




The state of enterprise AI

OpenAI

Foreword

At OpenAI, our mission is to ensure that artificial intelligence benefits all of humanity, and helping enterprises solve problems is central to this mission.

The majority of economically valuable activity takes place inside organizations, where innovation translates directly into improved outcomes for workers, customers, and other stakeholders. Enterprise problems also present the hardest technical challenges for frontier intelligence, requiring reliability, safety, and security at scale. The revenue generated from solving these problems can help fund broad, free access to powerful AI for hundreds of millions of people worldwide.

For much of the past three years, the visible impact of AI has been most apparent among consumers. However, the history of general purpose technologies—from steam engines to semiconductors—shows that significant economic value is created after firms translate underlying capabilities into scaled use cases. Enterprise AI now appears to be entering this phase, as many of the world’s largest and most complex organizations are starting to use AI as core infrastructure.

More than 1 million business customers now use OpenAI’s tools. This report brings together evidence from de-identified and aggregated enterprise usage data and a variety of other sources to provide a grounded view of how AI is being deployed inside organizations today.

Four key findings stand out

- 01 Enterprise usage is scaling, with deeper workflow integration.** ChatGPT message volume grew 8x and API reasoning token consumption per organization increased 320x year-over-year, demonstrating that more enterprises are using AI and their intensity of usage has increased.

- 02 Enterprises that leverage AI are experiencing measurable productivity and business impact.** Enterprise users report saving 40–60 minutes per day and being able to complete new technical tasks such as data analysis and coding. Case studies indicate AI is contributing to important outcomes such as revenue growth, improved customer experience, and shorter product-development cycles.

- 03 Enterprise growth is global and rapidly accelerating across industries.** Over the past six months, international adoption has surged as organizations worldwide deepen their use of AI, complementing continued strong momentum in the U.S. In the past 12 months, the median sector grew by more than 6x, with the technology sector leading the pack at 11x.

- 04 A widening gap is emerging between leaders and laggards.** Frontier workers are sending 6x more messages and frontier firms are sending 2x as many messages per seat than the median enterprise. There's a substantive gap in the likelihood to utilize the most capable AI tools today, despite broad availability of these tools. Models are capable of far more than most organizations have embedded into workflows, and this presents an opportunity for firms.

Looking ahead, the next phase of enterprise AI will be shaped by stronger performance on economically valuable tasks, better understanding of organizational context, and a shift from asking models for outputs to delegating complex, multi-step workflows. As these capabilities mature, we expect organizations to not only improve efficiency, but discover new ways to serve customers and deliver value.

The findings in this report represent early signs of how AI is beginning to reshape the modern enterprise. As enterprise AI evolves, OpenAI will continue to share real-world evidence on how AI is influencing firms, workers, and the broader economy.

Ronnie Chatterji

Chief Economist
OpenAI

Introduction

Over the past three years, enterprises have integrated AI systems across a wide range of use cases and operational workflows.

These deployments provide insights on how AI is shaping work, particularly in environments where accuracy standards are high, workflows are complex, and improvements in productivity or decision quality have direct economic outcomes. Because much of the world's economically valuable activity occurs inside firms, enterprise adoption patterns provide a clear signal of where AI is delivering value today and where it will likely do so in the future.

The scale and diversity of OpenAI's more than 1 million business customers provides a distinctive view into this shift. This report summarizes key findings from across OpenAI's enterprise customer base, and what those patterns suggest about the current state and trajectory of enterprise AI. By examining how adoption varies across industries and functions, the analysis also highlights where AI is becoming deeply embedded in firms, and where gaps are emerging.

Findings are based on two primary data sources

- 01 Real-world usage data from enterprise customers of OpenAI.
- 02 An OpenAI survey of 9,000 workers across almost 100 enterprises documenting patterns of AI adoption.

All analyses in this report are based on de-identified, aggregated enterprise usage data. Message content was classified using automated systems, and no OpenAI employee reviewed individual enterprise, business, or API customer data as part of this analysis.

