization does day to day—its institutional knowledge, communication and processes.²

A close-up portrait of a man's face, looking slightly to the left. The image is bathed in a vibrant blue and red light gradient, with the blue being more prominent on the left and the red on the right. The man has short dark hair and a light beard. He is wearing a light-colored shirt. The background is blurred, showing some architectural lines.

**Consume or
customize:
Generative AI
for everyone**

Consume or customize: Generative AI for everyone

Easy-to-consume generative AI applications like ChatGPT, DALL-E, Stable Diffusion and others are rapidly democratizing the technology in business and society. The effect on organizations will be profound. The ability of LLMs to process massive data sets allows them to potentially “know” everything an organization has ever known—the entire history, context, nuance and intent of a business, and its products, markets and customers. Anything conveyed through language (applications, systems, documents, emails, chats, video and audio recordings) can be harnessed to drive next-level innovation, optimization and reinvention.

97% of global executives agree AI foundation models will enable connections across data types, revolutionizing where and how AI is used.³

We’re at a phase in the adoption cycle when most organizations are starting to experiment by consuming foundation models “off the shelf.” However, the biggest value for many will come when they customize or fine tune models using their own data to address their unique needs:

Consume

Generative AI and LLM applications are ready to **consume** and easy to access. Companies can consume them through APIs and tailor them, to a small degree, for their own use cases through prompt engineering techniques such as prompt tuning and prefix learning.

Customize

But most companies will need to **customize** models, by fine-tuning them with their own data, to make them widely usable and valuable. This will allow the models to support specific downstream tasks all the way across the business. The effect will be to increase a company’s efficacy in using AI to unlock new performance frontiers—elevating employee capabilities, delighting customers, introducing new business models and boosting responsiveness to signals of change.



Companies will use these models to **reinvent the way work is done**. Every role in every enterprise has the potential to be reinvented, as humans working with AI co-pilots becomes the norm, dramatically amplifying what people can achieve. In any given job, some tasks will be automated, some will be assisted, and some will be unaffected by the technology. There will also be a large number of new tasks for humans to perform, such as ensuring the accurate and responsible use of new AI-powered systems.

Consider the impact in these key functions:

Advising. AI models will become an ever-present co-pilot for every worker, boosting productivity by putting new kinds of hyper-personalized intelligence into human hands. Examples include customer support, sales enablement, human resources, medical and scientific research, corporate strategy and competitive intelligence. Large language models could be useful in tackling the roughly 70% of customer service communication that is not straightforward and can benefit from a conversational, powerful and intelligent bot, understanding a customer's intent, formulate answers on its own and improve the accuracy and quality of answers.⁴

Creating. Generative AI will become an essential creative partner for people, revealing new ways to reach and appeal to audiences and bringing unprecedented speed and innovation in areas like production design, design research, visual identity, naming, copy generation and testing, and real-time personalization. Companies are turning to state-of-the-art artificial intelligence systems like DALL·E, Midjourney and Stable Diffusion for their social media visual content generation outreach. DALL·E, for example, creates realistic images and art based on text descriptions and can process up to 12 billion parameters when transforming words into pictures. Images created can then be shared on Instagram and Twitter.⁵

Coding. Software coders will use generative AI to significantly boost productivity — rapidly converting one programming language to another, mastering programming tools and methods, automating code writing, predicting and pre-empting problems, and managing system documentation. Accenture is piloting the use of OpenAI LLMs to enhance developer productivity by automatically generating documentation — for example, SAP configuration rationale and functional or technical specs. The solution enables users to submit requests through a Microsoft Teams chat as they work. Correctly packaged documents are then returned at speed — a great example of how specific tasks, rather than entire jobs, will be augmented and automated.

Automating. Generative AI's sophisticated understanding of historical context, next best actions, summarization capabilities, and predictive intelligence will catalyze a new era of hyper-efficiency and hyper-personalization in both the back and front office—taking business process automation to a transformative new level. One multinational bank is using generative AI and LLMs to transform how it manages volumes of post-trade processing emails—automatically drafting messages with recommended actions and routing them to the recipient. The result is less manual effort and smoother interactions with customers.

