# EU Economic Blueprint 2.0

January 2026

OpenAI

# Foreword

## Deploying AI of, for and by Europe

AI capability has advanced at extraordinary speed. In 2022, frontier models could reliably complete tasks that take human experts about one minute to complete; today, they can complete tasks that take human experts more than 30 minutes. But Europe's ability to translate that capability into broad-based economic and social impact is uneven. Left unaddressed, it risks concentrating productivity gains in a small number of countries, sectors, and firms, while others fall behind.

We refer to this gap as the capability overhang: the widening distance between what frontier AI can do and the extent to which those capabilities are being adopted by users and absorbed across the economy. Worldwide, the typical power user uses 7x more thinking capabilities than the typical user. And across more than 70 countries with the highest ChatGPT users, leading countries use 3× more thinking capabilities per person than users in lagging countries. This capability overhang, ultimately, is about the opportunity for individuals, companies, and countries to participate in the Intelligence Age.

**Europe uses 17% more thinking capabilities on average than the rest of the world, however, there are big differences between Member States**, with the most thinking-intensive country using approximately 40% more thinking capabilities than the least intensive. Although most Member States are above the global average, nine EU countries still fall below the global average. Understanding and addressing the overhang in Europe, including through **our new SME AI Accelerator campaign supported by Booking.com, to help 20,000 SMEs across Europe better use AI**, is the core focus of this year's Blueprint report.

We see the capability overhang as the primary challenge for Europe moving forward: the countries that have a strategy to better utilize AI will be the nations that ensure that their people have the opportunity to participate in an AI global economy. But it's worth also reflecting on the progress that has been made since our first Blueprint was published last year. In that time, the contours of a European AI ecosystem — spanning infrastructure, skills, and governance — began to take shape with EU AI gigafactories, Union of Skills, and AI Act Codes of Practice underway.

At the same time, Europe continues to contend with weak productivity growth, ageing populations, and rising fiscal constraints, alongside intensifying global competition. Against this backdrop, AI is one of the few levers available to boost productivity, strengthen resilience, solve hard problems, and sustain Europe's long-term economic model. A recent study from the European Investment Bank found that AI is already increasing labour productivity in Europe by around 4%, underscoring why adoption at scale and with intention is vital moving forward.

On the regulation front, the European Commission has taken welcome steps to ease compliance timelines for local players that want to compete on the world stage. It's also acting on the EU Inc. 28th regime to support startups, a proposal we supported in last year's Blueprint. Looking ahead, continued discipline around regulatory simplification will be essential — ensuring that new rules meet a high bar, reduce fragmentation, and support innovation rather than slow it down. Europe must not stand in its own way.

For our part, we were the first US-based AI developer to sign the EU's AI Act General Purpose Code of Practice and want to help set the standards for youth protections in the AI era. Today's young people will be the first generation to grow up with AI as a part of daily life. It will support their learning, become a foundational skill for their careers and offer new ways to create and express themselves. Parental controls, strong content protection and new tools such as age prediction, combined with guidance from physicians and our Wellbeing Council, and **a new €500,000 EMEA youth and wellbeing grant programme for local NGOs**, reflect our belief that youth safety is non-negotiable.

OpenAI has been an active partner in Europe's efforts to reap the benefits of AI over the past year and looking ahead, we are keen to do our part to help countries close their capability overhang, too. Across Europe, we have already worked with governments and partners on a broad set of AI priorities — from data-centre initiatives in Germany and with Stargate Norway, to education, startup, and skills partnerships spanning Estonia and Greece through to Ireland and Slovakia.

In 2026, we will expand our work through the **OpenAI for Europe** umbrella to more governments and across additional policy areas, supporting countries as they pursue their national AI priorities. This will include new initiatives focused on education and health, AI skills training and certifications, disaster response and preparedness, cybersecurity, and startup accelerators.

This second edition of the EU Economic Blueprint sets out how Europe can build on recent progress, avoid costly missteps, and — most importantly — close the capability overhang. The choices made now will determine whether AI becomes a broad-based driver of European competitiveness, or whether its benefits accrue elsewhere.

Europe has what it takes to succeed and we look forward to working alongside local businesses, governments, startups, schools, and citizens on this journey.