Enterprise AI usage is accelerating and deepening

Over the past year, enterprise AI adoption has increased substantially as organizations incorporate AI into repeatable, multi-step workflows across functions and business units. OpenAI now serves more than 7 million ChatGPT workplace seats, and ChatGPT Enterprise seats have increased approximately 9x year-over-year.

Since November 2024, weekly Enterprise messages have grown approximately 8x in aggregate, with the average worker sending 30% more messages. This growth reflects both more frequent use of ChatGPT and a deepening in the intensity of use.

Two shifts underscore the deepening integration of AI into core enterprise workflows.

Custom GPTs and Projects are enabling deeper workflow integration

GPTs and Projects are configurable interfaces built on ChatGPT that can be tailored with instructions, knowledge, and custom actions, enabling workers to execute repeatable, multi-step tasks.

Weekly users of Custom GPTs and Projects have increased by approximately 19x year-to-date. In recent months, approximately 20% of all Enterprise messages were processed via a Custom GPT or Project. The most widely deployed GPTs either codify institutional knowledge into reusable assistants or automate workflows through integrations with internal systems. Some organizations have built a culture of developing and sharing Custom GPTs at scale. For example, BBVA regularly uses more than 4,000 GPTs, indicating that AI-driven workflows are increasingly implemented as persistent tools embedded in daily operations.

19x

Year-to-date increase in weekly users of Custom GPTs and Projects

20%

of all Enterprise messages were processed via a Custom GPT or Project

Developer and API workflows are rapidly scaling

Companies build on the API to integrate models directly into their products and systems with a high degree of control and customization. As firms transition from experimentation to production deployments, API consumption has rapidly increased. More than 9,000 organizations have now processed over 10 billion tokens, and nearly 200 have exceeded 1 trillion tokens.

Average reasoning token consumption per organization has increased by approximately 320x in the past 12 months, suggesting that more intelligent models are being systematically integrated into expanding products and services. Codex, while still early in its enterprise lifecycle, is gaining rapid traction as teams adopt it for end-to-end software tasks: code generation, refactoring, testing, and debugging.

In the past six weeks, Codex engagement indicates growing penetration of AI-assisted development inside enterprises.

2x

Increase in weekly active users

50%

Approximate increase in weekly messages

Workers report measurable value from using AI

In most settings, AI enables workers to produce higher quality work faster. However, productivity alone does not fully reflect how AI is reshaping work. Survey data from almost 100 enterprises highlights key operational gains across functions, and shifts in who performs specialized and technical work.

Enterprise workers report time saved and improved outcomes across functions

Seventy-five percent of surveyed workers report that using AI at work has improved either the speed or quality of their output. On average, ChatGPT Enterprise users attribute 40–60 minutes of time saved per active day to their use of AI, with data science, engineering, and communications workers saving more than average (60–80 minutes per day). Time saved per message varies by function: accounting and finance users report the largest benefits followed by analytics, communications, and engineering.

These gains translate into broad operational improvements across functions

87% of IT workers report faster IT issue resolution

85% of marketing and product users report faster campaign execution

75% of HR professionals report improved employee engagement

73% of engineers report faster code delivery

These results indicate that productivity benefits are already materializing across core enterprise functions, not only in early-adopting technical roles.

Technical work expands beyond traditional role boundaries

AI is not only accelerating existing work; it is also expanding the tasks and skills workers can perform. Several studies find that AI has an equalizing effect, disproportionately aiding lower performing workers.¹ Consistent with these findings, 75% of workers report being able to complete tasks they previously could not perform, including programming support and code review, spreadsheet analysis and automation, technical tool development and troubleshooting, and custom GPT or agent design.

The broadening of individual capabilities is particularly apparent in technical settings, where non-technical teams are increasingly engaging in coding and data-analysis work that was previously confined to specialized roles. Among ChatGPT Enterprise users, coding-related messages have increased across all functions, and outside of engineering, IT, and research, coding-related messages have grown by an average of 36% over the past six months.

