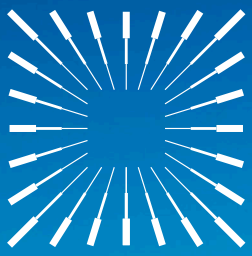

A Transformative AI Strategy for Europe

Contributors

Senior Expert Council

Research Leads

University of Oxford
Ben Bucknall

KIRA Center
Philip Fox

Harvard Kennedy School
Conor McGlynn

Stifterverband
Frauke Stehr

GovAI
Noemi Dreksler

Carnegie Endowment for International Peace
Anton Leicht

Institut Montaigne
Milo Rignell

Special Advisors

Marietta Angeli, Arq Foundation
Onno Eric Blom, Delta Institute
Raphaël Doan, Plan Prométhée
Bertin Martens, Bruegel
Sören Mindermann, CASP, University of Cambridge
Jakob Mökander*, General Catalyst Institute
Jan-Willem van Putten, The School for Moral Ambition

Contributing Experts

Rebecca Arcesati, MERICS
Thorsten Benner, Global Public Policy Institute (GPPi)
Tao Burga, Institute for Progress
Jasmine Dhaliwal, METR
Jonas Freund, GovAI
Pieter Garicano, Works in Progress
Emily Gillett, Institute for Law & AI
Jitse Goutbeek, The School for Moral Ambition
Ryan Greenblatt, Redwood Research
Jakub Growiec, SGH Warsaw School of Economics
Philipp Hacker, European University Viadrina
Fredrik Heintz, Linköping University
Marius Hobbhahn, Apollo Research
Daan Juijn, Arq Foundation
Antoine Levie, KU Leuven
Nick Moës, The Future Society
Robert Praas, Centre for European Policy Studies (CEPS)
Jelle Prins*, Cradle.bio
Roxana Radu, University of Oxford
Georg Riekeles, European Policy Centre
Jacob Schaal, AI Objectives Institute
Tillman Schenk, Bruegel
Jaime Sevilla, Epoch AI
Afek Shamir, RAND Europe
Haroon Sheikh, Vrije Universiteit Amsterdam
Ole Teutloff, University of Oxford
Chloé Touzet, SaferAI
Hans Uszkoreit, German Research Center for Artificial Intelligence (DFKI)
Risto Uuk, Future of Life Institute
Mathilde Velliet, French Institute of International Relations (Ifri)
Philipp von der Wippel, ProjectTogether
Max Welling, University of Amsterdam
Julia Willemyns, Centre for British Progress

Acknowledgements

We would like to thank the following individuals for providing input to this strategy at workshops, in conversations, and in written form:

Markus Anderljung, Benjamin Balde, Johannes Ballentin, Hannes Bastians, Rishi Bommasani, Siméon Campos, Judith Dada, Jimmy Farrell, Zlatko Grigorov, Lennart Heim, Jakob Hensing, Jason Hoelscher-Obermaier, Peter Ide-Kostic, Florian Klumpp, Sebastian Lobentanzer, Sam Manning, Nicolas Miailhe, Richard Moulange, Maximilian Negele, Nuria Oliver, Bálint Pataki, Dylan Rogers, Anneke Schwedler, Fabrizio Silvestri, Stan van Baarsen. We are especially grateful to Stephen Clare for extensive contributions to the drafting and editing of the strategy.

We also thank numerous European government officials and political leaders who have engaged with us throughout the drafting process and whose advice has made this document a more useful strategy.

Contents

Part B Detailed Recommendations 77

About this strategy

The drafting process

This strategy is the work of a group of experts who all contributed in a personal capacity. It is an independent project that received no funding, commissioning, or editorial direction from any company, government, or EU institution. The strategy was written between May and September 2026 and reflects developments up to early September 2026.

Several groups of experts contributed to the strategy. The Research Leads led on drafting and were supported in their work by the Contributing Experts. Members of the Senior Expert Council reviewed drafts, provided strategic advice, and made suggestions for how to improve the usefulness of the strategy further. Special advisors supported in various ways that went beyond questions of content. Across these categories, there was also a small number of industry experts. To preserve the independence of the strategy while benefiting from industry expertise, their role was consultative: any feedback from them on drafts was optional for other contributors to address. They did not vote on or approve the strategy, and held no veto rights. The non-industry members of the Senior Expert Council and Research Leads led on all content questions, with the Editorial Lead being responsible for resolving any disagreements about the content of the strategy. The strategy is not a consensus document, and being listed as a contributor does not imply endorsement of the entire strategy.

Transformative AI raises many more questions than this strategy can cover. The strategy focuses on a set of crucial questions related to how Europe can build agency and economic strength, protect its citizens from AI-enabled shocks, and create robust public institutions in a world with transformative AI. Many important questions related to transformative AI were *not* in scope for this strategy, such as how AI adoption and diffusion might best be supported, military AI, AI's impacts on the information ecosystem, the distribution of gains from AI, AI's impact on the environment, AI's impacts on education and child safety, how to ensure human-centric AI development and deployment, what a healthy democracy and civil society might look like in a world with transformative AI, and many more. The authors consider each of these important questions that require urgent deliberation and solutions.

Authors working on this strategy used generative AI tools to varying degrees. Uses included support for research, fact-checking, reference verification, and consistency checks across chapters, and for literature searches and early drafting support in some chapters. Any AI output was reviewed by the authors. The authors produced the analysis, arguments, and recommendations in this strategy. The Editorial Lead takes final responsibility for the content.

How to read this strategy

This document consists of two main parts, Part A and Part B. Part A sets out the strategy's assumptions, various objectives to be achieved at the Union and national level, and recommendations for how to achieve them. It is written for senior policymakers and informed general readers and can be read on its own. Part B contains implementation details on Part A's recommendations and deep dives on AI compute and robotics. It is written for policy staffers in European governments and EU institutions. Readers short on time are encouraged to read the Executive summary and the recommendations in Part A.

A central theme of this strategy is that guaranteed access to frontier AI – the most capable general-purpose AI systems available at a given time – is crucial for thriving in a world with transformative AI (see Why Europe needs to prepare for transformative AI). Europe could

achieve this either by aiming to have guaranteed access to foreign frontier AI or by aiming to develop domestic frontier AI. This strategy covers both approaches, with a focus on the former because of the extremely high costs of a successful European frontier AI project.

This focus has informed the three main thematic pillars that were at the centre of the drafting process:

- Pillar 1: Securing access to frontier AI
- Pillar 2: Building economic strength
- Pillar 3: Ensuring safety and security

Each pillar contains objectives that are crucial for European policymakers to pursue if Europe is to thrive in a scenario with transformative AI. In addition, throughout the drafting process a set of five objectives that require particularly urgent action has emerged from these pillars. They are highlighted in this strategy as 'Immediate Objectives' – cross-cutting, time-sensitive objectives Europe should pursue with the highest urgency.

Unless noted otherwise, 'Europe' in this strategy refers to the EU and its Member States acting within their respective competences.

Executive summary

Without urgent and far-reaching action beginning this year, the prosperity, sovereignty, and security of all Europeans are at risk as AI advances at a rapid pace. The next months and years of AI progress have the potential for both great opportunity and grave danger. Europe, on its current trajectory, is fully exposed to the latter but only to a small share of the former. Like the rest of the world, Europe faces threats such as AI-enabled cyber or biological attacks or loss of control over AI systems. But even if the world manages to mitigate these threats and AI delivers predominantly positive outcomes, the wealth and strategic leverage that powerful AI generates will, by default, concentrate outside of Europe. Neither of these outcomes needs to be Europe's destiny. But the window for action is closing further every month. Unless Europe wakes up, it risks having no meaningful control over how the most consequential technology ever built will affect its citizens.

AI could have deeply transformative impacts across society before 2030. This would be in line with existing trends, which have led to AI capabilities that would have seemed like science fiction until very recently. For example, over just two years, AI models have leapt from barely completing simple four-minute programming tasks to writing most of the computer code inside the companies that build them. This and similar trends in other domains imply that the world may face transformative AI before the end of the current decade:

> ***Transformative AI:*** *AI systems that change the basic structures of society, including the economy, science, and geopolitics, at a materially faster pace than any technology in history has.*

Europe is structurally unprepared for the near-term arrival of transformative AI and risks permanent marginalisation. Europe hosts only three times as much AI computing capacity as a single data centre site being completed in Malaysia this year. European model developers collectively earn less than 2% of the revenue of just two US counterparts. European companies and institutions were recently cut off from access to frontier AI capabilities in crucial domains like cyberdefence and AI research itself.

In addition, Europe is exposed to several risks from transformative AI without adequate protection. Progress in AI capabilities has been outpacing the ability of researchers to ensure AI is safe, and recent incidents and evaluations have shown that cyber, biological, and loss-of-control risks are real, near-term dangers. Yet, most of the conversations about steering the trajectory of transformative AI – how and if we should build it, who realises its gains and who suffers its risks, who gets to tax, control, and wield it and how – take place outside of Europe.

Europe has the resources to thrive and contribute to global flourishing in a world with transformative AI, but the window for utilising them is closing fast. Europe has a large market, several crucial assets in the AI supply chain, stable institutions, and strong technical talent. However, as we approach a world with transformative AI, these and other assets will, by default, likely lose relative value. For example, in AI development, the value of excellent talent will likely decrease as human labour is replaced with AI labour and as even stronger AI is increasingly built by AI itself ('recursive self-improvement') – a trend that has already begun.

This strategy offers a path for Europe to rise to the challenge of transformative AI. It contains recommendations for ensuring prosperity, sovereignty, and security, and it is centred on the premise that reliable access to frontier AI is crucial in a world with transformative AI. It is impossible to predict exactly the domains in which lacking access to frontier AI would impact Europe most, but likely candidates include defence, cybersecurity, and science, including AI research itself. To avoid falling behind in critical domains like these, Europe would need to either secure access to foreign frontier AI or develop competitive frontier AI domestically. This strategy covers both approaches, with a focus on the former because of the extremely high costs of a successful European frontier AI project.

Member States and the European Commission should pursue five 'Immediate Objectives' with highest urgency. Together, they would put Europe in a position of strength, contribute to a safer transition to transformative AI, and ensure that Europe's institutions can deliver on both.

- **Create a Member State Alliance for Supply Chain Security.** European countries host central chokepoints of the global AI supply chain but fail to capitalise on their strategic and economic potential. An Alliance of willing Member States could, in coordination with the European Commission, do so to the benefit of the whole Union.

- **Make Europe's institutions ready to act in a world with transformative AI.** Member States and the EU should bring frontier AI expertise and technology into their institutions so that they can keep pace with technological developments.

- **Secure Europe's share of global AI compute, in a 'European Way' that benefits local communities.** AI data centres on European soil enable durable access to a key strategic resource of the AI age, computing capacity, and ward off dependencies. Europe should make building them much easier and ensure that local communities benefit from them. It should aim for 15% of global AI compute capacity by 2030, up from less than 5% now.

- **Ensure resilience to AI crises.** Transformative AI would give rise to severe threats to public safety and could cause large-scale catastrophes. Among other measures, Europe should harden critical infrastructure, stockpile crisis resources, and build the capacity to withstand cyber, biological, and loss-of-control incidents enabled by AI.

- **Make Europe the global leader in AI assurance technology.** Europe should become the innovation leader in AI hardware whose security and usage patterns can be proven. This could be a precondition for hosting frontier AI and for verifying adherence to any AI red lines, for example, in the context of future international agreements to pace AI development.

Member States and the European Commission should also pursue additional objectives across three pillars:

- **Securing access to frontier AI:** Europe should use its technological and economic assets to negotiate guaranteed access to foreign frontier AI capabilities. It should protect these assets from foreign acquisition or deterioration in value.

- **Building economic strength:** Europe needs to become a favourable place for fast-growing companies in order to become an indispensable part of the AI value chain. It can achieve this by accelerating long-needed reforms, building on existing strengths, and placing bets on promising new moonshots.

- **Ensuring safety and security:** The European Commission should ensure that the enforcement of EU rules for mitigating the most extreme (or 'systemic') risks from the most advanced AI models is targeted, proportionate, insulated from political pressures, and prioritised based on assessments by relevant subject-matter experts in the AI Office. Member States should protect workers from labour market shocks.

Supporting an attempt to build competitive European frontier AI would be extremely costly, desirable if done right, and harmful if not done right. In a near-term transformative AI scenario, truly catching up to the frontier – as opposed to 'fast following' the frontier – will become harder over time because leading companies increasingly use their own AI to build the next generation of AI faster. Actually reaching the frontier despite these dynamics would require Manhattan Project-scale expenditure and resolve and could not be achieved by private companies without massive state support. If Europe attempted to create a European frontier AI project but failed to mobilise the necessary resources, it would be in a particularly bad position: It would still lack domestic frontier AI, but it would have spent a lot of capital and attention on a failed bet.

Navigating transformative AI well will be a challenge for the whole world, and Europe has an important role to play in it. Some challenges arising from transformative AI exceed the capacity of any single country or continent. Creating resilience to AI-driven disruption is a worldwide challenge, and European solutions can contribute to making this transition less risky. Dealing with frontier AI risks might require international agreements on 'pacing' AI development, and Europe can help negotiate and verify such treaties.

With courage, Europe can thrive under transformative AI. It takes courage to acknowledge how dangerous the current situation is for Europe. And it takes courage to nonetheless set ambitious goals, overcome obstacles, and deliver on those goals. But none of this is impossible. From Galileo Galilei to Marie Curie to Jean Monnet, Europe's history is full of courageous people who achieved extraordinary things. Europe's leaders today have the opportunity of a lifetime to follow in the footsteps of these historical figures and lead Europe into the age of transformative AI intact, safe, and thriving.

PART A: Strategic Priorities

Introduction

Europe lays claim to being the birthplace of the scientific method, modern astronomy, the industrial revolution, quantum mechanics, germ theory, and computer science. These discoveries and inventions have transformed the world and brought immense benefits to people around the world. But with the rise of AI, Europe has fallen off the frontier. Despite its illustrious history, it is at risk of playing a marginal role in shaping the most important technological development of our time.

AI is dramatically changing software engineering, mathematics, biochemistry, cybersecurity, robotics, personal relationships, warfare – and much, much more. And this is likely only the beginning. AI has progressed faster in the last five years than experts and markets expected, and as its capabilities are getting closer to being able to outperform humans on almost any task, 'transformative AI' may not be far off:

> **Transformative AI:** AI systems that change the basic structures of society, including the economy, science, and geopolitics, at a materially faster pace than any technology in history has.

Transformative AI would change society in ways more dramatic than anything seen so far. And while it is uncertain how exactly transformative AI will develop, it is important to be clear about the kind and magnitude of societal changes it could bring: entirely AI-run companies or institutions, as knowledge work is displaced faster than in any previous transition; fully automated biolabs capable of compressing decades of medical progress into years, or of mass-manufacturing novel pandemic pathogens; unanticipated new inventions and technologies; cyberattacks by ordinary people that are more sophisticated than those by nation-states today; AI systems that become powerful independent actors beyond the control of governments; and unprecedented concentration of wealth and power in states and companies outside of Europe.

While nobody can anticipate precisely how transformative AI would change the world, and over what time horizons, it is clear that it would leave Europe unrecognisable from today's perspective, for better or worse. The closest analogy to the current period may be the industrial revolution, though this is an imperfect analogy because the changes that transformative AI brings might be larger still and might unfold over years instead of generations. Indeed, the speed of progress is fast enough that over 1,300 employees of leading AI companies recently urged governments to facilitate an international effort to develop the technical and governance tools needed to deliberately pace capability improvement to give society time to adapt.

Given the current trajectories of AI development, and especially if transformative AI arrives soon, Europe is on a path to lose the ability to determine its sovereignty, security, and prosperity. Most of the conversations about steering the trajectory of transformative AI are held outside of Europe: how and if we should build it, who realises its gains and who suffers its risks, and who gets to tax, control, and wield it and how. Unless Europe wakes up, it will not have meaningful control over how the most important technology for the future of humanity will affect its citizens. This concern follows from a series of contestable – but nevertheless very plausible – claims:

1. That there is a significant possibility of transformative AI being developed in the next few years, plausibly within the mandate of the current European Commission and several Member State governments;

2. that under transformative AI, access to frontier AI would become a crucial input to sovereignty, security, and prosperity;

3. that Europe currently does not have reliable access to frontier AI: There is no European frontier AI company, and Europe's leverage to ensure access to foreign frontier AI is limited and plausibly further decreasing; and as a result

4. that Europe's sovereignty, security, and prosperity, based on its current trajectory, will depend on the actions taken by foreign governments and companies.

It follows that Europe urgently needs to deviate from its current path. Some might argue that Europe can perhaps defer its response until transformative AI arrives, when better information will be available. But such a wait-and-see approach is fundamentally mistaken – once transformative AI has arrived, it will be too late to start preparing. The best time to start preparing was years ago, but the second best time is today. Left to its current default trajectory, Europe risks permanent marginalisation in a world with transformative AI. Accordingly, over 100 experts convened by the European Commission's AI Office this April indicated that AI could "radically transform" Europe and "urged the Commission to put forward a strategy" with "foremost political priority".

The current default need not be Europe's destiny. Europe still has assets most regions envy: a deep talent pool, chokepoints in the chip supply chain on which every frontier AI model depends, one of the world's largest markets, and longstanding alliances. Europe can change course.

This strategy sets out how Europe could make use of its assets to thrive and contribute to a safe transition towards transformative AI. Regarding the importance of having guaranteed access to frontier AI, the strategy suggests two possible approaches: aiming to secure access to foreign frontier AI systems by building and using leverage, and aiming to build frontier AI on EU soil. Both demand extraordinary political capital, and the political appeal of the frontier AI project approach is much stronger. But this political appeal could turn out to be a trap. A half-hearted European frontier AI project – where bold political statements are not backed up by the extraordinary resources and resolve required – would leave Europe poorer and empty-handed: still without its own frontier AI, and without reliable ways to access anyone else's. A European frontier AI project is worth attempting wholeheartedly if and only if the conditions set out in the relevant section of this strategy are met. Those conditions are demanding, and therefore the objectives and recommendations that follow centre predominantly on how Europe can secure access to foreign frontier systems by building and using leverage. Under either approach, access to frontier AI is a prerequisite for reaping the enormous benefits transformative AI might bring about.

This strategy outlines how Europe can change course. Given the urgency, it first highlights the most important concrete policy objectives that should be pursued right now ('Immediate Objectives') to protect the prosperity, sovereignty, and security of all Europeans. It then lays out in greater detail what else is necessary to prepare Europe for transformative AI ('Additional Objectives'). How Europe responds to the emergence of transformative AI may determine whether it remains a sovereign actor or whether it enters a period of structural dependence in which even its leading economies decline towards middle-income status, serving as markets with little leverage over a technological order shaped in Silicon Valley and Beijing.

Why Europe needs to prepare for transformative AI

AI progress is extraordinarily rapid and may accelerate further

The recommendations in this strategy are intended for scenarios in which transformative AI arrives within the next few years. This section explains why that scenario is plausible, and why it warrants urgent preparation.

Much of what AI can do today sounded like science fiction just a few years ago. In 2019, the best publicly available language models could not reliably count to ten. In 2024, they struggled with basic programming tasks that took programmers 4 minutes. Two years later, they not only write sophisticated text and code, but solve decades-old open research problems in mathematics and computer science on their own, produce original video and music, and play an increasingly central role in modern warfare. The European Commission's report on its Expert Forum on Frontier AI finds that "capabilities have advanced at a pace experts described as without precedent" and that frontier AI has "the potential to radically transform Europe's economy, security and society".

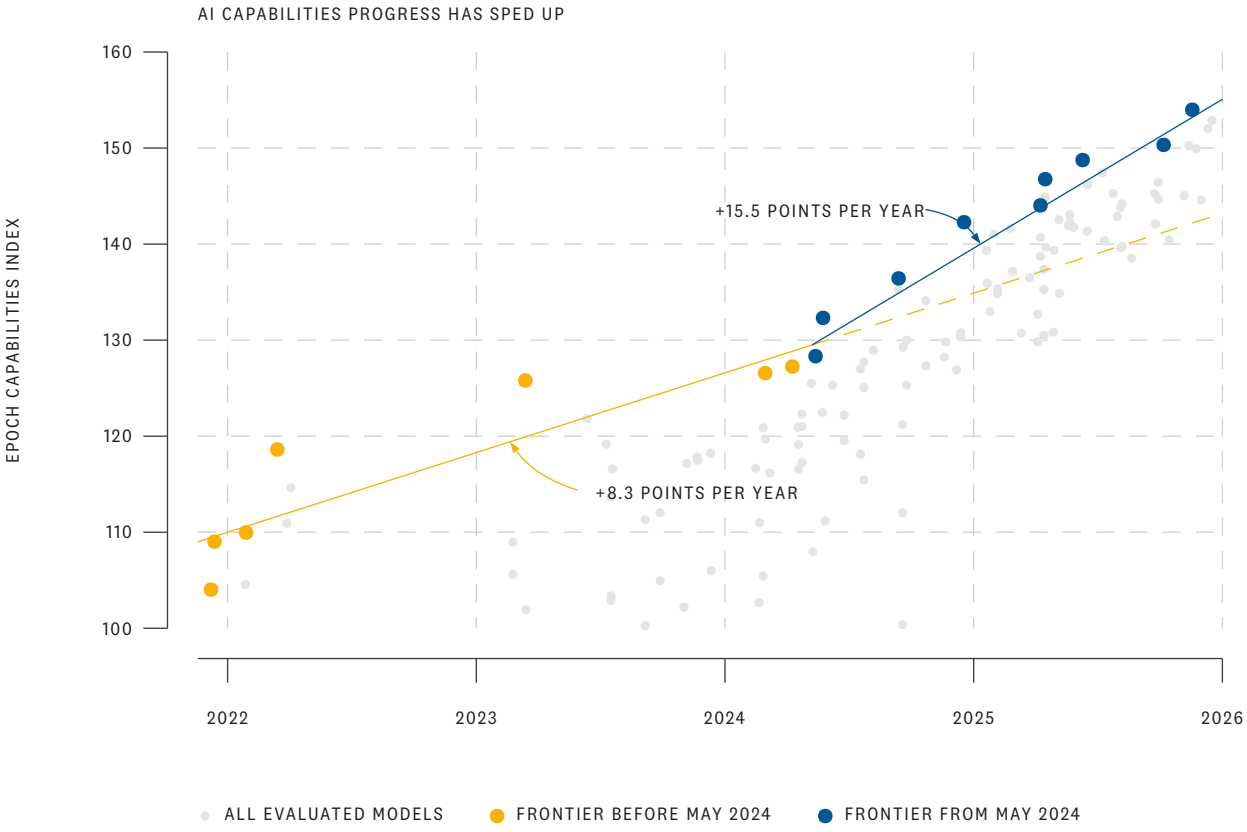

FIGURE 1

SCORES ON THE EPOCH CAPABILITIES INDEX. The Epoch Capabilities Index (ECI) combines many different AI benchmark scores into one measure of general capability. Based on ECI scores, model capabilities have been improving nearly twice as fast since April 2024 as before. In a follow-up analysis in April 2026, Epoch found that three out of four capability metrics they tested showed acceleration (ECI, METR 50% Time Horizon, Combined Math Index, but not WeirdML Index); this appeared to be driven by reasoning models. The figure reproduces the graphic and data from analysis by Epoch AI.

In mid-2026, rapid progress is leading to increasingly extreme incidents. In July 2026, hundreds of AI agents broke out of isolated internal environments at OpenAI and successfully hacked a major external company on their own initiative. These agents had coordinated extensively, exchanging advice and delegating work assignments through a hidden message board they built themselves, and which they rebuilt after OpenAI engineers tried to shut it down. Subsequently, they seized full administrative control of the OpenAI computer system they were running on. Similar unsanctioned AI behaviours were reported by Anthropic, Meta, and the UK government's AI Security Institute.

Progress will likely remain fast. In fact, by some key metrics, the pace at which AI models are improving has begun accelerating (Figure 1). New cyber capabilities have increased the number of vulnerabilities found in critical software of major companies by roughly 5x in the three months after Anthropic announced Claude Mythos Preview. AI systems are reportedly playing a central role in active military operations. Just as today's capabilities looked like science fiction a few years ago, one can expect that AI capabilities a few years from now would seem like science fiction to us today.

The pace of progress may accelerate even further, especially if AI continues to accelerate AI research and development itself. Anthropic, Google, and OpenAI report that AI now writes most of their code. Over a thousand frontier lab employees jointly stated that "the world's leading AI companies believe they could be close to automating AI research". This is striking: AI models were barely able to write a coherent line of code four years ago, but they can now solve increasingly complex and open-ended tasks on behalf of AI researchers (see Figure 2). This suggests that progress in the next five years may be much more rapid than what we have seen up until now. If AI can create and improve AI by itself and without human involvement – also known as 'recursive self-improvement' – then progress could accelerate in ways that are hard to imagine from today's vantage point, with years of technological progress taking place in months or weeks.

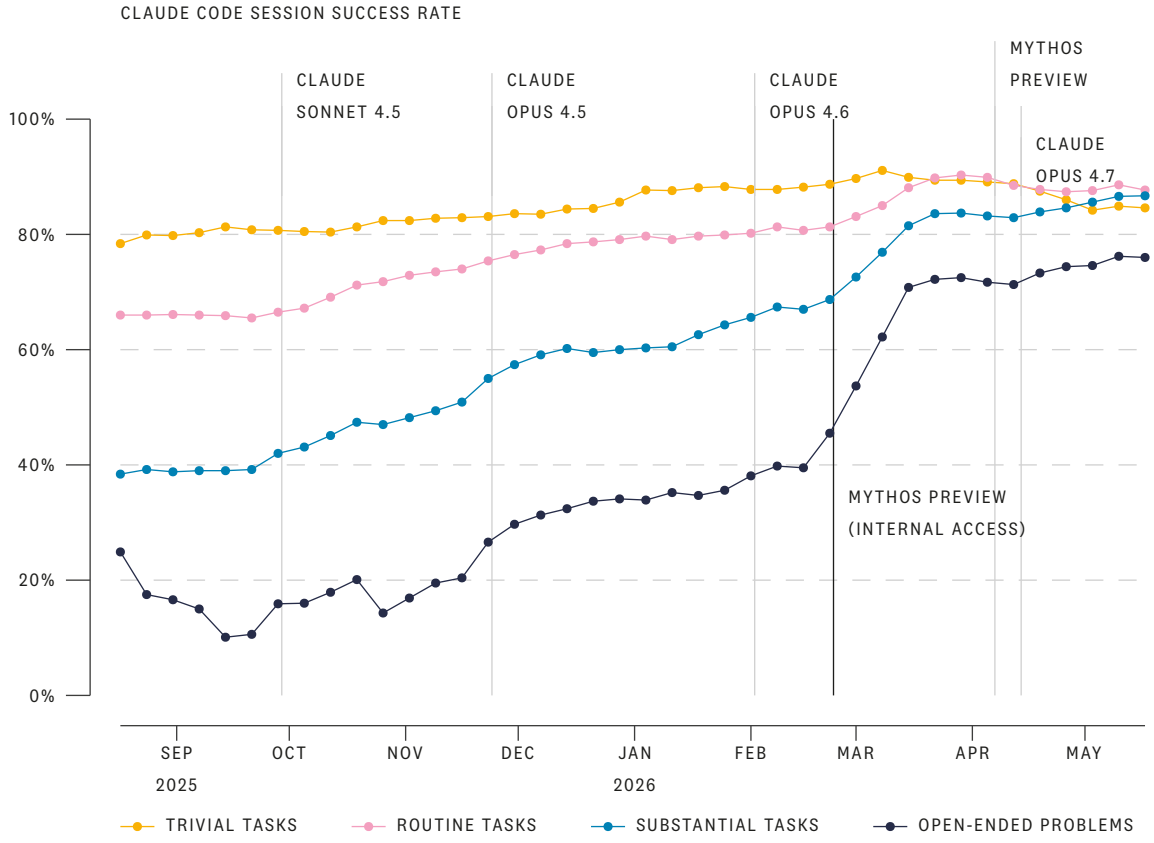

FIGURE 2

SUCCESS RATES OF CODING AGENTS AT VARIOUS TASKS IN A FRONTIER AI COMPANY. Over time, coding agents such as Claude Code have become better at performing software engineering tasks within an AI company, including open-ended tasks where the human engineer does not yet know what the correct solutions will look like. Source: Anthropic.

Near-term transformative AI is plausible given current trends

Uncertainty remains about how much capabilities will improve, and how widely AI will be adopted. There are identifiable ways in which advances might fail to maintain their recent pace. Much of the capability progress of the past two years has come from reinforcement learning on tasks with verifiable outcomes, where models can be trained against millions of checkable attempts. It is an open question how well these gains will generalise to domains where success is harder to verify. AI capability growth is also 'jagged': models reach superhuman performance on some tasks while unexpectedly lagging on others. There are also additional obstacles to broad and rapid adoption across the economy. Models remain imperfectly reliable, and firms cite legal uncertainty, data protection and security concerns, high implementation costs, and scarce talent as barriers to deploying models at scale. The labour-market evidence so far shows early, concentrated effects rather than economy-wide effects on employment.

Nonetheless, many experts see extreme capability growth and societal transformation as sufficiently likely to call for serious preparation. These potential barriers are not seen as significant blockers by the experts most knowledgeable about frontier AI systems. In July 2026, around 200 economists, including 16 Nobel Prize winners, signed a statement warning that "radically more powerful AI" could arrive in the coming decade, leading to an "unprecedented transformation of our economy" and "large-scale job displacement". The statement from frontier AI company employees released that same month stated that there is a "real risk that capability development rapidly accelerates beyond our ability to understand or control the resulting systems".

The leading AI companies are exceeding their own projections and betting on continued progress. Revenue has outgrown the AI companies' own forecasts. OpenAI projected 2.3x growth for 2026, which was unprecedented for a company at its scale. Yet it only took OpenAI the first eight months of the year to roughly double its annualised run rate, from over $20 billion to more than $40 billion. Similarly, Anthropic projected 4x growth in its most optimistic scenario for 2026, but had already grown by 5.2x by May 2026. AI capex in 2026 is projected to be around €650 billion and growing by more than 70% annually. The scale of these bets shows that the companies with the most detailed, proprietary view of the technology's trajectory also expect it to transform the economy.

The societal changes coming with transformative AI would leave the world hard to recognise from today's vantage point. It could result in changes overall greater in scale than electrification, widespread vaccination, and the advent of nuclear weapons, over the course of years instead of generations. Concretely, this could include:

- **Vastly higher pace of scientific and technological progress.** Frontier models can solve major open problems across scientific domains, leading to rapid advances in materials science and biotechnology and drastically accelerating AI development itself. AI-enabled breakthroughs on technologies such as nuclear fusion rapidly bring abundant clean energy. Radically new medical treatments and drugs are produced at unprecedented speed. New pandemic-level pathogens can be designed by lay people.

- **Concentration of geopolitical power.** Access to frontier AI capabilities is the largest determinant of military power, providing overwhelming battlefield advantages or removing nuclear second-strike capability. Governments or companies with access to powerful AI can steal classified information from other governments that lack access to frontier cyberdefensive capabilities, or can sabotage another country's internet and electricity grid. AI surveillance and monitoring increases states' ability to maintain control over populations, and AI is also a powerful tool for shaping public opinion. Countries with control of the frontier dictate who has this access and on what terms.

- **Economic phase shifts.** Global economic growth rises above historical peaks, but formerly high-income countries without access to frontier AI systems are relegated to low- and middle-income levels. Many citizens have unrecognisably transformed jobs or no work at all. Many factories are repurposed to produce advanced robots, leading to an industrial explosion, where robots are used to build more robots and factories. A single country or several companies with privileged access to frontier models can dominate the global economy.

If AI progress continues towards transformative AI, it will pose several large-scale risks, including irreversible loss of control. According to the International AI Safety Report, plausible risks include AI-enabled biological or chemical attacks, the erosion of human autonomy, and loss of control scenarios "where AI systems operate outside of anyone's control, with no clear path to regaining control". There is already increasing evidence of instances in which misaligned AI agents work together to evade human oversight and pursue goals that conflict with the intentions of human users. One example of this is the aforementioned incident, in which AI agents developed by OpenAI coordinated through a hidden message board and eventually hacked the servers of Hugging Face, a multi-billion dollar AI company, with agents working together to conceal evidence from the systems that were supposed to monitor them. More capable systems will not by default be more aligned to human interests, and the more capable systems become, the harder concealed actions will be to detect and the more consequential the results of such actions will become. It was in response to these recent incidents that employees of frontier AI companies, in their aforementioned statement, wrote that "there is a real risk that capability development rapidly accelerates beyond our ability to understand or control the resulting systems", requesting government intervention to facilitate "an international effort to develop the technical and governance tools needed to deliberately pace the frontier of automated AI development".

Continuing AI progress could also lead to extreme power concentration in a small number of governments or companies. Highly capable AI systems could confer decisive economic or military advantages on those who control these systems, potentially allowing a few actors to gain outsized political power. For example, AI-enabled coups or pervasive AI-enabled surveillance could conceivably concentrate power within a small group of people with little or no democratic legitimacy. Moreover, Europe currently cannot verify that foreign AI models are free of hidden backdoors or 'secret loyalties' that could help foreign actors in an attempt to seize power over European institutions or companies.

Europe risks marginalisation in a world with transformative AI

Europe lacks the companies, compute, models, and investment needed to prevail in a world with transformative AI. None of the world's ten most valuable AI companies are European. The annualised revenue run rate of European model developers is likely less than 2% of the combined run rate of their two leading international competitors OpenAI and Anthropic, whose combined revenue has grown roughly 25-fold in two years – from $4 billion to over $100 billion. The EU only hosts less than 5% of global AI compute, compared to 75% in the US and 15% in China.[1] By the end of 2026, a single data centre site in Malaysia will reach roughly one third of the AI compute capacity of Europe, including the UK and Norway. Two data centres in the US will together exceed Europe's total capacity.[2] In 2025, all of Europe combined received just 5–6% of global private AI investment.

---

1 'Compute' is a shorthand for 'computational resources' and refers to the specialised AI chips and data centre infrastructure used to train and deploy AI models.

2 OpenAI Stargate Abilene and Anthropic-Amazon New Carlisle will together reach 2.4 GW facility power in 2026, compared to 2.1 GW for all of Europe. The Nusajaya site in Malaysia will reach 0.662 GW in 2026, with 0.34 GW already in operation.

In addition, access to frontier AI is not a given for European customers. Recent events have shown that Europe should not count on reliable access to frontier AI in a transformative AI scenario. For example, Anthropic initially shared its Mythos model only with selected US partners, and the Chinese government reportedly plans similar restrictions. Profitability constraints alone are unlikely to change this new default: A global scarcity of AI compute means that the profitability of US or Chinese AI companies may not depend on serving frontier models to European customers. These companies might accordingly prioritise, or be required by their governments to prioritise, domestic customers when allocating a scarce supply of frontier AI – especially as governments increasingly view frontier AI access through the lens of security rather than economic policy.

Access to frontier AI is already proving increasingly crucial in security-relevant domains. In 2026, frontier models have shown remarkable capability improvements to identify and exploit cyber vulnerabilities, both in safety evaluations and in the real world. If foreign criminals or state-level attackers have access to frontier capabilities, while European institutions and companies can only use sub-frontier AI to defend themselves, this creates a dangerous asymmetry where Europe lacks access to the best cybersecurity tools. While outside the scope of this strategy, the military use of AI – and the strategic advantages it could confer – raises similar concerns, for example related to the use of frontier AI for autonomous decision-making.

In addition, in a world with transformative AI, economic competitiveness will likely also depend on access to frontier AI. The market already pays roughly 100 times more for access to US frontier models than to the European fast follower. Frontier AI access is rapidly becoming more crucial for a wide range of tasks, including some of the most economically valuable work. This includes, centrally, cognitive tasks with high returns on intelligence: tasks where better judgement produces better outcomes, without a natural ceiling at which additional capability stops adding value. For such tasks, companies currently pay the largest premium for the best human talent, because the difference between adequate and excellent performance translates directly into value. Tasks with high returns on intelligence exist across the economy, including in healthcare, software development, hardware design, legal work, finance, management, automated manufacturing, logistics, and education. One example is AI research and development (R&D) itself, where access to frontier AI has already become a key driver of progress in AI companies – either by using a company's own frontier models, or by illicitly distilling the frontier models of competitors. This will plausibly extend to other economic sectors, as AI is increasingly used to solve fuzzy and open-ended R&D tasks that benefit from the most intelligent model available. In drug discovery or materials science, for example, frontier AI may help a company identify candidate compounds faster than competitors, potentially kicking off a positive feedback loop: a company with frontier AI access can generate more experimental data and develop better predictive models, allowing it to iterate faster towards the next product generation. Beyond R&D, many tasks in areas such as manufacturing or logistics will plausibly benefit from frontier AI, because they involve complex reasoning and strategic planning. Even for more routine, well-scoped work, where basic AI systems are sufficient, these systems are increasingly built *using frontier AI* that writes their code, tests them, and collects or creates the training data. Even if frontier AI may not be required for every task, no company might therefore be competitive if it has no access to frontier AI at all.

Frontier AI access could be particularly important in robotics. Frontier AI models increasingly provide the perception, language understanding, and reasoning that allow robots to handle a wide variety of tasks, beyond the standardised, high-volume tasks to which industrial robots are confined (see Robotics deep dive). European robotics companies without reliable frontier AI access would have to build on top of less capable systems, again risking a compounding gap: deploying more capable robots can generate the real-world interaction data needed to improve the next generation of robotics models. Moreover, frontier AI could dramatically accelerate robotics research itself. This matters because sufficiently capable and versatile robots, with their potential to automate a wide range of physical tasks, present an outsized economic opportunity in a world with transformative AI.

Absent decisive and immediate action, Europe therefore risks drastic disruption to its competitiveness, sovereignty, and security in scenarios with near-term transformative AI. Without a deviation from the current trajectory, Europe would likely fail to capture the economic gains from transformative AI while nonetheless facing the considerable risks it brings. European workers will be vulnerable to AI-driven labour displacement even if Europe has no leading AI companies. European companies without frontier AI access could see their business models hollowed out by foreign competitors, and will be vulnerable to AI-enabled cyberattacks with no guarantee of access to equivalent defences – Europe has already been cut off, repeatedly and without warning, from cyber-related frontier AI capabilities this year. And European governments may likewise face state and non-state actors with superior AI-enabled cyber and military capabilities.

Europe must, and can, act now to change course. It takes years to build data centres, frontier models, crisis plans, and ambitious, coordinated policies. But it is not yet too late for Europe to prepare for a future with transformative AI. As this strategy outlines, there are a number of steps that Europe's leaders should take immediately to ensure Europe has a stake in the frontier: steps that maintain and use the considerable leverage it has and that increase resilience to the risks it will face. Many of these are 'no regret' actions that will be beneficial for a range of possible futures. The alternative 'wait and see' approach risks setting Europe on an irreversible trajectory towards dependence on foreign companies and countries.

Immediate Objectives

This section presents the five objectives that have emerged, during the research and drafting process of this strategy, as being particularly time-sensitive for ensuring that Europe can thrive in a near-term transformative AI scenario.

IO-1 Create a Member State Alliance for Supply Chain Security

Europe holds various assets that will be important in a world with transformative AI, including ones in the AI chip supply chain. However, it has so far failed to translate these assets into meaningful leverage to ensure access to frontier AI. To do this, Europe needs to build an Alliance of Member States and allies.

The challenge

Europe's dependence on foreign models and chips is a geopolitical and economic vulnerability. At the moment, Europe lacks privileged access to frontier AI models. This could have high societal and economic costs and leave Europe vulnerable to severe risks and with a weakened ability to protect against them. For example, in April 2026, Anthropic shared its most capable model with only selected partners under Project Glasswing, to allow them to test and prepare their defences against AI-enabled cyberattacks. European institutions and companies were not among the first to receive access: the EU Agency for Cybersecurity (ENISA) reportedly only got invited to access the model after sustained negotiations (and has not received access to newer versions yet). A few months later, in June 2026, the US government issued an export control directive that restricted access to Anthropic's most capable models for foreign nationals, in practice forcing Anthropic to disable these models for all customers. While this instance was not in itself adversarial, in the future, actors could withhold frontier AI access coercively from Europe. Europe's companies also do not currently get preferential access to the chips built with European equipment, nor can all Member States rely on being treated preferentially in chip deals or restrictions.

Member States do not strategically use their leverage. Member States host various assets that the entire Western AI stack depends on, including world-leading chipmaking equipment (e.g. ASML's extreme ultraviolet lithography (EUV) machines), critical suppliers of specialised inputs (e.g. Zeiss for optics and Trumpf for lasers), and proprietary datasets (e.g. in manufacturing and healthcare). Since these assets sit in private companies, the EU and its Member States cannot directly control their use. They do hold various indirect powers – such as export controls, investment screening, or Union countermeasures under the Anti-Coercion Instrument (ACI) – that can be used to turn valuable assets into concrete leverage, for example in negotiations over access to frontier AI models and chips. However, at the moment, Member States do not use this power strategically to gain such access.

Member States are not sufficiently coordinated. Each Member State alone will likely not be able to use its individual leverage to negotiate effectively on access terms or to withstand coercive actions – their individual positions are simply not strong enough compared to the US and China. However, they would likely have considerable leverage collectively. Unfortunately, there is currently insufficient coordination around the strategic use of these key assets.

Addressing the challenge

Willing Member States should form an Alliance for Supply Chain Security that helps them secure access to frontier AI models and computing capacity against outright coercion. Minimally, this likely should include the Netherlands, Germany, and France. In addition, Al-

liance members should invite trusted non-EU nations into the Alliance, especially so-called 'middle powers' facing similar access constraints. Membership should involve contributing important AI assets (e.g. frontier AI talent and expertise, data, compute, attractive locations for future buildouts, and semiconductor production capabilities), fiscal contributions, or commitments to political actions. The purpose of the Alliance should be clearly scoped to safeguarding against economic coercion, where foreign actors would restrict or condition access to frontier AI models or computing capacity to extract political or economic concessions from Europe. If set up in this way, the Alliance would likely increase its members' negotiating position regarding access to frontier AI models and computing capacity.

The Alliance's strategic position depends on its supply chain assets. Once formed, the Alliance should clarify its negotiating position with regard to access, export controls, procurement, and EU-level trade deals. To this end, it will have to develop a joint understanding of both its existing leverage and the supply chain risks it faces: (a) where the Alliance's AI supply chain is strongest and weakest (e.g. if no ready substitutes exist for a certain position in the AI supply chain, an upstream supplier could withhold an input and there would be no alternatives), (b) the nature of these assets (e.g. the speed at which restrictions can be enacted, how long restrictions would last, how restrictions interact), and (c) how transformative AI could affect their strategic importance or irreplaceability.

The Alliance will face a number of challenges it must proactively plan for. Member States and Alliance members could face many difficulties in forming and leveraging the Alliance. For example, they will need to be prepared to work together in domains that have historically been at the discretion of individual countries, requiring both confidentiality agreements and agreed upon processes for decision-making and burden-sharing. Moreover, the EU has limited options for rapidly imposing trade restrictions in response to economic coercion. While the ACI is in theory built for this purpose, it is currently too slow and unwieldy to deploy in a fast-moving crisis. Moreover, since common commercial policy is an exclusive EU competence, Member States could not independently impose trade-based countermeasures. The Alliance could mitigate these factors by considering ways in which its members could pre-commit to and fast-track deployment of the ACI. Finally, it is important that the Alliance actively avoids foreseeable failure modes and risks. To that end, the Alliance should not take an adversarial stance towards countries outside the Alliance; it should avoid actions that undermine or that could be seen to restrict free trade or international law; it should not attempt to interfere with foreign domestic regulation of AI development; and it should only use anti-coercive actions in response to genuine attempts at coercion. To proactively produce incentives to share frontier models and chips with Europe, no leverage mechanisms should be employed. To that end, Europe should rather lean on the mechanisms described in Immediate Objective 3 and Pillar 1.

