MICROSOFT RESPONSE

to the European Commission's call for evidence for the Apply AI strategy

Microsoft welcomes the opportunity to participate and provide feedback to the call for evidence on the Apply AI Strategy (AAIS). We appreciate the European Commission's efforts to address the critical need for AI adoption in EU strategic sectors and beyond to make Europe an AI continent.

This contribution will touch upon the investments and conditions needed to rapidly accelerate the development and diffusion of AI throughout the European economy, and as such address the four main challenges that the call for evidence identified, the three main objectives that the AAIS aims to achieve, as well as the need for regulatory simplification, as outlined in the fifth pillar of the AI Continent Action Plan.

1. Investing in infrastructure and people to accelerate AI adoption throughout Europe

Microsoft fully supports the European Commission's main objectives for the AAIS. Wide integration of AI throughout the EU's industrial and public sectors and further unlocking the EU's innovation potential with AI will be essential for its competitiveness and enable its companies to become global AI frontrunners. Importantly, broad diffusion of AI will be critical to address other critical challenges facing the continent, including demographic developments, energy, climate and biodiversity.

AI is rapidly becoming a "General Purpose Technology", or GPT. Development and subsequent wide adoption of transformative GPTs, like electricity and the internet, have not only driven economic growth but sparked new discoveries and inventions, changing the way we live and work. Notably, the most important long-term determinant of a country's economic growth during an industrial revolution is not whether it is at the forefront of innovation in a "leading sector" of the time. Instead, it's whether the country "diffuses"—or spreads—the adoption of a critical GPT broadly across its economy. **This key insight has profound implications for the impact of AI over the next 25 years and should be appropriately considered in the further development of the AAIS.** While there are important policy issues to be addressed around leadership in key sectors and AI technology development, it's equally, if not more, important to address what it will take to ensure the widespread and effective adoption of AI across all the societal sectors that can benefit from it.

We recognize that Europe wants and needs a world class and broad AI and cloud ecosystem. The Commission's proposed Cloud and AI Development Act, AI Factories and AI Gigafactories place an important emphasis on the need for large-scale AI and cloud infrastructure. They can create real-world capabilities at scale to fuel business and manufacturing innovation, run national health systems, enable secure government services, and support digital tools in education. **Microsoft is proud of its contribution in this space, and is stepping up further to help Europe meet its AI ambitions.** This spring, we announced plans to increase our European datacenter capacity by 40% over the next two years, expanding datacenter operations in 16 European

countries and thereby more than doubling our European datacenter capacity between 2023 and 2027. It will result in cloud operations in more than 200 datacenters across the continent.

These large-scale investments call on suppliers that are manufacturing critical components throughout the EU, including Poland, Finland, Ireland, Germany, France and Italy. In addition, **we are developing and operating our AI infrastructure and platform services with a constant focus on Europe's needs.** Microsoft's AI infrastructure is accessible, open, and available on fair terms to the entire European economy, as set out in our AI Access Principles that govern our operations.

Another critical element of AI adoption at scale is investments in people, both through advanced AI skilling and vocational training. Building and operating Europe's AI and datacenter infrastructure will depend on a skilled workforce, from electricians and pipefitters to AI engineers and system designers. We also recognize that AI fluency must extend across sectors and job roles, not only among developers, but also frontline workers, policy makers and organisational leaders. **Advanced AI skilling infrastructure is therefore essential in growing the professions that will design, develop and deploy AI-powered solutions throughout the economy.** This will require EU Member States, coordinated by the European Commission, to develop and invest in national AI skilling strategies, built on strong collaboration between the public and private sectors. These strategies should build on existing disciplines like computer and data science, anticipating how these fields will evolve into future careers and jobs for AI engineers, AI systems designers, applied AI practitioners, and others. Equally important, they must integrate AI competencies into key sectors, such as education, manufacturing, healthcare, public administration, etc, while scaling access to training across all parts of the workforce. That's why we are investing heavily in AI skilling across Europe, including through our AI Skilling Initiative. We are partnering with government, education providers, industry, and civil society to support both foundational and advanced AI competencies. Through our strategic partnerships, we have helped skill 2.9 million Europeans and are on track to engage 8 million people by the end of the year.