75%

of users report being able to complete new tasks


36%

Average increase in coding-related messages outside of engineering, IT, and research

Workers report greater productivity from more intensive AI use

At the individual worker level, impact increases as workers deepen their use of AI. Across a large sample of workers, time saved is correlated with the use of more advanced ChatGPT features, including Deep Research, GPT-5 Thinking, and Image Generation. Workers consuming the most intelligence (as measured by credits used²) report higher time savings. Workers who save more than 10 hours per week are not just using more intelligence, they are also using multiple models, engaging with more tools, and using AI across a wider range of tasks.

Productivity gains increase with intensity of AI use



¹Noy & Zhang (2023), Dell'Acqua et al. (2023), Schwarcz et al. (2025), and Brynjolfsson et al. (2025)

²Credits map to usage, with more advanced features like Codex and Deep Research consuming an higher number of credits

Pace of acceleration varies based on industry and geography


Over the last year we've seen overall rapid adoption as companies move from AI pilots to full deployments, and there are notable differences based on industry and geography.

Growth is rapid across most industries

OpenAI customer growth is broad-based across industries, with the median sector expanding more than 6x year-over-year and even the slowest-growing sector exceeding 2x.

AI adoption by industry: enterprise scale vs. year-over-year growth

Technology, healthcare, and manufacturing show the fastest growth, while finance and professional services operate at the largest scale.



In absolute terms, ChatGPT Enterprise customers are most concentrated today in professional services, finance, and technology, sectors that were early adopters and continue to lead in their scale of AI usage. Healthcare and manufacturing started from a much smaller base but are now among the fastest-growing sectors, rapidly closing the gap.

Fastest-growing sectors	Year-over-year customer growth
01. Technology	11x
02. Healthcare	8x
03. Manufacturing	7x

The API is most commonly used to build and scale customer-facing applications (e.g., in-product assistants, search, and automation), particularly by technology companies. But usage is diversifying: customer service and content generation now represent approximately 20% of API activity, and non-technology firm API use has grown 5x year-over-year. Taken together, this pattern suggests adoption is expanding beyond technology-led product embedding toward a broader set of operational and workflow deployments across industries.

Technology companies are using the API at a rate 5x higher year-over-year as they scale external, customer-facing applications. They also lead in coding workflows, where frontier models such as Codex are accelerating software development.

Top API use cases

- 1. In-app Assistant & Search
- 2. Agentic Workflow Automation
- 3. Coding & Developer Tools
- 4. Customer Support
- 5. Data Analysis, Summarization & Extraction

Professional services concentrate API spend in coding and developer tools to build custom tooling that accelerates delivery, improves the customer experience (often via personalization), and enables assistant applications.

- 1. Coding & Developer Tools
- 2. Content & Creative Generation
- 3. In-app Assistant & Search
- 4. Customer Support
- 5. Agentic Workflow Automation


Finance organizations often start with customer-support because support is a large, scalable cost center with proven ROI. Coding and developer tools rank second as firms invest in system migration and custom applications for trading, risk, and compliance.

- 1. Customer Support
- 2. Coding & Developer Tools
- 3. Agentic Workflow Automation
- 4. In-app Assistant & Search
- 5. Data Analysis, Summarization & Extraction


Enterprise growth is global and accelerating

While early AI adoption was primarily U.S.-based, international growth is now accelerating rapidly:

- Among the largest markets, Australia, Brazil, the Netherlands, and France show the fastest growth in business customers, increasing more than 143% year-over-year
- ChatGPT usage among business customers continues to scale globally, with the United States, Germany, and Japan among the most active markets by message volume.
- The United Kingdom and Germany now rank among the largest ChatGPT Enterprise markets outside the U.S. by number of customers.
- International API customer growth has exceeded 70% over the last 6 months, with Japan having the largest number of corporate API customers outside of the U.S.