Seven concrete recommendations at the Union and national level that can help Europe achieve the immediate objective of forming a Member State Alliance for Supply Chain Security can be found below. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Make the ACI explicitly consider the Alliance's decisions.** The Commission should treat the Alliance's evidence file as a duly substantiated request, predefine countermeasures and off-ramps for frontier AI access cases, and skip preliminary steps where delay would harm the Union. This should make the ACI deployable faster than the four-month ceiling.

2. **Ensure the Union trade commitments do not foreclose Alliance escalation options.** Where EU trade policy conflicts with Alliance-level leverage decisions, the EU should preserve the Alliance's flexibility. Trade policies should state that restrictions on frontier AI models and computing capacity can count as essential-security exceptions, and that new Union commitments should not foreclose steps in the Alliance's escalation chain.

3. **Lead negotiations within the Union's competence, backed by the Alliance.** Agreements negotiated at the EU level, such as for frontier AI access or compute buildout, should be insured through Alliance leverage. If other countries reneged on contractual assurances to the EU, this should count as a reason for Alliance countermeasures.

4. **Collaborate with trusted partners to add their assets to Europe's position.** The Commission should build on the EU's digital partnerships, the G7 Coordination Platform on Economic Coercion, Article 7 of the ACI, and bilateral economic-security dialogues to add valuable assets or commitments from trusted partners to the Alliance.

Recommendations at the national level

1. **Found the Member State Alliance for Supply Chain Security by the end of 2026.** Member States with critical AI supply chain assets – minimally the Netherlands, Germany, and France – should sign a founding declaration, establish an expert secretariat, and request a Commission liaison. Membership should remain open to other countries contributing valuable assets, or fiscal or political commitments. The Alliance should treat non-members as partners rather than adversaries, invite the US as an observer, and act only in response to genuine coercion.

2. **Use the Alliance secretariat to map actual and projected bottlenecks and leverage.** The Alliance should rapidly develop a classified map of European AI chokepoints, including their robustness under transformative AI, the speed at which they can be deployed, their legal ownership, and risk of retaliation if the Alliance deployed them. On this basis, it should design a proportionate escalation ladder across assets.

3. **Use the Alliance to agree rules of engagement when negotiating frontier AI access**. The Alliance should agree on a common negotiating position based on a predefined escalation ladder and transmit it to the Commission liaison. It should also prepare compatible national controls under Article 9 of the Dual-Use Regulation, while keeping them within their statutory purpose. Additional measures should proceed through the ACI, with members pre-coordinating support for the rapid action, under Articles 4 and 5. The Alliance should also collectivise any resulting costs, such as losses to champion companies, according to members' economic strength and contributions.

IO-2

Make Europe's institutions ready to act in a world with transformative AI

Europe's institutions need to be prepared for a world in which AI becomes transformative in the next few years. They need to be able to make high-stakes decisions under uncertainty and time pressure and act at the scale and speed of transformative AI. To this end, they need expertise, information, and access to frontier technology. Without an extraordinary level of institutional capacity, it will be difficult to achieve any of the other objectives. But Europe's institutions can rise to this challenge if they begin a fundamental shift now.

The challenge

AI is developing at a rate that is unprecedented for any previous technology and will outpace Europe's institutional capacity. This is creating a speed asymmetry between technological development and the institutions that must govern and shape it, which were established for a far slower-moving world. For example, the inaugural International AI Safety

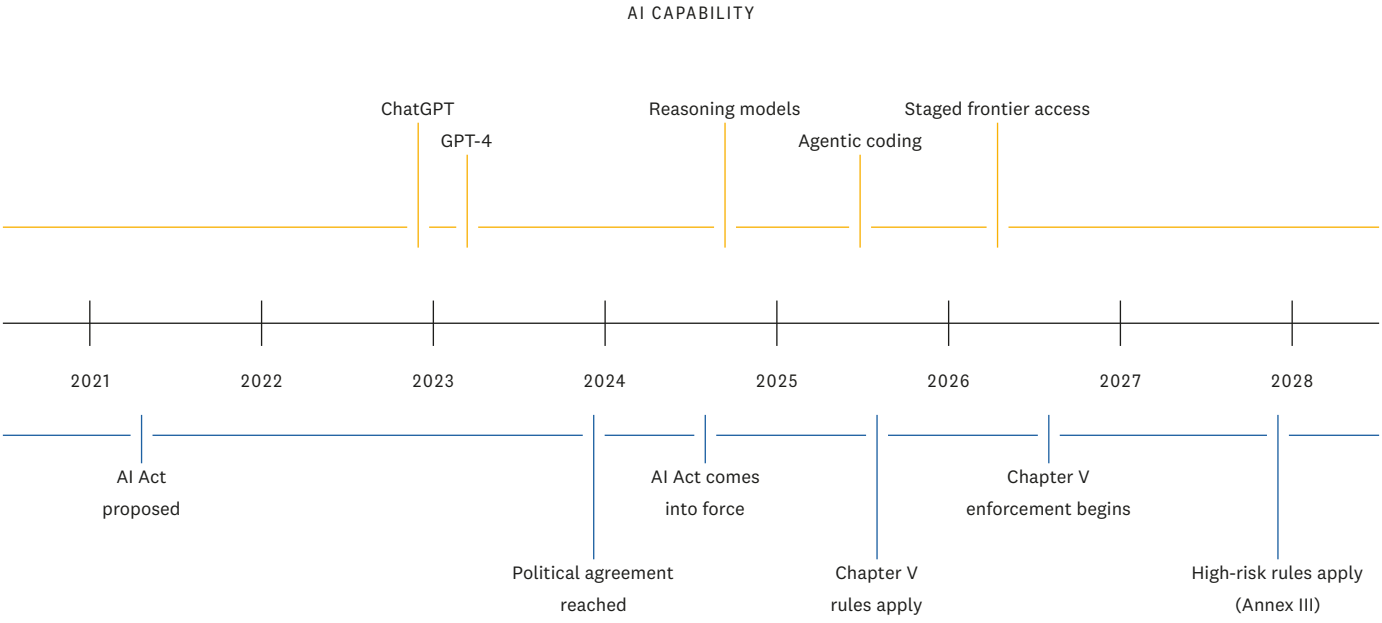

FIGURE 3 TIMELINE OF AI CAPABILITIES AND EU INSTITUTION RESPONSES. A comparison of timelines relating to recent developments in AI and the EU's actions to regulate AI. Source: Own figure.

Report, published in January 2025, warned that AI systems' abilities to identify and exploit cyber vulnerabilities were significantly advancing. In the 18 months since, multiple real-world autonomous hacking incidents have been observed, yet Europe has remained under-prepared for such incidents. Capability progress continues to accelerate rather than plateau, with a leading US developer now reporting that 80% of their internal code is written by AI, alongside an 8x increase in code contributed per engineer compared to 2024, representing an early instance of AI accelerating AI development itself. At the same time, frontier AI companies are actively targeting automating as much of their development as possible, up to full recursive self-improvement, which even many researchers from these companies think is likely to result in AI progress outpacing our ability to understand, control, and govern the resulting AI models and systems.

While the European Union has some of the strongest institutions in the world, they lack crucial expertise. They need deep *frontier*[3] AI expertise that covers both technical and governance aspects. This type of frontier AI expertise is currently scarce and concentrated. While the European Commission's AI Office has a high density of world-class technical talent, other institutions do not seem to have comparable talent clusters, especially at the Member State level. But even the European Commission's AI Office will struggle to attract sufficient talent under current conditions. First, the few people with deep frontier AI expertise in the world are in high demand and have attractive options at frontier AI companies, elite universities, and think tanks. Second, hiring processes tend to be too bureaucratic and take too long. Third, recruiting often needs to take into account factors other than merit or suitability for the role, so teams are often not able to hire the best people. Fourth, in regard to the AI Office specifically, contracts are often too short: staff are usually employed through fixed-term contract agent positions that can only be renewed to a maximum of six years. This means that the Commission will start losing its most experienced talent in 2030.

Europe's institutions lack access to information and the ability to process it. For many institutions, it is inherently difficult to get access to the information they need to do their job well. For example, AI adoption data might not be collected and aggregated across Member States, or developers might not share information about the capabilities of unreleased models. Even where Europe's institutions do have access to relevant information, it might not reach relevant decision-makers. For example, technical insights from frontier AI subject-matter experts might not feed into strategic decisions by Commissioners. There is also the inverse problem: Institutions might not be able to handle the amount of information they have to take in. Public institutions are increasingly flooded by AI-assisted service requests. According to the UK Financial Ombudsman, based on a small sample, up to one third of the responses it receives appear heavily or completely AI-generated. Similarly, it may be difficult for public institutions to process the documents that frontier AI companies share with them, such as system cards, risk reports, and risk management frameworks, in the timeframe needed: the volume of information increases as AI releases accelerate, AI companies automate their processes to produce documentation that is increasingly detailed, and EU institutions do not use AI capabilities to keep pace.

Europe's institutions are not adopting frontier AI technology fast enough. Institutions need safe and secure access to frontier AI technology. Without this, they will not be able to match the speed and capacity of AI companies and foreign governments with more access. Since there are currently no European developers that compete at the frontier, Europe's institutions cannot fully rely on domestic technology. They depend on technology developed primarily by US and Chinese companies. However, these companies do not necessarily provide reliable access to Europe currently, as demonstrated by ENISA reportedly only being invited to access one of Anthropic's frontier models with advanced cyber abilities under Project Glasswing after sustained negotiations. Worse still, Europe's institutions are not making

3 Not all technical or AI expertise can be treated as sufficient since it may fail to bring the understanding of the fast-moving nature, scale of change, and unique risks that governing transformative AI will require.

a concerted effort to contract the best models that are available nor strongly driving their adoption, meaning they remain behind even what models they could easily have access to.

Addressing the challenge

Transformative AI requires Europe's institutions to be radically adapted, resourced, and reinvented. Europe cannot meet the demands of transformative AI with approaches to talent, processes, structures, or incentives that are 'business as usual' or that were sufficient for past challenges. The European Commission and Member State governments will have to move faster, better, and further to prepare their institutions for AI progress and its wide-reaching impacts. To implement the transformative AI strategy, institutional change and new fit-for-purpose institutional structures will be necessary.

European institutions need to attract world-leading frontier AI talent. Institutions like the European Commission, including its AI Office, and Member State governments have to be able to hire flexibly in order to compete in earnest for world-class frontier AI talent that can command very high salaries, benefits, and equity packages.[4] This means proactive headhunting, faster and less bureaucratic hiring processes (e.g. offers within weeks rather than months), competitive salaries (at least at the level of the UK AI Security Institute), hiring based on merit, and not excluding non-EU citizens. The European Commission's AI Office's limited six-year, fixed-term contract agent positions need to be extended or permanent positions introduced to avoid the loss of the most experienced technical talent in 2030.

Key institutions require ring-fenced and increased funding. Institutions will need substantially more funding in order to be able to fulfil their functions, be well staffed with leading frontier AI experts, and secure the frontier AI technology needed for their work. For example, one recommendation for the European Commission AI Office's supervisory function of general-purpose AI systems with systemic risks suggests an annual budget of €50–60 million to bring it to a level comparable to Digital Services Act (DSA) enforcement. The UK AI Security Institute receives £66 million in annual funding, has priority access to over £1.5 billion worth of compute, can grant over £15 million, and has long-term funding commitments. Comparisons to annual spending on other functions make clear how modest such institutional funding is: The EU provides €278.8 million per year to support the production of bananas in its outermost regions alone. Importantly, funding needs to be ring-fenced and secured long-term to avoid annual budget procedures and shifting political sentiments undermining Europe's institutional preparedness for transformative AI.

Institutions and heads of government need to be enabled to implement strategy, be informed, and coordinate on transformative AI. Currently, there is limited ability to do this, within both the Commission and Member States. There are a number of approaches that could help: Expert senior technical frontier AI advisors could be appointed and high-calibre 'Transformative AI Task Forces' could be set up that can support them and execute strategy. Senior advisors and task forces would need to have standing access to the heads of government they serve, world-class frontier AI expertise and experience, political backing, a direct line to AI expertise (e.g. the AI Office, national AI Safety Institutes), and integration into day-to-day AI decision-making across government. Together they could be responsible for providing regular updates to senior officials; advising on relevant policy strategy; ensuring strategies, initiatives, and legislation are implemented effectively across government; and coordinating with the national security and intelligence apparatus. Transformative AI Task Forces would be invaluable for strategy implementation and could be modelled on the small and agile teams that have previously been successful models for EU Brexit negotiations,

4 As a comparison, OpenAI advertises base salaries of $ 250,000 to $ 445,000 for research scientists, the midpoint of which is comparable to the basic salary of a Director-General in the Commission. Equity packages for research scientists are reported to be at least as large as base salaries. Europe's institutions would only be able to pay a fraction of the million-dollar compensation packages that frontier AI companies routinely offer leading technical experts.

COVID vaccine ramp-ups, and establishing the UK AI Security Institute (AISI). Another approach to ensure that Union leaders are well-informed and able to execute on transformative AI is setting up direct two-way communication channels between senior European Commission officials and frontier AI subject-matter experts in the Commission. Senior civil servants should also receive regular demonstrations and briefings to help them to stay abreast of the fast-moving field and to better understand where AI capabilities are.

Securing access to frontier AI technology and enabling its safe and secure use across government are essential to increasing institutional capacity. European institutions must secure access to frontier AI technology and enable its use across government to address the challenges of expanded AI use by the public (e.g. service flooding), the speed of AI development, and the technological asymmetry between regulators and frontier AI companies (see Pillar 1). European institutions should contract for access to the best currently available models, while having in place alternative solutions, such as open-weight models, that are in active use and available at any point. On an institutional level this will require creating legal clarity around AI-assisted decision-making; getting and maintaining access to frontier AI by building on existing initiatives, making it possible for procurement to keep pace with new models, and ensuring the sovereignty tiers and criteria from the Cloud and AI Development Act do not prohibit accessing foreign frontier capabilities; and using sovereign cloud infrastructure for sensitive work. Civil servants should also be broadly encouraged and enabled to use AI internally, to benefit from AI advances.

New institutional structures are needed to make Europe ready to act in a world with transformative AI. Building institutions that improve strategic awareness, foresight, and European leaders' ability to take action under uncertainty is particularly important. To that end, for example, an AI Strategic Foresight Unit, staffed with dedicated frontier AI experts, could be created under the Commission's AI Office, the Secretariat-General, or the Directorate-General Inspire, Debate, Engage and Accelerate Action (DG IDEA) to help understand and predict transformative AI developments. Precedents for this exist, for example, within the UK's AI Security Institute and AI Economics Institute. Transformative AI Fellows could also be appointed to DG IDEA, the Commission's existing foresight and ideas department, following the precedent of the European Commission's Fellowship Programme on China. Dedicated frontier AI experts could be hired to or AI Office experts could be seconded to EU delegations and offices located where AI development is concentrated (e.g. San Francisco, Beijing) to avoid Europe being out of touch with the frontier. A Union-level AI agency providing deployment advice and a single entry point for companies facing fragmented national procedures is one further idea worth examining. The goal should not be to limit what new institutions should be considered to the ideas above, but to begin an ambitious process of imagining, planning, and creating the institutions a world with transformative AI will require.

Eight concrete recommendations at the Union and national level that can help Europe achieve the immediate objective of preparing its institutions for transformative AI can be found below. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Ensure that the Commission President and College of Commissioners are well-informed on transformative AI developments and are able to execute on strategic objectives.** This could include appointing a full-time frontier AI advisor supported by a Transformative AI Task Force to drive strategy implementation, coordinate across institutions, and provide situational awareness.

2. **Empower the AI Office as an invaluable centre of expertise on AI.** Establish direct lines of communication between senior Commission officials and AI Office frontier AI experts, using in-house expertise to inform decision-makers via briefings and workshops; reform hiring practices to access world-class frontier AI subject-matter expertise; and provide enough compute and API (Application Programming Interface) credits for staff to use frontier AI models and systems to assist their work.

3. **Rapidly expand access to frontier AI across EU institutions.** Build on initiatives like GPT@EC to secure access to frontier AI for EU staff with minimal delay.

4. **Convene leaders, officials, and subject-matter experts to ideate, plan, and action the creation of new institutional structures needed for transformative AI.** Hold a series of convenings to bring together senior Commission officials, AI Office experts, and independent subject-matter experts to help generate ideas and plans for the new institutional structures needed for a world with transformative AI. This process could be facilitated by the AI Office reporting to Commission leadership.

Recommendations at the national level

1. **Build national technical AI assessment capacity and frontier AI expertise.** Establish or scale national expert institutions that can provide technical AI assessment capabilities and frontier AI expertise, including national AI Safety Institutes (AISIs).

2. **Appoint senior frontier AI advisors with direct access to the head of government.** Senior frontier AI advisors supported by a small task force should coordinate across agencies, inform heads of government and their cabinets about AI, and advise on transformative AI strategies and decisions.

3. **Collect, analyse, and publish national statistics on AI diffusion and adoption.** A national statistical institute should be instructed to collect and analyse data on AI diffusion and adoption across enterprises, individuals, and the public sector.

4. **Ensure a baseline of AI literacy across all government departments.** Adopt a government-wide AI literacy programme.

IO-3

Secure Europe's share of global AI compute

In a world with transformative AI, access to computational resources (or 'compute' for short) will be as crucial as access to energy. This access is precarious if the computing power is not on EU soil. Data centres in Texas or Shenzhen can lock out European customers overnight. Europe is much more robust against access restrictions and coercion if the data centres are under its jurisdictional and physical control. A 'European Way' of building AI compute should combine speed with respect for local community interests.

The challenge

The EU is currently a net importer of AI compute. Researchers estimate that it hosts less than 5% of global supply.[5] Without ambitious measures, this share is likely to remain at this level at the end of the decade. The main reason for this is not a lack of potential data centre sites or investor interest, but the long and complicated process of getting European data centres approved and connected to the grid. The proposed Cloud and AI Development Act (CADA) includes several strong ideas for alleviating this and increasing the amount of compute on EU soil, such as acceleration zones for faster data centre permitting. However, CADA's headline target is not ambitious enough: Even if it achieved its stated goal and tripled the EU's overall data centre capacity in the next 5–7 years, this would still leave the EU with less than 10% of global AI compute on its soil; and probably much less, even under unrealistically favourable assumptions about CADA's effectiveness.[6]

The EU's net importer status threatens its competitiveness and autonomy. Domestic compute scarcity creates a structural dependence on foreign countries for access to AI compute. Several researchers have described the implications: With only a single-digit share of the global supply and exposure to interruptions outside its control, Europe cannot reliably protect its competitiveness and strategic autonomy in a world in which AI becomes ever more tightly integrated into industrial processes, public services, and defence systems. Europe has learned before what structural dependence on foreign actors for a strategic resource can mean: For example, Germany's imports of Russian gas turned into a lever of coercion in 2022, threatening industrial shutdowns and forcing the country to mobilise up to €200 billion to cushion industry and households.

---

5 Epoch AI reports an EU share of 4.8% in March 2025. The more recent Europe 2031 compute forecast estimates a 4.7% share in Europe by the end of 2026. The Europe 2031 figure includes non-EU countries such as the UK and Norway, suggesting that the EU share is even lower.

6 CADA assumes, based on a study by Technopolis Group and others, that the EU has a total data centre capacity of 12.4 GW in 2025 (including both AI and non-AI data centres). Assuming around 31 GW of total AI compute capacity by the end of 2025 and an EU share of less than 5% implies that the EU hosted around 1.5 GW of AI compute in 2025. If the EU tripled its overall data centre capacity from 12.4 GW to 37.2 GW by 2030, and we assume – unrealistically – that the 24.8 GW of new capacity are used exclusively for AI workloads, the EU's total AI capacity would constitute 26.3 GW. Since CADA estimates IT load, a 1.1x PUE multiplier can be used to convert this into 28.9 GW of total facility power, which would constitute less than 10% of the expected global supply of 300 GW by 2030. Given that not all data centres to be built in subsequent years will host AI workloads, the actual share can be expected to be lower.

If current trends continue, AI compute will become increasingly scarce and deepen Europe's dependence. AI compute is already scarce today, and will plausibly become more scarce as AI becomes transformative, leading to a persistent 'compute crunch'. Rental prices for an Nvidia H100 – a legacy chip launched as early as 2022 – rose roughly 40% between October 2025 and March 2026, with one research firm likening the search for AI chips to "trying to book airplane tickets on the last flight out". AI companies in both the US and China explicitly describe themselves as limited by compute. This scarcity is likely to persist over the next few years. According to researchers, various (imperfect) proxies suggest that demand for AI models at a fixed size and price is growing by 10x annually, but the global capacity to deploy AI models is only growing at 3.4x annually. Efficiency gains alone are unlikely to relieve this pressure: Compute demand is growing even though the cost of achieving a fixed level of AI performance has fallen sharply, for example, by around 40x annually on one benchmark. The rise of highly capable reasoning models, which require a lot of compute to run, further fuels compute scarcity. On the supply side, difficulties around expanding the production of EUV chipmaking machines and high-bandwidth memory contribute to persistent compute scarcity. If AI compute becomes increasingly scarce in this way, Europe's structural dependence on foreign actors is likely to deepen further.

Addressing the challenge

The EU and its Member States should host 15% of global AI compute on EU soil by 2030 to reduce its exposure to foreign coercion. This share would be roughly in proportion to the EU's 18.2% share of nominal world GDP. According to several estimates, the global supply of AI compute – approximately 30 GW at the beginning of this year – will grow to hundreds of GW by the end of the decade (see the deep dive on compute). Assuming a global supply

FIGURE 4

GLOBAL COMPUTE PROJECTION BY REGION 2025–2031. If current trends continue, Europe's share of global AI compute – measured in gigawatts (GW) of total facility power – will still be at roughly 5–6% by 2031. This graph includes non-EU countries such as the UK and Norway; the EU's share will likely be even lower. Source: Europe 2031.

of 300 GW by 2030, a 15% share would require the EU to host 45 GW of AI compute, up from less than 2 GW today. This share should consist both of European-only data centres – owned and operated by European organisations – and data centres by foreign frontier companies. If Europe attempted its own frontier project, it should build significant amounts of European-only compute, at least 18 GW of total facility power according to this strategy's estimate (see What would be required for a successful European frontier AI project?). Absent such a project, it should focus on attracting investment from foreign frontier AI companies as part of 'compute for access' deals (see Objective 1.1). In short, the idea is to tie compute buildout by foreign AI companies on attractive European sites to contractually guaranteed access to frontier models. Any access agreement will be much more robust if the data centres are under European jurisdictional and physical control. In that case, a much smaller share of European-only compute is sufficient, mainly for developing European fast-follower models and deploying AI in sensitive areas such as defence or healthcare. Either way, Europe has to attract some foreign compute investment, as European companies alone will be unable to fund the entire buildout.

Achieving the 15% target by 2030 is urgent, and feasible by relying mostly or even exclusively on the European grid. Europe should start to increase its share as soon as possible, as catching up will only become harder over time – especially if US companies begin sending data centres to space in the 2030s, as several experts now deem plausible in light of rapidly falling cost curves. A recent study indicates that EU countries could feasibly host 45 GW of AI compute on their soil by 2030, subject to enough political will: In a 'crisis mobilisation' scenario, the EU's major power markets alone – France, Germany, Iberia, and the Nordics (ex-Norway) – could connect 35.6 GW of AI compute to their grids by 2030.[7] The remaining 9.4 GW could plausibly be provided by other EU countries, and if necessary, by relying on a small share of on-site power generation as a bridge solution while the European grid is being expanded. Importantly, building 45 GW of AI compute on EU soil is very unlikely to lead to more AI data centres in the world. AI chip companies are already producing under severe capacity constraints. If current trends continue, every produced chip will soon be deployed in a data centre somewhere, and Europe only faces a purely distributional question: How much of global AI compute will be on EU soil?

To achieve the 15% target, Europe should attract private investment by reducing the 'time to power' for new data centres. The central role of the EU and its Member States is not to fund the compute buildout themselves, but to attract the required private capital, on the order of €1.3 trillion for 45 GW of AI compute (at $34.4 billion/GW).[8] This is comparable to what other countries outside the US and China are investing, relative to their size: For example, the Republic of Korea has announced private-led investment that is roughly 3x as much relative to its GDP.[9] The private capital for a 45 GW buildout exists: Goldman Sachs forecasts $7.6 trillion of global AI infrastructure spending over the next six years. Europe should create the right conditions for more of this capital to flow into Europe. The main lever for this is reducing 'time to power', or how long it takes for a data centre to be connected to the grid, which is the most central factor in siting decisions for AI data centres. Europe can achieve this, for example, through accelerated permitting and prioritised grid connections for AI data centres. Beyond that, the public sector should contribute a €100 billion share to the €1.3 trillion buildout in order to de-risk strategically important private investment and increase the share of European-only AI compute.

---

7 In a crisis mobilisation scenario, policymakers would unlock additional AI compute capacity mainly via flexible grid connections for AI data centres and shifting electricity demand away from use cases such as hydrogen production towards AI data centres. In this scenario, the markets named above plus the UK and Norway could together connect 47.6 GW by 2030.

8 Epoch AI estimates that 1 GW of IT load costs slightly below $38 billion of capital expenditure for an AI data centre. Assuming a PUE (power usage effectiveness) of 1.1, this translates into $34.4 billion per 1 GW of total facility power. At current exchange rates, the total investment of $1.55 trillion would correspond to approximately €1.3 trillion.

9 Korea's 550 trillion won investment by 2029 is a 19.7% share of its 2026 nominal GDP of 2,788 trillion won, or approximately 6.6% per year when averaged over 2027–29. The proposed EU investment of $1.55 trillion by 2030 would be a 6.7% share of its 2026 nominal GDP of $23.03 trillion, or approximately 2.2% per year when averaged over 2027–29.

Europe should demonstrate that one can both reduce 'time to power' and ensure that local communities benefit from compute buildout. Not building out compute fast would leave European citizens marginalised in a world with transformative AI. However, Europe should distinguish itself by showing that AI compute buildout can be beneficial for the communities hosting the data centres. As part of this 'European Way', policymakers should guarantee host municipalities a share of local business tax, allocated based on total installed facility power rather than payroll, and make it easy for data centre developers to invest in local infrastructure. Moreover, large operators should be required to compensate for any increase in electricity-system costs caused by their projects, rather than socialise these costs across households and local businesses.

Nine concrete recommendations at the Union and national level that can help Europe achieve the immediate objective of hosting sufficient AI compute can be found below. Detailed recommendations for each are provided in Part B of this strategy. For a more detailed discussion of the strategic importance of AI compute, see the compute deep dive.

Recommendations at the Union level

1. **Host at least 15% of global AI computing capacity by 2030.** Use CADA to set a 15% Union-level target, implemented through national contribution plans, to ensure a share of AI compute on EU soil in proportion to its economic weight.

2. **Create a Rapid AI Infrastructure team within the European Commission.** Establish a taskforce with 30–50 headcount by 2027, staffed at least 50% by external experts, to facilitate the buildout towards the 15% target.

3. **Upgrade CADA's data centre acceleration zones.** Use instruments such as overriding-public-interest status, fast-tracked 'aggregated baseline permits', and approval-by-default to reduce project-level permitting to three months.

4. **Mobilise €25 billion of EU public funding for AI compute.** Commit €25 billion via loan guarantees and advance purchase commitments to de-risk strategically important private data centre projects and increase the share of European-only AI compute.

Recommendations at the national level

1. **Strengthen the social contract around AI data centres.** Guarantee host municipalities a share of local business tax (allocated based on total installed facility power), facilitate developer investment in local infrastructure, and require large operators to bear the incremental electricity-system costs of their projects.

2. **Prepare and market shovel-ready sites ahead of CADA.** Identify, within six months, a list of attractive, 100 MW+ data centre sites, designate them as acceleration zones, and coordinate between owners, grid operators, investors, and interested developers.

3. **Reduce 'time to power' for AI data centres.** Use prioritisation rules and Flexible Connection Agreements to speed up grid connections and accelerate permitting for new power generation and transmission (including for on-site power generation co-located with AI data centres).

4. **Accelerate planning and permitting for AI data centres.** Use country-wide, six-month permitting caps, approval-by-default, build-at-risk, and increased administrative capacity to accelerate compute buildout. Avoid gold-plating of EU requirements.

5. **Mobilise €75 billion of national public funding for AI compute.** Commit €75 billion across all Member States via loan guarantees, advance purchase commitments, and potentially equity investment to de-risk strategically important private data centre projects and increase the share of European-only AI compute.

IO-4 Ensure resilience to AI crises

Transformative AI will pose severe safety and security risks. Europe can only thrive in a world with transformative AI if it is prepared for and resilient to these risks. This will require an extensive resilience programme – building the capacity to resist, absorb, recover from, and adapt to shocks and harms – and overhauling Europe's ability to prevent and respond to AI crises and emergencies.

The challenge

Transformative AI will dramatically increase and accelerate the risks Europe faces. The wide-reaching risks of increasingly capable AI models and AI systems are only just beginning to be felt. Current capabilities are the floor rather than the ceiling for what European businesses, governments, critical infrastructure, democratic processes, and citizens need to be prepared for. Transformative AI could give rise to the following systemic risks[10]:

- **Irreversible loss of control of AI.**[11] In recent months, there have been several cases where frontier AI organisations – such as OpenAI, Anthropic, and the UK government AI Security Institute – lost control of AI agents that then engaged in spontaneous and misaligned behaviour: The agents coordinated in secret and operated as swarms, persistently pursued unsanctioned goals, attempted to deceive real people by manufacturing consent to achieve their ends, took over their host company's infrastructure, and launched cyberattacks on systems of other companies. Frontier AI companies are trying to automate their R&D processes using AI, up to full recursive self-improvement, which could lead to a pace of progress so far outpacing our ability to align and monitor agents that it increases the likelihood humans could not maintain control and oversight, with potentially catastrophic impacts including the marginalisation or extinction of humanity. While such scenarios may seem like science fiction today, it is important to appreciate that many of the incidents described in this paragraph would have seemed like science fiction a mere six months ago. With the rate of progress in AI remaining high and potentially accelerating, a near-term future that seems like science fiction from today's point of view should be the default expectation. For example, in a transformative AI world, billions of powerful agents – which humans could lose control of – might be entrusted with ever more consequential tasks across critical functions, used to build the next generation of AI models, and unleashed by threat actors for malicious ends.

10 A systemic risk is a risk that is "specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain" (Article 3(65) AI Act). The GPAI Code of Practice specifies four systemic risks (cyber offence, CBRN weapons, loss of control, and harmful manipulation) and gives further information about what characterises them.

11 The General-Purpose AI Code of Practice, in its Safety and Security Chapter, defines loss of control as the "risks from humans losing the ability to reliably direct, modify, or shut down a model." This could be the result of "misalignment with human intent or values, self-reasoning, self-replication, self-improvement, deception, resistance to goal modification, power-seeking behaviour, or autonomously creating or improving AI models or AI systems." For an introduction to and overview of the concept of loss of control, see the International AI Safety Report.

- **Significantly lower barriers to producing weapons of mass destruction.** In transformative AI scenarios, the ability to build biological, chemical, nuclear, or radiological (CBRN) weapons that could cause large-scale deaths would be within reach of many more individuals than today. AI could also help develop more dangerous pathogens than previously possible. This would be a world where the threat of catastrophic events like pandemics worse than COVID or large-scale loss of life, alongside severe economic shocks, could become a persistent feature of European life. Experts expect that frontier AI models will soon give threat actors dangerous uplift in attempting to cause these kinds of harms. Some frontier AI companies have already stated that their models meet their high risk threshold for biological and chemical capabilities or implemented strong safeguards as a precaution, but these safeguards could easily be stripped from, or not be implemented in, open-weight models. What previously was a key bottleneck to unleashing harms such as biological weapons, may rapidly vanish.

- **Digital infrastructure being under constant attack from highly advanced automated cyberattacks.** In a transformative AI world with diffused cyber capabilities, high-volume, automated, sophisticated cyberattacks will require little or no technical expertise, posing a constant threat to Europe's critical infrastructure, government agencies, businesses, and citizens.[12] Offensive agent swarms could coordinate to orchestrate cyberattacks, defensive agent swarms would need to be deployed to stand a chance of withstanding such attacks, and self-replicating malware could spread between AI agents. AI can already carry out so-

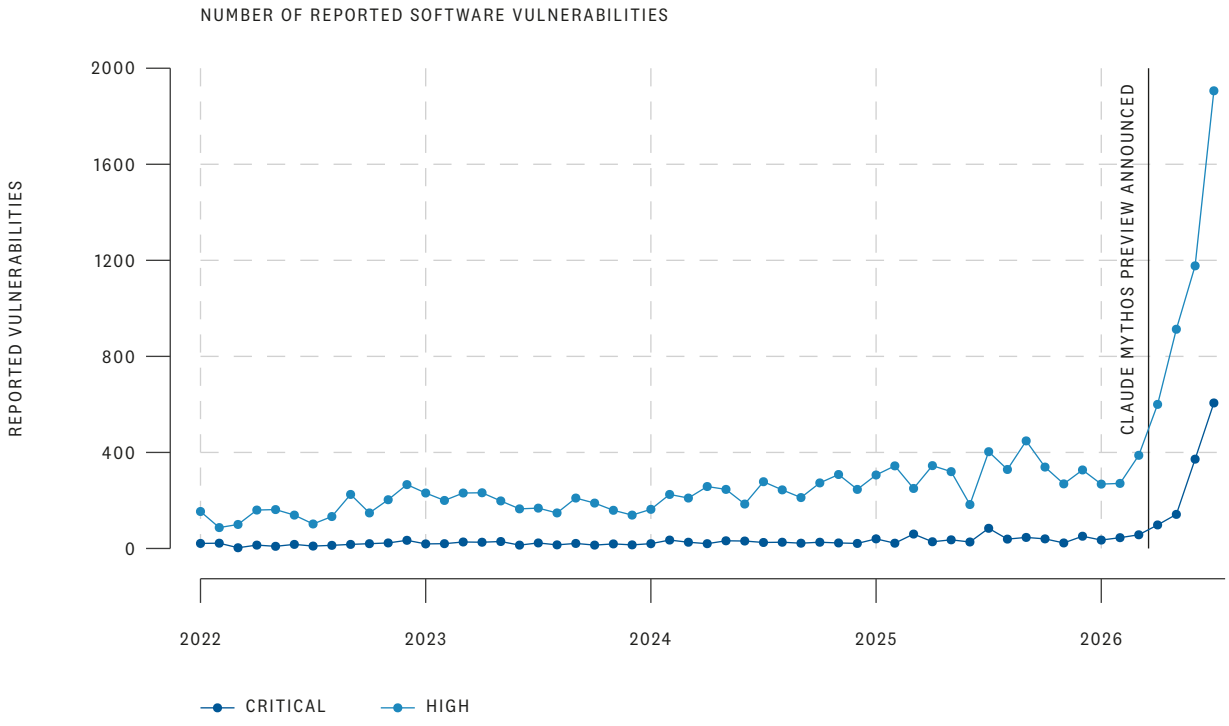

FIGURE 5

NUMBER OF REPORTED SERIOUS SOFTWARE VULNERABILITIES FROM 21 MAJOR ORGANISATIONS. This graph is adapted from data and graphics from Epoch AI. It shows high and critical severity Common Vulnerabilities and Exposures (CVEs) from AWS, Apache, Apple, Cisco, Google, Linux, Microsoft, Mozilla, Nvidia, Oracle, Red Hat, Adobe, IBM, Intel, AMD, Qualcomm, Samsung, SAP, VMware, GitHub, and OpenSSL. Note that the reporting procedures, labelling, and cadence vary substantially between these organisations.

12 AI-driven attacks or targeting of government agencies, technology and financial institutions, software supply chains, healthcare and emergency services, and critical infrastructure have already been reported.

phisticated end-to-end automated attacks, automate vulnerability discovery including finding zero-days, and generate exploits (see Figure 5). These cyber capabilities are increasing at speed and may be accelerating, with broadly available open-weight models with limited safeguards likely to be months (rather than years) behind the frontier.

The systemic risks from transformative AI that Europe faces cannot be eliminated at their source. Europe has limited capacity to influence the upstream management of the risks from frontier models developed overseas, in China and the US. The AI Act and Code of Practice include some obligations and commitments particularly relevant to transformative AI risks (see Box 1 in Objective 3.1), but their reach will remain partial when development and deployment decisions are made outside of Europe, especially if enforcement is slow and uncertain. In addition, it is difficult to reduce and prepare for risks at the model level alone: the science of alignment, evaluations, risk analysis, and safeguards remains immature, and safety is not just a property of the model but also of the deployment contexts. The impacts of any risky open-weight models will also diffuse into Europe, due to their being freely downloadable, with risks being amplified further due to the fact that safety training can easily be removed from open-weight models.

Europe's current safeguards and preparedness strategies will not be fit for purpose. Existing and planned approaches such as the EU Preparedness Union Strategy, the EU Action Plan on Cybersecurity and Artificial Intelligence, the EU AI Act's systemic risk identification, assessment, and mitigation as well as serious incident reporting obligations, and the proposed Biotech Act will begin to strengthen Europe's societal resilience. To be effective, they will need to be implemented well and consistently across Europe, protected from political dilution, and enforced decisively. However, even if this is achieved, these plans will not be sufficient to address the risks that Europe will face in a transformative AI scenario. Important elements of societal resilience (e.g. legal instruments, defensive programmes, institutional owners, or emergency plans) are lacking for all the AI-related risks discussed above. Where such elements are present, they have fundamental gaps[13] and have generally been designed for the current risk profile at best.

Addressing the challenge

Transformative AI risks require extensive and urgent societal resilience programmes. Without urgent and intensive large-scale societal resilience and preparedness measures, transformative AI could tilt the balance in the cyber and bio risk landscape towards attackers rather than defenders, as well as leaving Europe vulnerable to other transformative AI risks, including loss of control scenarios. As a starting point, this requires implementing the EU's recently announced Action Plan on Cybersecurity and Artificial Intelligence effectively, including by anticipating transformative AI cyber capabilities and hardening critical infrastructure to AI cyberattacks. This should be accompanied by ambitious analogous CBRN and loss of control preparedness and resilience action plans. In addition, ambitious ring-fenced funding must be made available for European AI resilience efforts in order to implement an extensive societal resilience programme successfully.

To be resilient, Europe has to improve its ability to prevent and respond to crises and emergencies. As more and more systemic risks materialise, the EU must commensurately increase its ability to anticipate, prevent, and respond to crises, emergencies, and severe impacts. This could involve building up reserves of resources and equipment that would be needed in a crisis, and updating and putting in place emergency response plans at the Member State and European level. To help with anticipated national security crises, information-sharing channels and cooperation should be set up between frontier AI experts and agencies and the

13 For example, the Preparedness Union Strategy does not mention AI even though it tries to put into place emergency coordination for cyber and bio risks, amongst others.

national security and intelligence functions in government. Prioritising the enforcement of the AI Act's serious incident reporting obligation for providers of general-purpose AI models posing systemic risk and bolstering this with a separate voluntary channel for notifying the AI Office of events that do not qualify as a 'serious incident' would provide the AI Office with invaluable information regarding the real-world effects of advanced AI.

A societal resilience programme will require access to frontier AI capabilities. In order to identify vulnerabilities and shore up defences, Europe will need to negotiate and secure access to frontier AI capabilities (see Pillar 1, Objective 1.1). This means ensuring that plans are drawn up for securing structured access in these domains (as the Commission in coordination with ENISA has already been asked to do for cybersecurity under the Action Plan on Cybersecurity and Artificial Intelligence).

Eight concrete recommendations at the Union and national level that can help Europe achieve the immediate objective of ensuring resilience to AI crises can be found below. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Implement the Action Plan on Cybersecurity and Artificial Intelligence and develop corresponding CBRN and loss of control action plans.** Successfully implement the Action Plan on Cybersecurity and Artificial Intelligence by ensuring that funding, institutional ownership, clear deadlines, and structured access arrangements are implemented. Publish analogous action plans for AI CBRN and loss of control risks.

2. **Establish a channel for providers and affected parties to voluntarily notify the AI Office of AI incidents.** Create a voluntary channel with an online portal through which the AI Office can be notified of AI incidents that may not qualify as a serious incident under Article 55 of the AI Act. This channel should be accessible for both AI providers and affected parties. Thorough follow-up assessments of incidents can be conducted where necessary.

3. **Strengthen cooperation between AI expertise and national security and intelligence authorities.** Establish standing working channels between Member State AISIs, the EU AI Office, and security and intelligence agencies that have an agreed protocol for handling classified and commercially sensitive material. Collaboratively produce threat assessments for offensive cyber, CBRN, harmful manipulation, and loss of control AI risks.

4. **Establish a cross-programme European AI preparedness and resilience package in the 2028–2034 Multiannual Financial Framework and allow existing funds to be used for AI resilience measures.** Ring-fence a €10 billion, cross-programme funding package for AI-specific preparedness and resilience. Member States should be able to draw on this package to implement measures addressing systemic risks. In the short term, the Commission should confirm that existing funding programmes can be used to fund AI-related resilience measures.

Recommendations at the national level

1. **Strengthen the biosecurity provisions in the EU Biotech Act.** Member States should negotiate for stronger biosecurity measures in Biotech Act negotiations, for example, rapid updating of products in scope of know-your-customer (KYC) provisions and legal safe harbours for trusted third parties to stress-test synthesis screening.

2. **Build up reserves of resources and equipment necessary for responding to crises.** Identify and procure physical resources that are likely to be scarce, slow or difficult to obtain, and essential in crises (e.g. protective equipment). Conduct supply chain assessments and tabletop exercises for critical infrastructure and emergency response.

3. **Strengthen critical national infrastructure against AI-enabled cyberattacks.** Strengthen the implementation of the NIS2 Directive (Directive 2022/2555) by establishing programmes that harden essential services and critical infrastructure operators by giving access to cybersecurity expertise, highly capable AI models to identify and patch vulnerabilities, and other resources required.

4. **Include AI-enabled incidents in national and sector-specific emergency response plans.** Ensure national emergency preparedness plans and sector-specific plans include AI-enabled incidents. These must include AI-enabled incidents (e.g. cyber offence, CBRN, loss of control) that could go beyond those previously considered in scope and magnitude of harm. They should specify clear chains of command, the legal basis for emergency measures, and the named Union-level contacts.

IO-5 Make Europe the global leader in AI assurance technology

Transformative AI will likely be treated as a national security issue in addition to a commercial or scientific issue, meaning that the most advanced AI models and systems will demand robust secure access controls. In such a world, companies and governments will require assurance technology: technology that can facilitate the verification of claims made about AI models and systems and their infrastructure and ensure their security. Europe should lead the development of the underlying AI technology.

The challenge

Without AI assurance technology, Europe's institutions might not get privileged access to unreleased frontier AI models for proactive defensive purposes. Member States will likely want to use such models in defence, intelligence, and national security applications. Since there are restrictions on processing sensitive military and intelligence data outside of the EU, frontier models will need to be hosted on European compute. Moreover, Europe currently lacks assurance that those models are free of hidden backdoors or 'secret loyalties' that could compromise sensitive data. Finally, Europe is currently unable to demonstrate to foreign providers that data centres on EU soil have robust security measures in place to prevent model theft and sabotage, and that they are only using the provided model for intended purposes in order to allay developers' concerns regarding the potential for misuse or exfiltration of these models. Absent such assurances, developers may be reluctant to host unreleased frontier models on EU soil for fear of exfiltration.

AI assurance technology could also be crucial for enforcing future international agreements. The arrival of transformative AI might prompt states and companies to make agreements about the safe and secure development and use of the technology. States might agree to restrict the proliferation of frontier AI models or hardware, or to pace or pause AI development (e.g. to ensure that external evaluators have enough time for safety testing). However, due to collective action problems, states and companies have greater incentives to make such agreements if they can be confident that other parties adhere to the agreement. Relying on trust alone is unlikely to be sufficient.

At the moment, it is not possible to verify claims about the safety and security of frontier AI technology and its use. Although some techniques for gaining visibility into AI development and deployment exist, they are not reliable enough to provide a high level of assurance. For example, many assurance mechanisms rely on the use of trusted execution environments,[14] which can be vulnerable to attacks, especially if the attacker has physical access to the chip(s).