2. Continuing Europe's path towards a Data Union

Data is the fuel that powers AI and Europe can achieve better outcomes from AI with better access to high quality data. The quality, quantity, and accessibility of data directly determines the strength and sophistication of AI models. We therefore welcome the Commission's announcement to propose a Data Union Strategy. While a large amount of data remains untapped, **Europe should find ways to incentivize more data sharing, including the application of incentive structures to publish more open government data and more open research and science data.** Yet today, despite deliberate efforts at national and EU level, many of these datasets remain either inaccessible or not usable for AI development. By making more data publicly available for AI training, the EU can significantly accelerate the advancement of AI capabilities, driving innovation and discovery. **To enable broad access to data and promote a thriving data ecosystem, Microsoft is supporting the establishment of open data commons.** Opening access would allow for the analysis of themes, patterns, and insights across broad datasets, ensuring Europe's cultural heritage and diversity is reflected in the development process of AI, while importantly at the same time leveling the playing field. Because accessible

public data empowers not only large companies, but also startups, academic institutions, and nonprofits to train and refine AI models. This fosters a more competitive and inclusive AI ecosystem, where innovation is driven by ideas and ingenuity – not just proprietary data.

3. Ensuring a streamlined regulatory framework that builds trust and spurs innovation

Widespread adoption will only materialize in the face of public or social acceptance, which in turn requires both usefulness and trust: Technologies must solve real-world problems and improve people's lives. At the same time, they must be trustworthy, with safeguards in place to protect a country's societal and ethical values. The AI Act, the EU's centerpiece of its regulatory framework for AI, aims to promote the uptake of human centric and trustworthy AI, while ensuring a high level of protection of health, safety and fundamental rights. At Microsoft, we are building our products and services to comply with our obligations under the EU AI Act and working with our customers to help them deploy and use the technology compliantly. As the Commission is developing further guidelines and secondary regulations, it should **ensure this is done in a balanced and straightforward way, that fosters trustworthy and secure AI while allowing companies of all sizes to innovate and grow**. Moreover, the Commission and national authorities should ensure that enforcement approaches and interpretations are closely aligned and pragmatic, in order to allow development and adoption throughout Europe to flourish.

In addition, in view of the rapid developments in the field of AI, and to maintain Europe's competitive edge in the global AI landscape, **we welcome the Commission intention to include the AI Act as a whole within its digital regulatory simplification exercise.** In this context, it is concerning that European harmonized standards related to the application of high-risk requirements are currently delayed and yet crucial for companies to facilitate compliance. The focus should therefore be on prioritizing such standards, rather than providing the possibility for the Commission to adopt common specifications. In addition, more flexibility is also needed in how providers of high-risk AI systems design their post-market monitoring plans.

Ahead of a possible legislative streamlining proposal, key implementation efforts for GPAI models currently underway ahead of their legal application date in August should also be considered in the simplification context, including through significant streamlining of the final version of the GPAI models Code of Practice, announcing a grace period for implementing this Code beyond the (imminent) legal deadline, and ensuring the Commission's upcoming GPAI guidelines avoid burdensome interpretations of model fine-tuning and modifications, as well as overly detailed reporting requirements in the transparency template, both leading to significant trade secret concerns. With regard to designations of GPAI models with systemic risk, the AI Office should urgently provide more clarity on how it will use Annex XIII criteria for ex officio designations, announce it will work towards introducing a capabilities-based approach to assessing systemic risk over the mid- to long term, and accelerate the issuing of a draft delegated act to increase the compute threshold for systemic risk GPAI models before model requirements start to apply on 2 August. This will give a clear signal to AI developers about the intent to capture only the most highly capable models as GPAI models with systemic risk. Further detailed proposals can be found in the annex to this paper.

4. Addressing dependency and sovereignty concerns in Europe

As the EU charts its path towards digital transformation of its economies and societies, and in the face of continued Russian aggression on the continent and increased global geopolitical volatility, it is understandable that questions around dependency and sovereignty of digital technology become more prominent.

As a multinational company, we believe in Transatlantic ties that promote mutual economic growth and prosperity. Since the launch of our first Microsoft product localized for Europe forty-two years ago, our economic reliance on Europe has always run deep. We recognize that our business is critically dependent on sustaining the trust of customers, countries, and governments across Europe. We respect European values, comply with European laws, and actively defend Europe's cybersecurity. Our support for Europe has always been–and always will be–steadfast.