Monthly business message volume




The growing divide in AI adoption

There are clear differences emerging in how AI is used across industries and among individuals within firms. Whether this gap widens or contracts will depend on how organizations approach change management and their ability to build the systems, skills, and operating models required to successfully deploy AI.


To understand this growing divide more deeply, it is useful to compare frontier workers (defined as those in the 95th percentile of adoption intensity) to the median worker. Frontier workers generate 6x more messages than the median worker. Even among those who work in data analytics, frontier workers use the data-analysis tool 16x more than the median.

Worker usage gaps widen with more advanced tools

Frontier (95th percentile) vs median users




The gaps are widest between frontier and median workers for writing, coding, and analysis. Coding exhibits the largest relative gap in message volume, with frontier workers sending 17x as many messages as the median.




These differences matter. Usage data matched to survey results show that users who engage across roughly seven task types report five times more time saved than those who use only about four. In other words, the benefits users realize from AI scale directly with depth of use.

Time savings increase as users engage across more distinct tasks




Even among active ChatGPT Enterprise users, many have not tried some of the most capable tools. Of monthly active users, 19% have never used data analysis, 14% have never used reasoning, and 12% have never used search. Among daily active users, those shares drop to 3%, 1%, and 1%, respectively.



There’s significant headroom for firms to increase their AI maturity

At the firm level, gaps in adoption intensity mirror those observed at the individual level. Frontier firms (95th percentile) generate approximately 2x more messages per seat than the median enterprise and 7x more messages to GPTs, indicating markedly deeper organizational integration and workflow standardization. These firms invest systematically in the infrastructure and operating models required to embed AI as core organizational capability rather than a peripheral productivity tool.

Firm usage gaps widen with more advanced tools
Frontier (95th percentile) vs median firms



AI adoption and business impact: case evidence

The following case studies illustrate how AI is generating measurable business outcomes across a range of organizational contexts. Rather than a one-size-fits-all solution, their impact reflect the application of AI to specific operational and strategic challenges. Across these examples, AI is associated with revenue growth, improvements in customer experience, automation of manual processes, and accelerated product development.

These effects are not confined to a small set of firms, and external research shows that AI adoption is beginning to influence core financial performance indicators.

A 2025 Boston Consulting Group (BCG) study found that over the past three years, AI leaders achieved 1.7x revenue growth, 3.6x greater total shareholder return, and 1.6x EBIT margin. They also outperformed on non-financial measures such as patent output and employee satisfaction, linking AI maturity to both financial and organizational strength. While this evidence is still early, it suggests that AI adoption is correlated with improved financial performance and organizational outcomes.

Case studies	
p. 17	Intercom
p. 18	Lowe's
p. 19	Indeed
p. 20	BBVA
p. 21	Oscar Health
p. 22	Moderna



Intercom used OpenAI’s Realtime API for Fin Voice, delivering a low-latency, enterprise-ready voice AI Agent

Challenge

Fin, Intercom’s AI agent for customer service, delivers industry-leading resolution rates across chat, email, and social channels, resolving millions of customer queries each month. But extending Fin to a new channel — the phone — introduced a new, critical challenge: latency. In phone support, where issues are often urgent, even brief pauses can destroy the customer experience, and lead callers to abandon the interaction or escalate to a human.

Solution

Intercom built Fin Voice on OpenAI’s Realtime API to dramatically reduce latency and enable natural, interruption-friendly phone conversations. The Realtime API’s low time-to-first-token, strong instruction-following, and dependable tool-calling capabilities allow Fin Voice to navigate complex, multi-step requests with high quality and reliability.

Impact

Latency has decreased by 48% since March through Fin Voice’s use of the Realtime API for answer generation.

With faster responses enabled by the Realtime API, **customers are seeing Fin Voice resolve 53% of calls end-to-end on average** — a significant result given that phone calls are typically far more complex than chat.