Protecting. In time, generative AI will support enterprise governance and information security, protecting against fraud, improving regulatory compliance, and proactively identifying risk by drawing cross-domain connections and inferences both within and outside the organization. In strategic cyber defense, LLMs could offer useful capabilities, such as explaining malware and quickly classifying websites.⁶ In the short term, however, organizations can expect criminals to capitalize on generative AI's capabilities to generate malicious code or write the perfect phishing email.⁷

**A look ahead at the
fast-paced evolution
of technology,
regulation and
business**



A look ahead at the fast-paced evolution of technology, regulation and business

Moments like this don't come around often. The coming years will see outsized investment in generative AI, LLMs and foundation models. What's unique about this evolution is that the technology, regulation, and business adoption are all accelerating exponentially at the same time. In previous innovation curves, the technology typically outpaced both adoption and regulation.

The technology stack


The complex technology underpinning generative AI is expected to evolve rapidly at each layer. This has broad business implications. Consider that the amount of compute needed to train the largest AI models has grown exponentially – now doubling between every 3.4 to 10 months, according to various reports.⁸ Cost and carbon emissions are therefore central considerations in adopting energy-intensive generative AI.

“The hottest new programming platform is the napkin.”

Paul Daugherty, Accenture Group Chief Executive & Chief Technology Officer

Referring to the use of OpenAI to generate a working website from a napkin drawing

Figure 1: Each layer of the generative AI tech stack will rapidly evolve



The risk and regulatory environment

Companies will have thousands of ways to apply generative AI and foundation models to maximize efficiency and drive competitive advantage. Understandably, they'll want to get started as soon as possible. But an enterprise-wide strategy needs to account for all the variants of AI and associated technologies they intend to use, not only generative AI and large language models.

ChatGPT raises important questions about the responsible use of AI. The speed of technology evolution and adoption requires companies to pay close attention to any legal, ethical and reputational risks they may be incurring.

It's critical that generative AI technologies, including ChatGPT, are responsible and compliant by design, and that models and applications do not create unacceptable risk for the business. Accenture was a pioneer in the responsible use of technology including the responsible use of AI in its Code of Business Ethics from 2017. Responsible AI is the practice of designing, building and deploying AI in accordance with clear principles to empower businesses, respect people, and benefit society — allowing companies to engender trust in AI and to scale AI with confidence.

AI systems need to be “raised” with a diverse and inclusive set of inputs so that they reflect the broader business and societal norms of responsibility, fairness and transparency. When AI is designed and put into practice within an ethical framework, it accelerates the potential for responsible collaborative intelligence, where human ingenuity converges with intelligent technology.

This creates a foundation for trust with consumers, the workforce, and society, and can boost business performance and unlock new sources of growth.

Figure 2: Key risk and regulatory questions for generative AI

- ☐ **Intellectual property:** How will the business protect its own IP? And how will it prevent the inadvertent breach of third-party copyright in using pre-trained foundation models?
- ☐ **Data privacy and security:** How will upcoming laws like the EU AI Act be incorporated in the way data is handled, processed, protected, secured and used?
- ☐ **Discrimination:** Is the company using or creating tools that need to factor in anti-discrimination or anti-bias considerations?
- ☐ **Product liability:** What health and safety mechanisms need to be put in place before a generative AI-based product is taken to market?
- ☐ **Trust:** What level of transparency should be provided to consumers and employees? How can the business ensure the accuracy of generative AI outputs and maintain user confidence?
- ☐ **Identity:** When establishing proof-of-personhood depends on voice or facial recognition, how will verification methods be enhanced and improved? What will be the consequences of its misuse?


The scale of adoption in business

Companies must reinvent work to find a path to generative AI value. Business leaders must lead the change, starting now, in job redesign, task redesign and reskilling people. Ultimately, every role in an enterprise has the potential to be reinvented, once today’s jobs are decomposed into tasks that can be automated or assisted and reimaged for a new future of human + machine work.