**Christopher Lehane**
Chief Global Affairs Officer
OpenAI

# The EU's AI Capability Overhang

## Bringing Europe's capability overhang into focus

Volume of thinking capabilities used per monthly active user, Plus plans, by EU country

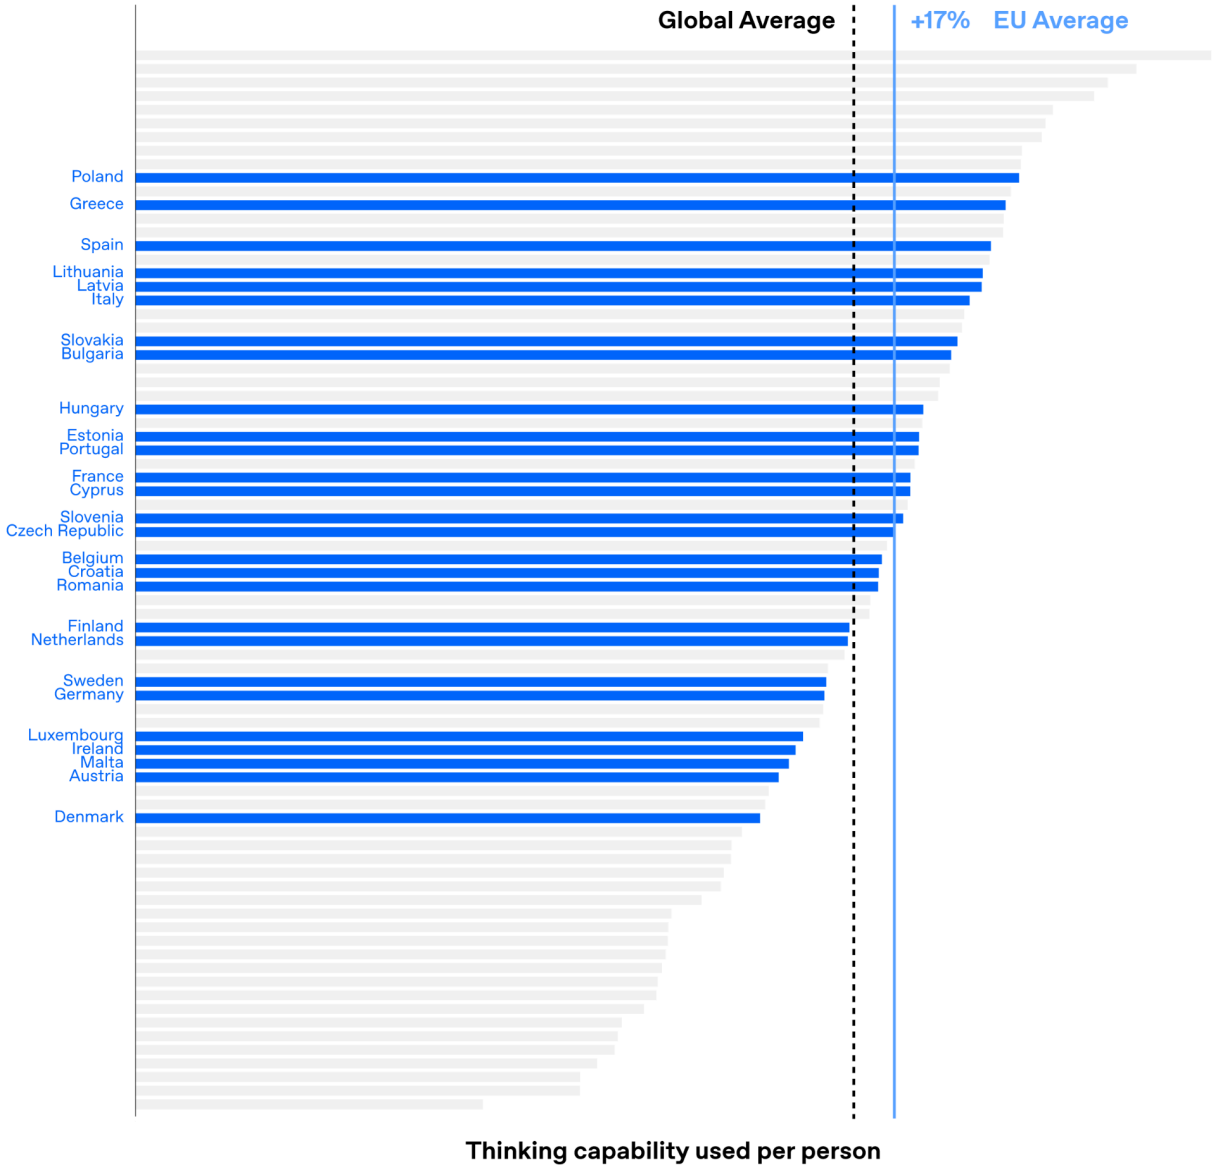

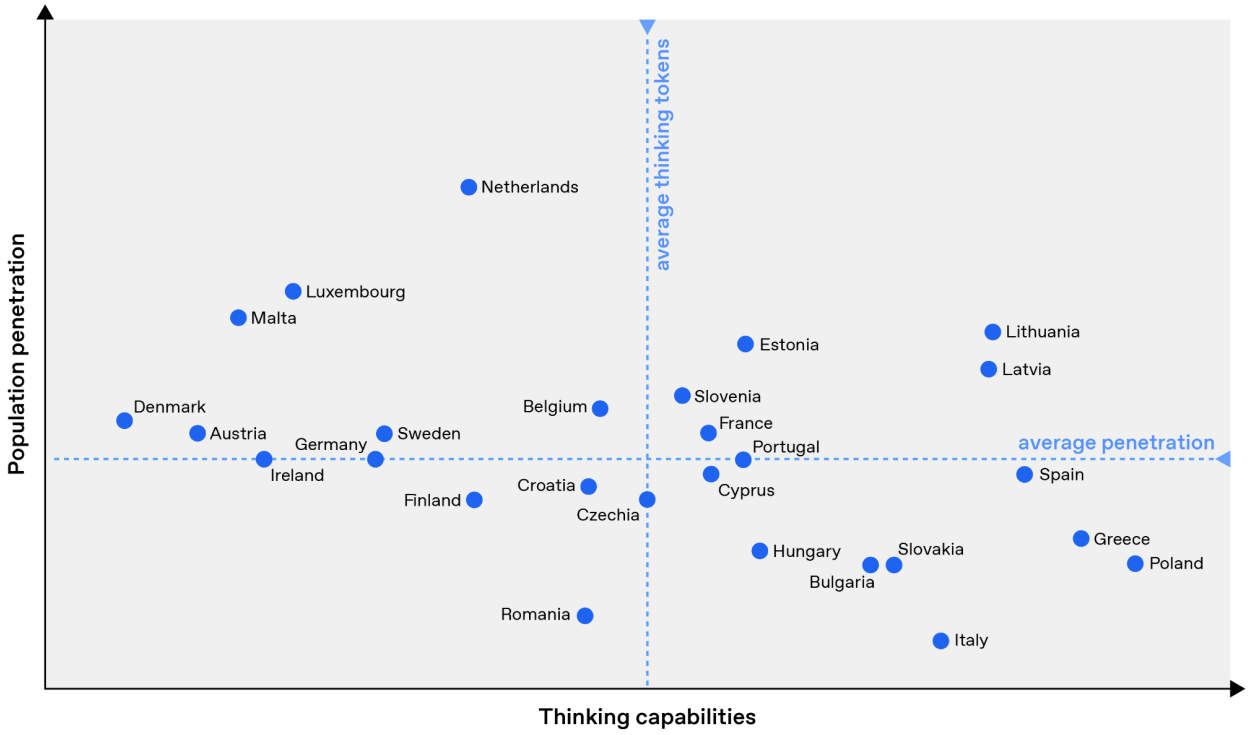


## Methodology

Population penetration measures the share of total population in a country who use ChatGPT across all plans (including Free), reflecting the total breadth of usage.

Thinking capabilities measure the depth of usage and reflect the volume of reasoning tokens used by the median monthly active Plus user. See the global Ending the Capability Overhang report for more information about methodology.

AI capabilities are advancing at unprecedented speed, but their economic impact depends on how fully they are adopted in the real world. While frontier AI systems can now reliably perform complex, multi-step tasks that would take human experts more than 30 minutes, most users still rely on far simpler uses. This creates a growing capability "overhang" or "gap" — the difference between what AI tools are capable of and how they are typically used in practice. Closing this gap matters because deeper, more advanced use of AI is associated with higher productivity gains, time savings, and the ability to tackle more complex problems. Without broad uptake of advanced capabilities, countries risk leaving significant economic potential untapped.