Customers report that calls that ultimately require human agents are **resolved 40% faster** once Fin Voice completes the initial steps, improving efficiency on higher-touch calls.

Given that human-handled support conversations typically cost upwards of \$5–\$20 (depending on region and industry), **Fin is already saving customers hundreds of millions of dollars annually.**



Lowe’s deployed Mylow and Mylow Companion to scale expert home-improvement guidance to every online visitor and in-store associate

Challenge

Lowe’s needed to scale expert home improvement guidance to online shoppers and help store associates, especially new hires, answer complex questions consistently across more than 1,700 stores.

Solution

Lowe’s deployed Mylow on Lowes.com to provide customer project and product advice, and Mylow Companion for store associates in every store.

Impact

Mylow and Mylow Companion answer **nearly 1 million questions per month** about everything from product specs to project know-how to the status of a customer order since launching in March of this year.

Mylow is available on Lowes.com and in the award-winning Lowe’s mobile app. When customers engage with Mylow during their online visits, **the conversion rate more than doubles.**

Mylow Companion is deployed in 100% of stores and answers hundreds of thousands of associate questions each week. **Lowe’s is seeing customer satisfaction scores increase 200 basis points** when associates use Mylow Companion to help customers shopping in the aisle.



Indeed uses GPT-powered job matching and career coaching to improve hiring outcomes for job seekers and employers

Challenge

Indeed’s mission is to help people get jobs. Job seekers can face friction when searching, evaluating fit, and applying for roles, while employers want more qualified applicants for their open roles. Both sides benefit from deeper personalization and clearer context about what makes a strong match.

Solution

To address this friction, Indeed launched a suite of AI-powered products, using its proprietary AI to match job seekers and employers coupled with GPT-powered explanations that help explain why they’re a good fit. Indeed Invite to Apply uses AI to generate and send contextual, personalized job invitations at scale, helping candidates understand why a role is a strong match and improving employer reach. Indeed Career Scout acts as an AI career coach, accelerating job discovery and streamlining the application process for job seekers.

Impact

In experiments, Invite to Apply with LLM-generated explanations **increased started applications by 20%** and **improved downstream success (interviews and hires) by 13%** versus traditional matching.

Early results show job seekers using Career Scout **find and apply to relevant jobs 7x faster** and **are 38% more likely to be hired**, with **84% rating it valuable**.



BBVA deployed a legal AI chatbot to instantly validate corporate signatory authority and unblock branch commercial operations

Challenge

In Mexico, BBVA must perform a legal check (also known as bastanteo) to confirm that a company representative has the authority to sign and act on behalf of the company before key transactions can proceed (e.g., opening accounts, signing contracts, issuing credit). Historically, this process relied on a specialist legal team responding to repetitive branch queries, creating delays, bottlenecks, and high demand for scarce legal capacity.

Solution

BBVA built a generative AI chatbot that provides instant access to standardized, pre-validated legal FAQs and documentation guidance for common signatory-authority questions. The content was developed and reviewed by BBVA’s Legal Services team, reducing manual handling of daily inquiries and making approved legal guidance consistently available.

Impact

The solution built with ChatGPT Enterprise **automates more than 9,000 queries annually** and has enabled BBVA to redeploy the **equivalent of 3 FTE's** toward producing **over 11,000 bastanteos per year, delivering 26% of the Legal Services division's annual savings KPI.**

Oscar

Oscar Health deployed member-facing chatbots to answer benefits, cost, and general health questions in real time and help members navigate the complexities of the healthcare system

Challenge

For many people, the healthcare system can be challenging to understand and navigate. Understanding benefits, finding the right doctor, estimating care costs, and getting clear answers to questions are often challenging and time-consuming. This is partially because the data needed to make the right decisions often lives in different places, including portals, benefits documents, and doctors' notes from past visits. Oscar wanted to create a single, trustworthy entry point that helped members better understand and navigate the healthcare system.