Generative AI will disrupt work as we know it today, introducing a new dimension of human and AI collaboration in which most workers will have a “co-pilot,” radically changing how work is done and what work is done. Nearly every job will be impacted – some will be eliminated, most will be transformed, and many new jobs will be created. Organizations that take steps now to decompose jobs into tasks, and invest in training people to work differently, alongside machines, will define new performance frontiers and have a big leg up on less imaginative competitors.


Nearly 6 in 10 organizations plan to use ChatGPT for learning purposes and over half are planning pilot cases in 2023. Over 4 in 10 want to make a large investment.⁹

Figure 3: Generative AI will transform work across industries



Work time distribution by industry and potential AI impact

Based on their employment levels in the US in 2021




40% of working hours across industries can be impacted by Large Language Models (LLMs)

Why is this the case? Language tasks account for 62% of total worked time in the US. Of the overall share of language tasks, 65% have high potential to be automated or augmented by LLMs.

Source: Accenture Research based on analysis of Occupational Information Network (O*NET), US Dept. of Labor; US Bureau of Labor Statistics.

Notes: We manually identified 200 tasks related to language (out of 332 included in BLS), which were linked to industries using their share in each occupation and the occupations’ employment level in each industry. Tasks with higher potential for automation can be transformed by LLMs with reduced involvement from a human worker. Tasks with higher potential for augmentation are those in which LLMs would need more involvement from human workers.

Embrace the generative AI era: Six adoption essentials



1

Dive in, with a business-driven mindset

Even when new innovations have obvious advantages, diffusing them across an organization can be challenging, especially if the innovation is disruptive to current ways of working. By experimenting with generative AI capabilities, companies will develop the early successes, change agents and opinion leaders needed to boost acceptance and spread the innovation further, kick-starting the transformation and reskilling agenda.


Organizations must take a dual approach to experimentation. One, focused on low-hanging fruit opportunities using consumable models and applications to realize quick returns. The other, focused on reinvention of business, customer engagement and products and services using models that are customized with the organization's data. A business-driven mindset is key to define, and successfully deliver on, the business case.

As they experiment and explore reinvention opportunities, they'll reap tangible value while learning more about which types of AI are most suited to different use cases, since the level of investment and sophistication required will differ based on the use case. They'll also be able to test and improve their approaches to data privacy, model accuracy, bias and fairness with care, and learn when "human in the loop" safeguards are necessary.

98% of global executives agree AI foundation models will play an important role in their organizations' strategies in the next 3 to 5 years.¹⁰

A bank uses enhanced search to equip employees with the right information

As part of its three-year innovation plan, a large European banking group saw an opportunity to transform its knowledge base, empower its people with access to the right information, and advance its goal of becoming a data-driven bank. Using Microsoft's Azure platform and a GPT-3 LLM to search electronic documents, users can get quick answers to their questions — saving time while improving accuracy and compliance. The project, which included employee upskilling, is the first of four that will apply generative AI to the areas of contract management, conversational reporting and ticket classification.




2

Take a people-first approach


Success with generative AI requires an equal attention on people and training as it does on technology. Companies should therefore dramatically ramp up investment in talent to address two distinct challenges: creating AI and using AI. This means both building talent in technical competencies like AI engineering and enterprise architecture and training people across the organization to work effectively with AI-infused processes. In our analysis across 22 job categories, for example, we found that LLMs will impact every category, ranging from 9% of a workday at the low end to 63% at the high end. More than half of working hours in 5 of the 22 occupations can be transformed by LLMs.

Figure 4: Generative AI will transform work across every job category



Work time distribution by major occupation and potential AI impact

Based on their employment levels in the US in 2021



In 5 out of 22 occupation groups, Generative AI can affect more than half of all hours worked

Source: Accenture Research based on analysis of Occupational Information Network (O*NET), US Dept. of Labor; US Bureau of Labor Statistics.