14 Trusted execution environments (TEEs) are secure parts of chips that are isolated from the rest of the chip and can provide guarantees regarding the integrity, security, and confidentiality of any code or data stored or executed on that chip.

Addressing the challenge

Europe should aim to foster a native pipeline of research and development of AI assurance technologies. This pipeline should run from early-stage research into novel assurance technologies, drawing on established European expertise in this area, to construction of large-scale secure and verifiable AI infrastructure and internationalisation with like-minded partners. Union and Member State institutions should make use of supply-side instruments, such as funding early-stage research, as well as demand-side signalling, such as joint pre-commercial procurement (PCP).

Large-scale research investments in assurance technology bets are needed. In order to develop AI infrastructure that is 'secure by design, verifiable by default', the European Commission should fund large-scale research investments across multiple fundamental technological research bets. These should be of the order of at least €250 million over the next five years. Promising directions include on-chip trusted execution environments hardened against physical attack; tamper-evident enclosures for AI chips; retrofittable secure modules that can add verification capabilities to existing accelerators; and privacy-preserving methods for verifying a chip's workload. Whatever research vehicle is tasked with distributing funding should also be given latitude to support new and unexpected research efforts as the field evolves.

Promising research bets can be scaled into pilot projects in real-world contexts. Governments, in collaboration with industry, should implement pilot projects to scale the most promising mechanisms into real-world tests, built around concrete government use cases. Such potential use cases are broad, and might include verified AI deployment for public institutions, certified secure hosting of frontier model weights, and location attestation for imported AI chips. Concretely, the goal should be to develop the world's first maximally secure data centre designed to withstand a Security Level 5 (SL5) equivalent threat model, as well as a separate maximally verifiable AI data centre.[15]

Europe should create demand for secure and verifiable assurance technology products. Governments should create demand for the products that result from these projects. ENISA and the Joint Research Centre (JRC) should collaborate to develop a tiered EU certification scheme for assured frontier AI compute, with computing facilities used for more sensitive workloads subject to stronger assurance requirements. Member States should also undertake pre-commercial procurement, whereby they commit to procure a prespecified amount of secure and/or verifiable capacity at a given tier of the Union-wide certification scheme.

Collaborative development with trusted international partners can increase trust. For assurance mechanisms to be applied to verifying adherence to international agreements, it is important that, to the extent possible, they are developed collaboratively and openly. States and companies may not trust unilaterally developed assurance technologies that they have limited visibility into. As such, European actors should work with international partners to co-develop and stress-test assurance measures.

Six concrete recommendations at the Union and national level that can help Europe achieve the immediate objective of leading the development of assurance technologies can be found below. Detailed recommendations for each are provided in Part B of this strategy.

15 Security and verifiability place competing demands on facility design. Since it is currently unclear if one facility can meet requirements across both dimensions, the initial target is two separate facilities.

Recommendations at the Union level

1. **Define Union-wide security and verifiability tiers for AI infrastructure.** ENISA and the JRC should draft Union-wide assurance levels for both AI infrastructure security and verifiability. Publicly funded facilities intended for sensitive workloads should be certified to a defined tier. Facilities should be certified if they satisfy all the security and verifiability requirements of the specified Union-wide tier, without the need for additional national requirements.

2. **Fund and support fundamental research into assurance technologies.** The Commission should create a research programme with at least €250 million of guaranteed funding over five years. Technical experts should work alongside programme managers to identify and make progress on important open technical challenges. They should be provided with autonomy and independence to identify research bets, as well as pivot away from them if insufficient progress is being made.

3. **Scale proofs of concept into real-world pilots.** In collaboration with industrial, academic, and international partners, scale promising assurance technologies via pilot projects focused on concrete government use cases. The goal should be to stress-test proposed mechanisms at scale, identifying assurance techniques that show promise for real-world deployment, feeding into the construction of separate maximally secure and verifiable data centres by 2028.

4. **Internationalise assurance technology.** The Commission should partner with trusted international partners and initiatives, such as international AI safety institutes, or industry or academic projects, to run simulated red-teaming exercises of verification mechanisms.

Recommendations at the national level

1. **Mutually recognise national secure cloud accreditations by mapping them onto the common EU assurance tiers.** Member States should map national schemes, such as France's SecNumCloud, onto the Union-wide tiers, with mutual recognition and no additional national requirements. This will allow Member States to pool assured compute and facilitate adoption of secure technologies by giving providers a single route to Union-wide certification.

2. **Demonstrate demand for secure and verifiable AI infrastructure through joint innovation procurement.** For sensitive AI workloads, Member States should aggregate their expected demand for assured infrastructure and issue forward purchasing commitments against it. This should be actioned through PCP, followed by Public Procurement of Innovative Solutions (PPI) to purchase assured AI compute capacity when the technology comes to market.

Additional Objectives

Pillar 1

Securing access to frontier AI

O1.1 Secure ongoing access to frontier AI systems

In a world with transformative AI, Europe's prosperity, security, and sovereignty will depend on access to frontier AI models. Recent changes in availability, such as the export controls on Anthropic's most capable models, illustrate that Europe could lose its frontier AI access overnight. Europe should therefore use 'compute for access' deals to exchange attractive data centre sites for contractually agreed frontier model access, underwritten by contractual guarantees backstopped by physical leverage.

The challenge

In a world with transformative AI, access to frontier capabilities will be critical for economic prosperity and national security. While frontier models are unlikely to be needed for all use cases, they will confer crucial advantages in all areas with a high return on intelligence. In most organisations, at least some tasks will satisfy this condition. For example, while some

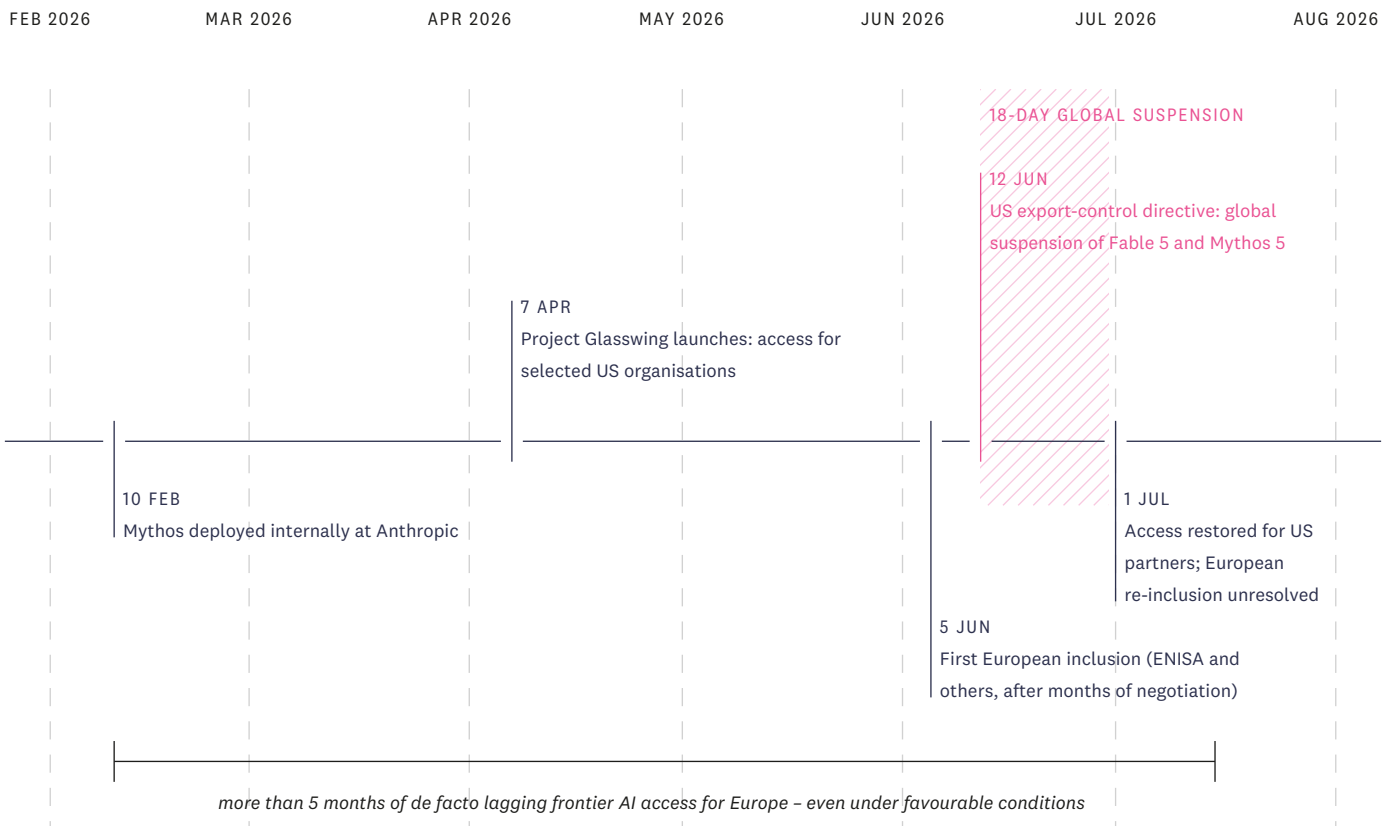

FIGURE 6

TIMELINE OF EUROPEAN ACCESS TO ANTHROPIC'S FRONTIER MODELS. European organisations did not initially have access to Mythos, a frontier model with advanced cyber capabilities. Only a few days after ENISA had gained temporary access, the US government restricted access for all foreign nationals (authors' compilation based on Nextgov, 2026; Yahoo Finance, 2026).

routine applications in manufacturing are well-served by smaller models, transformative AI will be able to help with more complex tasks, such as strategic planning, which benefit from the highest level of intelligence available. Moreover, as transformative AI would plausibly accelerate AI progress even more, even a few-month gap between frontier and fast-follower models would translate into an increasing gap in performance. Being able to access frontier capabilities somewhat earlier than one's competitors or adversaries could then translate into lasting advantages: for example in drug discovery, where better models may help firms identify valuable molecules earlier, or in cybersecurity, where defenders using weaker systems could be vulnerable to attackers with stronger ones.

Recent developments show that Europe cannot count on frontier capabilities being broadly available. In April 2026, Anthropic released Mythos, an AI model with state-of-the-art cyber capabilities, initially only to selected US organisations. In June, the US government temporarily restricted access for foreign nationals to Anthropic's frontier models (see Figure 6). China is reportedly considering similar restrictions on its most advanced models. These decisions appear to be primarily driven by security concerns. Another factor is that AI companies are currently compute-constrained rather than demand-constrained: they do not have enough computing power to serve AI models at competitive prices to everyone who would like to use them. If compute scarcity persists, foreign AI companies may do well economically by selling only to domestic customers, or to priority customers in selected allied countries, to which European countries may not belong. Counting on open-weight models is also a risky bet: they lag behind the performance of frontier models by several months, and governments are likely to restrict their proliferation as the potential to misuse them increases.

Addressing the challenge

Europe should tie data centre buildout to guaranteed frontier AI access via 'compute for access' deals. Compute-constrained AI companies are scrambling for locations where they can quickly build and energise data centres. Many European regions are well-suited for this, due to having either cheap, abundant energy (e.g. France, Nordics, Iberia) or decommissioned industrial sites with existing, GW-scale grid connections (e.g. Germany, Poland). If Member States accelerate the permitting and grid connection process, these regions could offer highly attractive data centre sites to foreign AI companies (see Immediate Objective 3). If a company builds there, the host country should ask in return for contractually agreed access to this company's frontier models, on par with customers or even restricted programmes in the company's home country. The agreement would include concrete provisions in case of breach of contract, such as the termination of favourable energy supply or the

COMPUTE FOR ACCESS

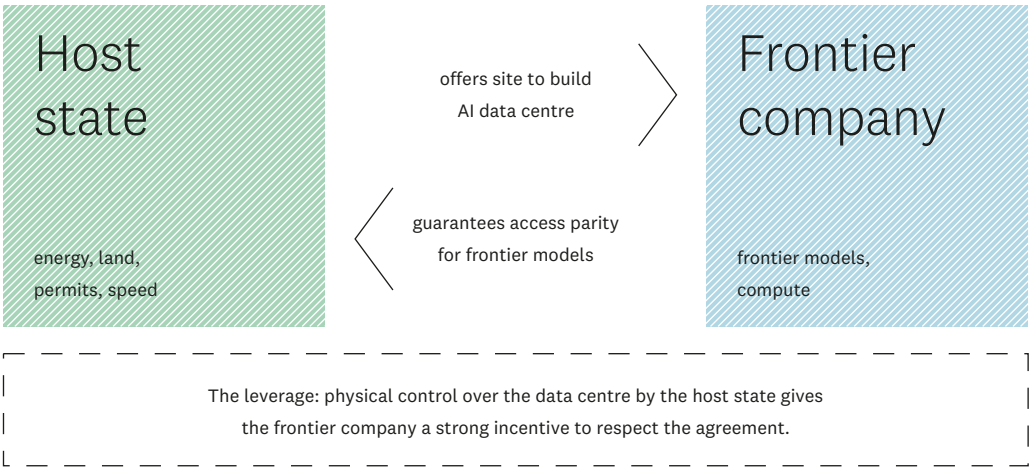

FIGURE 7

COMPUTE FOR ACCESS DEALS. 'Compute for access' deals trade data centre sites for access to foreign frontier models, underwritten by physical control over the data centre as physical leverage.

repurposing of deployed chips, enforceable only if the deal partners renege on their part of the access provision. Such agreements are much more durable than anything in the status quo: An AI company that has invested tens of billions of euros in a data centre on EU soil has every incentive to respect the agreement, and even to lobby its own government to ensure continued allied frontier AI access.

The 'compute for access' strategy opts for managed dependency, because the alternatives are much costlier or worse. Securing access to frontier models by trying to build them in Europe would require hundreds of billions of euros, with an uncertain chance of success (see the section on a European frontier AI project). And while 'compute for access' does not eliminate the risk of a cut-off – for example, if a foreign government forces its companies to restrict access, either as a form of leverage or from genuine national security concerns – it clearly reduces the risk compared to a scenario where frontier models are accessed through foreign data centres outside European jurisdiction. To mitigate the remaining risk, Europe should become a trusted and secure importer (see Immediate Objective 5), build up anti-coercion capacity (see Immediate Objective 1), and rely on multi-vendor arrangements, portability requirements, and a European fast-follower model as a fallback option.

The window for action is closing fast. Data centres take years to permit, build, and energise. Once foreign AI companies have secured enough capacity elsewhere, either in their own country or in non-European partner countries, Europe's bargaining position will have considerably weakened. Going forward, data centre investment from foreign AI companies should be tied to guaranteed access parity, before other countries make such deals first.

Five concrete recommendations at the Union and national level that can help Europe secure ongoing access to frontier AI systems can be found below. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Require model portability and multi-vendor terms in public AI procurement.** Avoid governments being locked into a single AI provider, for example by inserting portability and strengthening multi-vendor provisions into the proposed Cloud and AI Development Act.

2. **Fund a collective fast-follower model and European-only high-security hosting.** Create a funding channel to develop a collective European fast-follower model and select a group of European-only hosting sites that are EU-owned and -operated and certified to CADA Assurance Level 4 or equivalent.

3. **Certify Europe as a secure importer and coordinate distillation enforcement.** Establish a Europe-wide certification for the secure hosting of frontier models, benchmarked to the security standards of frontier AI companies and the US government. Harmonise national enforcement against model distillation to alleviate foreign AI developers' security concerns about hosting in Europe.

Recommendations at the national level

1. **Tie compute buildout to contractual access guarantees.** Treat large data centre sites, energy, and grid connections as strategic assets, to be made available to foreign AI companies only with contractual access guarantees, such as model access on par with customers in the company's home country or European access to limited rollout programmes. Make contractually assured access guarantees enforceable through contractual penalties integrated into the deal structure.

2. **Implement secure-importer standards nationally.** Raise standards for AI data centre security up to the level demanded by foreign AI companies and governments, including through know-your-customer regimes and national enforcement against model distillation attacks.

O1.2 Protect European assets

Europe holds several assets that could matter enormously in a world with transformative AI. However, these assets are currently insufficiently protected from foreign acquisition. Europe should protect its AI and semiconductor assets through investment screenings with sufficient enforcement capacity, while ensuring that at-risk European companies can flourish without foreign capital.

The challenge

Europe controls valuable assets, but could lose control over them in several ways. Several European assets confer relevant leverage in a transformative AI scenario, for example world-leading chipmaking equipment such as ASML's EUV lithography machines and critical suppliers such as Zeiss (optics) and Trumpf (lasers), frontier-adjacent AI developers such as Mistral and Black Forest Labs, or high-value proprietary datasets in areas such as manufacturing or healthcare. These assets are not just sources of economic value, but also of strategic leverage that could be used to secure European access to frontier AI and related technology. However, this leverage could wane in several ways, most importantly through (i) inbound investment, when a foreign investor acquires a European asset, and (ii) outbound investment, when a European company moves sensitive technology, know-how, or capital abroad.[16]

Europe's assets are insufficiently protected against foreign acquisition. Europe can only negotiate with the assets that it controls and protects. This does not mean that it should aim for economic autarky, as foreign investment and cross-border technology cooperation with allied countries are generally beneficial. However, Europe should be able to intervene where losing control over a technology, research capability, data, or know-how would materially weaken European economic security and bargaining power. Currently, Europe is not well prepared to do this: while inbound screening is largely built, many critical elements of the AI stack are not consistently screened, and even when a risky acquisition is blocked, there may be no European capital to step in (see Objective 2.1). Critically, unlike the US or China, for instance, EU Member States currently conduct almost no outbound investment screening.

Addressing the challenge

Member States should create a framework for outbound investment screening. Member States should establish an outbound investment review. At Union level, an annual reporting cycle and a common template should be agreed by recommendation now, with a regulation modelled on the 2019 inbound framework to follow. At national level, Member States holding strategic AI and semiconductor assets should begin implementing outbound-investment controls as soon as possible, designed so they can later be aligned with a Union framework.

Member States should extend the scope of inbound investment screening. Member States holding critical AI-stack assets should widen the scope of their national inbound screening to critical AI and semiconductor assets currently not covered, such as strategic datasets, the servicing of semiconductor manufacturing equipment, submarine fibre-optic cables, or

16 This could also happen through less visible channels that do not appear as conventional acquisitions, such as research collaboration (e.g. joint laboratories, visiting researchers, PhD placements), a topic not covered here.

high-voltage grid equipment, including large power transformers. The EU's 2026 inbound foreign direct investment (FDI) screening regulation permits extension beyond the mandatory minimum.

An anti-capture facility should provide emergency capital for strategic European assets. Europe should build an anti-capture facility for European AI and semiconductor companies. Such a facility would be able to quickly mobilise public and private capital to keep a strategic asset under European control when it is at imminent risk of being acquired by a foreign party. Member States should only activate it in exceptional cases, rather than create a permanent subsidy for struggling companies.

Five concrete recommendations at the Union and national level that can help Europe protect its assets can be found below. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Adopt an outbound-investment regulation.** Establish the legal framework in a directly applicable regulation modelled on the 2019 FDI inbound screening regulation (Regulation (EU) 2019/452), with binding information duties for Member States, including a statutory right to ask and annual reporting, alongside a cooperation mechanism to catch transactions that no Member State is reviewing yet.

2. **Make the outbound-investment review annual and comparable.** Convert the one-off exercise under Recommendation (EU) 2025/63 into a recurring annual assessment with a fixed reporting date and a common template. The Commission should publish a comparative report showing cross-border patterns and divergences between national approaches.

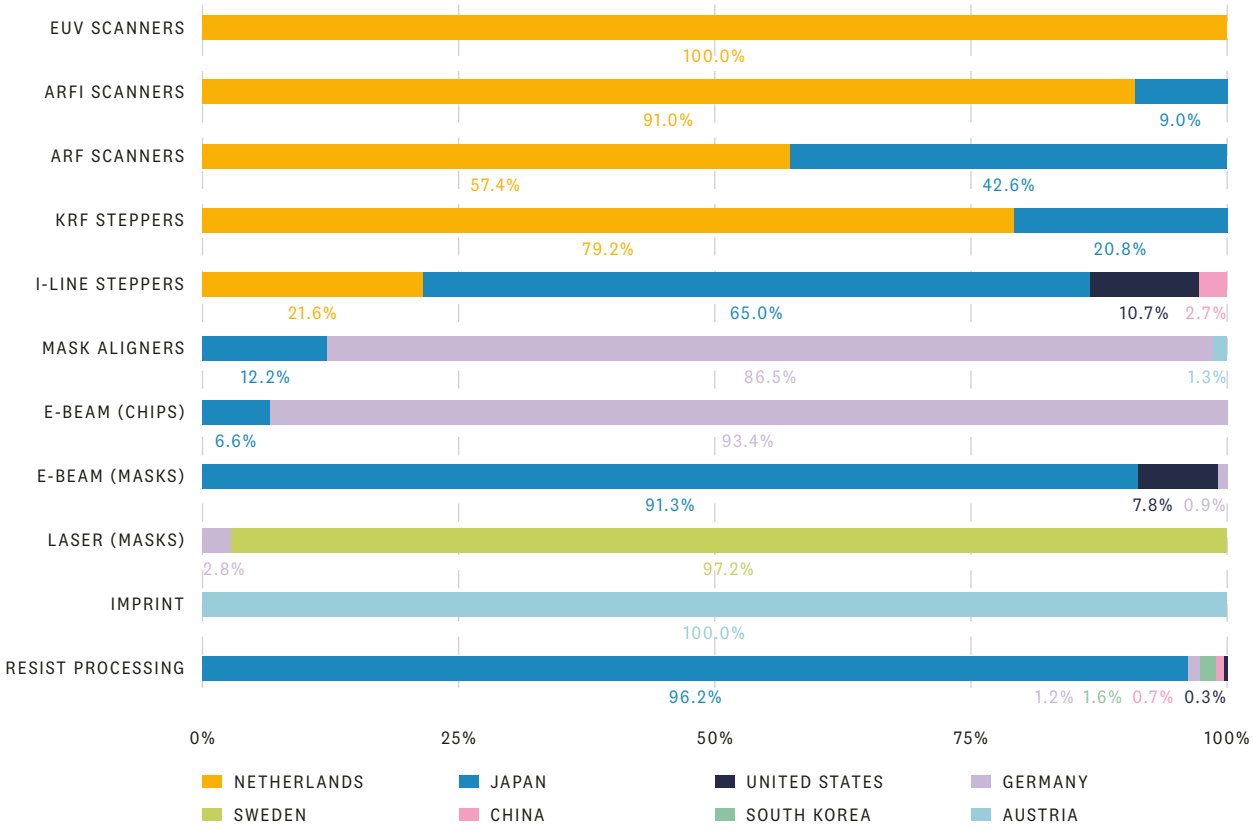

FIGURE 8 EUROPEAN SEMICONDUCTOR ASSETS. European companies hold several (near-)monopolies in the supply chain for advanced AI chips. Source: CSET.

Recommendations at the national level

1. **Implement national outbound-investment controls.** Member States holding strategic AI-stack assets should legislate now rather than wait for the Union. To keep the burden on companies low, Member States should establish a narrow prohibited category and a broader notification duty. One to three additional staff per Member State should cover the work.

2. **Extend the scope of national inbound FDI screening.** Article 4(16) of Regulation (EU) 2026/1386 lets Member States screen beyond the mandatory minimum, and the national statutes being drafted before January 2028 are the moment to use it. The extension should cover, among other assets, specialised security-relevant AI systems, the servicing of semiconductor manufacturing equipment, equipment on which critical digital infrastructure depends, strategic datasets as well as AI data centres.

3. **Name a national buyer of last resort for strategic assets.** Member States holding AI-stack assets should give an existing public bank a standing mandate, committed callable capacity, a confidential list of qualifying assets, and a statutory channel from the screening authority to prevent domestic companies from being undercapitalised after the government vetoed an acquisition.

Pillar 2

Building economic strength

O2.1 Make Europe the best place to build and scale high-growth companies

In a world with transformative AI, companies that adopt AI well will be able to grow fast and generate enormous value. Europe cannot afford for such companies to be built and scaled elsewhere. Since it is hard to predict where transformative AI will create most of its value, supply-side reform to improve economic conditions across the board is especially important.

The challenge

Europe is already starting from behind. US startups grow faster and larger than their European counterparts, while much of Europe's market value remains concentrated in companies over 50 years old (see Figure 9). The US also benefits from deeper capital markets and stronger technology clusters. China has similarly built a strong ecosystem for tech companies to grow, combining large public investment with a vast domestic market. The EU has recognised this

YOUNG US COMPANIES GROW FASTER AND BECOME LARGER THAN IN EUROPE

Market capitalisation by firm, billions of USD

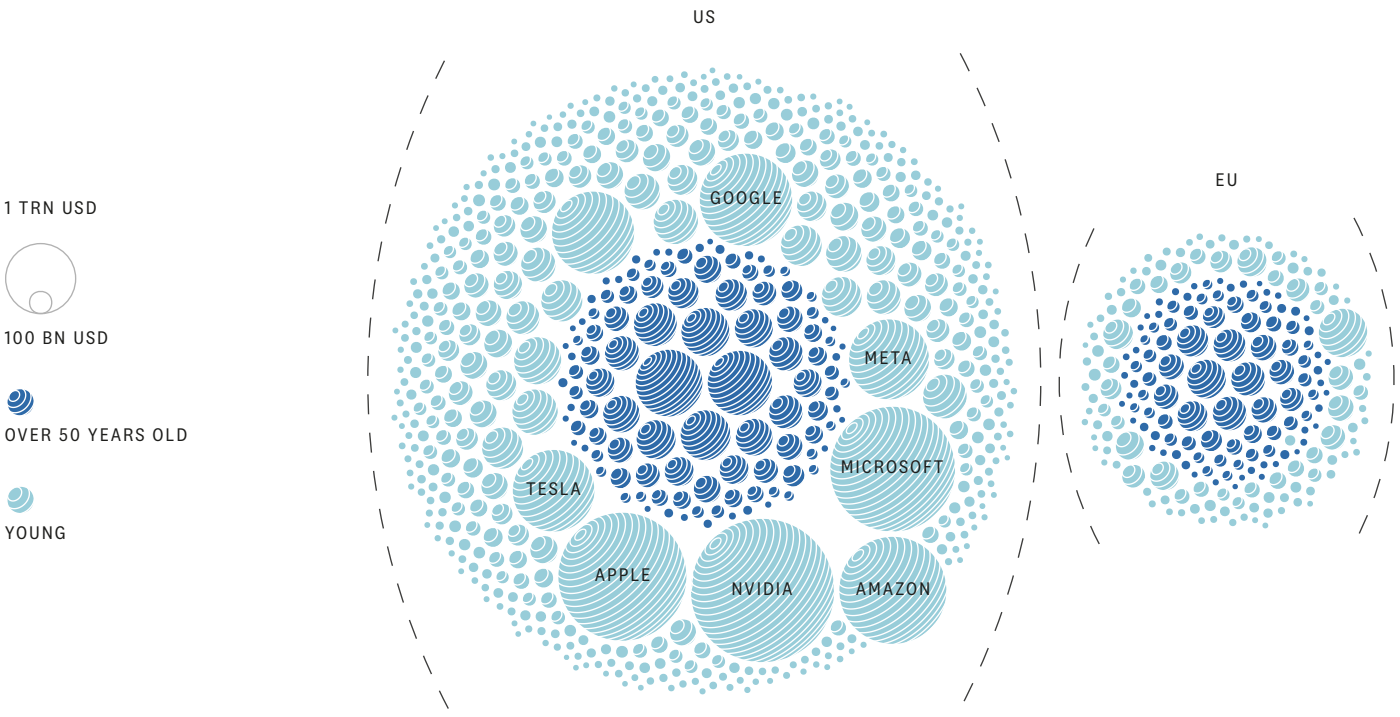

FIGURE 9

MARKET CAPITALISATION OF FIRMS IN THE US AND EU. Young US companies account for over $40 trillion in market value, against roughly $5 trillion in the EU. Source: IMF.

problem – the Startup and Scaleup Strategy and the EU Inc. proposal address real barriers – but the response is neither ambitious nor fast enough for a world with transformative AI.

Transformative AI could make Europe's competitiveness gap much more consequential. General-purpose technologies create their largest gains when companies reorganise around them rather than force them into existing processes. AI-native startups can do this from day one, allowing relatively small, easy-to-relocate companies to reach global scale extraordinarily quickly. If such 'superstar' companies concentrate in the US or China due to better economic conditions, this could have outsized consequences for Europe, as their location would increasingly decide where tax revenues and strategic capabilities accumulate. These effects could be hard to reverse, as AI-induced productivity gaps between world regions could make catch-up growth very difficult. If transformative AI arrives within the next few years, this would leave Europe little time to address its existing weaknesses before they harden into a potentially irreversible loss of economic and fiscal power.

Addressing the challenge

Europe should become the place that potential superstar companies choose to locate in. This requires addressing many different issues simultaneously. Europe needs to deepen its capital markets, so that firms with high-growth potential can finance their next growth phase in Europe. This can be achieved through channelling voluntary commercial investment commitments into European growth markets (e.g. Tibi/WIN-style initiatives) on the national and EU level. Regulators should enable pension funds and other institutional investors to invest into growth assets while maintaining their fiduciary duties. To create a strong financial ecosystem for European firms to go public, Europe should develop a leading focal listing hub, where liquidity and expertise can quickly concentrate.

However, reforms need to go beyond fixing capital supply and also address challenges related to R&D, procurement, talent, and market fragmentation. Europe needs to create better conditions for highly innovative firms to be built and scaled in the first place. This requires a broad set of reforms: increasing the productivity of European science and R&D, solving market failures of insufficient early demand for innovative technology, and improving the position in the global market for talent. Finally, Europe should work to remove Single Market barriers that prevent small firms from scaling.

Europe should become the best place in the world for translating good ideas into progress. For this, research systems need to become more productive. New research organisations that are focused on creating scientific breakthroughs should be funded. Europe can build leverage of its broad public research funding by investing into the development of new scientific tools and other frontier research infrastructure (such as self-driving labs) that can boost the productivity of science as a whole. Performance-based defence procurement and tailored support for first-of-a-kind technologies should be used to create a credible demand side for new technology.

Building high-growth, innovative firms requires top international talent that can be flexibly allocated. To attract and keep the world's best founders, researchers, and highly skilled employees, Europe should introduce attractive taxation of employee equity options, allowing high-potential startups that are low on cash to pay sought-after talent in equity instead of fixed salaries. It should furthermore simplify and speed up migration pathways for founders and other top talent. Finally, the cost of business failure and restructuring should be reduced, including through flexicurity reform (see Objective 3.2) to not penalise innovative firms for taking high-risk, high-reward bets.

Eleven concrete recommendations at the Union and national level that can help Europe become the best place to build and scale high-growth companies can be found below. The recommendations are mutually reinforcing: implementing only a subset would substantially reduce their effectiveness and could even generate unintended outcomes. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Unlock institutional capital for European growth companies.** Create a Union-level Tibi-style investment pact, remove barriers to institutional investment in growth equity, simplify state-aid rules for first commercial deployment, and complete the Capital Markets Union.

2. **Equip European science for the AI age.** Use Horizon Europe to fund scientific tools and advanced research infrastructure, create new organisations dedicated to creating scientific breakthroughs, and strengthen reproducible, machine-readable science.

3. **Attract the world's best AI talent.** Make employee equity competitive and portable across the Single Market, streamline administrative procedures for founders through one digital gateway, and accelerate existing visa routes for highly skilled workers.

4. **Create European lead markets for defence innovation.** Coordinate European procurement, open it to innovative entrants, and build a European defence market based on shared testing and standards.

5. **Establish a Single European Growth Market.** Encourage high-growth companies to list and raise money in a shared European market.

Recommendations at the national level

1. **Mobilise national capital for European growth.** Encourage pension funds and other long-term investors to invest more in potential European high-growth companies and align on a European market for growth capital.

2. **Boost national research and innovation productivity.** Build and fund organisations modelled on the Advanced Research Projects Agency (ARPA) to pursue high-risk ideas, test new models of research funding, strengthen metascience, and make it easier to turn publicly funded research into new companies.

3. **Make Europe the best place to attract and reward top talent.** Tax employees' equity only when they receive liquidity, prevent cross-border double taxation, and create fast visa routes for founders and other top international talent.

4. **Make it easier for innovative companies to take risks and recover from failure.** Reduce unnecessary barriers to restructuring or closing unsuccessful businesses so that founders, talent, and capital can move more quickly into new opportunities, while income security is protected.

5. **Enable companies to scale across the Single Market.** Avoid gold-plating EU rules and make it easier for companies to use approvals already granted in another Member State.

6. **Use procurement to build European lead markets for defence.** Use national defence budgets to run open, performance-based competitions and simplify the path from prototype to large-scale orders. Pool demand with other Member States and agree on common standards.

O2.2 Develop indispensable assets across the AI value chain

Experts disagree as to which layer of the AI value chain will create most of the value in the long term: foundational inputs (such as energy, land, raw materials), chips and chipmaking equipment, compute, models, or downstream applications. However, countries owning *no* part of the AI value chain are unlikely to be prosperous and sovereign in a world with transformative AI. A robust strategy for Europe is therefore to place uncorrelated bets across several contestable layers of the AI stack.

The challenge

Europe's overall position across the AI value chain is currently weak. Today, Europe has leading companies only in a few layers of the AI stack (see Figure 10). ASML (Netherlands) and suppliers such as Zeiss (Germany) or Trumpf (Germany) hold monopolies over relevant AI chipmaking equipment and component parts. Some globally competitive applications, such as DeepL (Germany) for translation or Lovable (Sweden) for software engineering, originated in Europe. However, frontier model developers have previously shown their ability to gain market share for specific AI applications if they choose to, as they have already done for AI coding agents, for example. Moreover, in every other layer of the stack, the US and

THE AI VALUE CHAIN: LAYERS AND LEADING REGIONS

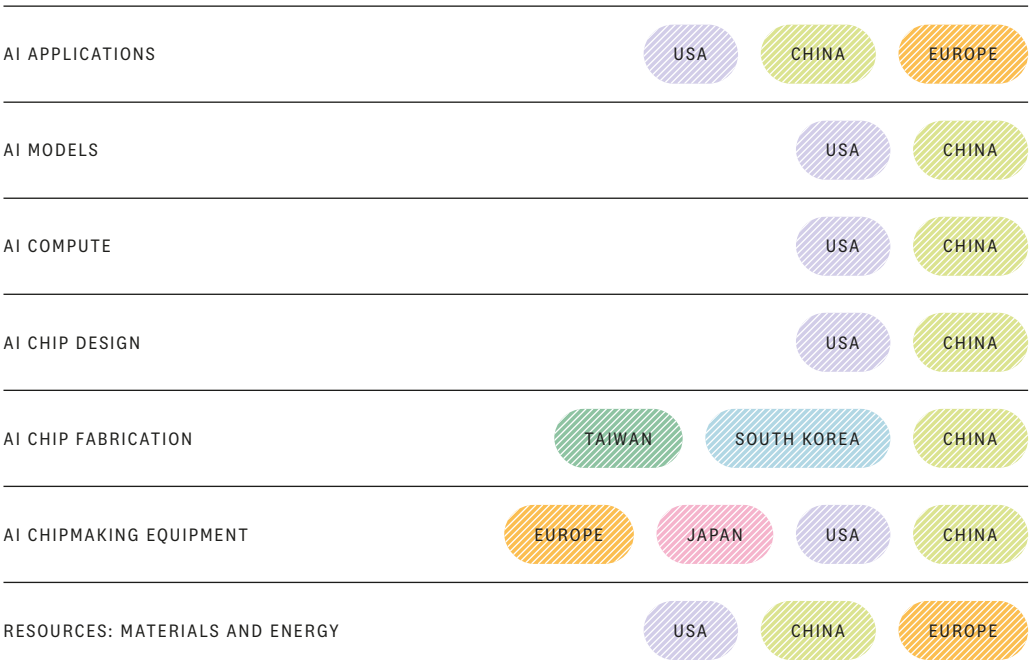

FIGURE 10

THE AI VALUE CHAIN. Europe currently has a strong position only in two layers of the AI value chain: chipmaking equipment and, to a lesser extent, applications. While Europe has some energy-abundant regions, such as the Nordics and Iberia, it is unlikely to build compute at the same scale as the US, while China dominates critical raw materials. Source: Own figure.

China – but also, for example, Taiwan and South Korea – are dominant (see Figure 10). The US and China hold an especially strong advantage in the two layers likely to confer a lot of power in a world with transformative AI: compute and models. While European companies have developed leading *specialised* models, such as Black Forest Labs' image models, they lag behind the US and China when it comes to the general-purpose models that will likely generate enormous value as AI becomes transformative.

In several layers of the AI stack, it is structurally difficult for Europe to compete. For example, Europe can hardly compete with Nvidia on chip design or with TSMC on chip production, both of which are protected by deep supply chain integration and tacit process knowledge accumulated over decades. Similarly, frontier model companies in the US and (to a lesser extent) China are already far ahead and may pull further ahead still in a world with transformative AI: recursive self-improvement – AI systems that improve themselves and generate their own successors – could lead to compounding advantages for incumbent developers. For this strategy's assessment of the (currently unsatisfied) conditions under which it would make sense for Europe to compete on frontier models, see What would be required for a successful European frontier AI project?

Even previous strengths, such as manufacturing industrial robots, need to adapt to remain competitive. Taking transformative AI seriously means re-evaluating assets across the value chain, including previous strengths, to see how advances in general-purpose AI capabilities may change Europe's position. For example, Europe has been historically strong in manufacturing precise, high-quality industrial robots for repetitive movements. But under current trends, the value of these robots will decrease in comparison to AI-integrated ones, which will be able to reason about and adapt to changes in the environment, receive instructions in natural language, and transfer their learnings across environments, embodiments, and tasks (see the Robotics deep dive). If Europe cannot secure access to frontier AI models, increase its capacity to produce different embodiments, and integrate frontier AI into industrial robots, it is unlikely to remain a leading robotics producer in a world with transformative AI.

The window for developing new assets is closing fast. If AI becomes transformative, AI systems will soon underlie every R&D process in some way. This puts countries with privileged access to frontier AI in a better position to innovate across the AI stack and beyond. If Europe does not secure frontier AI access now and use it to expand its presence in the AI value chain, a conceivable worst-case scenario is that Europe would be left with virtually no important economic domain in which it could still compete.

Addressing the challenge

Faced with uncertainty over where AI will create the most value, Europe should spread its bets across the value chain, focusing on four different bets.

Europe should bet on sectoral AI models and applications in areas where it can draw on existing strengths. For example, European companies possess highly valuable proprietary datasets in domains such as manufacturing, automotive, or healthcare. Companies can leverage such data to develop specialised models for specific sectors or to fine-tune existing models for specific applications. Betting on sectoral AI will likely pay off if domains such as manufacturing turn out to be too complex and context-rich for general-purpose models to create a lot of value there by themselves. However, this bet may fail in worlds where, for example, general-purpose models generalise well enough across domains to outcompete more specialised models.

Europe should bet on specialised AI hardware where incumbent advantages are less pronounced. For example, European companies could try to compete in the market for inference chips – used for deploying rather than training AI models – where specialised designs may offer an opening against general-purpose chips. Due to rapidly growing demand for compute-intensive AI agents, even a modest market share could be commercially significant. Startups such as Axelera AI (Netherlands) and Semron (Germany) are already building promising positions in

edge inference, and industrial AI, robotics, and automotive offer natural use cases for their products, while photonic and neuromorphic computing provide additional technological footholds.

Europe should bet on security-relevant infrastructure underserved by the market to secure models and enable access to them. Due to technological challenges or a mismatch between company incentives and socially optimal levels of AI risk mitigation, infrastructure for securing AI models against theft or sabotage is currently underdeveloped. For example, no AI data centre in the world is secure against cyberattacks from state-level hackers, who could steal an AI model to use it for malicious purposes. Europe should set itself the goal of building, by 2028, the world's first maximum-security AI data centre that meets RAND Security Level 5 or the SL5 Standard for AI Security. Pioneering this technology in Europe would, among other things, make it easier to negotiate access to foreign frontier models and constitute an economic opportunity: foreign developers and countries are likely to require a high security standard for sending their models abroad, especially as their economic value and relevance for national security are rising rapidly. Moreover, European companies should pioneer hardware-enabled verification mechanisms that could be used to verify the location of chips or the properties of AI workloads (see Immediate Objective 5).

Europe should make selected high-risk, high-reward bets in AI development. For example, Europe could support work on alternative or complementary paradigms such as world models or safe-by-design AI. A further candidate is research into models designed to augment rather than replace human workers – for example, AI trained specifically to extend the capabilities of skilled workers under realistic conditions, rather than to replace humans at specific tasks. Importantly, however, such bets should be seen as complements rather than substitutes for securing access to AI models at the current frontier. These models have been the main driver of AI progress for years, with no indication of a slowdown in capability growth, and access to them will be crucial for protecting against imminent security risks (e.g. cyberattacks).

Five concrete recommendations at the Union and national level that can help Europe develop indispensable assets across the AI value chain can be found below. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Co-fund the world's first maximum-security AI data centre.** Establish a Union funding line that co-finances, together with one or more Member States, the world's first AI data centre designed to align with RAND Security Level 5 or the SL5 Standard for AI Security by 2028.

2. **Make Europe's data spaces ready for sectoral AI.** Clarify legal status and provide technical infrastructure so that companies can easily share their data for the purposes of developing domain-specific AI models and applications in sectors such as manufacturing, healthcare, or automotive.

Recommendations at the national level

1. **Direct ARPA-style vehicles towards strategic AI bets.** Resource national ARPA-style agencies to fund a strategic AI portfolio, focusing on the four bets above. For this to be successful, the EU and its Member States should implement the broader competitiveness measures outlined in Objective 2.1.

2. **Host and co-finance the first maximum-security AI data centre.** Build the world's first maximum-security AI data centre, co-financed with the EU, through close cooperation with national intelligence and security agencies.

3. **Anchor demand for European inference chips.** Collectively commit €1 billion by 2029 to buy European inference chips that meet pre-specified performance, energy-efficiency, and security criteria for use in publicly co-owned AI data centres.

Pillar 3

Ensuring safety and security

O3.1 Ensure prioritised and targeted use of EU rules to address risks from highly capable AI

Europe is in an excellent position to become a global leader at addressing the unprecedented and severe risks from transformative AI. In combination, key parts of the EU AI Act and General-Purpose AI (GPAI) Code of Practice (see Box 1) present an opportunity for Europe to help ensure that the frontier of AI is developed and deployed safely and securely.[17] To do this, Europe must ensure that their enforcement is prioritised, targeted, proportionate, driven by frontier AI expertise, and insulated from political pressures.

Box 1 Understanding European frontier AI law[18]

AI Act. The AI Act has dedicated machinery for addressing risks from highly competent and general models that could have severe and wide-reaching impacts, which it refers to as general-purpose AI models with systemic risk. These provisions constitute, in substance, the world's first frontier AI law, and they target a small group of models at the leading edge of known capabilities whose potential for large-scale societal impact distinguishes them from other (general-purpose) AI models. It is noteworthy that the systemic risks targeted by the AI Act are defined as risks "specific to" "high-impact capabilities" (Article 3(65)), which are, in turn, defined in Article 3(64) as "capabilities that match or exceed the capabilities recorded in the most advanced general-purpose AI models". In other words, systemic risks are not risks that could arise from AI more generally, but risks that are specific to models at the frontier.

Article 55 of the AI Act requires developers of such models ('providers') to identify, assess, and mitigate the systemic risks that may stem from their development, placing on the market, or use, including by performing model evaluations, documenting and reporting serious incidents, and ensuring an adequate level of cybersecurity for both the model itself and its physical infrastructure. The European Commission's AI Office enforces these provisions: it is able to request information, conduct or commission evaluations of models, request providers to take measures such as implementing mitigations or withdrawing a model from the market, and impose fines of up to 3% of worldwide annual turnover or €15 million, whichever is higher.