To assess how deeply AI capabilities are being utilised, we analysed ChatGPT usage data and developed a proxy measure for usage depth: thinking capabilities - we measure the amount of effort the model uses to respond to user queries. When users ask more difficult and complex questions (such as long analysis or planning, data analysis, problem-solving, coding, etc.), the model typically spends more effort reasoning to produce a better answer, making thinking capabilities a useful indicator of advanced AI use. Our analysis focuses specifically on ChatGPT Plus users across the EU, allowing for a like-for-like comparison between users with the same level of access to frontier AI capabilities.

Globally, the data shows a large capability gap: the typical power user (95th percentile) uses roughly 7× more thinking capabilities than the typical user (50th percentile), and leading countries use around 3× more per-person. The majority of EU Member States outperform the global average, and EU countries on average use 17% more thinking capabilities per person than the global baseline. However, the most thinking-intensive country uses approximately 40% more thinking capabilities than the least intensive, and 9 EU countries still fall below the global average.

The EU data reveals three distinct patterns of AI use, reflecting different stages of adoption and depth rather than simple leaders and laggards.

1) **High Penetration / High Use of AI Thinking Capabilities:** First, a smaller set of countries combines both high population penetration and high-intensity use of advanced AI thinking capabilities. Lithuania and Latvia stand out in this category, indicating not only widespread access to AI, but also broad diffusion of deeper, more sophisticated use across users. This pattern points to ecosystems where advanced capabilities are moving beyond early adopters and into mainstream workflows.

2) **Low Penetration / High Use of AI Thinking Capabilities:** Second, some countries combine high-intensity use of advanced AI thinking capabilities with relatively lower population penetration. In these contexts — such as Poland, Greece, and Italy — Plus users are disproportionately using AI for complex, multi-step, and higher-value tasks. This suggests a more concentrated or specialist user base, where advanced capabilities are already embedded in demanding workflows, even if overall access remains narrower. Expanding broad access to AI technologies and AI literacy could unlock significant untapped potential across these countries.

3) **High Penetration / Lower Use of AI Thinking Capabilities:** Third, a different group of countries shows high population penetration, but lower average use of advanced AI thinking capabilities.

Examples include the Netherlands, Luxembourg, and Malta. Here, AI adoption is broad and mainstream, spanning a wide range of users and everyday use cases. However, there is ample opportunity to raise the bar for utility across these mainstream users by deepening value creation: moving from surface-level use cases toward higher-impact applications, integrating AI more systematically into business and industrial workflows, strengthening founders and start-up ecosystems, and anchoring applied AI research centres more directly to concrete industry and productivity needs.

While some countries lie near the margins, no countries are firmly positioned in the low-penetration, low-thinking-capabilities segment.

Taken together, these patterns reinforce a central insight: as AI adoption expands, average intensity often dilutes unless deeper forms of use spread alongside access. Europe's next productivity gains will come not from choosing between breadth and depth, but from enabling more users — especially SMEs and public-sector workers — to move from surface interactions to higher-value applications as adoption scales.

## Recommendations

- **Move from access to effective use:** Broad AI adoption alone is not enough. As penetration increases, average depth of use can dilute unless users are supported to apply AI to more complex, higher-value tasks. Policy should therefore prioritise measures that help mainstream users across the economy move beyond surface interactions and integrate advanced AI capabilities into real workflows.
- **Build on national strengths while closing critical gaps:** European countries display different AI strengths and weaknesses across their economies. Policy should reinforce areas of existing excellence—such as advanced thinking, coding, and technical use in countries like Poland—while ensuring that weaker sectors, regions, or firm sizes are not left behind. The objective is not uniformity, but balance: scaling what already works, and deliberately extending AI capability into parts of the economy where adoption or effective use remains limited.
- **Treat skills and adoption as productivity policy for competitiveness:** Closing the capability gap is ultimately an economic challenge, not a technical one. Deeper AI use is closely linked to productivity gains, time savings, and problem-solving capacity. Policies that invest in practical AI skills, sector-specific applications, and trusted deployment environments should be viewed as core components of Europe's competitiveness and growth agenda. We share more on how we are supporting this effort by working closer with SMEs, as an example, later in the report.

# Bridging the Gap

## Increasing adoption across European businesses

Among leading European enterprises, AI adoption has moved beyond experimentation and pilot projects and is increasingly reshaping how work gets done inside organisations. Rather than treating AI as a standalone tool, leading companies are integrating it directly into everyday workflows and repeatable processes. According to OpenAI's State of Enterprise AI Report, ChatGPT use in enterprises among these frontrunners is not only accelerating but deepening: workers report saving 40-60 minutes per day, with heavy users reporting more than 10 hours per week, while 75% of workers report that using AI at work has improved either the speed or quality of their output. Additionally, AI is helping people do new kinds of work with 75% of users reporting being able to complete new tasks they previously could not perform.