Notes: We manually identified 200 tasks related to language (out of 332 included in BLS), which were linked to industries using their share in each occupation and the occupations' employment level in each job category. Tasks with higher potential for automation can be transformed by LLMs with reduced involvement from a human worker. Tasks with higher potential for augmentation are those in which LLMs would need more involvement from human workers.


2

In fact, independent economic research indicates that companies are significantly underinvesting in helping workers keep up with advances in AI, which require more cognitively complex and judgment-based tasks.¹¹ Even domain experts who understand how to apply data in the real world (a doctor interpreting health data, for example) will need enough technical knowledge of how these models work to have confidence in using them as a “workmate.”

There will also be entirely new roles to recruit, including linguistics experts, AI quality controllers, AI editors, and prompt engineers. In areas where generative AI shows most promise, companies should start by decomposing existing jobs into underlying bundles of tasks. Then assess the extent to which generative AI might affect each task — fully automated, augmented, or unaffected.

Figure 5: Reinventing a customer service job, task by task

To assess how specific jobs will be reinvented with AI, an Accenture analysis decomposed one customer service job into **13 component tasks**. We found:



Importantly, new job tasks might also be needed to ensure the safe, accurate and responsible use of AI in customer service settings, such as providing unbiased information on products and pricing.

3

Get your proprietary data ready

Customizing foundation models will require access to domain-specific organizational data, semantics, knowledge, and methodologies. In the pre-generative AI era, companies could still get value from AI without having modernized their data architecture and estate by taking a use-case centric approach to AI. That's no longer the case. Foundation models need vast amounts of curated data to learn and that makes solving the data challenge an urgent priority for every business.

Companies need a strategic and disciplined approach to acquiring, growing, refining, safeguarding and deploying data. Specifically, they need a modern enterprise data platform built on cloud with a trusted, reusable set of data products. Because these platforms are cross-functional, with enterprise-grade analytics and data housed in cloud-based warehouses or data lakes, data is able to break free from organizational silos and democratized for use across an organization. All business data can then be analyzed together in one place or through a distributed computing strategy, such as a data mesh.



Read more on the practices data-mature companies are using to maximize enterprise data value: [A new dawn for dormant data: Unleash the intrinsic value of enterprise data with a strong digital core on cloud.](#)

4

Invest in a sustainable tech foundation

Companies need to consider whether they have the right technical infrastructure, architecture, operating model and governance structure to meet the high compute demands of LLMs and generative AI, while keeping a close eye on cost and sustainable energy consumption. They'll need ways to assess the cost and benefit of using these technologies versus other AI or analytical approaches that might be better suited to particular use cases, while also being several times less expensive.

As the use of AI increases, so will the carbon emissions produced by the underlying infrastructure. Companies need a robust green software development framework that considers energy efficiency and material emissions at all stages of the software development lifecycle. AI can also play a broader role in making business more sustainable and achieving ESG goals. Of the companies we surveyed that successfully reduced emissions in production and operations, 70% used AI to do it.¹²





5 Accelerate ecosystem innovation

Creating a foundation model can be a complex, compute-intensive and costly exercise. And for all but the very largest global companies, doing it entirely on their own will be beyond their means and capabilities. The good news is that there is a burgeoning ecosystem to call on, with substantial investments by cloud hyperscalers, big tech players, and start-ups. Global investment in AI startups and scale-ups is estimated to exceed \$50 billion in 2023 alone.¹³ These partners bring best practices honed over many years, and can provide valuable insights into using foundation models efficiently and effectively in specific use cases. Having the right network of partners—including technology companies, professional services firms and academic institutions—will be key to navigating rapid change.

6 Level-up your responsible AI

The rapid adoption of generative AI brings fresh urgency to the need for every organization to have a robust responsible AI compliance regime in place. This includes controls for assessing the potential risk of generative AI use cases at the design stage and a means to embed responsible AI approaches throughout the business. Accenture's research suggests most companies still have a long way to go. Our 2022 survey of 850 senior executives globally revealed widespread recognition of the importance of responsible AI and AI regulation. But only 6 percent of organisations felt they had a fully robust responsible AI foundation in place.