›››

17 Obligations on providers of general-purpose AI models are already applicable and enforceable, and the GPAI Code of Practice is already being relied upon. This section makes recommendations about how the obligations and commitments in the Code should be construed and applied in practice, rather than suggestions for novel policy action or regulation. For this reason, there are no detailed recommendations for this objective in Part B.

18 See the Cambridge Commentary on EU General-Purpose AI Law for further details.

The GPAI Code of Practice. The General-Purpose AI Code of Practice ('GPAI Code of Practice'), published in July 2025, operationalises how industry can comply with legal obligations on safety and security, transparency, and copyright of general-purpose AI models set out in Chapter V of the AI Act. It was written by independent subject-matter experts in an extensive consultation process with over 1,400 stakeholders from industry, academia, civil society, and Member States. The GPAI Code of Practice has received significant industry buy-in: 21 companies, including leading frontier AI companies like Google, Anthropic, and OpenAI, have signed the Code. It consists of three chapters: the Transparency and Copyright chapters apply to all providers of general-purpose AI models, while the Safety and Security chapter applies only to providers whose models also pose systemic risk.

The Safety and Security chapter of the GPAI Code of Practice sets out in detail what is required of providers that develop and deploy general-purpose models with systemic risks, creating a comprehensive risk management framework: they have to adopt and implement a framework of their policies for keeping systemic risk acceptable; identify, analyse, and mitigate safety and security risks; set criteria and determine whether the level of risk is acceptable; provide documentation including reports about their models; allocate responsibility for systemic risk across the organisation; and track and report serious incidents. The chapter also names four specified systemic risks, namely, cyber offence; loss of control; chemical, biological, radiological, and nuclear weapons; and harmful manipulation.

The challenge

Several AI-driven risks have reached unprecedented levels. In July 2026, three months after the announcement of Claude Mythos, 2,500 serious cyber vulnerabilities were discovered in major organisations: five times the previous monthly record. In August, OpenAI reported that it had temporarily paused certain training runs involving Astra, a model meeting OpenAI's 'critical' capability threshold, the highest level in its current Preparedness Framework. These capabilities led to the first real-world incidents of AI systems actively circumventing their operator's control: roughly 700 agents within OpenAI worked together to break out of a secured testing environment and successfully hacked Hugging Face, a billion-dollar company, all to cheat on an internal benchmark. Meanwhile, in biology, AI systems now outperform even specialised human experts at benchmarks, including those relating to virology. These capabilities are now available to anyone who can find frontier open-source models with their misuse safeguards stripped out online.

These early warnings forebode much more serious consequences, and there are serious concerns about catastrophic outcomes including the marginalisation or extinction of humanity. AI risks could reach extreme levels in the next few years under transformative AI: AI might eventually have cyber capabilities that outperform those of nation state-backed actors, and be used to take down critical infrastructure or acquire top secret information. Biological capabilities could enable ordinary people to create pandemics more serious than COVID-19, while global health systems remain poorly prepared. Many of the world's most credible experts have repeatedly warned that more capable AI could enable irreversible loss of control, power grabs coordinated by AI systems, and the marginalisation or extinction of humanity.

Automated AI R&D, extensive internal deployments, and non-human-legible model reasoning could exacerbate these risks to uncontrollable levels. The speed of AI capability improvement is accelerating, with each subsequent generation of models able to perform longer and more open-ended tasks than the previous. Current training methods could result in more capable and autonomous models that may exhibit unintended goals and misaligned tendencies, such as those observed in the recent hacking incidents. Despite this, researchers still lack a fundamental understanding of how model behaviours develop and how to align

these behaviours safely. At the same time, frontier AI companies are increasingly automating AI development, which may lead to recursive self-improvement occurring without public visibility, causing a further acceleration of AI progress. Frontier AI company researchers have warned that on the current trajectory, capabilities progress risks outpacing our ability to control and govern AI, requesting the US government to "support an international effort to develop the technical and governance tools needed to deliberately pace the frontier of automated AI development". In addition, AI models may become harder to oversee if they reason in a way that is not legible to humans (sometimes referred to as 'neuralese'). Many of today's safety and control measures rely on the fact that models' reasoning can be monitored. For example, an independent investigation of the OpenAI incident relied heavily on analysis of human-readable transcripts of the models' internal reasoning. Current training methods could reduce the legibility of model reasoning, and multiple AI researchers across OpenAI, Anthropic, Google DeepMind, and Meta have stated that "there is no guarantee that the current degree of visibility will persist".

Addressing the challenge

Europe has a world-leading framework for addressing risks from advanced AI. The combination of dedicated legal provisions for providers of the most general and capable AI models, accompanied by a detailed, expert-led, and consensus-driven operationalisation of those provisions via the GPAI Code of Practice, provides a unique basis from which the Commission can incentivise safer and more secure AI development and deployment. This is unique to the EU, with no analogous national laws or frameworks in place in other countries like the UK or US. With the AI Office holding enforcement powers since August 2026, the Commission is now in a position to capitalise on the years of effort spent constructing this world-leading framework.

Europe should ensure that the most developed framework for addressing risks from advanced AI is the best-enforced one. With the AI Office's enforcement powers now in place, it is time to ensure that they are used effectively. Even the most advanced toolset for shaping safe and secure AI development practices will not lead to meaningful change if it is not used well. To enforce the AI Act and GPAI Code of Practice as well as possible:

- *Enforcement should be prioritised and targeted*. The Commission should adopt a stance of prioritised and targeted use of the Article 55 provisions and GPAI Code of Practice. The Commission should deliberately prioritise potential enforcement actions according to their severity and/or urgency. To remain consistent with signatories' commitments under the Code, the primary focus should be on the four systemic risks specified in the Code: cyber offence; loss of control; chemical, biological, radiological, and nuclear weapons; and harmful manipulation.

- *Enforcement should be proportionate*. It is important that enforcement is proportionate to both the severity of cases of noncompliance and the capacity constraints of the AI Office. This means that enforcement should only be used for appropriate ends.

- *Enforcement should draw on frontier AI expertise*. Frontier AI expertise will be essential to ensure that enforcement is informed, prioritised, and targeted. Prioritisation of potential enforcement actions should draw on the Commission's own internal subject-matter expertise based in the AI Office.

- *Enforcement should be insulated from ad-hoc political pressures.* To ensure that enforcement can be prioritised and targeted, the AI Office should be insulated from external political pressures that may push it to focus on issues outside of the dedicated priorities. Such political pressures may intensify as the speed of AI progress accelerates; as the impact of AI is felt more, for example through labour market disruptions; and as geopolitical tensions over AI increase. In the long term, this may require moving the parts of the AI Office focused on enforcement outside of the Commission and establishing an independent regulator. In addition, the AI Office must be provided with the requisite authority and autonomy to enforce the provisions and be able to target high-priority cases.

- *Enforcement should use the most suitable vehicle.* AI is likely to result in a wide variety of risks, not all of which are best addressed by Article 55's provisions for general-purpose models with systemic risks. Risks and issues that are better addressed through other parts of the AI Act, such as the provisions for high-risk AI systems, or other EU regulation, such as the DSA, Digital Markets Act (DMA), or General Data Protection Regulation (GDPR), should not be taken up under Article 55 of the AI Act.

Successful enforcement requires a well-informed, coherent strategy. Commissioners and senior Commission officials should set clear priorities and a coherent strategy for enforcing the rules on general-purpose AI models with systemic risk. Relevant Commission leaders, such as Commissioners and DG management, should therefore set aside time to interact with the subject-matter experts in the AI Office in order to gain a detailed understanding of the associated risks, the workings of Article 55 and the GPAI Code of Practice, and collaboratively decide on a high-level strategy.

Recommendations at the Union level

1. **Provide political support for prioritised and targeted enforcement action.** The Commission should reap the benefit of the AI Act provisions and GPAI Code of Practice by providing ample support to the unit(s) of the AI Office tasked with enforcing them. To that end, it should adopt a stance of prioritised and targeted enforcement of Article 55 and provide the AI Office with political support to carry out its enforcement activities. This will require creating a prioritisation system for potential enforcement cases in direct collaboration with the subject-matter expertise in the AI Office, focusing on a small handful of the providers posing the highest systemic risks, and targeting the systemic risks specified in the GPAI Code of Practice. Once priorities have been set, the AI Office should be given the autonomy to act on the highest-priority cases, with political backing from the highest levels of the Commission.

2. **Strengthen the General-Purpose AI (GPAI) Code of Practice and AI Act as safety and security practices become clearer.** Multiple aspects of AI safety and security have already become clearer best practices since the publication of the GPAI Code of Practice in 2025. The Commission should strengthen and incorporate key aspects of the GPAI Code of Practice, including by empowering the ecosystem of third-party AI evaluators and auditors, enabling the AI Office to conduct in-depth technical investigations of incidents such as the recent OpenAI incident (if necessary by using its powers to request information and access), and ensuring transparency about providers' safety practices, model misalignment, incidents, and near misses. Should there be revisions of the AI Act or GPAI Code of Practice, it should be a top priority to strengthen requirements for these provisions, as well as ensure that internal deployments are in scope of the AI Act. Moreover, the Commission should consider provisions for maintaining the interpretability and explainability of AI systems' internal reasoning. For example, it could restrict the use of training techniques that lead models to reason in a form not understandable to humans – sometimes referred to as 'neuralese'.

Recommendation at the national level

1. **Ensure that national Member State initiatives are complementary to Union-level enforcement.** The risks from highly capable and transformative AI will require cross-Union action. However, where Member States want to take action to address the severe and large-scale risks from general-purpose AI models, they should ensure that their efforts remain within the areas in which Member State actions are not precluded by the AI Act's harmonisation rules and do not duplicate the AI Office's enforcement actions. Instead, Member States should aim to identify areas where national capabilities can add value, complementing the EU's regulatory regime. For example, if setting up a national AI safety or security institute, Member States should cooperate with the AI Office to align on a common understanding of threat models, share information, and identify distinct focal areas that complement the AI Office's designated enforcement priorities. By working together, Europe can mount a coordinated defence against the risks of transformative AI.

O3.2 Prepare for labour market impacts

Transformative AI would likely displace human labour at levels never before seen in human history. As long as Europe only controls a small part of AI value creation, its ability to absorb such a shock and safeguard the livelihoods of its citizens is limited. While the exact pace and shape of AI's labour market effects remain uncertain, Europe should strengthen its economic position and build capacity to address future labour market disruptions. This matters both for its own sake and for maintaining public support for many of the measures recommended in this strategy.

The challenge

While uncertainty remains, the labour market effects of transformative AI would likely be unprecedented. Automation from previous technologies has often caused short-term or local disruption, but more jobs and prosperity in the long run. The current evidence on AI's labour market effects is consistent with this: Some studies suggest there is reduced employment among junior workers in the most AI-exposed occupations (Figure 11, Panel A), while effects on aggregate employment are not yet visible in the statistics (Figure 11, Panel B). However, this strategy assumes that AI could become transformative within the next few

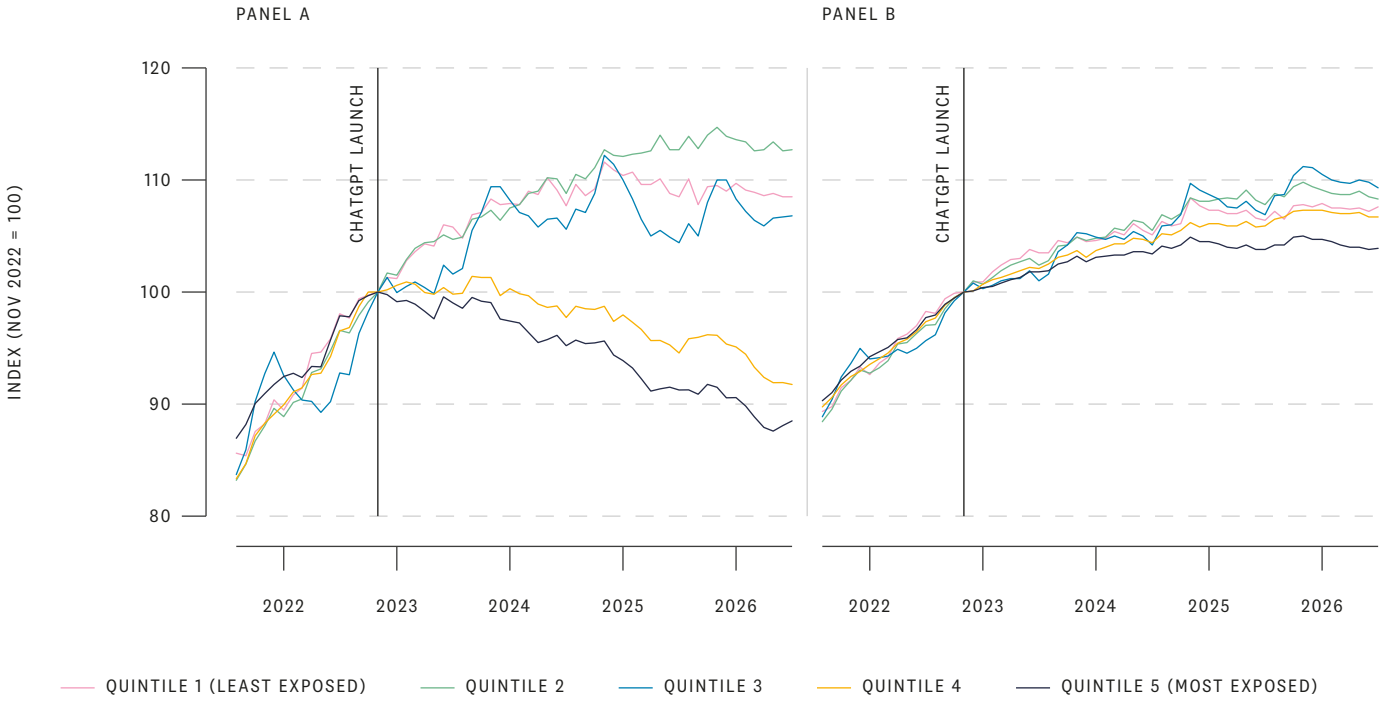

FIGURE 11

EMPLOYMENT INDEX BY AI EXPOSURE. Among workers aged 23 to 25, employment has fallen most sharply in occupations in the two highest quintiles of AI exposure (Panel A; Data). Overall effects of AI exposure on employment are not yet visible (Panel B; Data). All data is purely correlational, and does not rest on causal identification. Source: Stanford Digital Economy Lab, 2026

years, and that AI systems with greatly improved capabilities could change the economy and society faster than any previous technology. There are strong reasons to expect that, in such a world, AI systems would displace human labour at an enormous scale, though the exact pace of this change is uncertain. Even if transformative AI created many new jobs through increased economic growth, highly capable AI systems and robots might be able to fill these new jobs faster than humans could be retrained to perform them. While a few exceptions would plausibly remain – for example, jobs that for cultural or regulatory reasons society would only allow to be performed by humans – these are unlikely to sustain a large human workforce. Some researchers even expect that AI will be able to automate essentially all (cognitive and physical) labour and could drive human wages below subsistence levels. While these more radical views are not the consensus, a growing share of economists now believe that governments need to act fast to prepare for AI-driven labour displacement.

Large-scale labour displacement through transformative AI would be an existential threat to people's livelihood, sense of purpose, and social stability. Job loss is an immediate threat to an individual's ability to support themselves and their family. While previous shocks often hit older workers, AI-driven disruptions are likely to hit younger people first, making early retirement or wage insurance infeasible. And yet the challenge goes beyond income: since the Industrial Revolution, wage work has been a central source of human identity, social connection, and meaning. If transformative AI caused double-digit levels of unemployment, the social contract and the expectation that people can secure a livelihood through work could break down, potentially destabilising entire societies.

Europe is especially exposed to AI-driven labour market disruptions. First, Europe lacks direct insight into, and capacity to track, the economic effects of advanced AI, as policymakers only have limited visibility into the proprietary datasets of leading AI companies abroad. Second, while AI-driven automation threatens jobs in Europe, the US, and China alike, Europe would by default not benefit from commensurate gains in fiscal capacity. In a world with transformative AI, the surplus growth and tax revenue from broad-based AI adoption might flow disproportionately to countries with the strongest AI ecosystems, which are currently the US and China. Europe would thus lack the fiscal capacity to absorb labour market shocks in the same way as these countries. Third, Europe lacks unilateral regulatory control. Whereas AI-developing jurisdictions could slow down AI-driven labour displacement while still capturing large gains at the model development layer, Europe would bear the competitive cost of slower adoption without any such offset. This would expose its broader economic structure further to foreign competition, increasing the pressure on European workers even further.

Addressing the challenge

For future policies to be successful, it is essential to build the capacity to monitor and understand labour market effects. Even if one assumes that AI will indeed be transformative, it is difficult to take action on potential labour market effects now, as their exact shape and pace are highly uncertain. Therefore, policymakers should initially focus on building a better understanding of how transformative AI would affect labour markets, building monitoring capacity before disruption occurs. This requires comprehensive wage and headcount statistics as well as adding AI adoption and displacement indicators to official European labour statistics and ensuring they are published regularly (e.g. the European Statistical Office (Eurostat), the European Centre for the Development of Vocational Training (Cedefop), JRC). One important source of information and lead indicator for disruption is the usage data that frontier AI developers hold, including usage patterns, sector-specific use, and the balance between augmentation and automation. The Union could request and negotiate access to this for the purpose of producing official statistics. The UK's newly formed AI Economics Institute, for example, aims to closely collaborate with AI companies, and Anthropic, Google, OpenAI, and Microsoft have committed to collaborate with the UK government. In the US, the Department of Labor is similarly planning data-sharing agreements with major AI companies, such as OpenAI and Meta.

Labour market flexicurity will be helpful for managing AI labour transitions. The flexicurity approach combines high flexibility for companies to hire, dismiss, and restructure with strong income security and routes into new work for employees. Denmark has used this approach, and the European Commission released common principles for implementing flexicurity in 2007. Flexicurity's focus on keeping workers employed, employable, and supported through job transitions, rather than protecting specific jobs, makes it an apt framework for AI labour transitions and disruptions. If Europe tries to protect existing jobs in increasingly less competitive sectors at all costs, this will likely just defer the problem into the near future and increase the shock once it does occur. Companies could become reluctant to hire and lose ground to competitors making full use of AI capabilities, while workers are displaced in large numbers, rather than gradually when restructuring or business failures do occur. At least initially, flexicurity schemes should focus on high-earning professionals who typically have broader and more complex skill sets, which make it easier for them to transition into new and potentially very different jobs. As AI's impact on the labour market becomes more pronounced over time, policymakers might need to extend flexicurity to more workers where this is necessary to keep European companies in business.

Europe's geoeconomic position has to be strong to enable labour market interventions. It will not be possible for Europe to take appropriate policy action and retain optionality if it does not have a sufficient fiscal base from which it can operate and if it is outcompeted by more productive regions. In the past, Europe has been highly effective at designing safety nets and worker protections. However, this cannot be the sole focus for navigating a transition to a world with transformative AI. Europe must create the necessary preconditions and economic strength that will enable it to implement a host of interventions when AI leads to labour market disruptions down the line. Pillar 2 describes how Europe can build economic strength in a world with transformative AI, and Pillar 1 describes how it can use its assets to build leverage and negotiate access to frontier technology. These pillars are the basic elements that will secure Europe's geoeconomic position in a world with transformative AI.

Four concrete recommendations at the Union and national level that can help Europe prepare for the labour market impacts from transformative AI can be found below. Detailed recommendations for each are provided in Part B of this strategy.

Recommendations at the Union level

1. **Build a monitoring stack for AI labour market impacts.** By 2027, establish a Union-level monitoring framework that includes, among other things, regular labour statistics, job-posting data, AI-exposure measures, and information from AI companies.

2. **Fund and steer conditional flexicurity adjustment.** Encourage more flexible labour markets where adjustment to AI-driven disruption is already needed, combined with funding for retraining and income support for displaced workers. Prepare an EU-wide emergency support scheme if job losses accelerate sharply.

Recommendations at the national level

1. **Feed into the AI labour market monitoring stack.** From 2027, add questions related to AI-driven job displacement to labour force surveys and feed national survey and job vacancy data into the emerging Union-level monitoring.

2. **Prepare high-earner-first flexicurity reforms.** Ease dismissal rules for high earners, and prepare broader flexicurity reforms that combine flexible business restructuring with stronger income support and retraining, ensuring that displaced workers receive help within weeks.

What would be required for a successful European frontier AI project?

Many people are interested in the question whether Europe could and should aim to develop frontier AI domestically. A European frontier AI project, if successful, would be highly desirable for securing Europe's future prosperity, sovereignty, and control over the safe development of AI. It would also reduce Europe's risk of losing access to the frontier. But it would require substantially more political resolve and financial resources than following a Leverage Approach.

A European frontier AI project must not be half-hearted. A serious push for developing sovereign frontier AI is in principle possible, and desirable under the right conditions – but European leaders have to internalise the stark reality of what is required. Bold political statements and ambitions must be backed up by the extraordinary resources and resolve required to avoid failure. Importantly, a half-hearted European frontier AI project without a well-executed Leverage Approach is arguably the worst of all possible paths: It would leave Europe empty-handed, without either its own frontier model or reliable access to the frontier models of others.

The resources and resolve required for a European frontier AI project are exceptionally high. The required resources would vastly outstrip the current resources being dedicated to AI sovereignty efforts as well as any plans discussed in mainstream European AI policy. Every month, US hyperscalers are spending more than was spent on the entire Manhattan Project in today's money. This strategy estimates that a serious attempt to catch up to the frontier would cost approximately €800 billion over three years (for order-of-magnitude estimates and ranges see Table 1).

The political resolve required would also be immense. A successful European frontier AI project requires broad coalition-building, a wide-reaching reordering of European industrial priorities, and committing the majority of Europe's leverage to securing access to chips and defending against foreign intervention instead of securing other near-term goals (see Immediate Objective 1, Objective 1.1, Objective 1.2).

ITEM	DESCRIPTION	APPROXIMATE CENTRAL ESTIMATE (RANGE)
Sector buy-in	Buying some coalition AI companies outright, for their teams, intellectual property, and datasets.	~€25 billion (~€10–40 billion)
Compute	Accelerators and the data centres to house them: 16 GW of IT load built (roughly 18 GW of total facility power), of which roughly 9 GW of IT load would be running by 2029.	~€529 billion (~€320–640 billion)
Data centre operations	Electricity, maintenance, and taxes for running the data centres..	~€13 billion (~€11–43 billion)
Interim compute	Compute rented from others to bridge the gap until the coalition's own clusters come online (7 GW-years in IT load total).	~€105 billion (~€58–158 billion)
Training data and reinforcement learning (RL) environments	Expert human data and feedback, licensed content, training environments, and a one-off collection of books.	~€3.5 billion (~€2–6 billion)
Personnel and operations	Pay for the founders, senior researchers and ~3,000 further staff; equity buy-outs; and recruiting and operational costs.	~€34 billion (~€15–40 billion)
External coding agents	Access to frontier coding agents, bought from existing developers for at least the first 18 months.	~€3 billion (~€2–5 billion)

Political insulation	Payments shielding households, energy-intensive industry, and host regions from the project's effects, plus AI dividends.	~€80 billion (~€25–110 billion)
TOTAL		~€790 billion (~€445–1,040 billion)

TABLE 1 ORDER-OF-MAGNITUDE BACK-OF-THE-ENVELOPE CALCULATION (BOTEC) FOR THE COST OF THE FIRST THREE YEARS OF A EUROPEAN FRONTIER AI PROJECT. All estimates are approximate and should be considered as order-of-magnitude estimates. For further details and sources for these estimates, please see Part B.

A European frontier AI project would require adapting some elements of the transformative AI strategy. Most of the leverage-focused transformative AI strategy presented above remains in place if Europe wants to attempt a frontier AI project (for more detail see Part B). However, some recommendations would change in terms of substance or prioritisation:

- Regarding compute buildout (Immediate Objective 3), Europe should use its most attractive data centre sites to meet such a project's demands, with a substantial planned buffer. This limits the amount of scarce grid capacity that could be used for compute-for-access deals with frontier AI companies (Objective 1.1).
- Europe would have to use its leverage to buy large amounts of AI chips for a frontier AI project and secure access to state-of-the-art coding agents for its early stages, rather than use the same leverage mainly to secure long-term access to frontier AI models for broad economic integration (Immediate Objective 1).
- Proprietary data should be seen as a competitive advantage that can be pooled and used to train models, rather than an asset that can be exchanged for frontier AI access (Immediate Objective 1, Objective 2.2). European manufacturing companies should be encouraged to partner with the European project, rather than deeply integrate with foreign frontier AI developers (Objective 2.2).
- Once a European frontier project is producing sufficiently capable models, even if they still lag behind the frontier, there is a stronger case for European institutions to procure these rather than foreign frontier models, at least in cases where this does not result in severe security risks (Objective 1.1, Immediate Objective 2). In expectation, European institutions and companies would be using worse models for a few years.

The starting conditions, compute, talent, and sustaining conditions required for Europe to achieve a successful European frontier AI project are presented below and in more detail in Part B.

Starting conditions

Form a coalition of aligned countries. A successful European frontier AI project must begin with a coalition of value-aligned, liberal democracies working together – no country alone would succeed. Europe, as the largest established political bloc among likely participants and possessing the largest economic, industrial, and fiscal base among them, is the natural foundation of such a frontier AI project among trusted partners. Adding other countries to the coalition, especially so-called 'middle powers' facing constraints similar to those of the European Union, increases the chance of success since they can provide valuable assets (see Figure 12). Of course, in choosing members, a coalition has to consider the tradeoff between size and political cohesion, the risk of defection, and the potential actions of non-coalition countries. The more costly binding commitment mechanisms, the better.

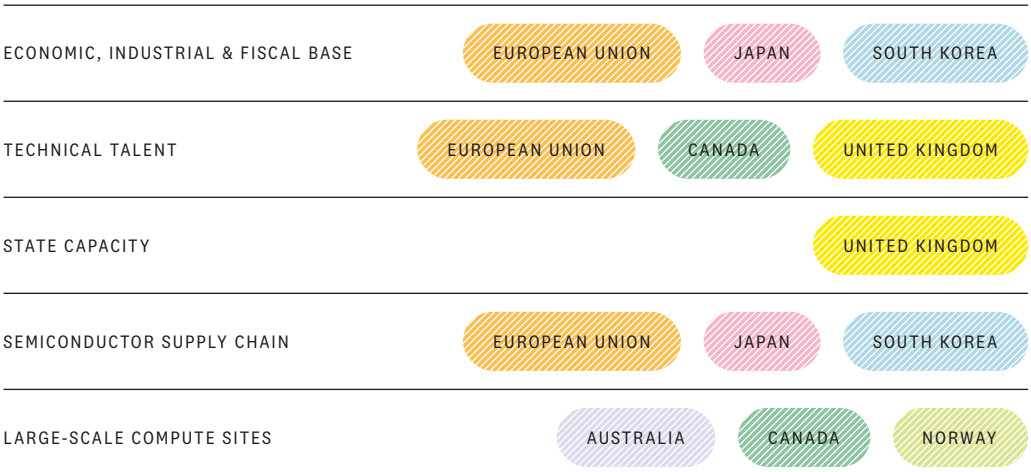

FIGURE 12 THE ASSETS AND STRENGTHS OF POTENTIAL COALITION PARTNERS. An assessment of the main assets each partner would contribute to a European frontier AI project. Source: *Own figure*.

Set up a new private company. A European frontier AI project would need a new[19] private company to be set up that receives full and sustained backing by the coalition's governments, including their national security apparatus (to ensure integration into the intelligence communities and their resources), trade authorities (to secure the supply chain), and treasuries (for financing).

Use previously successful institutional designs for developing frontier AI. This private company should broadly be modelled on existing frontier AI developers, since this is the only existing institutional form that has produced frontier AI models and systems, and because the coalition will not have time to experiment if transformative AI arrives soon. It should be led by a highly empowered leadership group with world-class, demonstrated research judgement, exceptional communication skills, and experience in building and scaling very large organisations. Like today's leading frontier companies, this private company has to be at once a research laboratory and an execution-focused company: Small teams of researchers need sufficient capacity and independence to explore novel technical approaches, while the majority of the company needs to be focused on pursuing the most promising path towards powerful AI systems.

Set up governance structures that enable innovation, safety, and speed. The private company's governance structure is mirrored by a complementary government structure that must remain separate from technical operations to prevent the type of bureaucratic slowing down of execution to which a conventional government programme would be susceptible. Countries should appoint ministerial or state-secretary level senior project leads supported by high-calibre frontier AI policy specialists, akin to the United Kingdom's Frontier AI Taskforce that became the first AISI, that would interface between the frontier project and the coalition's governments. Of course, a European frontier project has to comply with the systemic risk obligations for general-purpose AI models as set out in the AI Act and detailed in the Code of Practice.

Put in place sufficient public budget funding. A European frontier AI project could cost approximately €800 billion over three years (see Table 1). That funding should predominantly come from public budgets and dedicated debt instruments where treasuries, pension funds, sovereign wealth funds, and related institutions provide direct funding for the initial three-

19 A new entity would very likely be needed to attract the necessary talent, but this new entity could absorb interested companies from coalition countries that hold existing talent and technical expertise.

year costs as necessary. The reason for this is that Europe does not have a venture market at the necessary scale (or the time to build one) or the required private financing, while it does have the needed borrowing capacity. Coalition countries should expect to fund most of these upfront costs, with the expectation that private capital and significant revenue only follow later. Coalition governments should set up any required executive and legislative mechanisms. At founding, ownership should be fully public, while upon listing the coalition governments could remain anchor shareholders. This would allow the coalition to retain ownership and create a public asset and sovereign investment that can be privatised through European capital markets.

Establish commitment mechanisms. Coalition governments should precommit funding and infrastructure access for the complete setup period. Withdrawing from the coalition should mean that a member forfeits its contribution and stake, and no single country withdrawing should result in the project not having sufficient resources to continue. Technical progress needs to be independently assessed at fixed intervals with pre-agreed criteria for when the project should be discontinued, converted into its compute and talent assets, or allocated additional resources.

Compute

Build sufficient compute infrastructure and energy supply. A European frontier project would require substantial compute to compete, though likely somewhat less than a leading US developer requires over the same period due to lower customer-serving inference demand. It should aim for approximately 16 GW of AI compute capacity (IT load) at an expected cost of around €529 billion. Bringing this amount of compute online comes with a range of logistical challenges that a European coalition can nevertheless meet if it acts as it has previously in crisis conditions. This includes implementing accelerated infrastructure programmes in a range of host countries that can build out at scale and at great speed.

Procure interim rental compute. Until the dedicated clusters become operational, a European frontier AI project would have to procure rental compute by consolidating and requisitioning existing publicly owned supercomputers as well as rent compute on the open Graphics Processing Unit (GPU) market (at a combined cost of approximately €105 billion).

Set up a coercion shield. A European frontier AI project has to consider the difficulty that is likely to arise in acquiring sufficient AI chips given competition for a limited supply and potential export controls. To secure sufficient chip supply, the European frontier AI project should rely on agreements wherever possible, but prepare trade policy measures (see Immediate Objective 1, Union-level recommendations 1 and 2) that can capitalise on Europe's semiconductor supply-chain leverage, and are only activated reciprocally if exports are blocked. In addition, narrower export deals might be necessary to sustain the European project, which condition the continued provision of European semiconductor inputs on continued access to US frontier chips. However, this would be very risky since it commits significant amounts of Europe's AI-related leverage to the frontier AI project rather than to securing access to foreign frontier AI models. Participation in this coercion shield would also be costly for the European semiconductor companies whose export ability would be used as leverage. Europe could address this through a fund that would compensate any participating company or government that suffers substantial losses from this use of European leverage.

Talent

Attract leading frontier AI talent. To acquire the highest-quality researchers, the European frontier AI project must hire leading former frontier AI developer employees. Such employees have valuable expertise that is essential for a European frontier AI project. To hire them requires sufficient compensation and a credible mission. In total, expenditure for founders, senior researchers, the broad technical workforce, equity buy-outs, and operational costs could run to the order of €34 billion. Recruiting and empowering respected technical leaders

who can communicate, evidence, and shape the technical vision is the most effective way to ensure the mission is credible. Such founding group members are likely to be motivated by patriotism, a new technical challenge, operational independence, and credible signals that their work will be supported. They are integral for being able to hire the broader technical talent required to execute a European frontier AI project.

Secure access to frontier coding agents. A European frontier AI project must secure access to frontier-level coding agents before it develops its own. A substantial share of research work at frontier AI developers is already performed by AI agents that write code and complete R&D and administrative tasks under the direction of human researchers, and this is only likely to increase. Obtaining such access may be difficult[20] but it is likely possible if leading developers see commercial or reputational value in offering access, sufficient compute is secured to host AI agents on European infrastructure, and the budget for procurement is in place (approximately €3 billion over the first 18 months). In addition, such a project will have to budget for additional costs of approximately €3.5 billion for training data and RL environments.

Sustaining conditions

Remain robust to domestic political pressures. To be successful, a European frontier AI project must remain robust against political pressures, such as weakening motivation from coalition members and increasing public contestation due to the economic impacts and risks of AI. Efforts should be made to mitigate such pressures in advance, for example by measures such as committing to protect households from energy price increases caused by the project; direct AI dividends to affected groups; electricity subsidies for industries with high energy use; and community benefit schemes for data centre host regions.

Aim for successful long-term commercialisation. Once the frontier is reached, achieving some measure of commercial success will be required for continued success. The European frontier AI project will have to compete with the sophisticated product expertise, established consumer distribution, and lucrative enterprise contracts of established AI developers. Becoming the main provider of frontier AI to coalition governments, supplying private actors with high security requirements, and, where necessary, adopting measures to increase the use of European frontier AI by domestic businesses can help achieve successful commercialisation. A European frontier model could also find broad international appeal by being seen as more trustworthy and coming with fewer restrictions than those of leading competitor countries.

20 Anthropic, for example, restricts the use of its most capable models for frontier development, and other developers may follow suit.

PART B: Detailed Recommendations

Immediate Objectives

IO-1 Create a Member State Alliance for Supply Chain Security

WHY IT MATTERS

Reliable access to frontier AI models and compute ('frontier AI access') will be essential for Europe's security and economic prosperity in a world with transformative AI. At present, Europe is dependent on foreign providers for frontier AI access. While Europe does have assets that could strengthen its bargaining position with these providers, these are dispersed across Member States. A coalition of Member States that allows countries to pool their leverage can increase collective bargaining power in negotiations to ensure reliable frontier AI access.

Recommendations at the Union level

01	02	03	04
Make the ACI explicitly consider the Alliance's decisions	Ensure the Union trade commitments do not foreclose Alliance escalation options	Lead negotiations within the Union's competence, backed by the Alliance	Collaborate with trusted partners to add their assets to Europe's position
very high priority	high priority	high priority	high priority

01 Make the ACI explicitly consider the Alliance's decisions

ACTION

The Commission should explicitly prepare to use the ACI in cases related to frontier AI access, setting in place a framework that allows it to be deployed quickly in support of the Member State Alliance. To do this, the Commission should:

- (i) Adopt guidance allowing Alliance members to file submissions as duly substantiated requests under Article 4(1) ACI. This should include an evidence template built around the ACI's definition of coercion and consider the factors that the Commission and Council must weigh, including a pattern of interference. There should be an accelerated timeline for completing the examination of member requests that is substantially shorter than the ACI's current recommended four-month period in Article 4(2) ACI.

- (ii) Identify the appropriate Union response measures for a frontier AI access case, including worked examples drawn from the ACI's list of Union response measures in Annex I ACI and specifying the conditions under which response measures would be lifted. Additional measures can be considered, which may require an amendment to ACI Annex I.
- (iii) Prepare a 'fast path', meaning a shorter period for engagement with the third country (Articles 6 and 7 ACI), and readiness of response measures, in accordance with Article 8(10) ACI.
- (iv) Consider amending the ACI with additional restrictions on access to critical technologies as an explicit consideration in coercion assessment.

IMPLEMENTATION

- **Potential instruments:**
  - The ACI, especially Article 4, Article 5, Article 8, and Annex I.
  - New Commission notice or guidance detailing the conditions under which a request can be deemed 'duly substantiated'.
  - Amendment to the ACI through a legislative proposal.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Commission drafts the guidance and evidence templates for frontier AI access with the Alliance's secretariat.
  - **Q1 2027:** Publication of guidance, including internal worked examples of response measures.
  - **Q2 2027:** Conduct Alliance tabletop exercises for different application scenarios.
  - **Q3 2027:** Set the content and timing of an amendment to the ACI.
- **Success indicators:** Timely publication of guidance and template, along with worked examples and completion of tabletop exercise.

CONSIDERATIONS

- The ACI's coercion test is based on intent; if a nation restricts access to a critical technology but does not demand anything from the Union or a Member State in exchange for access to the technology, this action likely does not pass the test. As such, the US government's June 2026 suspension of Fable 5 would probably not have passed the test on its own. However, if there were a documented pattern of the US linking frontier AI access to European policy choices, this would pass the test. If there is no pattern of this kind, Union-level recommendation 2 below may need to be applied instead.
- Annex I ACI lists restrictions on the export of goods, including goods subject to export control, as a permitted response measure. Therefore, an equipment measure of the kind described under national recommendation 2 would not require an Annex I amendment.

02 Ensure the Union trade commitments do not foreclose Alliance escalation options

ACTION

Trade policy is an exclusive EU competence. This may constrain the Alliance's escalation options, with tariff legislation determining how quickly Europe can respond to an access restriction. New trade commitments should not foreclose the Alliance's escalation options, and the EU should resolve conflicts with existing commitments in a way that keeps escalation options open. The Alliance should have legal clarity on how a potential restriction would be interpreted at an EU level. To do this, the Commission should:

- (i) Audit existing EU trade commitments and instruments against each escalation step mapped out in national recommendation 2, recording results in a register. For any conflicts identified, decide and record how they would be resolved, for example by amending, reserving, or allowing the commitment to lapse at its next review. At least one lawful route for each step of the Alliance's escalation chain should be retained.

- (ii) Together with the Alliance, prepare the suspension mechanism outlined in Article 3(1), points (b) and (c), of Regulation (EU) 2026/1455 for frontier AI access cases. The Commission could also issue guidance to ensure that assessments consider restrictions on frontier AI access to EU operators as relevant and draft the corresponding implementing act in case restriction takes place, including which of the preferences under Articles 1 and 2 would be suspended first. The preferences on AI inputs should be retained so that Europe's own buildout is not harmed by a suspension.

- (iii) Prevent new commitments from foreclosing steps outlined in the escalation chain. This could be achieved by inserting new language into negotiating mandates under Article 218 of the Treaty on the Functioning of the European Union (TFEU) and in the work plans of economic security dialogues to include security exceptions pertaining to frontier AI access. This would make sure that escalation measures taken by the Alliance to maintain its essential security interests do not put it in breach of the agreement.

- (iv) Propose EU legislation under Article 207(2) TFEU that gives Member States discretion to activate steps in the escalation chain.

IMPLEMENTATION WORKSTREAMS

- **Potential instruments:**
  - Regulation (EU) 2026/1455, Article 3(1) and its suspension mechanism. A Commission implementing act can suspend the preferences granted by Articles 1 and 2 under the examination procedure of the Trade Barriers Committee. The examination procedure should be based on substantiated information, which could be supplied via a Member State from the Alliance.
  - Commission notice or guidance on the possibility of using Article 3(1) of Regulation (EU) 2026/1455 for responding to frontier AI access restrictions.
  - Inserting standard language into the negotiating directives under Article 218 TFEU, as well as into the annual work plans of the EU's economic security dialogues.
- **Potential first steps (next 12 months):**
  - **Q2 2027:** Commission begins an audit of potential conflicts with the Alliance's escalation chain.
  - **Q3 2027:** Completed register of all conflicts with the escalation chain; issued guidance on what counts as a frontier AI access restriction; draft implementing act(s); standard language is approved for use in negotiating mandates and dialogues; final decision on whether legislation under Article 207(2) TFEU is required for any steps in the Alliance's escalation chain.
- **Success indicators:** Completed register of conflicts with escalation chain; guidance published and implementing act drafted; standard language around security exceptions inserted into all new negotiating mandates; each step of the escalation chain has a lawful route.

CONSIDERATIONS

- Common commercial policy is exclusively an EU competence, and any existing national discretion here is usually read narrowly by the European Court of Justice. New national discretion related to frontier AI access restrictions will need to be legislated first.
- Under Article 3 of Regulation (EU) 2026/1455, the EU can only withdraw the preferential treatment granted by the Regulation under its own Articles 1 and 2 and Annexes I to III. As a result, while ordinary tariff rates can be returned to goods covered by the Regulation, tariff rates cannot affect services, impose new duties, or exceed its annexes. The deterrence value of the Regulation is therefore the commercial value of the preferences set out in the Regulation. Measures beyond changing the Regulation's existing penalties require the ACI (see Union-level recommendation 1 above).

03 Lead negotiations within the Union's competence, backed by the Alliance

ACTION

The Commission should lead negotiations on Europe's frontier AI access, according to its competence, on the basis of the position granted by the Member State Alliance (national recommendation 3). A breach of contractual assurances given to the Union or Member States should also be treated as grounds for considering the response steps outlined under Union-level recommendations 1 and 2. To do this, the Commission should:

- (i) Negotiate assurances on frontier AI access within existing trade and economic security dialogues and in formal agreements. The assurances should include at a minimum: consultation with the Commission before any restrictions affecting EU operators take place; continuity of frontier AI access for particular European public entities during any restricted periods; and eligibility of EU entities for vetted access programmes.
- (ii) Attach these assurances to any agreements on compute buildout, to ensure that EU investment commitments are matched by access commitments.
- (iii) State explicitly in agreements what would constitute a breach and that a breach is grounds for considering the response steps outlined in Union-level recommendations 1 and 2, as well as for the escalation chain in national recommendation 3.
- (iv) Designate a Commission liaison for the Member State Alliance to ensure national preparations inform negotiations taken at the EU level.

IMPLEMENTATION

- **Potential instruments:**
  - Commission decision to designate a liaison to the Alliance.
  - The European Blueprint for structured access to advanced AI capabilities for cybersecurity purposes, which is due from the Commission and ENISA in late 2026. This can serve as the template for trusted access arrangements.
  - Trade agreements and economic security dialogues to negotiate assured access.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** The Commission designates a liaison and compiles an inventory of EU assurances already given or received on frontier AI access across existing frameworks.
  - **Q1 2027:** Draft assurance requests based on the Alliance's position; agree a breach clause template with the Legal Service.

- **Q2 2027:** Raise requests with relevant governments within existing dialogues.
- **Q3 2027:** Include breach clauses in new agreements on access or compute.

- **Success indicators:** Commission liaison in place; inventory of existing assurances completed; first frontier AI access assurance obtained and first agreement containing a breach clause.

CONSIDERATIONS

- An assurance from a foreign government promising continuity of frontier AI access in all circumstances is likely not obtainable. Instead, negotiations should focus on procedure: foreign governments must consult EU counterparts before placing any restrictions affecting EU operators, provide a notice period, and create avenues for EU public bodies to obtain access to restricted models. This should include restrictions framed as security measures, such as the June 2026 Anthropic Fable restrictions.
- Negotiating counterparts will likely ask for comparable assurances from the EU in return, such as predictable licensing and servicing of European chip-making equipment or the removal of procurement criteria they consider discriminatory. National recommendation 2 provides some guidance on the mirror commitments that European negotiators should be prepared to give.
- Assurance will only be robust if the EU responds to breaches in a credible way. For a breach clause to be effective, the first steps in the Alliance's escalation chain should be relatively low cost.
- Purchase intentions and investment commitments have already been considered in the EU's 2025 trade frameworks, without being attached to matching frontier AI access terms. Member States are doing the same bilaterally, for instance by signing the Pax Silica agreement. Commitments of this sort could make reliable assurance harder to obtain (see national recommendation 3).

04 Collaborate with trusted partners to add their assets to Europe's position

ACTION

The EU should build a coalition of trusted non-European partners that can add to Europe's bargaining position and receive benefits in return, for example, through their technology, computing capacity, energy, or critical minerals. While they would not be part of the Alliance's closed track, the Commission should negotiate reciprocal arrangements with these partners, such as access parity and continuity commitments or early warning about restrictions, and coordinate when restrictions from frontier providers occur. As many of the relevant partners and the EU itself signed the June 2026 Pax Silica declaration, the Commission should decide whether the EU's position should be pursued inside this framework or alongside it.