Microsoft therefore supports the aims of the AAIS to address concerns identified around the large share of AI developments taking place outside the EU and reliance on technology from companies not headquartered in the EU, including perceived risks related to this, in particular for strategic sectors. As we will always be dedicated to creating jobs, promoting economic opportunities, and strengthening cybersecurity on both sides of the Atlantic, Microsoft recently announced additional commitments to take our support for Europe another step forward, including:

- Incorporating a "**European Digital Resilience Commitment" into all of our contracts** with European national governments and the European Commission, that specifies we will promptly and vigorously contest any government order to suspend or cease cloud operations in Europe, using all legal avenues available, including by pursuing litigation in court.
- Designating and relying upon **European partners with contingency arrangements for operational continuity** in the unlikely event Microsoft were ever required by a court to suspend services, using back-up copies of our code stored in a secure repository in Switzerland, together with the legal rights needed to access and use this code if needed for this purpose.
- Further cementing the nexus between Microsoft and Europe; going forward, our European datacenter operations and their boards will be overseen by a **European board of directors that consists exclusively of European nationals** and operates under European law.
- Further **expanding our European datacenter capacity by 40% over the next two years**, as touched upon above, using a diversified approach to public cloud datacenters (including our Cloud for Sovereignty) and supporting European cloud providers, to complement our Sovereign cloud datacenters. The tens of billions of Euros in datacenter investments across Europe are permanent structures, subject to local laws, regulations, and governments, that Microsoft is committed to adhere to.
- Further **strengthening and expanding solutions that allow European customers to control and protect their data**, on top of the robust, world leading capabilities across the entire cloud stack (including the EU Data Boundary and Confidential Compute) and strong contractual commitments that we offer our customers.
- Introducing new measures to protect against new cyberthreats, including a new European Security Program, the **appointment of a new Deputy CISO for Europe** as part of the

Microsoft Cybersecurity Governance Council and **engaging an independent auditor to verify and validate our commitments** to Europe.

- Further **enhancing our AI Access Principles** announced last year, to help strengthen Europe's economic competitiveness, including for open source.

5. Europe's path towards an AI continent

While measures to address the challenges outlined in the call for evidence for the upcoming Apply AI Strategy, and the AI Continent Action Plan more broadly, are pertinent, we would emphasize that the most important long-term determinant of a country's economic growth in the face of a critical General Purpose Technology is not whether it is at the forefront of innovation in a "leading sector", but rather how the country "diffuses"—or spreads—the adoption of such a technology broadly across its economy. Moreover, we encourage the Commission to reflect the importance of maintaining strong and trusted partnerships between public and private sector actors, irrespective of where they are headquartered, acknowledging mutual dependence and shared goals and interests, and taking into account continued commitment to constructively and effectively address concerns as they arise. Additionally, the Commission should endeavour to ensure the regulatory framework for AI is balanced and streamlined, and enforcement across the EU is coherent and pragmatic, in order to incentivize companies to develop and adopt AI. Finally, we would call upon the Commission to ensure that critical aspects such as security and sustainability remain key determinants for participation in initiatives and investments to help achieve Europe's AI ambitions.

ANNEX

EU AI Act simplification recommendations

GPAI model related recommendations

1. Revise the compute threshold to 10^26 FLOPs

We support continued efforts to capture only the most highly capable models as GPAI models with systemic risk. We welcome the AI Office's continued efforts to align the AI Act with technological developments, either by adjusting the value of the compute threshold and/or by specifying and updating the additional criteria for ex officio designation.

We furthermore appreciate the AI Office's acknowledgment that training compute is an imperfect proxy for generality and capabilities, and agree that compute currently provides the best metric for legal certainty, especially as scientific approaches to evaluation and measurement for the capabilities most relevant to highly capable models are still being developed.

The current compute threshold (10^25 FLOPs) risks being *too blunt*, bringing in a broader range of less capable models than the regulation originally intended.

We therefore recommend accelerating the issuance of the draft delegated act to raise the current compute threshold to **10^26 FLOPs**, before model requirements start to apply on 2 August 2025, to give a clear signal to AI developers about the continued intent to capture only the most highly capable models as GPAI models with systemic risk.

2. Introduce a capabilities-based approach to assessing systemic risk over the mid- to long term

As evaluation science matures and reliable benchmarks with lower saturation rates emerge, we recommend either complementing or replacing the compute threshold with a more direct assessment of capabilities associated with systemic risks, such as CBRN weapon development, offensive cyber capabilities, and advanced autonomy. This separate threshold based on a model's performance should be tested on state-of-the-art benchmarks for advanced general-purpose capabilities[1].