An organization's responsible AI principles should be defined and led from the top and translated into an effective governance structure for risk management and compliance, both with organizational principles and policies and applicable laws and regulations. Responsible AI must be CEO-led, beginning with a focus on training and awareness and then expanding to focus on execution and compliance. Accenture was one of the first to take this approach to Responsible AI years ago, with a CEO-led agenda, and now a formal compliance program. Our own experience shows that a principles-driven compliance approach provides guardrails while being flexible enough to evolve with the fast pace of changing technology, ensuring companies aren't constantly playing "catch up."

To be responsible by design, organizations need to move from a reactive compliance strategy to the proactive development of mature Responsible AI capabilities through a framework that includes principles and governance; risk, policy and control; technology and enablers and culture and training.

The future of AI is accelerating



This is a pivotal moment. For several years, generative AI and foundation models have been quietly revolutionizing the way we think about machine intelligence. Now, thanks to ChatGPT, the whole world has woken up to the possibilities this creates.

While artificial general intelligence (AGI) remains a distant prospect, the speed of development continues to be breathtaking. We're at the start of an incredibly exciting era that will fundamentally transform the way information is accessed, content is created, customer needs are served, and businesses are run.

Embedded into the enterprise digital core, generative AI, LLMs, and foundation models will optimize tasks, augment human capabilities, and open up new avenues for growth. In the process, these technologies will create an entirely new language for enterprise reinvention.

Businesses are right to be optimistic about the potential of generative AI to radically change how work gets done and what services and products they can create. They also need to be realistic about the challenges that come with profoundly rethinking how the organization works, with implications for IT, organization, culture, and responsibility by design.

Companies need to invest as much in evolving operations and training people as they do in technology. Radically rethinking how work gets done, and helping people keep up with technology-driven change, will be two of the most important factors in realizing the full potential of this step-change in AI technology.

Now's the time for companies to use breakthrough advances in AI to set new performance frontiers—redefining themselves and the industries in which they operate.

Glossary

ChatGPT is a generative AI chatbot interface built on top of OpenAI's GPT-3.5 large language model (see below). ChatGPT (and ChatGPT plus, which uses GPT-4) allows users to interact with the underlying AI in a way that seems remarkably accurate and feels surprisingly human. You can ask it to explain a subject, write an essay, run a calculation, generate some Python code, or simply have a conversation.

Generative AI is the umbrella term for the ground-breaking form of creative artificial intelligence that can produce original content on demand. Rather than simply analyzing or classifying existing data, generative AI is able to create something entirely new, whether text, images, audio, synthetic data, or more.

Foundation models are complex machine learning systems trained on vast quantities of data (text, images, audio, or a mix of data types) on a massive scale. The power of these systems lies not only in their size but also in the fact they can quickly be adapted or fine-tuned for a wide range of downstream tasks. Examples of foundation models include BERT, DALL-E, and GPT-4.

Large Language Models (LLMs) represent a subset of foundation models that are trained specifically on text sources. GPT-3, for instance, was trained on almost 500 billion words from millions of websites.¹⁴ Its successor, GPT-4, can take image as well as text as inputs.

Fine-tuning is the process by which foundation models are adapted for specific downstream tasks using a particular dataset. That can include everything from the hyper-specific (training a model to compose emails based on your personal writing style) to the enterprise level (training an LLM on enterprise data to transform a company's ability to access and analyze its core intelligence).

Data is the fundamental bedrock of generative AI. Not only in training foundation models themselves, but also in fine-tuning those models to perform specific tasks. In an enterprise context, examples might include everything from legacy code to real-time operational data to customer insights.