This reflects a broader transformation in how leading European businesses operate. AI is emerging as an operating layer for modern work, moving beyond chat-based experimentation toward integrated use across teams, departments, and sectors. European enterprises are central to this global adoption trend. France and the Netherlands rank among the fastest-growing business customer bases year-over-year, while Germany and the United Kingdom are now among the largest enterprise markets globally outside the United States.

At the same time, this level of adoption remains uneven across firm sizes and sectors, in line with the above-mentioned capability gap within the European economy. In 2025, according to Eurostat, around 20% of EU companies reported using AI, but uptake is significantly higher among large enterprises than among SMEs. While 55% of large firms now use AI (up from 41% in 2024), adoption among small enterprises stands at just over 17% (up from 11% in 2024). This growing gap highlights the risk of a new digital divide and underlines the importance of policies and initiatives that help smaller firms move from initial exposure to effective, scaled use, allowing them to fully capture associated productivity gains. This is critical for the EU economy, since AI can help small and medium-sized businesses operate with the resources and capabilities of much bigger rivals.

Even amongst large enterprises, OpenAI's research shows a gap across AI usage in workers and firms. Frontier workers (in the 95th percentile), send 6x more messages than the median employee and engage more intensively across advanced capabilities. Frontier firms send 2x more messages per seat and show deeper integration of AI across teams.

Across continental Europe, a growing number of organisations are already deploying OpenAI tools at scale to improve productivity, support employees, and modernise operations. These examples illustrate what is possible when the right capabilities, skills, and organisational conditions are in place, and how AI adoption is translating into practical, value-driven use across sectors:

**🇪🇸 BBVA (Banking, Spain)**
BBVA uses OpenAI tools to support employees across functions, including software development, documentation, and customer-facing operations. AI helps teams work more efficiently with complex information, improve response quality, and streamline internal processes in a highly regulated banking environment.

**🇳🇱 Booking.com (Travel, Netherlands)**
Booking.com applies OpenAI tools across customer service, internal knowledge management, and product development. AI supports teams in handling large volumes of information, improving customer interactions, and accelerating collaboration across a global workforce.

**🇸🇪 Scania (Manufacturing, Sweden)**
In the manufacturing sector, Scania integrates OpenAI tools into engineering and operational workflows. AI supports employees working with complex technical data, documentation, and problem-solving tasks across product development and production environments.

**🇩🇰 Lundbeck (Life sciences, Denmark)**
In life sciences, Lundbeck uses OpenAI tools to assist researchers and operational teams in managing scientific information, drafting technical materials, and supporting data-intensive workflows, enabling employees to focus more time on high-value research and decision-making.

**🇩🇪 Deutsche Telekom (Telecommunications, Germany)**
Deutsche Telekom and OpenAI have entered a multi-year collaboration to bring advanced AI capabilities to both internal teams and millions of customers across Europe. The partnership includes rolling out ChatGPT Enterprise across the company to streamline customer care, enhance workflows, and support innovation, while co-developing new multilingual AI experiences that improve communication, productivity, and network operations in a privacy-focused way.

**🇸🇪 Spotify (Entertainment, Sweden)**
OpenAI integrates Spotify through their app in ChatGPT to enhance music and podcast discovery, allowing listeners to connect their accounts and receive personalised recommendations, playlists, and suggestions via natural language prompts. This partnership brings conversational AI into the music streaming experience, helping users explore content more intuitively while respecting privacy and user control.

### Recommendations

- **AI Adoption Index**: Complement quantitative adoption metrics (e.g., adoption rate) with qualitative indicators (e.g., usage depth) that will help measure and bridge the capability overhang across countries and sectors.
- **AI Deployment & Training Vouchers for SMEs and Non-Profits:** Introduce targeted measures — such as AI deployment vouchers or upskilling programs — to support adoption, integration, and workforce uptake of AI tools by small businesses and non-profits organization, so as to help bridge the adoption gap.

More ideas to boost AI uptake in the Hacktivate AI report.