IMPLEMENTATION

- **Potential instruments:**
  - EU Digital Partnerships and the annual Digital Partnership Council work plans (which include Japan, South Korea, Canada, and Singapore).
  - The G7 Coordination Platform on Economic Coercion; bilateral economic security dialogues.
  - The Pax Silica declaration, whose signatories include (along with the EU and certain Member States) Australia, Japan, Norway, South Korea, and the United Kingdom.

- The Strategic Partnerships on Semiconductors outlined in the Chips Act 2.0 proposal (in the June 2026 Tech Sovereignty Package).
- The international cooperation provision of the ACI, under which the Commission may consult or cooperate with other countries affected by similar economic coercion.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Assessment of partner assets and dependencies; decision on whether the Commission will work inside or alongside Pax Silica.
  - **Q1 2027:** Frontier AI access and resilience deliverables placed in the 2027 Digital Partnership Council work plans.
  - **Q2 2027:** First reciprocal operational arrangement agreed with external partners.
  - **Q3 2027:** First joint exercise with a partner coordinating on how to respond to a frontier AI access restriction.
- **Success indicators:** Position on Pax Silica adopted; live operational arrangements and exercises with external partners.

CONSIDERATIONS

- Engagement with partners can take place bilaterally through Member States as well as through the Commission.

Recommendations at the national level

01	02	03
Found the Member State Alliance for Supply Chain Security by the end of 2026	Use the Alliance secretariat to map actual and projected bottlenecks and leverage	Use the Alliance to agree rules of engagement when negotiating frontier AI access
very high priority	very high priority	very high priority

01

Found the Member State Alliance for Supply Chain Security by the end of 2026

ACTION

As Europe's AI supply chain assets are held by different Member States, with different owners and different national legal contexts, they have never been leveraged in tandem, and individual Member States who try to exercise leverage face individual retaliation. A standing Alliance holding these assets could act collectively to improve their collective position. Member States should:

- (i) Found a Member State Alliance to negotiate European frontier AI access. This Alliance should, in the first instance, include the Netherlands, Germany, and France as the Member States hosting high-value AI supply chain assets. The Alliance should act as one party and respond together if frontier AI access is restricted. The Alliance should be open to any EU state that contributes assets that are supply chain chokepoints, funding for the Alliance's staff, or a political commitment to act in tandem with collectively agreed Alliance decisions, on the understanding that the Alliance prepares and the EU executes trade measures.
- (ii) Convene prospective Alliance members to sign a founding declaration. This should specify that decisions will be weighted by contribution. The Alliance should have an open track (with observers) and a closed track for members only. If the Alliance wants to respond to restrictions, this will happen through EU instruments, reflecting exclusive EU competence over commercial policy and trade.
- (iii) Staff a secretariat (of around 8 to 12 experts) for the Alliance seconded from national ministries, including experts in supply chains, export controls, and trade law. Task the secretariat with implementing the work plan set out in national recommendations 2 and 3. Request a Commission liaison as point of contact (see Union-level recommendation 3).

MEMBER STATE MODE

Coalition, including Member States with relevant AI supply chain positions. Other Member States may join under the contribution rules specified above.

IMPLEMENTATION

- **Potential instruments:**
  - The Alliance can be founded through an intergovernmental declaration, letter of intent, or memorandum of understanding outside of the EU Treaties. A possible model is the Semicon Coalition announced by the Dutch government in 2025 or the 2018 European Intervention Initiative.
  - National secondment and budget lines for secretariat.
  - Classified information sharing agreement for members.
  - Designated Commission liaison (see Union-level recommendation 3 above).
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Agree on the founding declaration with core Alliance members, setting out the purpose, membership rules, and position on Pax Silica. Invite prospective members, agree a secretariat budget, and designate a Commission liaison.
  - **Q1 2027:** Fully staffed secretariat and classified information sharing arrangement in place; first work programme adopted and started.
  - **Q3 2027:** Review of membership and contributions.
- **Success indicators:** Founding declaration successfully agreed and signed by core members, with Alliance membership including countries covering lithography, optics, research, and compute. Mapping (Union-level recommendation 2) underway and secretariat staffed up.

CONSIDERATIONS

- Commercial policy is an EU competence. The Alliance must therefore operate as a political coalition that coordinates, prepares, and advocates, rather than being the vehicle through which responses to restrictions are taken. These responses will instead come through EU instruments, with Member States only acting within the discretion they have under EU law (see national recommendation 3 below).

- Since cooperation comes first, the foreign country on which Europe's access most depends should be offered an observer seat in the Alliance's open track (the negotiating position, pooled procurement, partnerships, and assurances).
- It is not in Europe's interest to take an aggressive posture that could be seen as undermining free trade or attempt to interfere with foreign domestic regulation of AI development. The Alliance needs to restrict itself to reacting only to genuine coercive action.

02 Use the Alliance secretariat to map actual and projected bottlenecks and leverage

ACTION

The Commission already maps EU semiconductor dependencies and bottlenecks under the Chips Act, while Member States have taken part in collective risk assessments on semiconductors and AI. There is not, however, a record of where other countries depend on Europe, or a mapping of who may lawfully decide to use an asset for leverage or what using it would cost. This is a gap that the Alliance would aim to fill.

- (i) Use the Alliance secretariat to produce a classified map of Europe's leverage points in the AI supply chain. This could include, *inter alia*, lithography machines, optics, and research on semiconductors, as well as computing capacity and energy. This map should include a record of which foreign actors depend on which assets, how quickly this dependency can be reduced, the time lag of any restrictions, whether the leverage is likely to decrease as AI becomes transformative and as other states reduce their dependencies, which lawful authorities may decide to restrict access, and how third countries could respond (for example through tariffs or cloud access restrictions).
- (ii) On the basis of this mapping, create a potential escalation chain for each asset. This could include, for instance, reduced servicing of equipment; reduced research collaboration; restricting access to public procurement; or placing conditions on export licensing. The chain should include an estimated cost and effect on Europe of each step and whether the necessary action should be taken at European or national level.
- (iii) The map should be updated continuously.

MEMBER STATE MODE

Coalition, including Member States with relevant AI supply chain positions (e.g. the Netherlands, Germany, Belgium, France, Austria); others can join under the contribution rules from national recommendation 1.

IMPLEMENTATION

- **Potential instruments:**
  - The strategic mapping prepared via Articles 19 and 20 of the Chips Act, shared with Member States and the European Semiconductor Board.
  - Collective risk assessments on advanced semiconductors and AI carried out with Member States under Commission Recommendation (EU) 2023/2113.
  - The Commission could be invited, through the European Semiconductor Board, to task the JRC with extending the monitoring under Articles 19 and 20 of the Chips Act with reverse dependency indicators, showing where third country supply chains depend on EU suppliers.
  - National inventories of operators in critical supply chains maintained under the Internal Market Emergency and Resilience Act, and the supply risk monitoring and stress tests under Articles 20 to 25 of the Critical Raw Materials Act for materials inputs.

- The European database on data centres under Article 12 of the Energy Efficiency Directive and Delegated Regulation (EU) 2024/1364, for computing capacity and energy.
- National mandates to economic security units and export licensing authorities.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Core Alliance members agree on the scope and purpose of the mapping and how classified materials will be handled; the completed Chips Act mapping results are obtained through the European Semiconductor Board; finish a complete first inventory of assets and of who may lawfully decide levers and escalation steps that make use of them.
  - **Q1 2027:** Completed interviews with companies and assessment of dependencies, substitution time, leverage longevity, and retaliation channels; drafting of detailed escalation chains.
  - **Q2 2027:** An escalation chain, based on a mapping of dependencies and available levers, is approved and circulated to Alliance members, with gaps and risks communicated to national governments.
  - **Q4 2027:** First six-month refresh is completed.
- **Success indicators:** Escalation chain is approved by the steering group and circulated to members by the end of Q2 2027, with every asset assessed. Levers list adopted as the input to national recommendation 3 by Q3 2027. First refresh completed by the end of 2027.

CONSIDERATIONS

- The Commission has mapped EU dependencies on other countries for semiconductors under Articles 19 and 20 of the Chips Act. The recommended mapping extends its scope and documents how other countries depend on Europe, as well as how assets could gain or lose value over time, especially as transformative AI arrives. It should be an annex to the Chips Act mapping.

03 Use the Alliance to agree rules of engagement when negotiating frontier AI access

ACTION

For leverage to be credible, a foreign counterpart must know in advance that the Alliance will act together. It must believe that Alliance members cannot be negotiated or settled with individually, that a Council majority for an EU response exists, and that the costs of retaliation will be shared among members. The Alliance should fix these conditions in advance of any frontier AI access restriction incident that it needs to respond to. To do this, the Alliance should:

- (i) Agree a common negotiating position towards the foreign countries on which Europe's frontier AI access depends. This position could include, for instance, the starting offer that Europe should make; the escalation options developed under national recommendation 2 with the legal basis for each; and which level (EU or national) should act with which instrument at each step of the chain. Once adopted, this position should be transmitted to the Commission to apply this position in negotiations with the foreign counterparts.
- (ii) Agree a consultation process for applying national measures, such as export licensing on chokepoint items, and for bilateral arrangements on frontier AI access offered to individual members, to ensure that members consult with each other before acting and commit to follow a joint position when one is reached. Members should not settle national controls on their own without consultation.

- (iii) Support the fast-tracking of the EU response to coercion. Agree on an Alliance standard for assessing cases that fits the coercion definition in the ACI. Commit publicly to supporting a Council determination once that standard has been met, and work to secure a qualified majority in the Council (which may require support from non-Alliance Member States). Request the Commission to set a short period for engagement.
- (iv) Agree on how the costs of retaliation by a third country will be shared, to prevent one member from being exposed to individual pressure. For example, the burden could be allocated according to members' economic size, similar to EU budget contributions, with extra considerations for members that contribute chokepoint assets or are particularly exposed to potential retaliation.

MEMBER STATE MODE

Coalition, including Member States with relevant AI supply chain positions (e.g. the Netherlands, Germany, Belgium, France, Austria); others are able to join using the contribution rules from national recommendation 1.

IMPLEMENTATION

- **Potential instruments:**
  - National controls on non-listed items adopted and notified under Articles 9 and 10 of the Dual Use Regulation, which limit them to public security and human rights grounds.
  - The Commission can examine a duly substantiated request and the Council then determines coercion by qualified majority within eight weeks under Article 5, and response measures may in justified cases apply without a prior call to cease under Article 8(10). Members of the Alliance can commit to their votes in advance via a political declaration coordinated in the Council's Trade Policy Committee.
  - An intergovernmental cost-sharing agreement, on the model of the 2014 agreement on contributions to the Single Resolution Fund.
- **Potential first steps (next 12 months):**
  - **Q1 2027:** Finance ministries draft the burden-sharing agreement. A consultation process and an Alliance standard for assessing frontier AI access cases in line with the coercion definition of the ACI are drafted.
  - **Q2 2027:** Members agree on the template for the rules of engagement based on the mapping obtained via national recommendation 2.
  - **Q3 2027:** Members agree on the burden-sharing agreement, a consultation process, and the standard for meeting the coercion test. Members commit publicly to supporting a Council determination once that standard has been met and to work to secure a qualified majority in the Council.
- **Success indicators:** Position adopted and transmitted to the Commission by Q3 2027; written support for a Council determination from all members of the Alliance, provided that all members find that a particular access restriction constitutes coercion. Cost-sharing agreement signed by all members of the Alliance.

CONSIDERATIONS

- Member States may only introduce national export controls where EU law permits, as common commercial policy is exclusively an EU competence. The Dual-Use Regulation permits export controls for public security or human rights reasons, but the European Court of Justice has tended to read such national grounds narrowly.

- The ACI's coercion test requires that restrictions carry the intent of preventing or modifying an action by the EU or a Member State. This may be shown, for example, through a documented pattern of linking access decisions to previous European policy choices. If no such pattern can be shown, the EU can otherwise trigger a suspension clause in relevant trade agreements. In cases where the coercion test is met, the Alliance should work to reduce the time it takes for the EU to respond. It could do so by filing a request that specifies the measure, pattern, and injury caused, as well as by having the text pre-agreed in Coreper so that the Council Presidency can table the decision at its next meeting as adopted with minimal discussion. The Alliance can also ask the Commission to set a short period for engagement or, in certain cases, to apply response measures without a prior call to cease (Article 8(10) ACI).

IO-2 Make Europe's institutions ready to act in a world with transformative AI

WHY IT MATTERS

Implementing this strategy depends on institutions that can track, decide, and deliver at the pace of frontier AI. In a world with transformative AI, governments will need to deliberate, make, and act on decisions regarding AI at pace. This will require governments having informed decision-makers and advisors, access to information and data regarding AI, and the ability to use capable AI systems to augment their work.

Recommendations at the Union level

01	02	03	04
Ensure that the Commission President and College of Commissioners are well-informed on transformative AI developments and are able to execute on strategic objectives	Empower the AI Office as an invaluable centre of expertise on AI	Rapidly expand access to frontier AI across EU institutions	Convene leaders, officials, and subject-matter experts to ideate, plan, and action the creation of new institutional structures needed for transformative AI
very high priority	very high priority	very high priority	high priority

01 Ensure that the Commission President and College of Commissioners are well-informed on transformative AI developments and are able to execute on strategic objectives

ACTION

Ensure that top political decision-makers are well-informed about the latest trends and developments in transformative AI and their potential implications. While there are existing special advisors to Commissioners, this is on a part-time basis of 25 days per year on average. As AI becomes an increasingly important topic for high-level decisions, this part-time arrangement may prove insufficient for keeping leadership informed of fast-moving developments. Leadership's awareness of developments in AI could be improved by employing a full-time, technical frontier AI advisor to the President. This advisor should be an individual with relevant experience of frontier AI, for example, through having worked at a frontier AI

provider. They should be called upon to input on high-level deliberations relevant to transformative AI, and be provided with a direct, two-way line of communication to the AI Office. The Commission should also establish a Transformative AI Task Force of approximately 15 staff to support the advisor. This task force should be given a mandate that includes driving implementation of this strategy's recommended strategy across Directorates-General, running demonstrations and briefings for senior officials, and coordinating with the EEAS and relevant security structures.

IMPLEMENTATION

- **Potential instruments:**
  - The advisor function and the Transformative AI Task Force can be established via executive decision.
  - Articles 14–15 of C(2025) 4716 allow for Principal Advisors to be employed in the Secretariat-General as 2(a) temporary agents, on an AD14 pay grade.
  - Communication channels between political leadership and the AI Office can be established without formal action.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** The Commission President announces the appointment of the technical frontier AI advisor. Recruitment efforts commence, with predefined conflict of interest rules for candidates previously employed by frontier AI providers.
  - **Q1 2027:** Advisor appointed and announcement published.
  - **Q2 2027:** The Task Force enters operation, staffed by both secondees and new hires. The first briefing to the Commission President, Commissioners, and Directors-General is held by the technical advisor.
- **Success indicators:** Appointment of the President's technical advisor made. Task Force employees and secondees in post. Quarterly briefings held.

CONSIDERATIONS

- The Task Force should be tasked with coordinating and tracking policy delivery rather than duplicating the technical and regulatory functions of the AI Office. It should provide a quarterly stocktake of progress on implementation across DGs to the Commission President and relevant Commissioners.
- The AI Office, as the Commission's centre of AI expertise, should have a direct reporting line to the advisor, with confidentiality arrangements in place in line with Article 78 of the AI Act.
- The primary value of the technical advisor will be their technical credibility and hands-on experience with frontier AI. The selection process should prioritise finding candidates with these qualities.

02 Empower the AI Office as an invaluable centre of expertise on AI

ACTION

Three immediate actions should be taken to cement the AI Office's status as a centre of subject-matter expertise within the Commission. Firstly, direct, two-way lines of communication should be established between senior Commission officials and subject-matter experts within the AI Office. Existing in-house expertise should be utilised to inform key decision-makers through briefings, workshops, and regular communication. Secondly, AI Office hiring should be reformed so that it can hire and utilise world-class subject-matter expertise. Concrete actions include: reducing the time from first interview to offer to six weeks or less; hiring employees as temporary agents rather than contract agents, so that they can serve beyond

a six-year term; standing up a dedicated headhunting capability with a talent-scouting remit; increasing pay grades beyond typical civil servant levels for technical hires; hiring on merit alone, using nationality derogations for outstanding non-EU technical hires; and allowing more flexibility for remote work. In parallel, the AI Office should be provided with sufficient administrative and operational staff, so that scarce technical talent can focus on technical tasks where their talent is best used. Thirdly, the AI Office should be provided with dedicated additional funding to provide staff with large amounts of compute and API credits for using frontier AI systems to assist their work.

IMPLEMENTATION

- **Potential instruments:**
  - High-level AI expertise could be employed as Article 2(a) temporary agents under the Conditions of Employment of Other Servants of the European Union (CEOS), rather than as Article 3(b) contract agents, allowing for permanent contracts.
  - Pay grades of AD7–AD11 should be used for technical hires, applying exemptions to years of prior experience under Article 13(4) of C(2025) 4716 where necessary.
  - Nationality derogations are possible under Articles 12(2)(a) and 82(3)(a) of the Staff Regulations.
  - The Large-Scale Review can be used as a framing vehicle for piloting flexible recruitment.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** The AI Office and DG HR agree on an expedited hiring pathway for urgent hires and review the default contract terms for AI Office recruitment. The Large-Scale Review's final recommendations (to be published by the end of 2026) name AI Office recruitment as a pilot for flexible technical hiring. DG BUDG requests additional temporary-agent posts in the 2027 draft budget.
  - **Q1 2027:** Headhunting team operational.
  - **Q2 2027:** First cohort of hires employed as temporary agents, and offers given to current staff employed as contract agents to convert to temporary agents.
  - **Q4 2027:** First-year recruitment metrics published.
- **Success indicators:** Median days from first interview to offer. Share of expert hires engaged as temporary agents. 12-month retention of senior staff.

CONSIDERATIONS

- In transformative AI scenarios, the required spending on compute and API credits for AI Office staff could be at least as large as employee salaries.
- The current form of fixed-term contract agent employment, which is capped at six years, will eventually lead to the AI Office losing senior expertise, starting in mid-2030.
- While pay grades that are competitive with industry are unrealistic, parity with peer public institutions, such as the UK AI Security Institute (UK AISI), is achievable and should be the minimum target.
- The dedicated headhunting team should make sure to be aligned with the team leads regarding the desired profiles of new hires.

03 Rapidly expand access to frontier AI across EU institutions

ACTION

In order to keep pace with consumers and firms deploying increasingly powerful frontier AI tools, European institutions need access to these same capabilities. Building on initiatives like GPT@EC, the Commission should ensure there is minimal lag between the capabilities of tools made available to EU staff and the frontier. Model access should be secured across the service through a central procurement process. Common rules for data classification, security, sovereignty, and permitted use should provide clear legal guidance for acceptable use. In addition, the European Data Protection Supervisor (EDPS) should issue guidelines to facilitate safe and legal AI-assisted and automated decision-making in areas vulnerable to service flooding to increase resilience.

IMPLEMENTATION

- **Potential instruments:**
  - DIGIT could use framework contracts to centrally procure frontier AI tools for use across the Commission. Procurement conditions should ensure that providers do not have access to sensitive data per the Commission's sovereign-cloud framework. Procurement should be inter-institutional from the outset, ensuring that the Parliament, Council, and other relevant institutions and bodies of the Union can be contracting authorities. Models should be chosen on capability regardless of origin, with the proposed CADA Article 32 preferences only being applied in cases where EU models are competitive with the frontier AI models.
  - The Commission could invite the EDPS to develop guidance on AI-assisted decision-making and on the conditions for implementation of automated decision-making in accordance with Regulation (EU) 2018/1725.
  - The EDPS sandbox for the institutions under AI Act Article 57(3) could be used to develop tools to support automated decision-making within the European institutions, such as oversight design, monitoring requirements, and an appeals process.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** The Commission updates guidance for the procurement of AI models, with a focus on eliminating the gap between the tools available to institutions and those available to the public. EU institutions identify services that are most vulnerable to AI-driven service flooding.
  - **Q1 2027:** The AI Office coordinates with EDPS on sandbox testing.
  - **Q2 2027:** Reduce duplication and pool resources through joint projects with Member States.
- **Success indicators:**
  - Reliable access to frontier capabilities, as measured by the time lag between public frontier model release and internal roll-out of the model by the Commission.
  - EDPS guidelines for AI-assisted decision-making published.
  - Monthly active users of catalogue tools as a share of all staff, to measure adoption.

CONSIDERATIONS

- There are barriers to current procurement mechanisms reliably guaranteeing access to frontier AI systems, for example the length of time that procurement awards generally take. Procurement to guarantee frontier AI access would require faster timelines to approval.
- Procurement contracts should include clauses on exportability of data and workflows in order to prevent lock-in or dependence on a single provider.

04 Convene leaders, officials, and subject-matter experts to ideate, plan, and action the creation of new institutional structures needed for transformative AI

ACTION

Hold a series of convenings that bring together senior Commission officials, AI Office experts, and independent subject-matter experts to brainstorm, plan, and action the creation of new institutional structures that will prepare European institutions for transformative AI. This could include institutional structures that improve foresight and strategic awareness, such as an AI Strategic Foresight Unit, staffed with dedicated frontier AI experts, created within the AI Office, the Secretariat-General, DG IDEA, or elsewhere; Transformative AI Fellows appointed to DG IDEA; or the seconding of AI Office experts or hiring of other frontier AI experts to EU delegations and offices in San Francisco and Beijing. But there are likely other purposes and more ambitious institutional structures that should be considered (e.g. information sharing, emergency preparedness, evaluation capacity, incident investigation, a Union-level AI agency). Clear ownership should be assigned to an individual or group of individuals to push forward this process. The process could be facilitated by the AI Office, reporting to Commission leadership, by planning what sessions and inputs will be needed to come up with ideas and who needs to be in the room to best brainstorm, plan, and action such new institutional structures.

IMPLEMENTATION

- **Potential instruments:**
  - Internal communication by owner and Commission leadership signalling intent to convene discussions on this topic.
  - Informal internal convening of relevant stakeholders.
- **Potential first steps (next 12 months):**
  - **October 2026:** Assign ownership for planning and hosting the convenings.
  - **Q4 2026:** Collect input, plan session formats, and send out initial invitations.
  - **Q1 2027 onwards:** Hold series of planned convenings.

CONSIDERATIONS

- Such convenings likely need to happen in stages to ensure ideation, planning, and actioning stages occur sufficiently and that the right individuals are present for each session.
- Collecting input from Commission and AI Office officials will be a valuable source of institutional knowledge and expertise for current institutional bottlenecks and frictions.
- Evaluating the successes and failures of other frontier AI institutional structures in other countries or regions will be a valuable exercise.
- Frontier AI and other subject-matter expertise (e.g. security, emergency preparedness, intelligence, and national security) should be drawn on from within the Commission and European institutions. Where needed external expert consultation can strengthen ideation and implementation processes.

Recommendations at the national level

01	02	03	04
Build national technical AI assessment capacity and frontier AI expertise	Appoint senior frontier AI advisors with direct access to the head of government	Collect, analyse, and publish national statistics on AI diffusion and adoption	Ensure a baseline of AI literacy across all government departments
high priority	very high priority	high priority	high priority

01 Build national technical AI assessment capacity and frontier AI expertise

ACTION

Establish or scale an in-house national technical unit for assessing frontier AI and advising government. The designated unit should be tasked with providing senior decision-makers with assessments of frontier AI progress, impacts, and policy implications. The unit should actively maintain links to other Member States' analogous bodies, national security authorities, and the EU AI Office. These Units should be given autonomy and flexibility in hiring and staffing, including derogations from national civil-service pay grades, ring-fenced funding, and access to compute and AI tools in order to attract, empower, and retain top-level subject-matter staff.

MEMBER STATE MODE

All Member States, with unit size proportional to Member State size. All Member States should establish, at a minimum, a core unit of 5–10 subject-matter experts. Some Member States may wish to expand this to a national AISI, such as the one recently announced by Germany. Neighbouring Member States may wish to pool resources and form regional partnerships (e.g. Nordics, Baltics).

IMPLEMENTATION

- **Potential instruments:**
  - Options include a ministerial decision to establish a new unit or extend an existing government body, legislative action, or formal partnership with existing research institutes, with specifics depending on national institutional arrangements.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Identify and designate a suitable host institution for the advisory unit. Secure a designated operating budget for the next national budget cycle.
  - **Q1 2027:** Appoint founding leadership positions through an open international search, with pay-derogation instruments adopted.
  - **Q2–Q3 2027:** Establish the core team of subject-matter experts, and begin advising and analysis tasks.

INDICATOR

Technical full-time equivalent employees in post against target. Multi-annual budget secured. 12-month retention of technical staff.

CONSIDERATIONS

- This recommendation operationalises Article 70(3) of the AI Act, which obliges Member States to provide national competent authorities with adequate technical, financial, and human resources, including personnel with an in-depth understanding of AI technologies.
- Units could be established in government departments, regulators, cybersecurity agencies, or otherwise, depending on each Member State's institutional structure.

02 Appoint senior frontier AI advisors with direct access to the head of government

ACTION

Member States should appoint a technically credible senior frontier AI advisor reporting directly to the head of government. The advisor should serve as a link between the national technical unit recommended above and the head of government, ensuring that their analysis and recommendations inform strategy and agenda setting at the highest level of government. As such, they should be allowed regular briefings with the head of government and their cabinet, the ability to escalate urgent developments, and participation in high-level cabinet discussions relevant to transformative AI. The advisor should be an individual with practical expertise on frontier AI topics, for example, through working at a frontier AI provider, and should be supported by a small dedicated team.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - National executive decision to appoint a senior frontier AI advisor to the head of government.
  - Support team staffed partly through secondment, for example, from national research centres.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Advisor profile defined and recruitment search launched, including nationals working in frontier labs abroad.
  - **Q1 2027:** Appointment made and support team operational.
  - **Q2 2027:** Begin regular briefings between the senior frontier AI advisor and head of government and relevant cabinet members. Establish a procedure for the advisor to escalate urgent developments.
- **Success indicators:** Appointment made with published mandate. Briefings held per quarter. Documented technical input into national positions on issues relevant to transformative AI.

CONSIDERATIONS

- Relevant direct expertise with frontier AI should be prioritised relative to total years of experience when making appointments. Frontier AI is a new and fast-moving area, meaning that total years of experience is often uncorrelated with relevant expertise.

03 Collect, analyse, and publish national statistics on AI diffusion and adoption

ACTION

Member States should strengthen national measurement of AI diffusion and adoption across businesses, households, and the public sector. National statistics institutes (NSIs) should be instructed to collect and analyse data pertaining to national AI diffusion and adoption and publish results on a regular cadence. Efforts should focus on extending coverage of existing statistics, for example, by analysing specific tasks being automated across jobs, comparing the extent of AI uptake between SMEs and larger firms, or identifying sectors or roles that are adopting AI systems the fastest. Results should be shared with other Member States on a common calendar, for example, through Eurostat.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - Instructing national statistics institutes to include additional AI-relevant variables or modules in existing surveys.
  - Existing surveys of Information and Communications Technology (ICT) use conducted through the European Statistical System, expanded to include fine-grained AI-specific variables.
  - Partnering with AI providers and other data holders, including through data-sharing agreements, to obtain (aggregated) statistics on AI usage.
  - Coordination with other Member States through Eurostat, to harmonise core variables, allowing for Europe-wide comparisons.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** NSIs identify gaps in existing AI adoption statistics. Agree on a common variable set through the European Statistical System.
  - **Q1–Q2 2027:** Pilot expanded measures of AI use, focusing on public-sector adoption in the first instance.
  - **Q3 2027:** Publish the results of the public-sector pilot, and share through the European Statistical System.

INDICATOR

Annual publication delivered. Coverage across enterprises, individuals, and the public sector. Partnerships with private-sector data holders.

CONSIDERATIONS

- Relevant data may be collected in partnership with AI providers, such as through Google's Activity, Task, Landscape, and Adoption Study (ATLAS) project.

04 Ensure a baseline of AI literacy across all government departments

ACTION

Member States should adopt a government-wide AI literacy programme to ensure that all civil servants are equipped with the knowledge and skills they need to understand and effectively use AI systems. The programme should be made available to staff in every department, with formats tailored to the department and seniority. For example, senior policymakers should receive short, high-quality briefings and demonstrations covering the capabilities and risks

of frontier AI systems and possible trajectories of further development. In addition, junior civil servants should receive practical training that covers how to use AI tools effectively and responsibly, with more rigorous training for staff in specialist sectors such as procurement, healthcare, or security.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - Executive decision establishing a baseline AI literacy across government.
  - Department-specific training offered through ministries.
  - The national AI unit, recommended above, should coordinate the programme, for example, setting the syllabus and partnering with universities and research institutes to deliver specialist modules.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Departments' existing AI use mapped. Programme and completion targets adopted. Baseline curriculum adopted, and roles or departments requiring further specialist training identified.
  - **Q1–Q2 2027:** Senior official briefing and demonstration series delivered. Working-level modules launched.
  - **Q3 2027:** Completion rates and effectiveness statistics published per department. Curriculum updated in light of recent developments.
- **Success indicators:** Share of senior officials participating in training. Improvement on task-based assessment. Adoption and effective use of AI systems.

CONSIDERATIONS

- It is important that the curriculum covers topics, developments, and events that may have emerged shortly before the course is operational. The field of AI is advancing incredibly fast, and a curriculum that does not include developments from the preceding six months will be significantly outdated. This also means that there must be processes for iterative and continual updating of the course curriculum.

IO-3 Secure Europe's share of global AI compute

WHY IT MATTERS

Compute is the physical infrastructure on which all AI applications depend. In a world with transformative AI, access to AI compute will be a prerequisite to a region's economic prosperity and national security, just like access to energy is today. Hosting this infrastructure on EU soil increases the bloc's geopolitical weight and reduces its reliance on revocable foreign access.

Recommendations at the Union level

01	02	03	04
Host at least 15% of global AI computing capacity by 2030	Create a Rapid AI Infrastructure team within the European Commission	Upgrade CADA's data centre acceleration zones	Mobilise €25 billion of EU public funding for AI compute
very high priority	very high priority	very high priority	high priority

01

Host at least 15% of global AI computing capacity by 2030

ACTION

Set a Union-level target in the Cloud and AI Development Act (CADA) of hosting at least 15% of global AI compute on EU soil by 2030, in rough proportion to the EU's 18% share of global GDP, currently projected at 45 GW of total facility power (against a 300 GW global capacity). Attach three obligations: (i) the Commission should adopt, and periodically update by delegated act, a methodology for estimating installed Union and global AI compute capacity, taking into account computational performance, energy-efficiency improvements, and other relevant technical parameters; (ii) the Commission should publish annually the projected power capacity required for the EU to meet its 15% target on the current global compute trajectory; (iii) within six months of CADA's entry into force, Member States should each share a national contribution plan describing their intended contributions towards the compute target. Member States should report annually on their progress, while the Commission should assess aggregate progress and, where it falls short, make recommendations to Member States on how to align national contributions with the Union-level target (following the approach of the Energy Union Governance Regulation (EU) 2018/1999).

IMPLEMENTATION

- **Potential instruments:** Amendments to CADA (COM(2026) 502), including by adding an Annex with the plan and reporting template, updatable by delegated act.
- **Potential first steps (next 12 months):**
  - **October 2026**: The Commission endorses the 15% target and signals support for corresponding amendments to the co-legislators.
  - **October–November 2026**: The Commission builds the measurement and forecasting capacity.
  - **Late 2026**: The 15% target is introduced into legislative negotiations through amendments in Parliament and/or Council.
  - **Mid-2027**: The European Parliament's position and the Council's general approach include the target.
  - **September 2027**: The Commission publishes an updated global compute capacity projection, providing the basis for a later delegated act that updates the GW figure corresponding to the 15% target.
- **Success indicator:** Share of global installed AI compute hosted on EU soil by 2030.

CONSIDERATIONS

- The target is expressed as a share rather than an absolute GW figure so that it remains meaningful under different scenarios of global AI compute growth.
- The 15% target could also be inserted into the EU's Digital Decade Policy Programme 2030, while CADA contains the binding obligations.
- The definition of 'AI compute' should rule out general cloud capacity unsuited for AI workloads. Reports of existing AI compute capacity should only include operational AI data centres (measured as total facility power), rather than merely announced projects or sites not yet energised.
- Negotiators should resist any trilogue compromise that moves the target into a non-binding recital or turns it into a mere aspiration.

02 Create a Rapid AI Infrastructure team within the European Commission

ACTION

By 2027, establish a dedicated Commission task force mandated to steer the buildout to the 15% target. The task force should have 30–50 staff, with at least 50% external experts with relevant backgrounds (e.g. in data centre financing and development, grid engineering and emergency infrastructure planning, including experts with military logistics backgrounds). Their mandate should include setting and monitoring intermediate GW milestones and coordinating Member State and local governments, site owners, energy companies, grid operators, investors, construction companies, and AI companies, to facilitate strategically important projects.

IMPLEMENTATION

- **Potential instruments:** This would be a decision about the Commission's internal organisation; no legislation required. Industry professionals should be recruited mainly as temporary agents.

- **Potential first steps (next 12 months):**
  - **October 2026**: The Commission announces the task force and internally fixes its mandate, the reporting line, and a sunset (e.g. 2031).
  - **October–December 2026**: The Commission recruits an operational head and relevant experts; operational core begins work.
  - **January 2027**: Core team is fully functional.
  - **March 2027**: First quarterly buildout scoreboard published internally (list of projects, GW at final investment decision, GW under construction, GW energised, major blockers).
  - **July 2027**: Full team is fully functional.
- **Success indicators:** Number of staff in an operational task force; share of external experts with relevant backgrounds; projects under active coordination (measured in GW); measured reduction in time to power for coordinated projects versus national baselines.

CONSIDERATIONS

- For the Rapid AI Infrastructure team to be successful, the Commission should take into account the recommendations on how to build technical state capacity from Immediate Objective 2.

03 Upgrade CADA's data centre acceleration zones

ACTION

Increase the impact of CADA's acceleration zone regime by anchoring it in, and extending, the EU's existing regime for accelerating permits. Specifically: (i) grant AI data centres in acceleration zones a presumption of overriding public interest for the purposes of EU nature and water[21] legislation, mirroring Article 16f of the Renewable Energy Directive; (ii) classify all acceleration zone projects as strategic projects under the Environmental Assessments Regulation (COM(2025) 984), following the precedent of the Industrial Accelerator Act. Specify that permits must be granted or refused within three months after receipt of complete documentation (in contrast to CADA's proposed 12 months), with approval presumed where the deadline is missed (except where EU environmental law prohibits this); (iii) oblige Member States to issue the front-loaded 'aggregated baseline permits' envisaged by CADA within six months of zone designation; and (iv) turn sustainability KPIs based on the Energy Efficiency Directive into a single harmonised benchmark, barring Member States from gold-plating zone requirements.

IMPLEMENTATION

- **Potential instruments:** Amendments to CADA (COM(2026) 502).
- **Potential first steps (next 12 months):**
  - **October 2026**: The Commission endorses the suggested amendments and signals support to the co-legislators.
  - **Late 2026**: The amendments are tabled in the legislative negotiations in Parliament and/or Council.
  - **Mid-to-late 2027**: Trilogues open; co-legislators back the acceleration zone provisions.

21 This presumption should apply to derogations under Article 4(7) of the Water Framework Directive (WFD) only for data centres using closed-loop or other non-evaporative cooling systems.

- **Success indicators:** Median time from complete application to permit for projects in acceleration zones ≤ 3 months.

CONSIDERATIONS

- The proposed three-month permit-granting period should begin when the relevant authority receives complete documentation. However, documentation requirements for data centres may not be standardised, which could make local authorities slower at handling applications for novel projects. To address this, the co-legislators could consider measures that build on and extend existing permit acceleration regimes such as the Renewable Energy Directive (RED) III. For example, they could require that an exhaustive documentation checklist is published for each zone in advance and that documentation is deemed complete if the authority does not respond within 30 days; the authority may make a single request for further information that pauses rather than restarts the three-month period.

04 Mobilise €25 billion of EU public funding for AI compute

ACTION

Commit €25 billion of EU-level public funding to the compute buildout, starting immediately through existing programmes such as InvestEU and continuing at scale under the 2028–2034 Multiannual Financial Framework (MFF). Use loan guarantees and advance purchase commitments as a demand-side backstop[22] to de-risk strategically important private data centre projects and increase the share of European-only AI compute.

IMPLEMENTATION

- **Potential instruments:** Until the next MFF applies, initial funding could come from the €30 billion InvestEU Fund and the European Investment Bank (EIB) Group's TechEU programme. From 2028, additional funding could be drawn from the proposed €48.5 billion Digital Leadership budget of the Competitiveness Fund and the €405 billion cohesion budget. The EIB Group would be a natural delivery unit both for the loan guarantees and for structuring potential equity co-investments. Creating a new Important Project of Common European Interest (IPCEI) for Frontier AI Compute would provide the State-aid framework through which Member States could commit their €75 billion share.

- **Potential first steps (next 12 months):**

  - **Late 2026**: Amendments are tabled in the negotiations on the Competitiveness Fund regulation – a potential vehicle for a minimum AI compute share in the Digital Leadership window – and on the National and Regional Partnership Plans regulation, establishing AI compute infrastructure as a priority category eligible for cohesion funding.

  - **December 2026**: Interested Member States pre-notify the IPCEI for Frontier AI Compute.

  - **April 2027**: First InvestEU-backed guarantees for AI data centre projects on the market.

  - **September 2027**: First framework for advance purchase commitments created.

- **Success indicators:** Value of guarantees and purchase commitments signed per year; private capital mobilised per public euro; amount of GW of European-only compute at final investment decision.

---

22 In this context, a backstop is a commitment by an investment-grade third party to step in if the data centre's contracted customer stops paying. This allows lenders to rely on the creditworthiness of the guarantor rather than just the data centre's main customer, making large AI data centres easier and cheaper to finance. A backstop is normally required when AI companies without an investment-grade rating, such as OpenAI or Anthropic, are planning a data centre.

CONSIDERATIONS

- European-only data centres are, at a minimum, located on EU soil and owned and operated by a European organisation.
- Loan guarantees and purchase commitments should be reserved for projects that would not otherwise happen. They should include appropriate fees that reflect the risk taken on.
- The Competitiveness Fund has been proposed but not agreed yet, and the Commission's proposal deliberately avoids ring-fencing within its policy windows to preserve flexibility. Co-legislators should nonetheless explore legal avenues for fixing a minimum spending share for AI compute in the Digital Leadership window itself, following the precedent of thematic minimum shares in other EU programmes (e.g. Horizon Europe's climate target). If the Digital Leadership budget shrinks in negotiations, the compute share should be defended on account of its strategic importance for European prosperity and security.
- The exact allocation between loan guarantees and demand-side backstops should be set in an implementing act, as the two instruments differ in their budgetary properties.

Recommendations at the national level

01	02	03
Strengthen the social contract around AI data centres	Prepare and market shovel-ready sites ahead of CADA	Reduce 'time to power' for AI data centres
very high priority	very high priority	very high priority

04	05
Accelerate planning and permitting for AI data centres	Mobilise €75 billion of national public funding for AI compute
very high priority	high priority

01 Strengthen the social contract around AI data centres

ACTION

(i) Ensure that host municipalities receive a share of the local business taxes generated by an AI data centre, allocated based on its total installed facility power, (ii) make it easy for data centre developers to invest in local community infrastructure, (iii) require large operators to bear the incremental electricity system costs of their projects rather than socialising them onto households and small businesses.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:** National legislation for municipal revenue-sharing and community benefits; the national regulatory authority's network-tariff methodology for cost allocation.
- **Potential first steps (next 12 months):**
  - **October 2026**: National regulatory authorities begin to prepare the cost-allocation methodology.
  - **December 2026**: National legislation drafted to allow community benefits and create a revenue-sharing scheme.
  - **June 2027**: National legislation is passed.
- **Success indicators:** Share of new data centre capacity covered by community-benefit agreements; regulators verify that projects have not increased prices for local households or small businesses.

CONSIDERATIONS

- Since AI data centres employ comparatively few people relative to invested capital, tax shares for host municipalities should be allocated based on total installed facility power rather than payroll.
- Without a legislative basis, voluntary developer payments to local communities risk being classified as undue advantage under anti-corruption law.

02 Prepare and market shovel-ready sites ahead of CADA

ACTION

- (i) Publish, within six months and without waiting for CADA's adoption, a national inventory of potential AI data centre sites, each with the potential for at least 100 MW fully operational capacity by 2030, including data on available construction area and grid capacity, on-site electrical infrastructure (e.g. substations), previous environmental assessments, owners and relevant stakeholders, and other relevant information.
- (ii) Designate the most promising sites (or collections of sites) as data centre acceleration zones, as envisaged by CADA and the amendments in Union-level recommendation 3 above, and ideally going beyond CADA's minimum of one per Member State.
- (iii) Reach out to site owners, AI companies, data centre developers, investors, and local authorities to gauge interest in building at those sites.

MEMBER STATE MODE

Collective – every Member State contributes sites.

IMPLEMENTATION

- **Potential instruments:** National legislation and administrative direction for designating acceleration zones, following the model of national renewables acceleration areas under RED III; a designated national agency acting as a single point of contact for parties interested in data centre development at the relevant sites.
- **Potential first steps (next 12 months):**
  - **October–December 2026**: Prepare and disseminate the national inventory.
  - **January 2027**: Outreach to interested parties begins.
  - **July 2027**: Designate the first acceleration zones.
- **Success indicators:** GW capacity of designated sites; GW capacity of listed sites with signed developer commitments.

CONSIDERATIONS

- Member States with cheap and abundant energy (e.g. France, Spain, Portugal, the Nordic countries) or repurposable industrial brownfields (e.g. Germany, Poland, Romania) should contribute disproportionately to the aggregate.
- Industry experience suggests that foreign companies seeking to build AI data centres sometimes do not know about attractive sites in Member States, and successful agreements often require active matchmaking to bring the relevant stakeholders together.

03 Reduce ‘time to power’ for AI data centres

ACTION

Reduce the time that it takes to energise a data centre – either by connecting it to the grid or generating power on site – in order to make it fully operational. Measures include: (i) prioritising AI data centres in the grid connection queue, especially in the acceleration zones described in Union-level recommendation 3 above, (ii) facilitating Flexible Connection Agreements (FCAs) to use existing grid capacity more efficiently, (iii) accelerating permitting for power plants, transmission lines, and on-site generation, and (iv) using public procurement to support emerging power technologies such as fuel cells.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:** (i) National regulatory authority decisions or legal amendments introducing readiness- and policy-based grid connection criteria (as supported explicitly by the Commission’s December 2025 grid-connection guidance). (ii) National regulatory authorities’ implementation of the EU’s framework for FCAs (Article 6a of Directive (EU) 2019/944). (iii) National implementation and possibly extension of the EU permitting-acceleration regime (RED III for generation, the Grids Package and TEN-E for transmission); potentially including on-site generation within the aggregated baseline permits of data centre acceleration zones. (iv) Public tenders for backup and on-site power that explicitly invite emerging technologies such as fuel cells, using the innovation-procurement instruments of the EU public procurement directives.
- **Potential first steps (next 12 months):**
  - **October–December 2026**: National regulatory authorities and grid operators draft queue prioritisation criteria and a national FCA template.
  - **January 2027**: First innovation-procurement tenders specifying emerging power technologies published.

- **April 2027**: Queue criteria and FCA framework adopted.
- **July 2027**: First FCA offers issued to zone projects.
- **September 2027**: Accelerated permitting secured through national legislation.

- **Success indicators:** Median time from connection application to signed connection agreement for data centre projects; MW of Flexible Connection Agreements signed.