References

1. ChatGPT sets record for fastest-growing user base - analyst note, Reuters, February 2023
<https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
2. The Next Big Breakthrough in AI Will Be Around Language, Harvard Business Review, September, 2020
<https://hbr.org/2020/09/the-next-big-breakthrough-in-ai-will-be-around-language>
3. Accenture Tech Vision 2023
4. ChatGPT Is Coming to a Customer Service Chatbot Near You, Forbes, January 2023
<https://www.forbes.com/sites/rashishrivastava/2023/01/09/chatgpt-is-coming-to-a-customer-service-chatbot-near-you/?sh=730eeab97eca>
5. How AI Transforms Social Media, Forbes, March 2023
<https://www.forbes.com/sites/forbestechcouncil/2023/03/16/how-ai-transforms-social-media/?sh=739221ca1f30>
6. Large AI Models have Real Security Benefits, Dark Reading, August, 2022
<https://www.darkreading.com/dr-tech/large-language-ai-models-have-real-security-benefits>
7. OPWNAI: Cybercriminals starting to use ChatGPT, Checkpoint Research, January, 2023
<https://research.checkpoint.com/2023/opwnai-cybercriminals-starting-to-use-chatgpt/>
8. Accenture Technology Vision 2023
9. CXO Pulse Survey, conducted by Accenture Research, February 2023
10. Accenture Technology Vision 2023
11. [The Productivity J-Curve: How Intangibles Complement General Purpose Technologies - American Economic Association \(aeaweb.org\)](https://www.aeaweb.org/articles?id=10.3386/w30282)
12. Uniting technology and sustainability, Accenture, May, 2022
[Technology Sustainability Key to ESG Goals | Accenture](https://www.accenture.com/in/insights/technology-sustainability-key-to-esg-goals)
13. Pace Of Artificial Intelligence Investments Slows, But AI Is Still Hotter Than Ever, Forbes, October, 2022
<https://www.forbes.com/sites/joemckendrick/2022/10/15/pace-of-artificial-intelligence-investments-slows-but-ai-is-still-hotter-than-ever/?sh=853d8124c76c>
14. OpenAI's GPT-3 Language Model: A Technical Overview, Lambda, June, 2020
<https://lambdalabs.com/blog/demystifying-gpt-3>

Authors



Paul Daugherty

Group Chief Executive &
Chief Technology Officer



Bhaskar Ghosh

Chief Strategy Officer,
Accenture



Karthik Narain

Lead – Accenture Cloud First



Lan Guan

Lead – Cloud First, Data & AI



Jim Wilson

Global Managing Director – Thought
Leadership & Technology Research



The authors would like to acknowledge
Tomas Castagnino, Elise Cornille,
Ray Eitel-Porter, Linda King,
Amy Sagues, Ezequiel Tacsir and
Denise Zheng for their contributions.

About Accenture

Accenture is a leading global professional services company that helps the world's leading businesses, governments and other organizations build their digital core, optimize their operations, accelerate revenue growth and enhance citizen services—creating tangible value at speed and scale. We are a talent and innovation led company with 738,000 people serving clients in more than 120 countries. Technology is at the core of change today, and we are one of the world's leaders in helping drive that change, with strong ecosystem relationships. We combine our strength in technology with unmatched industry experience, functional expertise and global delivery capability. We are uniquely able to deliver tangible outcomes because of our broad range of services, solutions and assets across Strategy & Consulting, Technology, Operations, Industry X and Accenture Song. These capabilities, together with our culture of shared success and commitment to creating 360° value, enable us to help our clients succeed and build trusted, lasting relationships. We measure our success by the 360° value we create for our clients, each other, our shareholders, partners and communities. Visit us at www.accenture.com.

Contact us

For more information, contact the Accenture Generative AI/ Large Language Model Center of Excellence at: LLMCoE@accenture.com.

Copyright © 2023 Accenture. All rights reserved.
Accenture and its logo are registered trademarks of Accenture

Disclaimer: This content is provided for general information purposes and is not intended to be used in place of consultation with our professional advisors. This document refers to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.