## Boosting AI skills in education and work

AI adoption at scale depends not only on access to technology, but on people having the skills and confidence to use it effectively. To translate AI's potential into tangible economic and social benefits, individuals, organisations, and public institutions must be equipped to integrate AI into learning, work, and service delivery. Education and upskilling are therefore central to Europe's AI transition, underpinning productivity, competitiveness, and inclusive growth.

This is already reflected in how AI is used today. OpenAI's analysis shows that learning and education are among the most common use cases for ChatGPT, with millions of users relying on it for study support, skills development, and professional learning. AI is increasingly used as a practical learning assistant — helping people acquire job-relevant skills and adapt to changing workforce demands.

Recognising the central role of skills, OpenAI has made education a core pillar of its OpenAI for Countries initiative. In Europe, this has translated into partnerships with governments that place learning, workforce readiness, and support for research and innovation at the heart of national AI strategies.

In Estonia, OpenAI is partnering with the government through the AI Leap 2025 programme, which integrates AI tools across the national education system. From September 2025, around 20,000 secondary school students and 3,000 teachers gained access to ChatGPT Edu, alongside structured teacher training and pedagogical support. In Greece, OpenAI is piloting ChatGPT Edu in selected upper-secondary schools, with a focus on teacher training, inclusive implementation, and responsible classroom use, while also supporting talent retention and startup ecosystems in priority sectors.

Alongside country-level partnerships, OpenAI is investing in structured learning pathways and formal certification to help individuals and organisations build recognised AI skills. This includes the launch of OpenAI's first certification programmes, such as AI Foundations and ChatGPT Foundations for Teachers,

which are being piloted with universities, employers, and public-sector partners and provide credentials that signal job-ready capabilities.

Building on this, OpenAI is bringing OpenAI Academy to Europe, starting with the SME AI Accelerator, supported together with Booking.com. The programme will deliver free, practical training — both online and in person — across several Member States, including France, Germany, Italy, Ireland, the UK, and Poland, helping 20,000 SMEs turn AI into a driver of productivity and competitiveness.

Europe's AI ambitions will ultimately be judged by how widely their benefits are shared. Skills and education are the bridge between technological capability and real-world impact. As AI adoption accelerates, sustained investment in skills — across education systems, workplaces, and SMEs — will be essential to ensure that AI contributes to productivity, resilience, and inclusive growth across the Union.

## Recommendations

- **Portable AI Skills:** Standardize and certify job-level AI skills at EU level to ensure workers can apply them across employers and sectors.
- **Individual AI Learning Account:** Create for every EU worker an account with annual training credit which they can spend on approved AI-related courses or certifications.
- **National AI-in-Education Frameworks:** Adopt national frameworks for the safe and responsible use of AI across education systems, covering access, teacher training, responsible use, and curriculum integration. This would ensure AI supports skills acquisition, educational equity, and workforce readiness from school through lifelong learning.

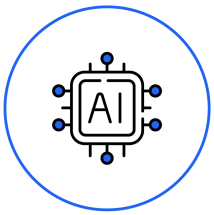


# Enabling Europe's AI Uptake

Closing Europe's capability overhang will not happen through skills and adoption alone. Even where access is expanding and advanced use is emerging, progress will stall without the right enabling conditions in place. Infrastructure, policy, and trust form the backbone that allows AI capability to move from experimentation to reliable, scaled adoption across the economy. This section sets out the foundational conditions required to translate Europe's growing AI skills and adoption into sustained productivity gains — ensuring that advanced capabilities can be deployed securely, at scale, and with public confidence.

## Infrastructure is destiny

AI represents a structural shift from previous waves of digital innovation. It is not simply software layered on existing systems, but a general-purpose capability that depends on a new class of infrastructure — combining compute, advanced chips, energy, data, and security controls — developed at unprecedented scale. We have seen that frontier AI capability is advancing faster than economies' ability to deploy and absorb it. Infrastructure is a key bridge between what AI can do and what organisations can reliably put into production workflows. Unlike earlier technologies, AI infrastructure must often be built before demand fully materialises, due to high upfront costs, long construction timelines, and strong lock-in effects. As a result, early, coordinated investment is now a precondition for adoption, competitiveness, and increased capability.