With this approach, models would be presumed to be in scope for systemic risk requirements under Article 55, if trained on 10^26 FLOPs and/or at the frontier in performance on state-of-the-art benchmarks for advanced general-purpose capabilities.

[1] Our FGF identifies the following advanced general-purpose capabilities as precursors to high-risk capabilities: 1) General reasoning; 2) Scientific and mathematical reasoning; 3) Long-context reasoning; 4) Spatial understanding and awareness; 5) Autonomy, planning, and tool use; and 6) Advanced software engineering.

3. Introduce a regular cadence to revising methods to assess systemic risk as model development techniques evolve

We also suggest implementing a regular cadence (e.g., annually) for reviewing the systemic risk threshold(s):

- A pre-training compute trigger should be subject to review regularly (e.g., every six months) given improvements in training efficiency and as new approaches to enhancing model capabilities outside of pre-training are further developed, including techniques leveraging test-time compute.
- When a capabilities-based approach is eventually adopted, we also recommend introducing a regular cadence to update performance benchmarks that might become saturated.

4. Clarify how the AI Office will use Annex XIII criteria for ex officio designations of GPAI models with systemic risk

We further suggest that the AI Office provide additional clarity regarding how it will ex-officio assess models below the compute threshold based on the systemic risk criteria in Annex XIII, ideally before 2 August 2025.

In order to determine whether a potential systemic risk model has high impact capabilities, the Commission and the scientific panel of independent experts can also use criteria listed in Annex XIII of the AI Act such as: number of parameters; quality or size of dataset; input and output modalities; benchmarks and evaluations of capabilities of the model; and, high impact on the internal market due to reach (e.g., based on the number of registered business users or end-users). There are no quantitative indicators or thresholds identified for these additional criteria, which makes it difficult to assess how decisions for ex officio designations of GPAI models with systemic risk would be made.

As with our previous recommendation, we also highly recommend the AI Office provides additional clarifications on the application of Annex XIII, focusing on capability-based criteria that are directly related to a model's capabilities, such as the following criteria as already mentioned in Annex XIII: "Benchmarks and evaluations of capabilities of the model, including the number of tasks without additional training, adaptability to learn, autonomy and scalability, the tools it has access to."

5. Provide a grace period for implementing the GPAI Code of Practice

The finalization of the GPAI Code is delayed despite the AI Act's intention to have a three-month gap between the Code's adoption (intended to be 2 May 2025) and the start of model requirements (2 August 2025). We recommend providing model providers that sign the Code with a sufficient amount of time beyond 2 August 2025, to adequately implement the Code's provisions.

High-risk AI-system related recommendations

1. **Remove the option for the Commission to adopt common specifications if standards aren't developed on time, instead prioritizing the development of harmonized standards in a timely and inclusive manner**

The development of harmonised standards is crucial to ensuring that regulatory requirements remain practical, industry-driven and reflective of technological realities. However, Art. 41 allows the Commission to adopt common specifications when harmonised standards are unavailable, delayed or deemed insufficient. Whilst this aims to address potential gaps, the mere presence of common specifications in the legislative framework discourages investment and engagement in the harmonisation process; Art. 41 should therefore be deleted. Developing harmonised standards within tight timelines already faces challenges, especially in the context of fast-evolving technologies such as AI, and introducing the option of common specifications creates parallel pathways that do not reflect technological evolution or practical feasibility. It is a shared responsibility between industry and the Commission to ensure that harmonised standards are developed in a timely and inclusive manner. Rather than relying on common specifications as a potential substitute, efforts should focus on strengthening the standardisation process itself.

2. **Provide flexibility in post-market monitoring**

Art. 72(3) requires providers of high-risk AI systems to follow a specific post-market monitoring plan, the template for which will be issued by the Commission through an implementing act. This approach limits providers' flexibility in developing monitoring plans tailored to their specific AI systems and risk contexts. Moreover, the process for drafting implementing acts allows very limited opportunities for industry consultation and co-design; as a result, companies are unlikely to meaningfully contribute to shaping the template, raising concerns about its practical feasibility. Art. 72(3) should therefore be revised in order to allow companies to develop plans that fit their organizational structure and technologies rather than requiring them to follow the upcoming template.