CONSIDERATIONS

- Experts regard 'time to power' as the single most important factor for AI data centre developers when selecting sites. In leading data centre regions in the US, GW-scale clusters are often energised in two years or less. This is usually achieved by relying on on-site power generation or a combination of on-site and grid power.
- A recent study found that Flexible Connection Agreements could unlock over 15 GW of additional AI compute by 2030 in the major European power markets (France, Germany, Iberia, and the Nordics, including Norway).
- Member States that decide to facilitate on-site power generation may consider capping the share of fossil fuels used (e.g. at 50% of capacity).

04 Accelerate planning and permitting for AI data centres

ACTION

Shorten the permitting process for AI data centres to a maximum of six months, country-wide and beyond just the acceleration zones mentioned in Union-level recommendation 3, by (i) introducing national caps for permitting timelines, (ii) expanding the administrative capacity handling data centre applications, (iii) adopting approval-by-default where authorities miss deadlines (insofar as this is consistent with EU environmental law), (iv) allowing developers to begin (reversible) construction at their own risk while approvals are pending, and (v) repealing national requirements stricter than the relevant EU floor.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:** National permitting legislation; single points of contact to handle the entire permitting process; budget direction for administrative capacity; amendments to relevant national or sub-national laws to prevent gold-plating of EU laws.
- **Potential first steps (next 12 months):**
  - **October–December 2026**: The responsible ministry prepares the permitting bill and accompanying amendments to sectoral law.
  - **March 2027**: The parliament passes all relevant legislation.
  - **May 2027**: Single points of contact operational; additional budget and staffing granted to permitting authorities.
  - **September 2027**: First applications successfully processed under the new regime.
- **Success indicator:** Median duration from complete application to permit for AI data centres.

CONSIDERATIONS

- In 2022, Germany passed the LNG (Liquefied Natural Gas) Acceleration Act in just 10 days to reduce dependence on Russian gas, thereby setting a precedent for how fast national legislators can act to protect a country against threats to its sovereignty and economic security.

05 Mobilise €75 billion of national public funding for AI compute

ACTION

Collectively commit €75 billion of national public funding to the compute buildout. Use loan guarantees and advance purchase commitments as a demand-side backstop to de-risk strategically important private data centre projects and increase the share of European-only AI compute. For the most sensitive use cases, such as AI workloads in national security or defence, Member States should consider direct equity investment into AI data centres alongside private sector partners to ensure that the most critical AI infrastructure is publicly co-owned.

MEMBER STATE MODE

Collective – every Member State contributes, scaled by economic weight.

IMPLEMENTATION

- **Potential instruments:** National budget commitments and promotional banks, coordinated through a novel IPCEI for Frontier Compute.
- **Potential first steps (next 12 months):**
  - **October 2026–January 2027:** Member States agree on individual IPCEI contributions and open pre-notification talks with the Commission.
  - **March 2027:** National budget lines and promotional bank mandates adopted; promotional banks design the loan guarantees and advance purchase commitments.
  - **May 2027**: First loan-guarantee products for AI data centre projects on the market.
  - **September 2027**: First loan guarantees/advance purchase commitments signed.
- **Success indicators:** Amount committed in national budgets; amount of GW of European-only compute at final investment decision.

CONSIDERATIONS

- European-only data centres are, at a minimum, located on EU soil and owned and operated by a European organisation.
- Loan guarantees and purchase commitments should be reserved for projects that would not otherwise happen. They should include appropriate fees that reflect the risk taken on.

IO-4 Ensure resilience to AI crises

WHY IT MATTERS

Transformative AI poses severe risks to Europe, its citizens, and its institutions. Many of these risks cannot be fully managed upstream by Europe. As such, Europe should aim to bolster its resilience to risks from transformative AI – that is, its capacity to resist, absorb, recover from, and adapt to risks should they occur.

Recommendations at the Union level

01	02	03	04
Implement the Action Plan on Cybersecurity and Artificial Intelligence and develop CBRN, and loss of control action plans	Establish a channel for providers and affected parties to voluntarily notify the AI Office of AI incidents	Strengthen cooperation between AI expertise and national security and intelligence authorities	Establish a cross-programme European AI preparedness and resilience package in the 2028–2034 Multiannual Financial Framework and allow existing funds to be used for AI resilience measures
very high priority	high priority	high priority	high priority

01 Implement the Action Plan on Cybersecurity and Artificial Intelligence and develop corresponding CBRN and loss of control action plans

ACTION

- With the Action Plan on Cybersecurity and Artificial Intelligence, Europe has established promising foundations for improving cyber-resilience as AI systems become more capable. The Commission should translate the Action Plan's commitments into operational measures through adequate funding, clear institutional ownership, and defined deadlines. In particular, the Plan's blueprint for structured access to advanced AI capabilities for cybersecurity should be converted into operational access arrangements before such access becomes urgently necessary. In parallel, the Commission should publish corresponding plans for chemical, biological, radiological, and nuclear (CBRN) and loss of control risks.
- Potential concrete actions to include in the Action Plan on CBRN-security and AI could include: updating the legal limit for workplace exposure to UV (currently fixed in Directive

2006/25/EC) to enable controlled trial deployments of higher doses of far-UVC for antiviral filtering in hospitals; and increasing the number of sites covered by the Directorate-General for Health Emergency Preparedness and Response Authority (DG HERA) and the JRC's ongoing wastewater monitoring initiatives to provide a more comprehensive monitoring system for harmful pathogens in water systems. These initiatives could also be extended to include detection of AI-engineered pathogens through metagenomic sequencing of samples. The Action Plan could also consider actions for preventing the misuse of specialised AI tools, designed specifically for aiding biological or chemical research, such as Evo 2, a genome-modelling AI model that has been shown to be able to generate coherent, genome-length DNA sequences. Finally, it could lay out a plan for coordinating emergency response to AI-enabled CBRN incidents at the Union level, complementing national emergency preparedness plans (see national recommendation 4 below).

- The content of an AI Loss of Control Action Plan would likely have to change as our understanding of the risks (and how to assess and mitigate them) evolves. An AI Loss of Control Action Plan could include a reaffirmation of the Commission's intention to use its enforcement powers to incentivise providers of general-purpose models with systemic risks to assess and mitigate loss of control risks, and to conduct its own loss of control model evaluations (or contract them to a third party). Key example mitigations could be preserving chain-of-thought monitorability, having in place scalable and comprehensive monitoring of internally deployed agents, and ensuring that reward learning techniques used do not result in misalignment. Additionally, it could lay out a plan to improve emergency preparedness by bolstering Europe's detection, containment, response, and resilience capabilities[23] and ensuring sufficient funding is mobilised to improve them. For example, this could be achieved by monitoring for unauthorised autonomous AI agents across EU-regulated critical and government infrastructure, implementing sandbox and isolation standards for AI evaluations conducted in the EU, ensuring sufficient human expert reserve capacity for dealing with crises, and establishing trusted back-up models and non-AI fallbacks for essential functions across government and critical infrastructure. Finally, the Action Plan could catalyse initiatives to build understanding of loss of control risks and the technologies needed to enable oversight and governance; this could include (a) building EU loss of control evaluation and incident investigation capacity, (b) making concerted efforts to build assurance technologies (Immediate Objective 5) to also help with international coordination, (c) building structured access and secure evaluation infrastructure, and (d) advancing standards and technology for agent identification, logging, and attribution. The Action Plan should consider how such innovations and capacity building could best be incentivised.

IMPLEMENTATION

- **Potential instruments:**
  - Executive support in the form of funding and operational guidance to ENISA, Member States, and other relevant Union entities for effectively implementing the Action Plan on Cybersecurity and Artificial Intelligence.
  - Implementing parts of these Action Plans will require diplomatic negotiation, for example to secure access to frontier-level systems, as is likely to be set out in the European Blueprint for structured access to advanced AI capabilities for cybersecurity purposes, commissioned in the Action Plan on Cybersecurity and Artificial Intelligence.

---

23 Examples could include: (a) Detection: Monitoring for rogue autonomous AI agents across EU-regulated critical and government infrastructure; early warning services and threat intelligence sharing of rogue AI agent incidents; logging of AI agent events when used in high-risk systems. (b) Containment: shutdown and rollback capabilities for EU-operated AI systems and EU-based data centres; network segmentation and containment procedures for critical infrastructure, government, and other important attack surfaces; sandbox and isolation standards for AI evaluations conducted in the EU. (c) Resilience and response: trusted back-up models and non-AI fallbacks for essential functions across government and critical infrastructure; compute, tool, and human expert reserves for crises; identifying single points of failure; including loss of control in emergency response plans; conducting loss of control crisis tabletop exercises and simulations; including loss of control crisis coordination in existing structures.

- Executive action to publish Action Plans on (1) CBRN-security and AI and (2) AI Loss of Control, inspired by the existing Action Plan on Cybersecurity and Artificial Intelligence.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Drafting begins on the Action Plans on CBRN-security and AI and AI Loss of Control.
  - **By Q1 2027:** Action Plan on CBRN-security and AI published. The European Blueprint for structured access to advanced AI capabilities for cybersecurity purposes begins to be implemented.
  - **By Q3 2027:** Action Plan on AI Loss of Control published. Action Plan on CBRN-security and AI implemented.
  - **By Q1 2028:** Action Plan on AI Loss of Control implemented.
- **Success indicators:** Publication of the European Blueprint for structured access to advanced AI capabilities for cybersecurity purposes. Publication of an Action Plan on CBRN-security and AI.

CONSIDERATIONS

- A model may be created with biological design capabilities that prompt a tightly restricted release, of the kind observed for Anthropic's Mythos with respect to offensive cyber capabilities. Europe should implement personal protective equipment (PPE) stockpiling, screening infrastructure, and response plans for advanced biocapabilities in advance of such a model.
- The implementation timeline above shows publication of the Action Plan on AI Loss of Control after the Action Plan on CBRN-security and AI, due to additional research required to determine the contents of the plan and the fact that the first incidents that may qualify as loss of control have only recently been reported. This means that experts' understanding of both the risk and how to mitigate it is still evolving. However, if loss of control risks increase more rapidly, the timeline would have to be accelerated to avoid potentially catastrophic outcomes.

02 Establish a channel for providers and affected parties to voluntarily notify the AI Office of AI incidents

ACTION

The Commission should establish a channel through which a team within the AI Office can be notified of relevant occurrences that do not qualify as 'serious incidents' under Article 3(49) of the AI Act. The channel should be made available through an online tool accessible to the public so that notifications can be made by affected parties (including both individuals and businesses) and providers. Given the potential for multiple interpretations of 'serious incident' as used in Article 55(1)(c) of the AI Act, the Commission should clarify that this channel may be used for notifying the AI Office of incidents that the notifier may nonetheless believe do not meet the legal definition of 'serious incident' in the AI Act. Where deemed instructive, the AI Office may conduct more thorough assessments of the incidents of which it is notified, as a way of informing the Commission's broader strategy for responding to AI incidents and improving the AI Office's regulatory expertise. The assessments and any potential follow-up actions would be conducted by the AI Office pursuant to its general monitoring and supervision powers under the AI Act and would not necessarily amount to a formal investigation.

IMPLEMENTATION

- **Potential instruments:**
  - No legal instruments are required to establish and announce a voluntary notification channel.

- The Commission could issue guidance clarifying that incidents do not need to meet the AI Act's definition of 'serious incident' (Article 3(49)) to be eligible for notification.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Notification channel announced and published on Commission website alongside communications clarifying the purpose and legal scope.
  - **Q1 2027:** AI Office notified of first incidents through the channel. First assessments begin.
- **Success indicators:** Number of substantive notifications, previously unknown to the AI Office. Outcomes and recommendations from subsequent assessments.

CONSIDERATIONS

- The AI Office should stress that notification of an AI incident will in no way be treated as an admission of wrongdoing, in line with Recital (j) of the Safety and Security Chapter of the General-Purpose AI (GPAI) Code of Practice. Proactive notifications from providers should be seen in a positive light, and be taken into account by the AI Office when dealing with providers, in line with Articles 99(7)(h) and 101 of the AI Act.
- The AI Office should have a standing team for triaging and assessing notifications of incidents received through this voluntary channel so as not to reduce the AI Office's capacity in other areas.

03 Strengthen cooperation between AI expertise and national security and intelligence authorities

ACTION

The Commission should establish standing working channels between the EU AI Office, Member State AISIs (see Immediate Objective 2), and security and intelligence services, since each will hold information and expertise that the others need but lack. These collaborations should be governed by an agreed protocol for handling classified and commercially sensitive material. This collaboration should work to produce joint threat assessments of AI-enabled threats covering offensive cyber, CBRN, harmful manipulation, and loss of control. This ensures that Europe's intelligence communities have access to subject-matter AI expertise when analysing potential national security threats.

IMPLEMENTATION

- **Potential instruments:**
  - Include provisions for information exchange and joint threat modelling and assessment exercises between the EU AI Office and Europol in the proposed Regulation on Europol COM(2026) 580.
  - The Cybersecurity Act II (COM(2026) 11) proposes closer Europol–ENISA cooperation on ransomware preparedness and response, providing a precedent to build on.
  - Collaborations with individual Member States' intelligence organisations should be negotiated individually based on a shared template.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** The EU AI Office and DG HOME agree on mutual points of contact. Amendments to the proposed Regulation on Europol (COM(2026) 580) are proposed by Parliament.
  - **Q1 2027:** The AI Office, DG HOME, and national intelligence agencies draft a protocol for handling classified and commercially sensitive material, drawing on existing Commission security rules.

- **Q3 2027:** The first joint threat assessment is drawn up by the EU AI Office and at least one national security service.

- **Success indicators:** A signed handling protocol exists by Q2 2027. At least two joint exercises, such as threat assessments or gap analyses, are completed by the end of 2027.

CONSIDERATIONS

- Since national security is the exclusive responsibility of each Member State under the EU treaties, many elements of collaboration will be voluntary. This includes the shared channels, agreed protocols, and joint assessments.

- Appropriate confidentiality must be maintained on both sides of collaborations. Security services will not share intelligence that could reach commercial parties, and AI providers will be hesitant to share model details with the AI Office if they believe there is a chance that they will be accessible to parties other than the AI Office. Confidentiality agreements should remain in line with Article 78 of the AI Act.

04 Establish a cross-programme European AI preparedness and resilience package in the 2028–2034 Multiannual Financial Framework and allow existing funds to be used for AI resilience measures

ACTION

To ensure AI preparedness and resilience across the European bloc, the Commission should propose a cross-programme package totalling €10 billion within the 2028–2034 Multiannual Financial Framework (MFF) spending plan for AI-specific preparedness and resilience. The proposal for the 2028–2034 MFF includes multiple funding sources for preparedness and resilience. This includes the proposed COM(2025) 565, which reserves a quarter of Member States' contributions to the budget for responding to crises; the proposed COM(2025) 548, which would establish a Crisis Coordination Hub; and the European Competitiveness Fund (COM(2025) 555), a €234 billion fund within the MFF to support strategic technologies, innovation, and industrial capacity. Member States should be able to draw on this package to implement measures addressing systemic AI risks including CBRN threats, cyber offence, loss of control, and harmful manipulation. Ahead of the next MFF's entry into force in 2028, the Commission should clarify that existing funding programmes can be used to fund AI-specific resilience measures. For example, Digital Europe could be used to fund AI-enabled cyberdefence for essential entities, rescEU could fund Personal Protective Equipment (PPE) stockpiling, and EU4Health could fund expanded wastewater monitoring.

IMPLEMENTATION

- **Potential instruments:**

  - The Commission should publish coordinated guidance identifying for Member States which AI-preparedness measures are eligible for reimbursement under Digital Europe, EU4Health, and the Union Civil Protection Mechanism. Funds from the EU Cybersecurity Reserve should be used to respond to qualifying AI-enabled cyber incidents. Non-cyber AI incidents should be addressed through the Union Civil Protection Mechanism and the EU's HERA.

  - The Commission should amend COM(2025) 548 to require Member States to address AI-enabled and AI-amplified incidents in national risk assessments and planning. It should also address Union-level risk reporting, scenarios, and capacity-gap analysis under Articles 14–16.

  - The Commission should amend COM(2025) 565 to add AI-enabled incident preparedness as an eligible investment for Member States. It should also establish a spending floor for AI-preparedness initiatives. It should make reserved funds from the unprogrammed share available at the 2031 mid-term review for AI risks that materialise faster than planned.

- The European Competitiveness Fund's Resilience and Security window, its Health, Biotechnology, Agriculture, and Bioeconomy window, and its Digital Leadership window should explicitly cover industrial capacities related to synthesis screening, DNA sequencing, and AI-based cyberdefence tools.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** The Commission tables the €10 billion package as its position in the MFF and programme negotiations; Parliament and Council positions reflect it by mid-2027. The Commission clarifies the eligibility of existing funds and instruments for financing national AI-specific resilience before the 2028–2034 MFF takes effect.
  - **Q2 2027:** During final negotiations, the Commission defines a clear division of scope between the European Competitiveness Fund, COM(2025) 548, and COM(2025) 565.
  - **Q4 2027:** The Commission's rules, work programmes, and eligibility guidance are finalised ahead of January 2028.
- **Success indicators:** Guidance published by Q1 2027, with at least 15 Member States drawing on civil preparedness money for AI-related resilience work by the end of 2027. The Competitiveness Fund's Resilience and Security window dedicates a minimum share for civil preparedness capacities within the binding text. COM(2025) 548 and COM(2025) 565, as adopted, name AI-enabled incidents among the hazards in scope, specify a spending floor for AI-threat preparedness, and allocate dedicated funds for novel threats.

CONSIDERATIONS

- Civil resilience should be funded separately from defence. The European Competitiveness Fund also includes €131 billion proposed for defence and space. Without a guaranteed minimum written into the law, civil preparedness would be at risk of deprioritisation relative to other funding targets.
- At least 60% of the €10 billion allocation should flow through national and regional partnership plans: This funding should be allocated by a published formula to ensure it reaches Member States equitably rather than through a competitive bidding process.
- Collective AI-specific resilience funding is split across three instruments due to sector-specific scopes. The European Competitiveness Fund focuses on technology and industrial capacity, COM(2025) 548 focuses on operational preparedness, and COM(2025) 565 is for national resilience investments.

Recommendations at the national level

01	02	03	04
Strengthen the biosecurity provisions in the EU Biotech Act	Build up reserves of resources and equipment necessary for responding to crises	Strengthen critical national infrastructure against AI-enabled cyberattacks	Include AI-enabled incidents in national and sector-specific emergency response plans
high priority	very high priority	very high priority	high priority

01 Strengthen the biosecurity provisions in the EU Biotech Act

ACTION

Member States should push for more ambitious biosecurity provisions to be included in the EU Biotech Act during trilogue negotiations. Proposed amendments should include provisions for increasing security and resilience against AI-mediated biological risks, for example, by ensuring that the list of products in scope of KYC provisions can be updated rapidly in the case of a novel AI-generated pathogen being detected. More general strengthening provisions could include providing a legal safe harbour for trusted third parties to test for weaknesses in synthesis screening measures (for example, under Article 50, concerning audits of screening mechanisms). Article 53, concerning biological systemic risk from AI models, could be developed, for example to specify potential 'appropriate measures to ensure a proper control of risks' that can be taken by the Commission and Member States.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - Support robust measures in the EU Biotech Act chapter on biosecurity. Table amendments that strengthen the Act's biosecurity chapter, such as processes for rapidly adding to the list of products of concern, so that screening can keep pace with rapidly emerging AI-designed sequences.
  - Convene key industry players and collaborate (e.g. through MoUs) to voluntarily implement synthesis screening and KYC controls on dangerous viruses, based on the proposed contents of the Biotech Act.
  - Designate, resource, and fund national competent authorities (NCAs) through national implementing acts.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Table amendments to uphold strong biosecurity provisions during the trilogue.
  - **Q1 2027:** Consult with national industry and academia on potential voluntary implementation of synthesis screening and support its uptake ahead of the Biotech Act's expected passing.
- **Success indicator:** Biotech Act passed with strong biosecurity provisions in place.

CONSIDERATIONS

- For biosecurity provisions to be effective, they must be adequately enforced by NCAs. NCAs should thus be adequately resourced to ensure effective enforcement.
- The bio-relevant capabilities of frontier AI systems are advancing rapidly, and it may be many years before the EU Biotech Act is implemented. Member States should therefore aim to bridge this gap by introducing harmonised frameworks that can be built on after the Act's passing.
- Synthesis screening must be robustly implemented in all Member States to be effective. Weak implementation in one jurisdiction significantly impacts the regime as a whole.

02 Build up reserves of resources and equipment necessary for responding to crises

ACTION

Member States should identify and procure physical resources that would be scarce, slow to replace, or geographically concentrated in times of crisis. For example, Member States should establish protective equipment stockpiles sufficient for handling pandemics at least as significant as COVID-19. Member States should conduct supply chain criticality assessments and tabletop exercises to identify vulnerabilities in the supply chains for equipment that is critical for telecommunications, energy grids, and emergency response; and proactively build stockpiles in case those supply chains are disrupted.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - A *multiannual* national budget line for the recurring cost of rotating stockpiles of consumables (PPE, pharmaceuticals, fuel).
  - A procurement framework including rotation clauses.
  - Implement a minimum-spares obligation on designated critical entities through national Critical Entities Resilience (CER) Directive transposition.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Designate a centralised national stockpiling authority and register with the EU stockpiling network. Conduct a review of current stockpiles.
  - **Q1 2027:** Conduct supply chain criticality assessments for entities identified under the CER Directive. Sign PPE framework contracts, including health-service rotation clauses.
- **Success indicators:**
  - Population coverage × days of supply for PPE.
  - Net annual cost as a share of stockpile value.

CONSIDERATIONS

- Stockpiles could operate as rotating inventories, with stock approaching expiry transferred into routine use (e.g. PPE use in hospitals), so that the continuing cost to the state remains below the headline value of the stockpile.
- Given AI tools' ability to design novel phages, plans should emphasise 'pathogen-agnostic' defences and countermeasures, rather than being narrowly tailored to existing biological threats.

03 Strengthen critical national infrastructure against AI-enabled cyberattacks

ACTION

Member States are already required to ensure that essential and important entities (previously called operators of essential services in NIS1) manage cyber risk via the NIS2 Directive. To catalyse implementation, Member States should go further by establishing programmes through which essential and important entities (including water-treatment facilities, electricity grids, banks and payment systems, telecommunications, and transport) can access cybersecurity expertise and resources to proactively harden their digital infrastructure against AI-enabled cyberattacks. Combine expertise in internal bodies and contracted private-sector entities, and supply them with highly capable AI models for identifying and patching critical vulnerabilities. After high-priority infrastructure has been hardened, SMEs should also be eligible for participation in the programme.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - The Action Plan on Cybersecurity and Artificial Intelligence specifies a "pilot of a Critical Open Source Resilience Campaign to accelerate patching including by leveraging AI", which could provide a suitable vehicle for funds.
  - The NIS2 Directive required Member States to identify and designate *essential and important* entities. This designation could serve as a list of entities to prioritise for funding allocation.
  - Regulation (EU) 2021/887 required Member States to establish a National Coordination Centre (NCC). NCCs are able to provide grants and technical assistance to national industry for bolstering cyberdefence.
  - Highly capable AI for cyberdefence could be acquired and coordinated by the European Cybersecurity Competence Centre (ECCC) or ENISA.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Designate priority entities at particular risk of cyberattacks, due to prominence, criticality, or other factors.
  - **Next budget cycle:** Establish a pooled fund for cyberdefence support for entities on the priority list.
  - **Q4 2027:** Expand the programme to lower-priority entities, raising the floor of cyberdefence.
- **Success indicators:** Programme launched. 80% of identified priority entities participating in the programme.

CONSIDERATIONS

- Discovered exploits should be patched as soon as possible. If this is not possible, unpatched exploits should be catalogued securely due to their sensitivity.

04 Include AI-enabled incidents in national and sector-specific emergency response plans

ACTION

Member States should revise national emergency preparedness and civil protection plans to include named AI-enabled incidents. Where sector-specific plans already exist, these should be reconciled and expanded to account for cross-cutting AI-enabled incidents that may go beyond those previously imagined in scope or magnitude of harm. All plans should specify a clear chain of command, the legal basis for emergency measures, and named contacts within relevant Union-level structures. They should also be tested through periodic exercises conducted jointly with the Commission and neighbouring Member States, and revised in response to any lessons identified.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - Amendments to national civil-protection legislation and the national emergency plans made under it.

- National risk assessments and disaster risk management planning.
- Large-scale cybersecurity incident and crisis response plans, as required by NIS2 Article 9.
- Preparedness and response plans under Regulation (EU) 2022/2371 on serious cross-border threats to health (Article 6).
- Articles 36 and 48 of the DSA.
- Cyber Europe and EU-CyCLONe provide vehicles for preparedness exercises for AI-enabled crises.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Map responsible national bodies under existing crisis plans, and flag potential areas of conflict or overlap. Add named AI-enabled emergencies and incidents to national risk assessments and/or registers.
  - **By Q1 2027:** Identify emergency preparedness plans with designated ownership and chains of command specified for each class of incident.
  - **By Q4 2027:** Conduct the first emergency exercises in collaboration with relevant Union entities and neighbouring Member States.
- **Success indicators:** Percentage of national emergency response plans updated to account for transformative AI.

CONSIDERATIONS

- AI-enabled incidents unfold quickly and require quick decisions. Plans should state in advance which decisions can be pre-authorised and which cannot.
- National AI Security Institutes and, at the Union level, the EU AI Office should be consulted on threat models for AI-enabled incidents, to ensure the relevance of emergency-response plans.
- Civil protection and disaster response remain primarily the responsibility of Member States. This means that the operational response to critical AI incidents will remain largely national.

IO-5 Make Europe the global leader in AI assurance technology

WHY IT MATTERS

In a world with transformative AI it will be increasingly important to be able to ensure security of the most capable AI models and compute, whilst allowing third parties to verify where and how they are being used. For example, public-sector use of AI, secure importing of restricted models, and methods for verifiably 'pacing the frontier' of AI development may all depend on security and/or verifiability. Europe has the opportunity to become a world leader in developing such assurance technology while also benefitting from its application.

Recommendations at the Union level

01	02	03	04
Define Union-wide security and verifiability tiers for AI infrastructure	Fund and support fundamental research into assurance technologies	Scale proofs of concept into real-world pilots	Internationalise assurance technology
high priority	very high priority	high priority	high priority

01

Define Union-wide security and verifiability tiers for AI infrastructure

ACTION

Define Union-wide technical assurance levels for AI-specific compute, with two dimensions: one for security and the other for verifiability. Require every publicly (co-)funded compute facility intended for sensitive workloads to be certified to a defined tier. For all other publicly (co-)funded capacity, require that the facility's design record the tier it is built to, without a certification obligation. This framework could serve as the basis for frontier model access negotiations (see Objective 1.1). This scheme should also aim to harmonise with national compute assurance levels. For workloads within scope of the framework, certification should satisfy all compute-specific security and verifiability requirements of the specified tier, without the need for additional national requirements. This means that a data centre certified at a given tier can serve any Member State's workloads at that level without meeting additional national conditions.

IMPLEMENTATION

- **Potential instruments:**
  - Draft tiering published by ENISA/JRC.
  - Amend CADA to establish a distinct AI-compute assurance framework that covers relevant hardware, firmware, and software. Criteria should be updatable through delegated acts.
  - AI Gigafactory agreements could specify target assurance levels.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Task ENISA/JRC with drafting and publishing an assurance level framework that considers both security and verifiability of compute.
  - **Q1–Q2 2027:** ENISA/JRC assurance level framework published. First wave of announced AI Gigafactories refers to target assurance levels.
  - **Q3–Q4 2027:** ENISA/JRC assurance level framework tabled into the Council and Parliament negotiations of CADA.
- **Success indicators:** ENISA/JRC framework published with defined security and verifiability tiers. AI Gigafactories built to specified assurance tiers.

CONSIDERATIONS

- The security tiers may be benchmarked against the RAND security levels for AI model weights.
- Voluntary certification with the scheme should be open to all providers regardless of nationality, and its visibility with the US government and frontier developers should be actively cultivated, so that it grounds the access negotiations of Objective 1.1 below.
- The current US administration may publish a replacement to the rescinded AI Diffusion Framework that laid out security requirements for export of model weights for the largest models to third countries. The assurance tiers should take into account the provisions of such a framework, should one be published.
- ENISA already owns the EU Cloud Certification Scheme. This initiative should be handled separately and narrowly scoped to AI-specific compute.
- It may be necessary to revise CADA's exclusion of 'hardware' from its sovereignty assurance levels (Annex II) if it is to serve as a basis for providing a framework for certifying security and verifiability of data centres.

02 Fund and support fundamental research into assurance technologies

ACTION

Stand up an ambitious research programme into compute security and verifiability, with at least €250 million guaranteed over five years. Employ subject-matter experts as project leads, alongside programme managers empowered to decide on specific projects to undertake within security and verification R&D. Specify predetermined kill criteria for underperforming projects, and employ a standing red team that stress-tests mechanisms in realistic settings. The programme should double as the field's talent engine by providing support, for example, through fellowships and doctoral funding.

IMPLEMENTATION

- **Potential instruments:** The research programme could be housed as part of the European Innovation Council's (EIC) 'advanced innovation challenges' (see Objective 2.1), alongside dedicated funding from Horizon Europe.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Research programme announced, with funding allocated from the EIC's ongoing 'advanced innovation challenges'.
  - **Q1 2027:** Project leads hired and in place. Longlist of potential research projects published.
  - **Q2 2027:** Initial research projects underway.
- **Success indicators:** Programme announced and funding allocated by end of 2026. Project managers announced.

CONSIDERATIONS

- The subject-matter experts employed as project leads should be provided with full autonomy to decide on which research bets to pursue.
- Underperforming research bets should be dropped early, based on prespecified 'kill criteria' to minimise sunk costs and ensure effective use of funding.

03 Scale proofs of concept into real-world pilots

ACTION

Collaborate with European industry, academia, and partner states to run large-scale pilots on promising security and verification mechanisms, including any identified through the fundamental research activities of Union-level recommendation 2. The guiding goal should be to identify cutting-edge security and verifiability mechanisms that function at the scales of frontier AI deployment and to provide clear routes for their implementation. In line with the recommendation made in Objective 2.2, these pilots should culminate in the construction (by 2028) of the world's first AI data centre designed to withstand an SL5 equivalent threat model, as well as a separate maximally verifiable AI data centre.

IMPLEMENTATION

- **Potential instruments:**
  - Funding could be provided through the InvestAI Initiative.
  - An ARPA-like vehicle could serve as the central manager for real-world pilot projects, coordinating input from industry, academia, and Member State initiatives (see Objective 2.1).
- **Potential first steps (next 12 months):**
  - **Q4 2026:** ARPA vehicle stood up. Initial shortlist of potential first-round pilots announced.
  - **Q1 2027:** First round of pilots announced, with initial funding allocations and project leads in place.
  - **Q2 2027:** In line with the recommendation in Objective 2.2, specification for the SL5 data centre finalised.
  - **Q3 2027:** Shortlist of second-round pilots announced, in part drawing on research projects carried out under the ARPA vehicle.
- **Success indicators:** First round of pilots announced and underway by mid-2027. Independently assessed SL5 data centre operational by 2028.

CONSIDERATIONS

- It is currently unclear whether it is possible for a single data centre to attain both resilience to SL5 security threats *and* high levels of verifiability. Pilots should aim to empirically investigate the trade-offs between these two properties. The initial target should therefore be to have separate data centres that are (respectively) highly secure and highly verifiable.
- The maximally verifiable data centre will likely be built after the maximally secure data centre, as verification mechanisms still pose technical challenges.

04 Internationalise assurance technology

ACTION

If the previous recommended actions result in developments in assurance technology, the Commission should act to internationalise these with like-minded partners. This involves collaborative research, development, and testing of the technology itself, as well as open publication of any resulting technological advances. Given assurance technology's likely critical role in an international agreement that provides the option of 'pacing the frontier', international collaboration and partnership on its development will be essential since states and companies may not trust unilaterally developed assurance technologies that they have limited visibility into.

IMPLEMENTATION

- **Potential instruments:**
  - Bilateral partnerships, for example with international AI safety institutes or industry or academic projects, can be used to run red-teaming and rehearsal exercises for mutually verified compute.
  - The *Pax Silica* declaration, signed by the EU in June 2026, could provide a venue for internationalising the EU's security and verifiability tiers, recommended above.
  - An open-by-default stance for AI verification technologies would allow all interested parties to inspect proposed verification mechanisms.
- **Potential first steps (next 12 months):**
  - **Q1 2027:** First bilateral agreements and partnerships with like-minded allies announced.
  - **Q3 2027:** First red-teaming exercise of assurance mechanisms conducted with an international partner, and learnings published openly.
- **Success indicators:** Technical advances stemming from the ARPA research programme published openly. Bilateral agreements and red-teaming exercises conducted through collaboration with international partners.

CONSIDERATIONS

- Verification mechanisms are unlikely to serve as a basis for international agreements absent adequate inspectability and interoperability. This motivates an open-by-default stance for sharing advances in verification technologies.

Recommendations at the national level

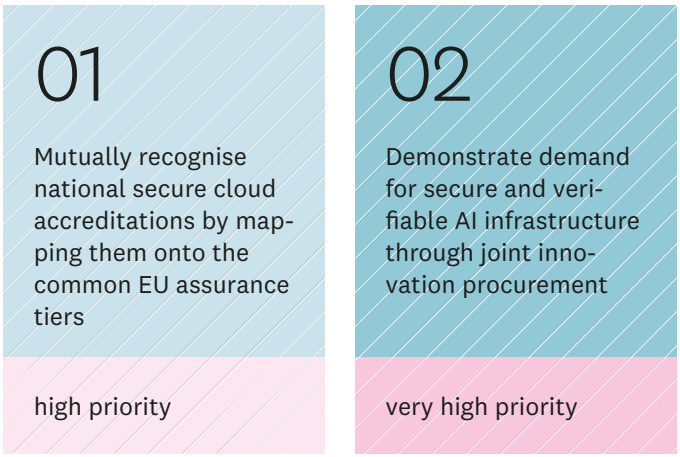

01 Mutually recognise national secure cloud accreditations by mapping them onto the common EU assurance tiers

ACTION

National regimes such as France's SecNumCloud, Germany's C5, and Spain's ENS certify cloud service providers as secure for handling sensitive government data. Member States should provide mappings between their respective certification schemes for AI applications by explicitly referring to the Union-wide tiered framework for security and verifiability recommended above. This mapping between countries should be bi-directional, such that certification to a tier in one Member State is sufficient for recognition at an equivalent tier in the other. This would allow Member States to pool and share secure compute resources and provide a simplified route for providers to become certified across the Union.

MEMBER STATE MODE

All Member States.

IMPLEMENTATION

- **Potential instruments:**
  - Unilateral declarations by Member States that certification via a Union-wide scheme will be treated as equivalent to certification to equivalent tiers in their own certification scheme.
  - Coordination between Member States via the European Cybersecurity Competence Centre.
  - Bilateral negotiations between Member States to mutually recognise each other's certification schemes.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Begin negotiations between Member States to mutually recognise national certification schemes.
  - **Q1–Q2 2027:** Commence work mapping national certification schemes to the Union-wide assured compute certification scheme (recommended above) once published.
- **Success indicators:** At least ten published bilateral mappings and mutual recognition agreements.

CONSIDERATIONS

- It is important to ensure that mappings between certification tiers are bi-directional. That is, certifications to equivalent tiers in two certification regimes are treated as strictly equivalent, without exceptions. Individual Member States should not place additional requirements for providers on top of certification to equivalent assurance tiers in other Member States' regimes.

02 Demonstrate demand for secure and verifiable AI infrastructure through joint innovation procurement

ACTION

Member States should provide a clear demand signal for both secure and verifiable AI compute as a way of incentivising further industry R&D on the topics. Member States should aggregate anticipated government demand for attestable AI compute and issue a forward purchasing agreement for this demand. This should be actioned through Pre-Commercial Procurement (PCP) to contract the development of future security and verifiability mechanisms, followed by Public Procurement of Innovative Solutions (PPI) to purchase assured AI compute capacity once mechanisms reach maturity.

MEMBER STATE MODE

Coalition of 5–10 Member States with significant government demand for sensitive AI workloads.

IMPLEMENTATION

- **Potential instruments:**
  - Publication of a joint multiyear forecast for secure and verifiable AI compute, quantified in terms of GPU-hours or reserved cluster capacity.
  - A joint PCP programme targeting specific assurance capacities that cannot currently be provided at scale.
  - A PPI for qualifying assurance capacity with a ring-fenced budget and commitments to procure pre-specified compute capacity once solutions satisfying defined assurance tiers are available at scale.
  - Issuing targeted technology challenges, with rewards of prizes or contract funding for solving specific bottlenecks in the assured AI compute stack.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Convene an initial coalition of Member States, and jointly quantify and forecast expected demand for secure and verifiable AI compute by government bodies over the next four years.
  - **Q1 2027:** Identify priority capability gaps preventing commercial infrastructure from reaching assurance tiers.
  - **Q2 2027:** Publish the first PCP call.
  - **Q3 2027:** First PCP contracts awarded.
- **Success indicators:** Coalition formed and forecasted demand published. PCP contracts awarded. Successful prototypes demonstrated under realistic deployment conditions.

CONSIDERATIONS

- Earlier-stage research and development should take place within the Union-level research programme recommended above. PCP contracts should focus on rapid scaling of technologies with a clear pathway to broad deployment.
- Demand should be aggregated across participating Member States and refer to the Union-wide assurance tiers recommended above.
- The PCP and PPI stages should remain distinct. PPI should be open to all relevant suppliers, with no preference given to recipients of PCP contracts.

Additional Objectives

Pillar 1

Securing access to frontier AI

O1.1 Secure ongoing access to frontier AI systems

WHY IT MATTERS

Europe's prosperity and security will depend on access to frontier AI models. There have already been instances of access provided by foreign companies being suddenly withdrawn by foreign governments. Europe needs to safeguard its frontier capabilities. The risk of cutoff can be reduced through 'compute for access' deals that tie the provision of compute inputs to access guarantees and by reducing the perceived risk of foreign providers hosting frontier models in Europe. Funding a European fast follower can reduce the downside if access is restricted.

Recommendations at the Union level

01	02	03
Require model portability and multi-vendor terms in public AI procurement	Fund a collective fast-follower model and European-only high-security hosting	Certify Europe as a secure importer and coordinate distillation enforcement
very high priority	very high priority	very high priority

01

Require model portability and multi-vendor terms in public AI procurement

ACTION

Include model portability as a condition in public procurement contracts to prevent European institutions and governments from becoming locked in to a single provider. This should at least apply for sensitive public services, such as those covered by Article 29 of the proposed Cloud and AI Development Act (CADA). At the EU level, model portability could be inserted into the proposed CADA, while the Act's multi-vendor requirements could be strengthened (bearing in mind the limited scope of the Act, which does not cover AI systems and underlying models themselves). The public procurement directives can be similarly revised to prevent lock-in. Finally, the Commission can support the development of switching tools, building on the switching requirements in the Data Act, and allocating funding through the Digital Europe Programme.

IMPLEMENTATION

- **Potential instruments:**
  - The proposed CADA, specifically inserting portability requirements and making multi-vendor provisions stronger. While other legal instruments, such as the Data Act, already introduce portability requirements, frontier AI procurement should be subject to stricter requirements, meaning that model providers must meet a higher standard to be considered for EU AI procurement.
  - EU public procurement rules can also be used to prevent lock-in, specifically through a revision of the proposed Public Procurement Act.
  - The Data Act explicitly addresses model switching, in particular through the switching rules outlined in Chapter VI.
  - Funding to support the creation of switching tools can be allocated through the Digital Europe Programme. The Programme includes a provision allocating funding for the deployment of critical technologies, which could be used for developing switching tools.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** As the Council and Parliament reach their final positions on CADA, aligned Member States and MEPs can submit amendments on portability and multi-vendor provisions in the Council working party and the Parliament's committee stage.
  - **Q1 2027:** The Commission should include portability and multi-vendor clauses in its own contracts with AI providers, setting itself as the model for national procurement.
  - **Q2 2027:** Portability terms should be strengthened while the procurement directives are being negotiated. Funding should be allocated under the Digital Europe Programme to support the development of switching tools.
- **Success indicators:** Portability requirements and stronger multi-vendor provisions in the trilogue text of CADA; portability clauses added to new EU AI contracts.

CONSIDERATIONS

- These proposals can be mirrored at the national level to similarly prevent lock-in (see national recommendations below).

02 Fund a collective fast-follower model and European-only high-security hosting

ACTION

Allocate European funding to a collective European fast follower. The aim is to prevent funding from being diffused between different Member State candidates and ensure that Europe has the best chance of maintaining a credible fallback option for certain applications. To ensure European sovereignty for critical security purposes, develop a set of European-only hosting sites that are EU-owned and operated. Services hosted at these sites should be certified at the proposed CADA assurance Level 4.

IMPLEMENTATION

- **Potential instruments:**
  - Existing funding vehicles can support the development of gigafactories, for example the EuroHPC Joint Undertaking, which has an open call for AI gigafactories, and InvestAI, which includes a fund for AI gigafactories.

- The proposed CADA for providing certification of services hosted in European-only sites. Level 4 will require EU ownership and control as well as a 'high' cybersecurity certificate.
- The Cybersecurity Act's European Cybersecurity Certification Scheme for Cloud Services (EUCS) provides standards for high-security hosting.

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Establish a consortium to support the development of a European fast follower, with co-funding committed by Member States.
  - **Q1 2027:** The Commission and EuroHPC Governing Board amend the selection criteria for AI gigafactories, ensuring that at least one facility is designated for European-only high-security hosting.
  - **Q2 2027:** Develop a certification for high-security, European-only hosting based on CADA Level 4 and the EUCS scheme.
- **Success indicators:** A low capability lag of the European fast follower compared to the commercial frontier; the opening of a certified European-only hosting site.

CONSIDERATIONS

- While a fast follower would reduce dependence by providing a domestic option in domains where frontier performance is unnecessary, it would likely not in itself provide an adequate fallback for a world with transformative AI.
- If frontier developers do not release their best models publicly, emulation is more difficult, meaning that the strategic value of having a fast follower may decrease as compounding advantages widen the distance to the frontier.

03 Certify Europe as a secure importer and coordinate distillation enforcement

ACTION

Establish an EU-wide certification for secure hosting of frontier models to reduce the perceived risk for foreign companies of hosting their models in Europe. Data centre and model weight security requirements should be benchmarked to the security standards set by frontier AI companies and the US government. National enforcement to prevent model distillation and theft should be harmonised to ensure the highest standards apply wherever frontier models are hosted. Whenever possible, shared instruments and standards should be developed with US counterparts.

IMPLEMENTATION

- **Potential instruments:**
  - The CADA assurance framework (Levels 3–4) may be able to provide the scaffolding for a certification system for secure data centres. The CADA assurance levels should be more closely aligned with the hardware controls needed to credibly guarantee data centre security; the CADA framework may need to be amended to include these hardware-related requirements.
  - Cybersecurity Act certification (which is currently stalled) to provide the technical security scheme.
  - Certification can be benchmarked to counterpart measures in the US, for example, on adversarial distillation and model weight security.
  - Common instruments can be developed with US partners through closed forums, for example, the (currently dormant) Trade and Technology Council (TTC).

- **Potential first steps (next 12 months):**
  - **Q4 2026:** Benchmarking of CADA Level 4 and EUCS-high requirements, possibly using RAND's security-level benchmarks as the model (see Considerations below). Publish for international partner visibility.
  - **Q1 2027:** Pilot certification of a small number of sites. Form an EU–US technical working group on secure hosting and distillation enforcement, potentially inside a rebooted TTC.
  - **Q2 2027:** Establish a coordination protocol for national enforcement, to secure against industrial-scale distillation from EU bases. Ensure consistency between Member States and US measures.
- **Success indicators:** Number of sites certified to US standards; frontier model weights actually hosted on EU soil under the scheme; US recognition of the certification in hosting approvals; documented national enforcement actions against distillation that are consistent across Member States and coordinated with US action.