Recognising this shift, OpenAI has focused on enabling AI infrastructure at the scale required by next-generation systems. This includes the launch of our Stargate initiative, designed to support long-term capacity, reliability, and secure integration across compute, energy, and deployment environments, treating AI infrastructure as a strategic, long-horizon investment. Crucially, capacity must translate into deployability: infrastructure should be planned not only for training, but also for widespread, low-friction inference and integration into day-to-day enterprise and public-sector services. In parallel, OpenAI launched OpenAI for Countries in response to growing demand from governments to treat AI as national infrastructure and to ensure that AI systems can be deployed at scale in ways that align with domestic ambitions in terms of AI-driven state modernisation and matching strong local legal, security, and governance requirements.

OpenAI for Countries is now being implemented in concrete national contexts in Europe. In Germany, OpenAI, SAP and Delos have launched OpenAI for Germany, a sovereign AI infrastructure initiative enabling large-scale AI deployment in the public sector within trusted, locally governed environments, backed by significant domestic compute investment. In Norway, Stargate Norway is under development, providing large-scale AI compute powered entirely by renewable energy and designed to support both sovereign and commercial workloads in a sustainable way. These partnerships reflect a broader European debate: many governments are seeking greater resilience and assurance over how critical digital infrastructure and data processing are governed and operated. OpenAI's approach is to engage directly and pragmatically —

listening to requirements and working with trusted local partners and governance models so that frontier systems can be deployed securely, at scale, and in line with European institutional realities.

In parallel, OpenAI now offers EU at-rest and inference data residency for ChatGPT Enterprise and Edu, allowing organisations to keep data stored in the EU and ensure prompts and responses are processed in-region, supporting resilient and compliant deployment at scale. Alongside infrastructure build-out, OpenAI is expanding its European presence by hiring public-sector and national-security experts to ensure AI systems reflect European institutional realities and public expectations.

As the EU advances initiatives such as AI gigafactories and EuroHPC expansion, the priority should be speed, cooperation, and effective use of infrastructure once built. AI infrastructure is now inseparable from economic, energy, and security policy, and Europe's competitiveness will depend not only on building capacity, but on enabling its deployment through coherent, workable regulation and strong public-private partnership. That also means reducing friction for Europe-wide roll-out — including procurement and compliance clarity, and interoperability — so that deployment is feasible not just for large enterprises, but also for SMEs and public administrations. Ensuring that Europe's governance framework supports, rather than constrains, the effective use of AI infrastructure will be critical to translating Europe's assets into lasting technological and economic leadership.

## Recommendations

- **AI Infrastructure Enablement**: Adopt policies that scale and democratise access to compute, specialised chips, energy, data, and security infrastructure essential for developing and deploying advanced AI systems.
- **AI Capacity Deployment Acceleration**: Prioritise rapid, coordinated deployment and effective use of AI infrastructure by aligning regulation, energy, and security policy and enabling strong public-private partnerships that turn built capacity into real economic and technological impact.
- **Define Sovereign-Compatible Operating Models:** Enable trusted, locally governed deployment environments (with clear specs on governance, assurance, and operational controls) that directly address European resilience expectations while enabling access to frontier capabilities.
- **Couple Compute Investments and Deployment Pathways:** Couple European and national compute investments with practical deployment support so public and private sectors, including SMEs, can move from access to effective use and close the capability gap.

## A rulebook that supports innovation

As AI systems grow in capability and adoption, safety depends not only on robust internal safeguards but also on constructive engagement with public authorities. Europe has positioned itself as a global leader in AI governance, and OpenAI's engagement with the EU framework reflects a commitment to approaches that are serious, practical, and forward-looking, focused on delivering real-world safety outcomes rather than formal compliance alone.

Over the past year, OpenAI has strengthened its internal safety architecture to better anticipate and manage frontier AI risks. This includes updates to the Preparedness Framework, with clearer capability thresholds, escalation pathways, and forward-looking risk categories, as well as the publication of an updated Model Spec that transparently sets expectations for model behaviour. Together, these tools reflect an approach grounded in early risk identification, transparency, and continuous iteration as models and use cases evolve.

This internal work is reinforced through close collaboration with public and private partners. OpenAI actively engages with regulatory frameworks, including as the first U.S. AI company to sign the EU AI Act's General Purpose Code of Practice, helping translate high-level legal obligations into workable, safety-enhancing practices. OpenAI works with the U.S. Center for AI Standards and Innovation and the UK AI Security Institute to support independent evaluation, testing, and research on advanced AI systems, and is a founding member of the Frontier Model Forum to advance shared safety research and best practices across industry.