Solution

Oscar developed a pair of member-facing chatbots to answer member benefits, costs and general health questions, on-demand and in real-time. Unlike general-purpose AI chatbots, these are integrated with Oscar systems and data, allowing them to draw from medical records, claims, and customer service interactions to personalize responses. Their chatbots can also assist with common tasks, including finding in-network doctors and refilling prescriptions.

Impact

The result is a platform that can address a wide array of questions and tasks, including understanding benefits, supporting symptom-related questions, preparing for visits, and explaining follow-up instructions, while also escalating members to providers or care guides as needed. Their platform **answers 58% of benefits questions instantly** and is able to handle **39% of benefits messages without any human escalation**. Today, they now have **the foundation for future capabilities**, including appointment booking, voice interactions, and condition-specific management.



Moderna used AI to significantly compress Target Product Profile development time

Challenge

Writing a Target Product Profile (TPP) is typically a multi-week, cross-functional effort involving teams across clinical, product, and marketing roles. Teams must review and process large evidence packs, sometimes up to 300 pages of information, to create these blueprints for product development.

Solution

Using ChatGPT Enterprise, Moderna has streamlined substantial parts of the TPP drafting and analysis workflow. The system helps extract key facts and assumptions from large data packages, generate structured draft sections, and flag important details or potential errors to the teams providing human oversight.

Impact

Delays or errors in TPPs can affect downstream activities such as research planning, cross-functional alignment, and product launch preparation. By reducing the time required to review, cross-reference, and integrate large evidence packages, **teams can spend more time pressure-testing trade-offs and making higher-quality decisions** earlier in the TPP creation process. Moderna reports that a core analytical step in this process has been **reduced from weeks to hours in some cases**, and believes that each day gained in early TPP planning can **help the company deliver for patients more quickly**.

In practice, leading firms consistently do several things

Deep system integration through enabling context	They turn on connectors to give AI secure access to company data inside core tools, enabling context-aware responses and automated actions. Roughly one in four enterprises still has not taken this step.
Workflow standardization and reuse	They actively promote the creation, sharing, and discovery of repeatable solutions for common tasks. GPTs often power this work, while the most sophisticated organizations embed API-powered assistants directly into core internal systems.
Executive leadership and sponsorship	They set clear mandates, secure resources, and align teams, and create space for experimentation, all of which enable deployment at scale.
Data readiness and evaluations	They codify institutional knowledge into machine-readable routines, build APIs for key data pipelines, and run continuous evaluations to track model performance on real-world outcomes.
Deliberate change management	They build structures that speed organizational learning, combining centralized governance and training with distributed enablement through embedded AI champions.

The AI landscape is evolving rapidly; OpenAI releases a new feature or capability roughly every three days. The primary constraints for organizations are no longer model performance or tooling, but rather organizational readiness.

Conclusion

Across OpenAI's more than 1 million business customers, AI is being embedded into an expanding range of workflows, products, and internal systems. Adoption is broad-based and accelerating across industries and regions, though depth of integration varies widely by organization.

The data suggest that depth of use matters. Workers and firms that make more consistent use of advanced tools, such as reasoning models, data analysis, Custom GPTs, Projects, and APIs, report larger productivity gains and broader task coverage than those whose use remains limited.

AI is also beginning to change who performs certain types of technical work. Coding and analytical tasks are increasingly showing up outside of traditional specialist roles, expanding what some non-technical teams are able to do. At the same time, industry patterns remain distinct, reflecting different operational needs across technology, professional services, finance, healthcare, manufacturing, and more.

Despite a growing divide in AI adoption, enterprise AI is still in the early innings. Firms have an opportunity to catch up by adopting the patterns of frontier workers and organizations. As enterprise AI matures, firms will increasingly translate AI capabilities into products and services that deliver new sources of value through faster iteration, deeper personalization, and new experiences. Organizations that succeed in bringing these capabilities into market-facing workflows will use AI not merely as a productivity tool, but as a durable engine of revenue growth and competitive advantage.