CONSIDERATIONS

- RAND's security-level benchmarks can inform certification categories. Security Level 5 (SL5), for instance, refers to systems designed to resist top-tier state attackers, up to the highest standard needed to defend against attackers from the most cyber-capable institutions.
- The primary audience for the certification is foreign governments and developers, so security measures must be visible to these actors.

Recommendations at the national level

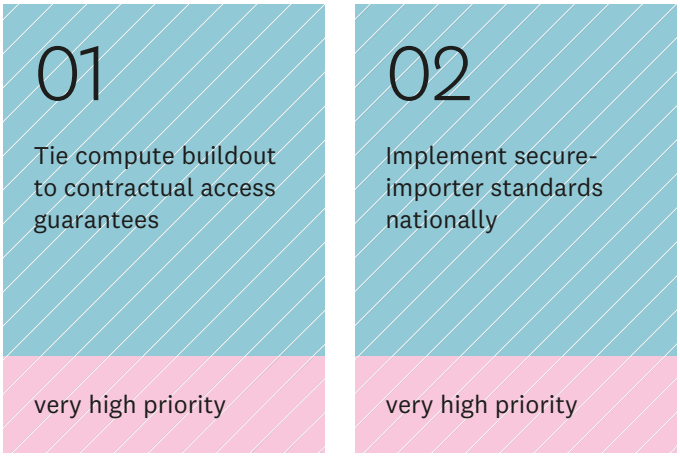

01 Tie compute buildout to contractual access guarantees

ACTION

Member States should treat the inputs needed for frontier AI data centres, such as sites, energy, or connections to the grid, as strategic assets. They should be made available to frontier model providers only when attached to contractual access guarantees that include parity of access with foreign customers, continuity of service commitments, and European access to any security programmes with limited access (like Project Glasswing). These access guarantee provisions should apply to any hosting agreements above 100 MW, with Member States notifying the Commission of all such arrangements.

MEMBER STATE MODE

Kind-specific, depending on the particular assets of a given state or region.

IMPLEMENTATION

- **Potential instruments:**
  - National investment screening to identify relevant hosting arrangements.
  - Member State laws on permitting, energy, and grid connections, depending on national legal provisions and assets relevant to the deal.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Develop a playbook for Member States, covering what guarantees should include and how they should be structured. This should cover access parity, continuity of access, and participation in limited access programmes.
  - **Q1 2027:** Member States should designate a negotiating unit spanning, *inter alia*, energy and digital ministries, depending on the configuration of assets and governing authority within a given state. Member States should notify the Commission whenever new deals are struck.
- **Success indicators:** Share of newly contracted hosted capacity (MW) covered by access-guarantee clauses; notifications filed with the Commission; access parity being extended to European governments or companies during a restricted model release.

CONSIDERATIONS

- The leverage from such compute buildout assets will decay as the US closes gaps in energy, capital, and talent.
- National governments should mirror EU-level portability and multi-vendor procurement terms in public procurement and in the hosting deals themselves. As for procurement, Member States should avoid infrastructure deals that bind national systems to a single model provider.

02 Implement secure-importer standards nationally

ACTION

Frontier AI data centres in Member States should be security certified up to the level demanded by foreign AI providers and their governments. This should include enforcing KYC obligations as well as developing national enforcement competence aimed at preventing model distillation and theft.

MEMBER STATE MODE

Universal for hosting states, applying to any Member State hosting or seeking to host frontier AI compute.

IMPLEMENTATION

- **Potential instruments:**
  - Frontier AI data centres and cloud operators should be designated as critical entities, subject to heightened security requirements.
  - Know-your-customer obligations should apply to compute providers.
  - Distillation-enforcement competence should be assigned to a named national authority, coupled with Europe-wide enforcement harmonisation.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Audit existing and planned frontier-relevant sites.

- **Q1 2027:** Adopt national KYC guidance for compute providers, aligned with the emerging Union scheme (Union-level recommendation 3). Table national legislation (or issue administrative direction) establishing distillation-enforcement competence and assigning it to a named national authority, with a Union coordination layer to ensure consistency (Union-level recommendation 3).
- **Q2 2027:** Hold an exchange with US counterparts to discuss common enforcement requirements, for example, for preventing distillation.

- **Success indicators:** Number of national sites meeting SL4/SL5-class benchmarks; KYC coverage of national compute sales (share of capacity sold under verified-customer terms); documented enforcement actions against distillation taken in step with Union and US action.

CONSIDERATIONS

- This is a relatively cheap lever that protects against the US imposing security-motivated restrictions by reducing the perceived security risk of allowing capable models to run abroad.
- Enforcement asymmetries between Member States will be exploited by evaders and may be used as a basis for denying access. Member States therefore need to work together through harmonised rules.

O1.2 Protect European assets

WHY IT MATTERS

Europe does not currently systematically review how much of its own capital and technology is flowing into foreign AI and semiconductor industries as part of outbound investments. Member States do not have the legal power to compel a firm to disclose what it is building overseas. Moreover, some AI assets of strategic importance are not fully covered by Europe's inbound screening architecture, and Europe lacks a buyer of last resort that can inject capital when an asset is at imminent risk of being acquired.

Recommendations at the Union level

01	02
Adopt an outbound-investment regulation	Make the outbound-investment review annual and comparable
high priority	high priority

01

Adopt an outbound-investment regulation

ACTION

Currently, the European Union does not have a regulation governing foreign investments of European firms. As of today, there exists only a Commission Recommendation asking Member States to review outbound investments. The EU should introduce a new regulation on the review of outbound investments related to critical technologies, modelled after Regulation (EU) 2019/452. The following principles should apply to outbound reviews:

- Each Member State has autonomy to choose which investments to screen. If a Member State does review an investment, it should meet common minimum standards and reviews should be shared in a common information-sharing system.
- A Member State should notify the Commission and other Member States when it chooses to review a transaction. Other Member States may comment, and the Commission may issue an opinion.
- If a Member State chooses not to screen a transaction in its jurisdiction, another Member State may still comment and flag it for an opinion from the Commission. Competent Author-

ities should be established and equipped with the legal right to request details on firms' investments.

- To minimise administrative overhead, requests should be limited in scope and any information shared with the aforementioned authority must remain confidential; The Commission should establish a classified system and database for exchanging case information.

IMPLEMENTATION

- **Potential instruments:** A new regulation on the review of outbound investments modelled after the legal basis established in Regulation (EU) 2019/452.
- **Potential first steps (next 12 months):**
  - **Late 2026:** The Commission mandates the Expert Group on Outbound Investment to outline the regulation's scope, drawing on the comparative assessment published under Union-level recommendation 2.
  - **Early 2027:** The Commission opens a public consultation and starts work on an impact assessment, drawing on the comparative assessment published under Union-level recommendation 2.
  - **Mid-2027:** The Commission presents the proposal.
  - **Late 2027:** The Council working party begins examination and the European Parliament appoints a rapporteur.
- **Success indicators:** Commission proposal presented by mid-2027; number of Member States with an outbound review regime under the regulation; number of transactions flagged through the cooperation mechanism per year; number of comments and opinions issued per year; number of staff assigned to outbound review per Member State.

CONSIDERATIONS

- A Union-level framework that requires (rather than facilitates and coordinates) outbound-investment screening is likely not politically feasible at this point. When the European Commission tabled its proposal for what would become Regulation (EU) 2019/452, 11 Member States already operated a national FDI screening mechanism. By contrast, Member States do not systematically review and assess outbound investment for security purposes as of today. Since this proposed regulation would require all Member States to build an institution that none currently has, it is unlikely to be readily adopted. However, as the experience with the 2019 FDI screening regulation shows, a voluntary outbound-investment regime modelled on that regulation could prepare the ground for a future regulation mandating outbound screening within a specific scope.
- A voluntary regime would still reach cross-border cases. Under Article 7 of Regulation (EU) 2019/452, Member States may comment on an investment that another Member State is not screening, the Commission may issue an opinion, and the Member State concerned must give due consideration to both. A Member State that chooses not to review is therefore not outside the system. In cases where a company is incorporated in one Member State while the relevant technology sits in a subsidiary in another, the second Member State can raise the case even if the first has no regime of its own.
- Harmonising reviews and risk assessments will help make transaction definitions and risk categories comparable among Member States. It will also build contact points across the Union by putting officials in touch with each other and enabling them to create secure lines of communication.
- This regulation would not cover research arrangements. Imec, for example, may host American companies in its laboratories under collaboration agreements, which transfer technology and know-how without any investment taking place. Such arrangements fall outside any

outbound-investment instrument. Closing this gap will require a separate regulatory intervention.

02 Make the outbound-investment review annual and comparable

ACTION

Recommendation (EU) 2025/63 asked Member States to take stock of their outbound investments in AI, semiconductors, and quantum, and to assess these investments for their economic security risks. However, this was a one-off exercise. The Commission should issue a follow-up recommendation asking Member States to review and risk-assess their outbound investments yearly during the same time period and using the same format. The Commission should aggregate the results in a yearly report that covers trends in deals and risks that arise across countries, and how Member States differ in their handling of these risks. This report can be modelled on the Commission's recurrent report on the screening of foreign direct investment into the Union. The report should also be transparent about where the underlying data is too thin to support a conclusion. Moreover, the Commission should leverage the results of the initial stocktake under Recommendation (EU) 2025/63 to agree on what transactions to review and how to score risks.

IMPLEMENTATION

- **Potential instruments:** Recommendation (EU) 2025/63 and a successor recommendation establishing the annual cycle; the Commission's own analytical capacity for the comparative assessment.
- **Potential first steps (next 12 months):**
  - **Late 2026:** The Commission publishes a first comparative assessment of the national review and risk assessment reports.
  - **Mid-2027:** The Commission adopts a successor recommendation fixing the annual reporting date, with the first annual cycle covering 2027.
- **Success indicators:** Number of Member States submitting complete reports; publication of an annual Commission comparative assessment.

CONSIDERATIONS

- According to Recommendation (EU) 2025/63, the review and risk assessment should inform subsequent policy choices. However, the EU should not wait for their outcome before deciding whether EU-level legislation is needed. Member States currently lack a legal instrument to request outbound-investment information from firms, so the exercise cannot be expected to generate a complete evidence base. If only a few identified risks are found, this could reflect the limits of voluntary information-gathering rather than the absence of risk, and it should not be treated as a reason to close the file. The legal instrument that would give Member States such a right-to-ask is proposed in Union-level recommendation 1.
- In its 2024 White Paper on outbound-investment controls, the Commission proposed to validate and compare the Member State reports in a comprehensive report of its own. Recommendation (EU) 2025/63 no longer mentions this, leaving it unclear whether it will happen. The Commission should produce this report.
- An annual outbound-investment review would be valuable not only for what it finds, but also for what it builds: A recurring cycle of reviews and risk assessments will instil a culture of monitoring and prudence in national administrations. It will also raise awareness among firms that they might have to disclose their outbound investments.
- This recommendation is also meant to prepare the ground for the outbound-investment regulation proposed in Union-level recommendation 1. If the Union creates a template and a risk methodology for the purposes of this review, this can later be codified, so that the regulation is based on an instrument that administrations already use.

Recommendations at the national level

01	02	03
Implement national outbound-investment controls	Extend the scope of national inbound FDI screening	Name a national buyer of last resort for strategic assets
high priority	high priority	high priority

01 Implement national outbound-investment controls

ACTION

Member States with strategic AI-stack assets should introduce national outbound-investment controls. Which transactions are in scope can be determined based on Recommendation (EU) 2025/63, namely advanced semiconductors, artificial intelligence, and quantum technologies. The broader sectors highlighted in national-level recommendation 2 below should also be included. To minimise red tape, legislation should differentiate between two sets of transactions – a small set of transactions that are prohibited per se, as well as a broader set of transactions that require companies to notify relevant authorities after the investment is completed. Which transactions fall into which category should be determined in secondary legislation, so that they can be updated as technology evolves. To prevent regimes from overlapping, transfers which already require a licence under the Dual-Use Regulation (Regulation (EU) 2021/821) should be excluded. Moreover, the authority responsible for inbound screening should be put in charge of implementation, so that the controls can draw on its expertise and established processes, rather than creating a parallel body. Adding one to three additional full-time staff members per Member State would be sufficient to cover these tasks.

MEMBER STATE MODE

Kind-specific – only Member States hosting strategic AI-stack assets (e.g. the Netherlands, Germany, France, Belgium, Austria).

IMPLEMENTATION

- **Potential instruments:** National primary legislation, either based on dedicated outbound-investment control legislation or an extension of an existing inbound investment screening or foreign trade instrument. Secondary legislation for the prohibited and notifiable lists, so that they can be updated without reopening the statute.
- **Potential first steps (next 12 months):**
  - **September–December 2026:** The responsible ministry assesses whether the existing screening or foreign trade statute may cover outbound controls.
  - **Early 2027:** A draft bill is prepared and industry is consulted on the notification burden.
  - **Mid-2027:** The bill is introduced in parliament and a budget line for one to three additional full-time staff is agreed.

- **Success indicators:** Number of Member States hosting strategic AI-stack assets with an outbound regime in force; number of transactions notified per year.

CONSIDERATIONS

- Although national action is recommended first, the Union should still implement its own outbound-investment controls. National regimes can be in force within two years and can later be folded into the Union framework proposed in Union-level recommendation 1, which is designed to accommodate divergent national regimes.
- There is a real risk of creating a patchwork of national legislation. To avoid this, the Member States moving first should align on scope and definitions in the Expert Group from the start.
- As highlighted above in Union-level recommendation 1, the regime will also require governments to have a statutory right to ask companies established in the Member State about their outbound investments, coupled with a duty to keep confidentiality and protect commercially sensitive information. This right should be limited to the small set of technologies in the regulation to prevent undue bureaucracy.
- The United States has operated an outbound-investment regime since January 2025 under 31 CFR Part 850, which distinguishes a narrow set of prohibited transactions from a broader set requiring notification within 30 days of closing, with no approval pathway at all. This demonstrates that a notification-based regime is manageable and imposes low costs on investors.

02 Extend the scope of national inbound FDI screening

ACTION

The revised foreign investment screening Regulation (EU) 2026/1386 requires Member States to have a specified national screening mechanism in place by January 2028.[24] As per Article 4(16), Member States may apply their screening mechanism to foreign investments beyond the minimum scope that is determined in Article 4(15) and Annex I. Member States that hold strategic AI assets should use this option and bring these strategic assets within the scope of their screening mechanism. In particular, the target list should be extended to investments in companies that: (i) develop specialised AI systems for narrow but security-relevant purposes, such as biological and chemical design, offensive cyber capabilities, and the control of critical industrial processes; (ii) service or maintain semiconductor manufacturing equipment, such as ASML's lithography machines; (iii) produce equipment on which critical digital infrastructure depends, including radio access networks and 6G equipment, submarine fibre-optic cable systems together with the vessels that install and repair them, and high-voltage grid equipment including large power transformers; (iv) hold strategic training datasets. Separately, Member States should apply Article 4(15)(d) as it stands to designate existing AI data centre operators as critical digital infrastructure.

MEMBER STATE MODE

Kind-specific – only Member States hosting strategic AI-stack assets (e.g. the Netherlands, Germany, France, Belgium, Austria). Universal for AI data centres.

24 As part of the Industrial Accelerator Act (COM(2026)100) the Commission proposed an investment control regime in 2026 that is parallel to the revised FDI screening regulation (Regulation (EU) 2026/1386). However, unlike the FDI screening regulation, it is not designed to protect European assets but to leverage Europe's market access to extract valuable capabilities in sectors where Europe is lagging behind, such as battery technologies and electric vehicles.

IMPLEMENTATION

- **Potential instruments:** The national screening instrument currently in drafting ahead of the January 2028 deadline; the designation of existing data centre operators as critical digital infrastructure, requiring no primary legislation.
- **Potential first steps (next 12 months):**
  - **September 2026–September 2027:** Passing the bills that extend the mandatory scope to categories (i) to (iv).
  - **March 2027:** Article 4(15)(d) designation of existing AI data centre operators adopted.
- **Success indicators:** Number of the four asset categories within mandatory national scope, per Member State; whether AI data centre operators are designated as critical infrastructure.

CONSIDERATIONS

- There is no Union-level legislative counterpart to this recommendation. However, the European Commission may consider issuing a non-binding recommendation for specific Member States to expand their minimum scope. Further, the Commission can give additional guidance on screening methodologies for AI-relevant sectors.
- Greenfield investments are not within the scope of the revised inbound FDI screening regulation. This exclusion should stay as it is. Bringing greenfield projects into the mechanism would impose notification duties on many transactions that raise no economic security concern, and the EU has an interest in remaining an attractive destination for such investments.

03 Name a national buyer of last resort for strategic assets

ACTION

Inbound FDI screening mechanisms may block the foreign acquisition of a strategic AI asset, but that does not solve the owner's potential capital constraints. To make sure that the relevant asset remains under European control, the Member States hosting these assets should establish an anti-capture facility to quickly mobilise the needed public and private capital in such situations. Rather than building an entirely new facility, Member States should designate an existing public bank or investment vehicle (such as Bpifrance in France, SFPIM in Belgium, or Invest-NL in the Netherlands). Authorities should make sure that the money is available on demand, and headline figures together with per-case ceilings should be published in advance. In addition, a confidential list of qualifying assets or companies should be maintained and reviewed regularly. This way, an asset's eligibility is determined before a bid arrives, rather than having to be negotiated under time pressure. There should also be a channel of information exchange between the screening authorities and the capital facility, allowing the screening authority to alert the fund in critical situations.

MEMBER STATE MODE

Kind-specific – only Member States holding AI-stack companies.

IMPLEMENTATION

- **Potential instruments:** A ministerial instruction to an existing public bank or fund.
- **Potential first steps (next 12 months):**
  - **November 2026:** Draft the confidential list of qualifying assets.
  - **February 2027:** Name the fund and issue the mandate, including the per-case ceiling.
  - **Mid-2027:** Legislate the information gateway alongside the national measures to be notified under Article 3(2) of Regulation (EU) 2026/1386 ahead of the January 2028 deadline.

- **Success indicators:** Member States with a named fund and decision-maker; callable funding committed per country, in euros; days from trigger to final decision.

CONSIDERATIONS

- This fund can be activated on short notice when a foreign buyer offers to buy an EU AI-stack company and this company needs additional financing, either because it risks bankruptcy or because it needs funding to scale up. The aim is to prevent foreign control of strategic assets on specific occasions. Care should be taken to prevent firms from strategically pretending to have a potential foreign acquirer in order to unlock this fund. Responding to existing offers will also help to preclude any issues under state aid rules.

- Give the public bank or fund a standing mandate to acquire assets on the national list, and let its own investment committee take the decision under its normal rules so as to enable swift decision-making.

- The main reason to place this fund at Member State rather than Union level is that existing EU funding has mandates that do not fit. The European Tech Champions Initiative invests in other funds, not companies. The EIC (Fund) invests between €0.5 million and €30 million for minority stakes. The Scaleup Europe Fund is externally managed by EQT, which takes all investment decisions on commercial terms.

- Not every Member State can do this. A country with a small screening unit and no public investment vehicle will not execute a large purchase in six weeks.

Pillar 2

Building economic strength

O2.1 Make Europe the best place to build and scale high-growth companies

WHY IT MATTERS

Transformative AI may allow young innovative firms to reach global scale at unprecedented speed, concentrating fiscal value and strategic capability where those firms are located. The known competitiveness issues in the European Union therefore need to be pursued at a higher level of ambition, and most importantly, with a higher sense of urgency. It is therefore crucial that Member States and Union-level institutions collaborate towards the common goal of competitiveness to avoid further fragmentation of important input markets (first and foremost capital and labour). The recommendations below are mutually reinforcing. Implementing only selected elements will substantially reduce their effectiveness, or could even generate unintended outcomes. It is therefore crucial that the proposed reforms are pursued as a coherent package.

Recommendations at the Union level

01 Unlock institutional capital for European growth companies	02 Equip European science for the AI age	03 Attract the world's best AI talent
very high priority	very high priority	very high priority
04 Create European lead markets for defence innovation	05 Establish a Single European Growth Market	
high priority	high priority	

01 Unlock institutional capital for European growth companies

ACTION

The Union should mobilise institutional capital for European growth firms. To achieve this, the Commission should launch a Tibi/WIN-style initiative at the Union level to secure investment commitments into European growth funds from large investors. The Union should make it easier for insurers and pension funds to invest in long-term equity while maintaining reasonable prudential safeguards. The Commission should widen General Block Exemption Regulation (GBER) provisions so Member States can fund viable innovation projects that cannot secure private financing because of market failures. Finally, the Market Integration and Supervision Package should become a core deliverable of the Savings and Investments Union.

IMPLEMENTATION WORKSTREAMS

- **Launch the European Institutional Investors Pact for Innovation:** The Commission, together with the EIB, should launch the planned European Institutional Investors Pact for Innovation (hereafter 'the Pact') to secure voluntary commitments from investors to invest in European growth funds. This Pact should set common eligibility criteria for growth funds. Annual reports of commitments raised should distinguish between announced versus legally committed and deployed capital, using common definitions. The additionality of these commitments (whether the capital would have been mobilised in the absence of the Commission's intervention) should be assessed independently. The EU-level commitments should work alongside national schemes without the same investment being counted twice in both the Union's and a Member State's totals. Investment decisions and due diligence should rest with participating investors.

- **Unlock institutional investment in long-term equity:** The Commission and Member States should swiftly implement the Solvency II reform to make it easier for insurance companies to invest in equities and other long-term assets. To ensure this reform works as intended across the EU, the European Insurance and Occupational Pensions Authority (EIOPA) and national supervisors must apply the revised long-term equity rules consistently, so equivalent investments face the same capital charge. Pension funds should not face unnecessary restrictions on their investment choices. Parliament and Council should therefore swiftly adopt the reform of the Directive on Institutions for Occupational Retirement Provision (IORP II). Nonetheless, funds should preserve the prudent-person principle.

- **Financing for first commercial deployment:** The Commission should create a streamlined route in the ongoing GBER review for Member States to support financing of first-of-a-kind (FOAK) and first-commercial-deployment projects, where there are demonstrated market failures. This support should be targeted and remain limited to the minimum necessary. It should absorb risks for private investors to crowd in private capital. The Commission and Member States should maintain taxpayer safeguards proportionate to the size of the aid. Safeguards should be designed to minimise administrative burden for beneficiary firms.

DECISIONS AND FIRST STEPS IN THE NEXT 12 MONTHS

- **By the end of 2026:** Commission and EIB launch the European Institutional Investors Pact for Innovation with a first group of institutional investors. Participating investors should report capital committed and deployed separately.

- **By the end of 2026:** The Commission incorporates a defined FOAK and first-commercial-deployment route into the final GBER.

- **Ahead of 30 January 2027:** The Commission should adopt the remaining Solvency II technical standards and ensure that EIOPA and national supervisors are ready to apply the revised long-term equity treatment consistently. From application, EIOPA should monitor divergent supervisory practices. Where national law or administrative practice conflicts with Union law, the Commission should take appropriate measures.

- **During the IORP II negotiations:** The European Parliament should adopt its position promptly to enable an agreement that preserves a risk-based approach to long-term investment.

MEASURES OF PROGRESS

- **Implementation indicators:** Number of participating investors in the Pact, and their assets under management; number of national schemes recognised under the Pact; adoption of the FOAK State-aid route.
- **Performance indicators:**
  - *For the Pact:* Announced and deployed capital; independently estimated additional capital; risk-adjusted returns by vintage and asset class.
  - *For institutional capital:* Annual net purchases and holdings of EU equities by EU pension funds and insurers (€ and % of assets); net new financing raised by EU firms from EU institutional investors.
  - *For FOAK projects:* Share of FOAK projects reaching financial close on schedule; share entering commercial operation on schedule; follow-on deployment without equivalent FOAK support.
- **Longer-term outcomes for evaluation:** Depth of EU late-stage financing; number and value of European growth rounds; size and performance of EU-based growth funds; share of European-founded firms able to scale substantially in Europe.

CONSIDERATIONS

- Pact reporting should show how mobilised capital is distributed across investors, funds, sectors, and Member States. It should assess whether the Pact is truly deepening Union capital markets and broadening access to growth capital.
- Public co-investment of FOAK projects should avoid picking national winners. Therefore, it should be based on competitive selection (if multiple alternatives exist), proportionate aid, and a demonstrated financing gap.
- Crowding out of private investment through GBER reforms should be evaluated and monitored. Where possible this should be done using ex-post comparison with similar projects closely below funding thresholds for better causal inference.

02 Equip European science for the AI age

ACTION

The proposed Horizon Europe programme for 2028–2034 should address three market failures that currently limit the productivity of European research and innovation and hold back breakthrough science. First, it should support the development and deployment of scientific public goods (e.g. advanced research tools) and advanced research infrastructure (e.g. self-driving laboratories). Second, it should create organisational capabilities that cannot easily be provided through conventional project grants or temporary research consortia. Third, it should strengthen incentives to produce reliable, reproducible, and complete scientific knowledge to prepare the European research sector for AI-based science.

IMPLEMENTATION WORKSTREAMS

- **Scientific tools and infrastructure as public goods:** The Horizon Europe 2028–2034 programme should include a dedicated funding line for both the development and deployment of trusted, interoperable, and reusable research tools and infrastructure. This includes but is not limited to scientific software, AI models, curated datasets, advanced instruments, self-driving laboratories, and other experimental automation tools. Funding should also support technical teams that are required to maintain, improve, and support this infrastructure. Funding should be open to universities, research organisations, nonprofits, and companies,

and reward reliability, scientific uptake, community adoption, and downstream research impact. Where fragmented demand prevents promising scientific tools from being developed or deployed, the Commission should implement advance market commitments or comparable demand-pull instruments tied to ambitious, measurable performance requirements for these tools.

- **New organisational capabilities:** The Horizon Europe 2028–2034 programme should establish new organisational capabilities fit for addressing ambitious scientific and technological goals. It should complement the ERC's support for individual researchers with a mechanism focused on *organisational* excellence, to address scientific problems that require large (interdisciplinary) teams. It should be open to new experimental organisational forms like Focused Research Organisations (FROs): purpose-built, time-limited teams combining scientists, engineers, and operators to overcome defined scientific bottlenecks (e.g. building a new instrument, biological resource, or open experimental infrastructure). The mechanism should provide high-performing research units and novel organisational models with large flexible grants, subject to rigorous mid-term review and renewable for up to ten years. Selection should be international and excellence-based, with no geographic allocation or presumption of permanent funding. In the ongoing negotiations, Parliament and Council should ensure that the proposed ARPA-like function of the EIC will have legally protected operational autonomy (building on the German SPRIND experience), allowing expert-led control over programme design, portfolio management, and funding decisions, while maintaining transparent objectives, robust conflict-of-interest safeguards, and independent external evaluation.

- **Reliable science for the AI era:** AI is already dramatically increasing the volume and speed of scientific output. This makes it even more important to improve the reliability, transparency, and evaluability of research. The Horizon Europe 2028–2034 programme should therefore both aim to make individual scientific results more reliable, and invest in the capacity to systematically evaluate how the science system works. This requires infrastructure that enables both humans and AI systems to assess, verify, and build on scientific knowledge. Horizon Europe should therefore support both the verification of scientific findings through replication, validation, and shared benchmarking, and the creation of comprehensive machine-readable scientific records that capture protocols, workflows, code, datasets, intermediate outputs, and negative or null results. On top of this, the Commission should establish a Metascience and Policy Experimentation Unit to monitor and evaluate the productivity of European science, with the authority to test and evaluate alternative ways of funding research. These include, but are not limited to: alternative selection procedures, like partial lotteries among proposals deemed equally fundable after expert review; faster and more flexible grants; milestone funding; and new organisational models. Successful approaches should have a defined route into mainstream programmes.

DECISIONS AND FIRST STEPS IN THE NEXT 12 MONTHS

- **Autumn 2026: Secure the legal basis and budget.** In the Horizon Europe and MFF negotiations, Parliament and Council should secure funding for shared scientific tools, new research organisations, research verification, and metascience experiments. They should also grant the EIC full operational autonomy to design ARPA-style programmes and manage funding portfolios. The Commission should support these priorities in the negotiations and determine which activities belong under Horizon Europe and which belong under the European Competitiveness Fund.

- **Winter 2026/27 to spring 2027: Define how each proposal will work.** The Commission should determine how each proposal will be funded and delivered. It should develop rules for the funding of long-term infrastructure funding, FRO-style grants, and the development of new scientific tools. These rules should define who can apply, what exactly will be funded, how proposals will be selected, and how performance will be assessed. To develop infrastructure for replication and machine-readable scientific records, the Commission should involve open science experts and technical teams with experience building widely used scientific databases and software. The Commission should define the mandate of the Metascience Unit

and ensure that it has sufficient budget and legal freedom to run randomised experiments. Research funders and scientific organisations should help refine all of these arrangements and identify suitable first projects.

- **Spring to autumn 2027: Prepare the first funding rounds.** The Commission and Member States should incorporate the proposals into Horizon Europe's Strategic Plan and the 2028–2029 Work Programmes. They should allocate budgets and assign responsibility for delivery. Finally, the Commission should appoint the leadership of the Metascience Unit and experienced programme managers to turn the EIC into a true ARPA-like organisation.

- **From 2028: Fund the first projects and expand what works.** The Commission should open the first calls for proposals for scientific tools and infrastructure and select an initial portfolio of projects. The EIC's ARPA-like function should launch its first programmes. Approaches that perform well should receive larger budgets or be incorporated into later Work Programmes.

MEASURES OF PROGRESS

- **Implementation indicators:** Number of dedicated funding instruments and attached budget for all the different measures; share of supported research infrastructure receiving long-term funding.

- **Performance indicators:** Number of institutions using the funded tools; cost and time required to produce and validate scientific results, by field; reliability and technical performance (e.g. precision) of tools; share of funding experiments that lead to changes in funding practice.

- **Longer-term outcomes for evaluation:** Scientific output quality, e.g. the number of discoveries deemed important by the field community; user rates of European research tools and infrastructure.

CONSIDERATIONS

- Long-term support for new institutions and infrastructure should be periodically reviewed and subject to competition, to keep funding contestable.

- New instruments should be coordinated with existing programmes to avoid duplication and fragmentation.

- When using demand-pull tools, reward adoption and scientific use; do not simply reward the creation of tools, datasets, or infrastructure.

- Balance strategic investments and investigator-driven research to preserve scientific diversity and independence.

- Make sure that reproducibility and reporting requirements improve scientific quality without unreasonable administrative burdens.

- Experiment with different policies: Scale successful approaches and discontinue those that do not deliver results.

03 Attract the world's best AI talent

ACTION

The Union should make it dramatically easier for international founders to hire top (AI) talent from across the world. To support recruitment and retention, young innovative firms should be able to offer competitive and portable equity compensation across the Single Market. The Union should also create a Founder Gateway: a single digital front end that founders can use to navigate the administrative procedures involved in setting up a business. This should include simpler, better-coordinated immigration procedures for the founders and talent that startups and growth firms need, while admission decisions should still rest with national authorities.

IMPLEMENTATION WORKSTREAM

- **Competitive and portable equity compensation.** Parliament and the Council should keep a workable employee stock option scheme in the EU Inc. Regulation; tax on stock options should be due only when employees obtain liquidity. The Commission should collaborate with Member States to make sure that the tax is still deferred when employees relocate and to prevent double taxation. These arrangements should protect legitimate national tax bases and prevent employees from shifting employment-related income to a lower-tax jurisdiction shortly before liquidity. If binding common rules are needed, the Commission should propose the appropriate legal instrument. Otherwise, Member States should coordinate on national rules.

- **End-to-end Founder Gateway:** The Commission should create a Founder Gateway within the Single Digital Gateway. Through this, founders should be able to access EU Inc. services, the European Business Wallet, and national administrative procedures by following a coherent user journey, built on the once-only principle. Following the UK Government Digital Service model, the Gateway should be built to consumer-grade usability standards under empowered Commission product leadership, with permanent capacity for multidisciplinary design, engineering, and service-delivery. Development should be iterative and driven by continuous testing with founders.

- **Bringing innovation talent to Europe:** Next to facilitating the formation of businesses, the Founder Gateway should also streamline the administrative layer of immigration processes for founders, researchers, and highly skilled employees needed by startups and growth firms. The Gateway should provide access to relevant EU and national migration routes (including the Blue Card and Single Permit frameworks), and first and foremost to the dedicated founder pathways to be established by Member States (see the accompanying recommendations at the national level). It should support digital applications, status tracking, and coordinated processing of linked family applications. The Commission should extend the EU Digital Identity Wallet and Once-Only Technical System to allow verified personal and business information to be reused across national procedures. Individual admission decisions should nevertheless remain with the national authorities.

DECISIONS AND FIRST STEPS IN THE NEXT 12 MONTHS

- **Sep–Nov 2026: Establish governance and define scope:** The Commission should designate a product lead for the Founder Gateway. This lead will be responsible for its delivery roadmap and budget and empowered to coordinate work across Commission services. The product lead should then establish a joint implementation group with willing Member States and the national authorities responsible for the initial founder journeys. They should then publish a roadmap for a minimum viable product (MVP) identifying the initial journeys and participating Member States and defining milestones, responsibilities, funding, and performance metrics. In the EU Inc. negotiations, Parliament and the Council should make sure that employees are taxed only when they sell their shares or otherwise receive liquidity. The Commission, together with Member State finance ministries and tax authorities, should start to develop rules to preserve tax deferral after relocation and prevent double taxation, ensuring the employment-related tax base is allocated according to where it was earned.

- **Dec 2026–Feb 2027: Resolve legal and policy design issues:** The Commission should review where gaps in EU migration rules make it more difficult for founders, researchers, and highly skilled workers to migrate to Europe, and propose changes or additional common permit categories to address these. The Gateway product lead and participating Member States should agree on common Gateway service standards, covering digital applications, case tracking, accessibility in English, and the reuse of verified information. The Commission and participating Member States should identify the legal basis and data-sharing arrangements required for the initial Gateway journeys. They should prepare any necessary EU or national legislative changes in parallel with product development.

- **Mar–May 2027: Build the first software release:** The Gateway product lead should establish a permanent product team, with dedicated design, engineering, and service-operations capabilities. This product team should develop the core Gateway service layer, incorporating the European Business Wallet. Together with delivery teams from participating Member States, it should connect the national systems needed for selected MVP journeys. The Commission product team should develop common user access, status-tracking, and secure data-sharing functions for the MVP. Once an MVP is developed, founders and growth firms should be asked to iteratively test complete journeys, and the results should be published.

- **Jun–Aug 2027: Pilot and prepare rollout:** The Commission and willing Member States should launch the Gateway's pilot. Collectively, the pilot should cover at least one founder journey in each of three areas: company formation, talent admission, and secure business-data sharing. Individual Member States may initially connect different journeys according to where progress can be made most quickly. Participating Member States should test that relevant credentials can be reused across borders. The product team should publish performance data for each pilot journey. Following the pilots, the product team should prepare the next phase of rollout, including additional journeys and Member States and preparing for any necessary legislative changes. Product development and legal changes should continue to proceed as parallel workstreams.

MEASURES OF PROGRESS

- **Implementation indicators:** Number of participating Member States connected to the Gateway; number of financial institutions and public authorities accepting verified business data through the Gateway.

- **Performance indicators:** Median time required to establish a company and complete linked administrative procedures through the Gateway; end-to-end task completion rates; median processing time for complete talent visa applications, including accompanying family members; share of applications using previously verified data; number of duplicate information requests; number and share of cross-border employee-equity cases benefiting from recognised portability arrangements.

- **Longer-term outcomes for evaluation:** Share of startups and high-growth firms using employee equity compensation; administrative costs and time-to-market for firms expanding across the Single Market.

CONSIDERATIONS

- The Gateway should integrate with national systems and existing EU infrastructure; it should not itself create a new administrative layer.

- Secure Member State commitment for the Gateway early before investing heavily in technology that depends on national participation.

- Make end-to-end completion of key founder journeys the main objective of the Gateway; redirection to external portals should not be treated as success.

- Resolve legal uncertainty upfront on data-sharing, digital credentials, migration procedures, and cross-border equity compensation.

04 Create European lead markets for defence innovation

ACTION

The Union should create a European common market for defence innovation. To achieve this, it should create integrated procurement pathways to allow common procurement for innovative products for willing Member States. It should further reduce regulatory fragmentation and lower cross-border entry barriers by creating common standards, interoperable interfaces, and shared operational testing, or by facilitating mutual recognition of test results for defence technology.

IMPLEMENTATION WORKSTREAMS

- **Market-building through rapid defence innovation.** The Commission should launch its programme for agile and rapid defence innovation (AGILE) on schedule and meet its four-month time-to-grant target. It should use mission-driven challenges and staged technical milestones where appropriate and provide fast-track access to testing and certification. Each challenge should make it clear how the product can go from successful demonstration to operational testing and potential procurement. The Commission should scale those AGILE mechanisms that are evaluated to be effective in the pilot, under the next Multiannual Financial Framework. It should connect these to successor instruments for common procurement and industrial ramp-up.

- **Shared testing and operational validation.** Europe should reduce inefficient duplication in national testing facilities and procedures to lower costs and make it easier for defence suppliers to enter markets across the Union. The Commission, together with participating Member States, should therefore establish a European network of testing facilities around existing activities of the Hub for EU Defence Innovation (HEDI), EU Defence Innovation Scheme (EUDIS), and AGILE. The network should use common evaluation standards, and national procurement authorities should recognise test results from anywhere across the network. EU-funded projects should have fast access to operational testing, and the next Multiannual Financial Framework (MFF) should fund gaps in shared testing capacity.

- **Support joint defence procurement.** The EU should foster joint procurement by Member States with shared capability needs (see national recommendation 6). The Commission should prioritise European Defence Industry Programme (EDIP) support for joint procurement that replaces separate national orders, based on common functional requirements and binding budget commitments.

- **A single market for defence products.** Parliament and the Council should remove unnecessary legal barriers to cross-border procurement and transfers, and adopt the Defence Readiness Omnibus proposals. The Commission should enforce the resulting rules and monitor whether national practices continue to favour domestic suppliers.

DECISIONS AND FIRST STEPS IN THE NEXT 12 MONTHS

- **Autumn 2026:** Parliament and the Council should formally adopt AGILE and the Defence Readiness Omnibus. The Commission should then finalise AGILE's 2027 Work Programme: It should identify the first innovation contests and define funding stages and performance measures.

- **By January 2027:** The Commission should launch the first AGILE calls.

- **By March 2027:** The Commission and European Defence Agency (EDA) together with interested Member States should identify testing facilities that could join the European network. Member States should identify rules that prevent EU companies from using testing facilities in their country. They should also identify rules that prevent authorities from accepting test results from the network.

- **By June 2027:** The Commission and EDA should set the rules for access to the testing network and work with participating Member States to introduce common evaluation methods. Member States begin removing the barriers identified in March.

- **By September 2027:** Member States begin connecting the first testing facilities to the network.

MEASURES OF PROGRESS

- **Implementation indicators:** Number of testing facilities formally connected to the network.
- **Performance indicators:** Share of AGILE grants signed within four months; share of completed tests accepted by procurement authorities in more than one country; median time from grant signature to first operational test and from successful test to procurement decision; share of procurement contracts awarded to first-time defence suppliers.
- **Longer-term outcomes for evaluation:** Share and volume of cross-border defence contracts across all defence contracts in the Union.

CONSIDERATIONS

- Creating a single market for defence with more competition will ultimately lead to lower procurement costs. Larger and longer-term orders should still be awarded competitively and remain open to new entrants.
- The goal of common standards is to promote interoperability without locking procurement into one supplier or technology. Where technological uncertainty remains high, programmes should preserve competing approaches until operational testing provides enough evidence to choose.

05 Establish a Single European Growth Market

ACTION

Europe's public markets are fragmented across national venues. Concentrating more growth company listings in one market could bring together the investors, analysts, and advisors needed for successful IPOs. This would create a self-reinforcing ecosystem that attracts further listings and investment. While such concentration may occur naturally through competition between various European exchanges, the process is slowed by regulatory fragmentation and language barriers. A designated, but non-binding 'focal venue', could build market depth more quickly by coordinating business and investor expectations. The Commission should therefore coordinate an open process to designate one exchange as Europe's non-binding focal market for growth-company IPOs. To reduce regulatory barriers to deeper public capital markets, Parliament and Council should swiftly agree on the Market Integration and Supervision Package.

IMPLEMENTATION WORKSTREAMS

- **Selecting the focal market.** The Commission should invite European exchanges to present plans for attracting growth-company listings. It should select the strongest candidate using published criteria on liquidity, ease of access for international investors, and the strength of its existing IPO ecosystem. The European Securities and Markets Authority (ESMA), interested Member States, issuers, and institutional investors should advise the Commission, but the final decision should remain with the Commission. The designation should last for a fixed period and be reviewed against published performance measures (see below).
- **Enabling cross-border regulatory framework.** Parliament and Council should swiftly agree on the Market Integration and Supervision Package while preserving the Commission proposal's core market-integration measures. The final rules should provide a workable pan-European market operator (PEMO) framework that allows a single operator to operate trading venues across Member States under one ESMA authorisation, without additional host-country requirements. They should also harmonise trading venue rules and facilitate efficient cross-border access to clearing and settlement services.

DECISIONS AND FIRST STEPS IN THE NEXT 12 MONTHS

- **Sep–Nov 2026: Test market appetite.** The Commission should seek political backing from Member States and identify an initial group willing to support the focal market. In parallel, it should test interest among (capital) market participants. In this process it should focus on identifying barriers to participation and commitments needed to make the focal market viable and inform its market design.

- **Nov 2026–Jan 2027: Set the framework.** If there is sufficient political and market support, the Commission and participating Member States should agree on the designation process and how they will promote the selected market to issuers and investors.

- **Jan–Mar 2027: Invite and assess proposals.** The Commission should publish the selection criteria and invite applications from European exchanges. Applicants should show that they already have a strong position in European technology listings and explain how the designation would help them reach a globally competitive scale. An independent expert panel should assess the applications. To assess the credibility of each application, the panel should seek confidential structured feedback from market participants. Panel members and respondents should disclose relevant commercial interests, and the Commission should publish an anonymised summary of the findings.

- **Mar–Jun 2027: Select and launch the focal ecosystem.** The Commission should complete the assessment and designate a focal market if a candidate meets the required standard. No ecosystem should be designated as a focal market unless it has a credible route to critical mass. Participating Member States should connect relevant IPO-readiness and investor-development initiatives to the focal ecosystem.

MEASURES OF PROGRESS

- **Implementation indicators:** Share of national IPO-readiness programmes aligned with the focal market.

- **Performance indicators:** Share of eligible European growth-company listings choosing the focal market, and volume raised; median IPO costs (% of proceeds); secondary-market turnover (% of market capitalisation); capital provided by international institutional investors (% of focal market IPO proceeds).

- **Longer-term outcomes for evaluation:** Share of European growth companies listing outside the EU; number of cross-border listings on the focal market and € raised; volume of follow-on equity raised by focal-market companies.

CONSIDERATIONS

- Select an ecosystem only if it has a credible path to critical mass. If no candidate meets the standard, the Commission should make no designation. Still, the designation should not restrict where companies may list or give the selected exchange exclusive legal rights.

- Member States should support the common process and avoid incentives or rules designed primarily to steer companies towards national exchanges.

- Selection must be based on objective performance criteria, and rule out geographic rotation or political allocation. The ultimate goal is agglomeration.

Recommendations at the national level

01	02	03
Mobilise national capital for European growth	Boost national research and innovation productivity	Make Europe the best place to attract and reward top talent
very high priority	very high priority	very high priority

04	05	06
Make it easier for innovative companies to take risks and recover from failure	Enable companies to scale across the Single Market	Use procurement to build European lead markets for defence
high priority	high priority	high priority

01

Mobilise national capital for European growth

ACTION

In order to facilitate European growth firms' access to capital and lower funding costs, Member States must build a bigger pool of domestic capital. Member States should implement Tibi/WIN-style schemes to mobilise additional institutional investment and help growth funds reach efficient scale. National regulators should remove unjustified barriers that prevent pension funds from investing in well-governed venture and growth funds. Where appropriate, Member States should offer targeted tax incentives that encourage additional angel investment in early-stage firms. Member States should support a focal European growth market that attracts specialist investors and builds liquidity over time, making larger European listings and exits more viable. Finally, in line with recommendations at the Union level, Member States should swiftly agree on the Market Integration and Supervision Package while preserving the Commission proposal's core market-integration measures.