As the AI Act moves from legislation to implementation, the EU should continue to work with stakeholders to frame compliance in a way that is common sense and fits with the EU's overall simplification agenda. Europe already has one of the most comprehensive AI governance frameworks globally; overlapping requirements and legal uncertainty risk diverting resources away from safety and compliance. Simplification efforts like the Omnibus agenda are essential to ensure coherence, proportionality, and workability in practice. New legislative initiatives should meet a high bar, addressing clearly identified gaps and adding demonstrable value. Europe's leadership in AI governance will ultimately be measured not by the number of rules adopted, but by how effectively they deliver safety, trust, and innovation.

### Recommendations

- **Regulatory harmonization mission:** A 1-year mission to align all definitions, taxonomies and thresholds across existing digital regulation as part of the Digital Fitness Check.
- **International cooperation:** Promote interoperability and mutual recognition for AI safety and security testing within the International Network of AI Safety Institutes so as to reduce fragmentation and foster regulatory convergence internationally.

More recommendations in our response to the European Simplification consultation.

## Responsible AI in Europe

As Europe moves from rule-making to implementation, trust becomes the central test of AI governance. Over the past year, OpenAI has implemented strong, age-appropriate protections for teens. This approach is set out in OpenAI's Teen Safety Blueprint, including using age prediction to apply the right protections to the right users, defaulting to protective safeguards when there is doubt about a user's age, identifying and mitigating risks to minors through strong under-18 safety policies, and empowering families with accessible parental controls. In practice, this has translated into stricter boundaries around sensitive content, protective model behaviour in situations involving distress or self-harm, and specific guidance for interactions with users under 18 embedded directly in our Model Spec.

OpenAI has also deepened its engagement with European civil society, researchers and experts, including through a dedicated Wellbeing Council translating research into real-world safety measures that guide product design and policy. Building on this work, OpenAI has announced a €500,000 grant programme for EMEA NGOs and research institutions working on youth safety and digital well-being, supporting organisations improving youth safety and wellbeing in the age of AI. Alongside these measures, OpenAI maintains long-standing, robust, and non-negotiable protections to prevent child sexual exploitation and abuse, which form a foundational baseline for all youth-related safeguards.

Trust also means empowering developers to build safer experiences. In partnership with ROOST, we've released a research preview of gpt-oss-safeguard—a family of open-weight safety reasoning models for safety classification that lets developers bring their own policies (e.g., scam detection, misuse prevention, conversation moderation, and emerging-risk flagging) instead of relying on fixed, one-size-fits-all rules.

As implementation progresses, the EU has an opportunity to reinforce trust-building practices that are operational, adaptive, and grounded in real-world use. This includes encouraging design-led protections for young people, supporting age-appropriate safeguards such as parental controls, fostering structured engagement between companies and civil society, and maintaining high standards for the protection of sensitive data. By recognising and reinforcing these approaches, the EU can help ensure that trust in AI is not only regulated, but actively built and sustained over time.

### Recommendations

- **Make safety safeguards the implementation standard, not just paperwork**: Policymakers should require providers to demonstrate real-world protections, including effective CSAM prevention tools, clear user and parental controls, and transparency, supported by meaningful evidence.
- **Implement age-appropriate protections without mandating hard ID checks**: Require a proportionate age-assurance framework that supports: (i) age-appropriate experiences by default, (ii) parental tools and family resources, and (iii)

age-estimation/prediction, with appeal pathways—rather than blanket requirements for upfront hard ID checks.

- **Treat safety tooling and civil-society partnerships as core infrastructure**: Create incentives for shared safety infrastructure—including open, auditable safety tools (e.g., policy-driven safety classifiers) and structured, ongoing engagement with researchers/NGOs (expert councils, grants, external testing) to translate evidence into safeguards and keep protections adaptive over time.

**About OpenAI**

Artificial intelligence is an innovation like electricity—it will change how we live, how we work, and how we engage with one another. OpenAI's mission is to ensure that artificial general intelligence benefits all of humanity. We're building AI to help people solve hard problems because by helping with the hard problems, AI can benefit the most people possible—through more scientific discoveries, better healthcare and education, and improved productivity. We're off to a strong start, creating freely available intelligence being used by more than 800 million people around the world, including 4 million developers. We believe AI will scale human ingenuity and drive unprecedented productivity, economic growth, and new freedoms that help people accomplish what we can't even imagine today.

***Cover image created with ChatGPT Images***