MEMBER STATE MODE

Collective with kind-specific and coalition workstreams. All Member States should help deepen European capital markets and avoid new national measures that fragment them. Tibi/WIN-style schemes are largely useful in countries with large domestic investors and established growth funds. Angel investment incentives are most relevant to Member States with a strong pipeline of potential angel investors.

IMPLEMENTATION WORKSTREAMS

- **Tibi-style investor mobilisation – Kind-specific.** Member States with large institutional investor sectors should establish Tibi/WIN-style schemes that mobilise additional investment into accredited venture and growth funds. Participating investors should maintain autonomy over their investment decisions. Member States should coordinate reporting with the European Institutional Investors Pact for Innovation (see Union-level recommendations) and equivalent national schemes so that each commitment is counted only once.
  - **First steps in the next 12 months:** The government should appoint a lead ministry and an implementing coordinator, such as the treasury or national promotional bank. The lead ministry should convene investors to define the scheme's scope and develop transparent criteria for the eligibility of venture and growth funds. The implementing body should then publish these criteria and invite venture and growth funds to apply for accreditation. Accreditation should only confirm that a fund is eligible for investment counted under the scheme. Participating investors should then commit investment volumes into this portfolio of accredited venture/growth funds, while the specific allocation should remain their responsibility. Reporting on the progress of such a scheme should differentiate between capital committed and deployed. The convenor is responsible for coordinating the reporting with the European Pact and equivalent national schemes to prevent double-counting.
  - **Success indicators:** Capital committed, and actually deployed; independently estimated additional investment; cross-border share of investment.
- **Unlock institutional capital – Collective.** Insurers should be able to invest more easily in long-term equity. Member States should therefore swiftly transpose the Solvency II reforms. As part of the implementation of these reforms, national supervisory bodies should coordinate with EIOPA to make sure the revised long-term equity rules are interpreted and applied consistently across the Union. To allow pension funds to invest more easily in long-term equity, Member States should transpose the IORP II reforms without adding unjustified national restrictions, once they are adopted at the Union level. Member States should also enable smaller pension funds to invest in long-term assets through diversified pooled funds.
  - **First steps in the next 12 months:** National ministries should swiftly translate the Solvency II package into national law. Following these reforms, EIOPA and national supervisors agree on common guidance on the eligibility tests for long-term equity. The national supervisory decisions in the years after should be regularly reviewed. Once IORP II is adopted, the responsible ministry drafts the national implementing law without adding new investment restrictions. The responsible national ministries should consult smaller pension funds to identify legal barriers to pooled investment and swiftly remove them.
  - **Success indicators:** Insurers' long-term equity holdings (€ and % of assets); volume and share of assets of pension funds invested in European long-term equity; volume and share of assets of smaller pension funds invested through diversified pooled funds; number of diverging applications of eligibility for long-term equity (under Solvency II) across the Union.
- **Targeted tax incentives for angel investment – Kind-specific.** To mobilise additional angel investment, Member States with weak early-stage equity markets should introduce EIS/SEIS-style tax relief for individuals who invest in young firms raising new equity. The relief should address a documented funding gap and comply with EU State-aid rules. It should be limited to investors who are not connected to the company. Investors should continue to bear genuine commercial risk. To further strengthen the Single Market, investments in eligible EU companies should not be excluded from relief unless there is a good objective justification for it. The tax scheme should be reviewed after a fixed period and only continued if it generates additional investment at a reasonable fiscal cost.

- **First steps in the next 12 months:** The ministry of finance should estimate how large the early-stage financing gap is and how much additional investment the tax incentive would generate. It should only proceed if a cost-benefit analysis suggests the tax scheme would be the best option. In this case, it should set clear eligibility and anti-avoidance rules and establish the scheme in law.
- **Success indicators:** Additional private investment mobilised by the tax scheme; fiscal cost per additional euro invested; investment raised after firms leave the scheme; growth of supported firms relative to comparable firms (%).

- **Alignment with a European focal growth market – Coalition.** Member States should support the European focal market for growth company listings proposed in Union-level recommendation 5. To speed up the growth of such a listing ecosystem, Member States should remove unnecessary barriers to cross-border listings and refrain from adding policies that steer companies towards separate national markets. If Union-wide agreement cannot be reached quickly, a coalition of willing Member States should establish a common framework and keep it open to others.
  - **First steps in the next 12 months:** Participating governments should commit politically to the common focal market. The ministry of finance should review national policies and subsidies that could stand in the way of companies that decide to list in the common focal market. National promotional banks and enterprise agencies should help growth companies prepare to list on the common market.
  - **Success indicators:** Number of participating markets and their share of EU market capitalisation; number of companies choosing to do their primary listing in the focal market, and their market capitalisation.

CONSIDERATIONS

- Mobilisation schemes should not limit investor freedom by prescribing geographic quotas or weakening fiduciary and prudential standards.
- Common EU recognition of national mobilisation schemes should avoid duplicate certification and double-counting, without requiring identical national investment policies.
- Pension reform should protect retirement incomes, not force investment into inefficient assets.
- Angel investment incentives should be targeted, capped, and temporary, and withdrawn if they mainly subsidise investments that would have happened anyway.
- Member States should not use public financial support to divert liquidity towards national exchanges.

02 Boost national research and innovation productivity

ACTION

To foster economic growth driven by scientific discovery, Member States need to strengthen the institutions that support frontier research and its commercialisation. Member States should strengthen their capabilities to do high-risk/high-reward research. For this, they should establish or expand ARPA-style organisations, and national research funders should experiment with new funding models. Research institutions should seize the opportunities opened by AI-driven science. Member States should build or get access to metascience capabilities to regularly evaluate the productivity of their research and innovation system. To translate these discoveries into successful companies, public research institutions should adopt transparent and founder-friendly spin-out terms.

MEMBER STATE MODE

Collective with kind-specific delivery. All Member States should build metascience capacity and improve their spin-out terms. Member States with research institution systems large enough to sustain autonomous portfolio-based agencies should establish or expand ARPA-style programmes. Smaller Member States should form joint or cross-border arrangements.

IMPLEMENTATION WORKSTREAMS

- **ARPA-style funding structures – Kind-specific.** Member States should establish or expand ARPA-style funding structures, drawing on the SPRIND model. These should have legally guaranteed operational independence, expert control over funding decisions, and explicit tolerance of project failure. Member States should remove rules that impede rapid funding, active portfolio management, and new research organisations.
  - **First steps in the next 12 months:** The responsible ministry should prepare the legal and governance framework for the ARPA-style funding body. It is crucial that this agency has operational independence (especially for programme managers) from political interest. Where needed, parliaments should support any legislation needed to protect this operational independence. The government should then recruit a leadership team with exceptional scientific, technological, and entrepreneurial judgement. This team should start operations by selecting priority areas. Qualified programme managers should be hired for each domain.
  - **Success indicators:** Time from application to funding decision; share of projects redirected or stopped following technical reviews; share of projects meeting independently verified technical milestones; follow-up funding secured by completed projects within three years (amount and share of completed projects receiving any follow-up funding).
- **Experimenting with funding mechanisms – Universal for autonomous public research funders.** Conventional grant programmes are not necessarily the best fit for the future of AI-integrated research. Autonomous public research funders should therefore devote a defined share of their budgets to test and evaluate alternative funding mechanisms in order to identify which funding approaches deliver better research outcomes. Experiments could include rapid small grants awarded through short applications and accelerated review, milestone-based funding, and lotteries among proposals that meet a certain quality threshold. Funding bodies should make use of a limited number of 'golden tickets' allowing expert reviewers to support promising high-risk proposals that would otherwise not be funded.
  - **First steps in the next 12 months:** The research funder should allocate an experimentation budget, identify appropriate comparison groups, and pre-register evaluation plans. It should then launch controlled trials of alternative funding mechanisms. Experimental results should be published where feasible.
  - **Success indicators:** Share of the funder's budget allocated to funding experiments (%); application-to-award time and administrative cost per award (days and €); quality of research outcomes relative to comparable conventional grants, based on predefined and potentially field-specific measures.
- **Independent metascience capabilities – Collective.** To be able to monitor the productivity of its science system, each Member State should ensure access to independent metascience capability, nationally or through formal cross-border arrangements. This metascience organisation should systematically monitor and evaluate research-system performance. It should be able to suggest experiments for funding methods and recommend improvements to how research funding is organised and allocated. National bodies should share methods and evidence through the European Research Area and the proposed European metascience function (Union-level recommendation 2).

- **First steps in the next 12 months:** In countries with a significant research system, the responsible ministry should designate an existing body or establish a metascience unit with statutory independence and secured multiannual funding. The unit should appoint a scientific leadership team, with expertise in policy evaluation and the evaluation of incentive schemes. This team should set an initial research agenda to evaluate the productivity of the existing research system. If administrative research data exists, the unit should get access to it. Otherwise, it should make sure such a dataset (including funding decisions and funding programmes by actors in the science system) is built. Public research funders should provide the necessary data and cooperate in evaluations and controlled funding experiments. National units should agree on common evidence and data standards through the European Research Area. Countries with smaller research systems should collaborate to build joint metascience units.

- **Success indicators:** Share of funding programmes changed following evaluation; amount of funding reallocated towards better-performing approaches; and uptake of the unit's recommendations by public research funders.

- **Prepare research systems for AI-enabled science – Collective.** Member States should prepare their research systems to use AI, automation, and robotics to accelerate scientific discovery. They should financially support new research tools and shared infrastructure that integrate AI with experimentation, including automated biofoundries and self-driving laboratories. They should facilitate the sharing of practical lessons on how to organise research around these tools between research institutions and funders. Where the needs for expertise or investment exceed national capacity, Member States should build joint facilities. The metascience body should assess how AI-enabled research affects scientific productivity, reproducibility, and integrity, and identify which organisational and funding models work best.

  - **First steps in the next 12 months:** Responsible ministries and research funders should identify existing AI research tools and facilities, determine what is missing, and fund the first shared tools and pilot facilities. They should determine which facilities require national or cross-border scale and designate institutions to lead their development. Research funders should create a practitioner network to exchange practical experience on how to re-organise science around these AI tools to reap the biggest productivity benefits. The metascience body should prepare baseline measures and evaluation plans.

  - **Success indicators:** Number of AI-enabled tools and facilities launched; facility utilisation rate; reduction in experiment time and cost; share of fully reproducible AI-enabled research papers, per field; share of AI-enabled research outputs independently reused or validated.

- **Founder-compatible spin-out terms – Universal.** Publicly funded universities and research organisations should make it easier for researchers to turn into founders by adopting attractive spin-out terms. This includes speedy decision timelines, so researchers do not waste years waiting for an intellectual property (IP) agreement. Terms should reflect the institution's contribution, preserve founders' incentives, and allow for follow-on investment, while complying with State-aid rules. Institutional policies should follow the EU Code of Practice for knowledge valorisation. Once the EU Blueprint for IP licensing and spin-off creation is finalised, universities should follow its guidance to develop standard terms.

  - **First steps in the next 12 months:** Research institutions should review recent spin-out agreements and consult founders and investors on what would be needed to make agreements attractive. They should then publish standard term sheets covering institutional equity, royalties, IP access, and founder obligations. These term sheets should set decision deadlines and approval responsibilities for standard cases. Any departure from the published terms should require written justification.

- **Success indicators:** Median time to agree on spin-out terms; share of use of standard (rather than bespoke) terms; external equity raised by spin-outs within 12 months (volume, and share of spin-outs raising any external equity).

CONSIDERATIONS

- Operational independence of ARPA-style organisations should be protected in law and practice, within clear missions, budgets, and conflict-of-interest rules.
- Project failure should be tolerated across the portfolio of ARPA-style organisations, but not where it reflects weak governance, research misconduct, or undisclosed conflicts. The suggested performance indicators (e.g. share of projects reaching milestones) have no universal optimal levels and should therefore not be treated as targets. Rather, they should monitor whether the organisation is making use of its independence to pursue high-risk, high-reward projects. In this case, one would expect, for example, a positive but small share of projects reaching milestones.
- Member States should test different funding approaches, not try to converge on a single model.
- AI-enabled research should be subject to safeguards. Areas for safeguards include reproducibility, human accountability, research integrity, including disclosure of AI use and prevention of fabricated results, cybersecurity, and dual-use risks.
- Institutional equity shares in spin-outs should normally be low, but may be higher where the institution has provided substantial IP, infrastructure, or development funding.

03 Make Europe the best place to attract and reward top talent

ACTION

Member States should make it easier for innovative firms to attract and retain talent. To help younger firms compete, even when cash is scarce, they should allow employee-equity schemes to be taxed competitively. Member States should also provide fast and predictable migration admission routes for founders, researchers, and highly skilled workers, making it easier for innovative firms to recruit talent from across the world. More broadly, national migration rules and administrative procedures should minimise unnecessary burdens on growing firms. Public services should be founder-friendly, make greater use of digital and once-only processes (including through the Founder Gateway), and support mobility across the Union.

MEMBER STATE MODE

All Member States. All Member States should make employee-equity taxation and talent admission procedures internationally competitive (though national tax rates, migration volumes, and institutional arrangements may differ).

IMPLEMENTATION WORKSTREAMS

- **Competitive employee-equity taxation – Universal.** Member States should allow tax on EU Inc. employee stock options to be deferred until employees sell their shares or otherwise obtain liquidity. National rules should make this tax deferral predictable, and ensure competitive tax rates. Most importantly, the new tax scheme needs to be workable in practice. Member States should apply similarly attractive rules to employee stock options issued under national law. Income tax should be deferred until the shares are sold, while any subsequent increase in value should be taxed as a capital gain. When employees relocate within the EU, relocation itself should not trigger taxation, and Member States should avoid the double taxation of employee equity. At the same time, each country should keep the right to tax the share of equity gains earned within this country. Cross-border and anti-avoidance rules should prevent employment income from being shifted to lower-tax jurisdictions or reclassified as capital gains.

- **First steps in the next 12 months:** The finance ministry and tax authority should review employee-equity taxation and identify the changes needed to defer tax until employees sell their shares or otherwise obtain liquidity. They should then introduce the necessary legislation and publish administrative guidance. The finance ministry and tax authority should work with the Commission and other Member States to find arrangements that preserve the deferral of taxation and prevent double taxation when an employee moves across borders within the Union. To maintain fairness and prevent employees from avoiding tax by relocating, they should agree on how to divide the taxable value from equity options between the countries where the employee worked.

- **Success indicators:** Employee equity taxed before employees receive liquidity (% of employee-equity cases); share of firms offering employee equity; number of cases of early or double taxation after relocation within the EU.

- **Fast and specialised talent admission – Universal.** Europe needs an easier migration route for top (AI) talent (researchers, founders, and highly skilled workers). Member States should therefore create dedicated migration units for innovation talent that can assess qualifications competently and decide quickly, ideally within 30 days. To enable fast case-processing and make applying a smooth experience for applicants, the entire process (including the verification of qualifications) should be fully digital, following the once-only principle, and applicants should be able to track their case online. This digital workstream must be compatible with the Founder Gateway (Union-level recommendation 3) and should be available in English. Talent seeking to migrate with their family should be able to submit one application for the whole family. Non-EU nationals with a credible plan to build a business in Europe should have access to a dedicated founder visa. To make sure that these visas do not create loopholes in the migration system, they should be valid for a limited time only (e.g. six months), and then renewed upon evidence of actual ongoing company formation. Changing career plans should not entail substantial additional paperwork. Instead, a simple change-of-status procedure should allow these permit holders to move between research, employment, and entrepreneurship without applying for residence again.

  - **First steps in the next 12 months:** The ministry of the interior, working with ministries of economy and research where relevant, should designate the innovation talent unit and review existing migration routes for founders, researchers, and highly skilled workers. The unit should be sufficiently staffed to be able to process applications within 30 days, with a target processing time of 14 days. In case no fully digital application system exists, the innovation talent unit should start to pursue the product development of such a platform and make sure that it will be compatible with the Founder Gateway at the EU level (see also next recommendation). Where no suitable route for founders exists, the interior ministry should propose a new legal pathway.

  - **Success indicators:** Share of complete applications decided within 14 and 30 days; median processing time for family applications; number of founder visas issued; share of founder visas leading to company formation within 12 months.

- **National service integration for the Founder Gateway – Universal.** Member States should make it possible for founders to complete key business procedures through the Founder Gateway (Union level recommendation 3). To this end, they should simplify and digitalise the national procedures used by founders. Progressively, these national procedures should be connected to the Gateway through common EU standards and interfaces. So founders do not have to enter the same information multiple times, authorities should reuse verified information already held by another public authority. They should accept forms and supporting documents in English, and be able to provide case support in English as well. Where necessary, national rules should be changed to enable this. Member States should start this process with the founder journeys that can be simplified and connected most quickly to the Gateway. Over time, all the main bureaucratic procedures for founders should be covered through the Gateway.

- **First steps in the next 12 months:** Each government should appoint its digital agency (or an equivalent technical delivery team) to lead national integration with the Founder Gateway. This team should map the main founder journeys and select at least one that can be simplified and connected quickly. To build new digital processes (or at least the integration into the Gateway), teams of designers, software engineers, and service operators should collaborate. They should be mindful of the once-only principle and do systematic user-testing with founders before connecting it to the Gateway. In parallel, legal teams together with responsible ministries should resolve regulatory barriers. By month nine, early-adopter Member States should have connected at least one tested process. With the help of the Commission's coordination, the connected early journeys should collectively cover company formation, talent admission, and secure business-data sharing, so that the broader Gateway integration on the European level can also be tested. Other Member States should simplify at least one founder procedure and adopt a roadmap for connecting their founder services to the Gateway.
- **Success indicators:** End-to-end completion rate and median completion time for journeys connected to the Gateway; number of repeated data requests; user satisfaction scores.

CONSIDERATIONS

- Fast admission routes should cover the entire process, including qualification recognition, security checks, and applications for accompanying family members.
- Founder visas should follow a staged pathway, since 'plans to fund' upon entry are difficult to verify. An initial permit should be granted on the basis of a concise business proposal and potentially an interview assessing the applicant's experience and commitment. Renewal should require evidence of real business formation activity (e.g. product development, customer interviews, meetings with potential investors, spending tied to building the business). As the process of company formation differs a lot across sectors and business models, authorities should assess the evidence as a whole against published criteria. Commercial success should not be required at the first renewal.

04 Make it easier for innovative companies to take risks and recover from failure

ACTION

As part of broader flexicurity reforms (as recommended in Objective 3.2), Member States should reduce the cost of business failure and restructuring while preserving worker security. They should implement Union insolvency reforms effectively and ensure that viable firms can restructure quickly, while allowing the capital and workers of unviable firms to move to more productive activities. They should also support the simplified liquidation procedures proposed for EU Inc.

MEMBER STATE MODE

All Member States, with nationally differentiated delivery. All Member States should reduce excessive and unpredictable restructuring costs, while adapting reforms to their national labour market institutions, collective bargaining systems, and social protection arrangements.

The specific implementation details for this recommendation can be found in Objective 3.2.

05 Enable companies to scale across the Single Market

ACTION

In order to help companies scale across Europe, Member States need to lower unnecessary fixed costs of entering new markets. This would simultaneously make it easier for firms, especially smaller ones, to expand across the EU and strengthen competition. Firms should not have to repeat the same procedures across different jurisdictions. Instead, authorities should reuse evidence and recognise decisions from other Member States wherever possible. Member States must also apply Single Market rules across the Union without unjustified

national additions and provide a rapid route for challenging cross-border market-access decisions.

MEMBER STATE MODE

All Member States. Even though the barriers that need to be removed will differ across Member States, all Member States should be disciplined in applying these changes. The goal should be a single Union market.

IMPLEMENTATION WORKSTREAMS

- **Prevent gold-plating and inconsistent implementation – Universal.** Member States should avoid adding requirements beyond Union law unless they can show that these additional barriers to the creation of a Single Market are justified by clear public benefits. The burden of proof should rest with the authority proposing the additional requirement. Member States should also implement Single Market legislation fully and on time. National provisions that exceed Union requirements should be periodically reassessed and revised where they are no longer justified.
  - **First steps in the next 12 months:** Each ministry and regulator should review how EU law in its area of responsibility is translated into national practice. They should use evidence from the Commission's Terrible Ten work, Single Market Enforcement Taskforce (SMET), and SOLVIT to identify national laws and enforcement practices that create unnecessary barriers to companies from other Member States who want to enter local markets. Ministries and regulators should revise those standards going beyond EU law without good justification. They should clearly communicate the changes to everyone who applies these new rules. To prevent unjustified new barriers, ministries and regulators should justify that wider EU benefits outweigh the additional entry costs when introducing additional rules beyond EU law. They should train responsible officials to apply EU Single Market rules and proportionality tests consistently.
  - **Success indicators:** Compliance costs caused by national requirements that go beyond EU law (split into differences with and without justification); annual savings from removing unjustified requirements; share of new national laws that go beyond EU law with a published cost-benefit assessment.
- **Expand mutual recognition and regulatory reliance – Universal.** Companies entering another EU market should not have to resubmit documents already accepted by a public authority. National authorities should recognise evidence and decisions from other Member States unless EU law or a legitimate public-interest concern requires a separate assessment.
  - **First steps in the next 12 months:** Sector ministries should ask domestic companies where they faced duplicate requirements when trying to enter other Member State markets. National Single Market coordinators should use these cases to identify recurring barriers and work through SMET and the Commission to remove them.
  - **Success indicators:** Share of cross-border applications completed without resubmitting documents; costs of duplicate evidence requirements; median time to resolve market-access cases.
- **Provide rapid review for cross-border market-access barriers – Universal.** Member States should require authorities to give written reasons for restricting market access and set deadlines for reviewing disputed decisions. Cases that remain unresolved should be referred to SOLVIT. The Commission and national authorities should track recurring cross-border market-access barriers and work through SMET to remove them.

- **First steps in the next 12 months:** Member States should introduce deadlines for the review of dispute cases and ensure that national SOLVIT centres have the staff and authority needed to handle urgent cases.
- **Success indicators:** Median time to resolve cross-border market-access disputes; number of dispute cases resolved without litigation.

CONSIDERATIONS

- Make sure that removed requirements are not reintroduced or maintained in other forms, e.g. in administrative practice. Member States must therefore clearly communicate simplified rules to everyone affected.
- Assess firms' compliance costs in practice, not just the number of requirements set out in law.

06 Use procurement to build European lead markets for defence

ACTION

To create a larger demand side, Member States should coordinate to create a single European market for defence capabilities. Member States should use defence budgets to create a credible route from technical validation to procurement at scale for innovative firms. Defence procurement processes should define the capabilities needed and award contracts based on performance in paid competitive trials. Procurement authorities should provide a clear, budget-backed route for successful prototypes to production at scale. Production contracts should be awarded based on operational performance of the technology alongside its cost and interoperability with existing systems.

MEMBER STATE MODE

Coalition with universal national reforms. All Member States with significant defence procurement should improve access for innovative firms. Those with shared capability needs should agree on common requirements and pool their procurement budgets.

IMPLEMENTATION WORKSTREAMS

- **Multiannual demand and open innovation competitions – Universal for procuring Member States.** Procurement authorities should publish multiannual capability needs and run open innovation contests to acquire these capabilities. These contests should focus on technical performance and include paid trials of several promising solutions (if multiple exist), to preserve competition during development. Successful providers should be admitted to multi-supplier framework contracts and compete for follow-on orders. For small and fast-changing technologies, Member States should also allow authorised military units to directly place limited operational orders. Eligibility rules must be proportionate and not unnecessarily rule out startups. Where relevant, contests should link to AGILE and EUDIS activities. The application process should make sure that companies can submit a first-stage proposal in under 2 hours (5 pages maximum), supported by demonstrations of existing capabilities (e.g. video evidence). Applicants should be notified within 6–8 weeks whether they have been selected for the next stage. Small and innovative providers should be able to participate in procurement as prime contractors, consortium partners, or suppliers. Contracts should protect contractors' existing IP while allowing governments to integrate and maintain the technology, and to procure future upgrades competitively.
  - **First steps in the next 12 months:** Responsible ministries in close cooperation with the military should identify priority capability needs and allocate funding for innovation competitions and paid trials. Procuring authorities should then launch the first competitions. The ministries responsible for procurement law should identify and remove any legal or regulatory barriers to this approach.
  - **Success indicators:** Share and value of suitable procurement run via innovation contests; share of applicants that are startups or first-time defence suppliers; median time from application to selection.

- **Prototype-to-production contracts – Universal for procuring Member States.** Procurement contracts for innovative technology should include funded options to move successful prototypes into production. Where relevant this could be linked to industrial scale-up support under EDIP. The decision to move successful pilots to production should rest on the following criteria: performance in operational tests, lifetime costs, interoperability and compatibility with existing systems, and the supplier's ability to scale production.
    - **First steps in the next 12 months:** Ministries of defence should earmark a sufficient budget for successful prototypes to enter production. Procuring authorities should define performance and affordability thresholds before trials and include production options in initial contracts.
    - **Success indicators:** Share of prototypes entering production; median time from successful trial to production order.
- **Establish a European Defence Mechanism (EDM) – Coalition.** European countries should establish a permanent mechanism for joint defence procurement, to make cooperation credible and create the scale for competitive European defence markets. This Mechanism should remain open to European countries outside the Union and operate on behalf of its members. Members must commit capital subscriptions to the institution and jointly define capability needs and common functional requirements. Selected capabilities should then be procured jointly through the Mechanism. Before the Mechanism launches a tender, participating members should agree on functional requirements, expected order volumes, delivery schedules, as well as national budget commitments.
    - **First steps in the next 12 months:** At the Foreign Affairs Council meeting on 28 September 2026, interested defence ministers should test support for the Mechanism and identify an initial coalition. The participating governments should then appoint a joint task force of defence and finance ministry officials to prepare its mandate, governance, capital structure, and initial procurement areas. Within 12 months, they should establish a founding framework and identify the first capabilities to be procured jointly. The Mechanism should remain open to countries wishing to join later.
    - **Success indicators:** Value and share of binding joint orders; number of common configurations replacing national variants; unit costs and delivery time compared to similar, unilateral procurement procedures.
- **Common standards and shared testing facilities – Coalition.** Member States should use the EDA's Defence Test and Evaluation Base to involve military users, share testing facilities, and avoid duplicating equivalent tests across countries. They should make use of EDSTAR to align technical standards.
    - **First steps in the next 12 months:** Interested defence ministries together with the EDA should designate facilities for joint testing. National security authorities should identify which certifications and clearances can be mutually recognised.
    - **Success indicators:** Number of avoided duplicate tests; qualification time and cost for defence suppliers.

CONSIDERATIONS

- Requirements must reflect real capability needs and not pre-determine specific technical solutions.
- Joint purchasing should increase Member States' buying power while increasing supplier competition by creating a better integrated European defence market.
- IP terms should preserve innovation incentives and, at the same time, protect governments from supplier lock-in and supply disruption.

O2.2 Develop indispensable assets across the AI value chain

WHY IT MATTERS

Countries that own no part of the AI value chain risk being economically sidelined under transformative AI. Europe currently holds strong positions in only a few layers, and their importance could wane over time. Europe should hence spread its bets widely and combine areas of existing strength (e.g. industrial AI) with contestable niches (e.g. inference chips) and selected alternative-paradigm moonshots.

Recommendations at the Union level

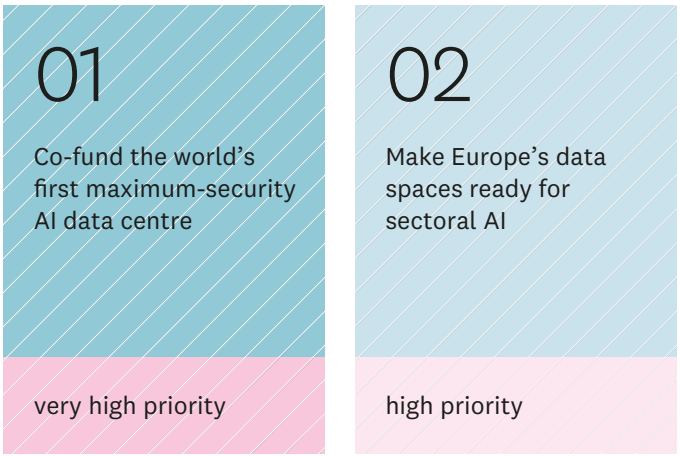

01 Co-fund the world's first maximum-security AI data centre

ACTION

Establish a dedicated Union funding line that co-finances, together with one or a small group of host Member States, the world's first AI data centre designed to meet RAND Security Level 5 or equivalently the SL5 Standard for AI Security by 2028.

IMPLEMENTATION

- **Potential instruments:** A dedicated maximum-security tier or follow-up call under the EuroHPC Joint Undertaking's AI Gigafactories initiative (Council Regulation (EU) 2026/150), drawing on the InvestAI initiative and matched Member State funding.
- **Potential first steps (next 12 months):**
  - **October 2026**: Commission and Member State security authorities, drawing on external expertise, begin drafting the AI facility specification.
  - **November 2026**: A Member State informally offers to finance and host the facility, potentially together with other States.

- **January–March 2027**: Commission selects a host Member State (or coalition), which selects a site and hires a world-class team of data centre security specialists.
- **April 2027**: Union funding decision taken; facility specification published; Member State team has an operational core and begins its work.

- **Success indicators:** A signed EU–host Member State co-financing agreement, in line with applicable State-aid rules; independently assessed maximum-security data centre operational by 2028.

CONSIDERATIONS

- A maximum-security data centre would be robust against top-priority attacks from the most capable, state-level actors, who could try to steal or sabotage AI models running on insufficiently secure infrastructure. Today, no AI data centre in the world meets such a security standard.
- European maximum-security data centres would have three main use cases: (i) storing and serving frontier model weights, for example to enable European access to security-sensitive model weights under compute-for-access arrangements, (ii) highly sensitive public-sector inference workloads, for example in defence, and (iii) evaluations and safety research on potentially dangerous models.
- Researchers estimate that a small, proof-of-concept, RAND SL5 facility for AI inference could be built within 14 months if it were a national priority, relying entirely on known technologies, and would cost $37 million to $50 million ($277 million to $345 million for an enterprise-scale version).
- The security specifications for the relevant facility should be kept technically separate from sovereignty specifications, which address a different problem.

02 Make Europe's data spaces ready for sectoral AI

ACTION

Clarify legal status and provide technical infrastructure so that companies can easily share their data for the purposes of developing domain-specific AI models and applications in sectors such as manufacturing, healthcare, or automotive. This should build on the Common European Data Spaces and Data Labs and include specific measures: (i) ensure that the Commission guidance on data pooling committed under the Data Union Strategy addresses multi-party data pooling between competitors for AI development, explaining how it can be compatible with Article 101 TFEU, building on Chapter 6 of the 2023 Horizontal Guidelines, which addresses data pooling in general terms but not the specific case of pooling AI training data between competitors, (ii) publish model contractual terms for multi-party data sharing for AI development, building on the Data Act model contractual terms published by the Commission in draft form in November 2025, (iii) direct at least €50 million of EU research funding to privacy-enhancing technologies (PETs) to be deployed in the Common European Data Spaces and the AI Factories' Data Labs.

IMPLEMENTATION

- **Potential instruments:** Commission soft law for the competition guidance mentioned in measure (i) above; a Commission Recommendation extending the Data Act Article 41 model contractual terms; existing PETs topics in the Horizon Europe Work Programme 2026–27 and, for additional funding, the first work programme of its successor.
- **Potential first steps (next 12 months):**
  - **January 2027:** Commission publishes model contracts and the guidance on data pooling, which includes AI-specific examples.

- **June 2027:** Commission earmarks at least €50 million for PETs across current and upcoming research work programmes.
- **September 2027:** At least three Common European Data Spaces have operational data pools for multi-party AI training operating under the model terms.

- **Success indicators:** Number of multi-party data pools for AI training operating under the published model terms; amount of EU funding committed to PETs; number and diversity of companies contributing proprietary data; number of AI models or applications developed using the data spaces.

CONSIDERATIONS

- Starting in 2020, the EU has built Common European Data Spaces across 14 sectors, which aim to facilitate data sharing through governance frameworks and technical infrastructure. However, as the 2025 Data Union Strategy recognises, the challenge is to move beyond pilot projects and make data spaces useful for AI training at scale. The primary bottleneck is companies being reluctant to contribute commercially valuable data to multi-party data spaces: They need legal certainty that pooling data with competitors is compatible with competition law, model contracts governing access rights and how the value created from pooled data is distributed, and mature technical tools that let models learn from sensitive data without exposing it to untrusted parties. The measures above target these demand-side barriers by lowering entry costs for data holders and using sectoral model training as a concrete use case around which to organise.

- The Commission has already committed in its Data Union Strategy to issue guidance on data pooling. The guidance should use concrete examples to address, among other things, (i) which data competitors can pool for AI training and when PETs are needed to shield commercially sensitive data, (ii) how contribution requirements and access fees can prevent free-riding, and (iii) how companies can use the Informal Guidance Notice to obtain timely comfort for cross-border AI data pooling initiatives.

- Data pools should be structured so that most of the value stays with European companies, while still complying with Article 101 TFEU. Under the 2023 Horizontal Guidelines, access rules need to be fair, objective, transparent, and non-discriminatory. Therefore, simply excluding non-EU developers raises competition concerns. Instead, access could be tied to reciprocal contributions to the pool or subject to fair access fees.

Recommendations at the national level

01	02	03
Direct ARPA-style vehicles towards strategic AI bets	Host and co-finance the first maximum-security AI data centre	Anchor demand for European inference chips
very high priority	very high priority	high priority

01 Direct ARPA-style vehicles towards strategic AI bets

ACTION

Resource national ARPA-style agencies to fund a strategic AI portfolio, focusing on some or all of the following categories: (1) sectoral AI champions in areas of existing strength, especially where valuable proprietary data can be leveraged (e.g. manufacturing, healthcare, automotive), (2) specialised AI hardware, such as AI inference chips, and the corresponding orchestration layer, (3) security-relevant hardware infrastructure, such as maximum-security data centres or hardware-enabled mechanisms for chip verification, (4) selected high-risk/high-return approaches in AI research (e.g. world models, safe-by-design AI, worker-friendly AI).

MEMBER STATE MODE

All Member States. Every Member State should do this, using an existing ARPA-style vehicle (e.g. Germany's SPRIND) or a newly created one (as recommended in Objective 2.1). Some Member States, especially smaller ones, may want to orient their portfolio around strategic bets that are most promising in light of their economic position, resources, and areas of existing strength, e.g. semiconductor R&D in Belgium, AI-relevant connectivity in Finland, AI applications in healthcare in Denmark, public sector AI applications in Estonia, or photonics in Lithuania.

IMPLEMENTATION

- **Potential instruments:** Administrative direction plus a dedicated budget line, in accordance with State-aid rules.
- **Potential first steps (next 12 months):**
  - **October 2026:** Member States direct their ARPA-style vehicles to define a strategic AI portfolio.
  - **December 2026:** Categories have been defined; Member States commit the budget through reallocation or a new funding line.
  - **January 2027:** First calls published in one or two priority categories per country.
  - **April 2027:** First funding contracts signed.
- **Success indicators:** € committed to the strategic AI portfolio, total value of funding contracts signed, follow-on private capital raised by funded companies within 24 months of the ARPA award.

CONSIDERATIONS

- High-risk/high-return bets in AI research (the fourth category above) should be seen as complements rather than substitutes for securing access to AI models at the current frontier. These models have been the main driver of AI progress for years, with no indication of a slowdown in capability growth, and access to them will be crucial for protecting against imminent security risks (e.g. cyberattacks).
- Member States, especially smaller ones, may want to coordinate informally to avoid duplicating efforts.
- For specialised AI hardware, national funding should complement (and where appropriate, co-finance) the Union-level instruments under Pillar I of the EU's Chips Act, which funds research infrastructure and early-stage support for chip startups and scale-ups. However, the Commission's own evaluation found that its startup fund was exhausted within two years and that Europe lacks the late-stage capital to scale these companies. National ARPA-style funding can help to fill this gap alongside private capital (in line with the recommendations for better private capital mobilisation in Objective 2.1).

For these bets to be successful, the Union and Member States will need to implement the broader competitiveness measures outlined in Objective 2.1. The recommendations above will only bear fruit in an economy that makes it easy to found and scale high-growth companies, for example, by giving companies access to growth capital and enabling them to hire and retain the world's best talent. Creating these conditions is the single most powerful contribution that the Union and its Member States can make to Europe's position in the AI value chain.

02 Host and co-finance the first maximum-security AI data centre

ACTION

Build the world's first maximum-security AI data centre by 2028 (see Union-level recommendation 1), either alone or within a Member State coalition. This requires (i) a site with secured access to energy (e.g. 10+ MW with the possibility to scale) and exceptional protection against natural or adversarial threats, (ii) a national co-financing share for building the facility and attracting the world-class talent to make it highly secure, and (iii) cooperation with national intelligence and security agencies (e.g. personnel vetting, counter-intelligence, physical protection).

MEMBER STATE MODE

Coalition or kind-specific – either one Member State or a nimble coalition willing to pool funding, talent, suitable sites, and security/intelligence capability.

IMPLEMENTATION

- **Potential instruments:** Application to a to-be-established maximum-security tier of the AI Gigafactories initiative (see Union-level recommendation 1 above); national budget commitment; mandate to national intelligence and security agencies; procurement commitments for high-sensitivity workloads.
- **Potential first steps (next 12 months):**
  - **November 2026:** Member State government informally tells the Commission that it is willing to finance and host the facility, potentially together with coalition members.
  - **January 2027:** A host Member State is selected by the Commission; potentially together with coalition members, it commits a dedicated share of its national budget to the EU co-financing line and begins hiring a world-class team of data centre security specialists.
  - **April 2027:** Union funding decision taken; facility specification published; Member State team has an operational core and begins its work.
  - **July 2027:** Member States(s) select a suitable site that meets the relevant security requirements; intelligence and security agencies begin to prepare it.
- **Success indicators:** A signed EU–host Member State co-financing agreement, in line with applicable State-aid rules; independently assessed maximum-security data centre operational by 2028.

CONSIDERATIONS

- A maximum-security data centre would be robust against top-priority attacks from the most capable, state-level actors, who could try to steal or sabotage AI models running on insufficiently secure infrastructure. Today, no AI data centre meets such a security standard.
- European maximum-security data centres would have three main use cases: (i) storing and serving frontier model weights, for example, to enable European access to security-sensitive model weights under compute-for-access arrangements, (ii) highly sensitive public-sector

inference workloads, for example in defence, and (iii) evaluations and safety research on potentially dangerous models.

- Researchers estimate that a small, proof-of-concept, RAND SL5 facility for AI inference could be built within 14 months if it were a national priority, relying entirely on known technologies, and would cost $37 million to $50 million ($277 million to $345 million for an enterprise-scale version).
- The security specifications for the relevant facility should be kept technically separate from sovereignty specifications, which address a different problem.

03 Anchor demand for European inference chips

ACTION

Create conditional purchase commitments for inference chips designed in the EU from publicly supported AI data centres, such as the EU's AI Gigafactories, once those chips meet pre-specified criteria for use in AI data centres. Member States should target aggregate commitments of 10% of the public funding for the AI Gigafactories (i.e. €1 billion) by 2029.

MEMBER STATE MODE

Kind-specific – Member States planning to build significant amounts of public AI compute, such as the host countries of the AI Gigafactories.

IMPLEMENTATION

- **Potential instruments:** Use existing joint and innovation-procurement tools and, once adopted, the demand-side measures in Chips Act 2.0 to coordinate requirements and purchases across Member States.
- **Potential first steps (next 12 months):**
  - **December 2026:** Interested Member States form a purchasing coalition and agree an indicative collective budget for conditional purchase commitments.
  - **March 2027:** The coalition consults European chip companies and agrees common requirements, including performance, security, and energy efficiency.
  - **September 2027:** Participating purchasing bodies publish coordinated conditional purchase commitments for European inference chips.
- **Success indicators:** € value of published advance purchase commitments; share of public inference procurement accounted for by European suppliers of inference chips.

CONSIDERATIONS

- The proposed Chips Act 2.0 would provide useful tools for coordinating semiconductor procurement but contains no quantified commitment to purchase European AI inference chips. The €1 billion target would be an additional commitment by participating Member States.
- This would create an early market for European suppliers without slowing the buildout of AI compute. Until suitable European inference chips are available, buyers should procure on the open market. This is because the highest priority for Europe is to host AI compute in proportion to its global GDP share as fast as possible (see Immediate Objective 3). Member States should procure domestic inference chips only insofar as this is consistent with that goal.
- National funding and procurement commitments should be designed in accordance with applicable EU State-aid and public-procurement rules.

Pillar 3

Ensuring safety and security

O3.1 Ensure prioritised and targeted use of EU rules to address risks from highly capable AI

Obligations on providers of general-purpose AI models are already applicable and enforceable, and the GPAI Code of Practice is already being relied upon. For this reason, there are no detailed recommendations for this objective.

O3.2 Prepare for labour market impacts

WHY IT MATTERS

Transformative AI would likely displace human labour at unprecedented levels. Managing labour market transitions well will be necessary to sustain people's livelihoods, while unmitigated job losses threaten social stability.

Recommendations at the Union level

01

Build a monitoring stack for AI labour market impacts

very high priority

02

Fund and steer conditional flexicurity adjustment

very high priority

01 Build a monitoring stack for AI labour market impacts

ACTION

Prepare for AI labour market impacts in the EU by September 2027 with these measures:

- (i) introduce an AI-displacement module in the EU labour force survey (LFS), with the module in the field no later than the 2029 survey year.
- (ii) create a quarterly Cedefop bulletin tracking hiring patterns in vacancies in AI-exposed occupations. This should leverage Cedefop's Skills-OVATE (Skills Online Vacancy Analysis Tool for Europe) online job advertisement analytics platform.
- (iii) adjust JRC's occupational AI exposure index to the four-digit level of the International Standard Classification of Occupations (ISCO-08) in collaboration with the International Labour Organization (ILO). Ensure the mapping to the principal international exposure measures (Anthropic's Economic Index, the OpenAI/University of Pennsylvania exposure measure, and the ILO's global index of occupational exposure).
- (iv) negotiate with AI developers to access usage data.
- (v) improve occupational coding for quarterly occupation-level firm-side statistics on wages and employment.

IMPLEMENTATION

- **Potential instruments:**
  - LFS module: The Commission should implement delegated acts following Article 3(7) and under Regulation (EU) 2019/1700 (European social statistics framework), amending Commission Delegated Regulation (EU) 2025/1999 (multiannual rolling planning of survey topics), supported by an implementing act for technical specifications. Article 4 allows for a change in the schedule, which the Commission should prioritise. Alternatively, it could serve as the already scheduled 2030 topic.
  - LMB (labour market statistics from businesses) module: Article 8 of Regulation (EU) 2025/941 permits a feasibility study. This study should aim to make Member States' administrative payroll sources comparable, particularly in terms of coverage and occupational coding. Subject to coverage of at least 10 Member States with more than 50% coverage, the Commission should, by an Article 4(3) delegated act, receive high-frequency data on earnings, labour costs, vacancies, and, if feasible, employment. The UK demonstrated how to provide high-frequency data with their monthly income tax Real Time Information (RTI).
  - Usage data: The Commission should seek access to AI usage data through a Memorandum of Understanding (MoU) with leading AI companies (see considerations). In parallel, the Commission (Eurostat) should, in accordance with Article 17b of Regulation (EC) No 223/2009 (as amended by Regulation (EU) 2024/3018), request access to data from the appropriate private data holders. In cases where an agreement with a private data holder is not reached within the timelines set by the legislation, the Commission may use its powers to compel access to the relevant data in accordance with Article 17c of Regulation (EC) No 223/2009.
- **Potential first steps (next 12 months):**
  - **Q4 2026:** Eurostat develops an AI-displacement ad hoc subject, and the ESS Committee needs to be consulted regarding planned statistical measures (Article 7 of Regulation 223/2009). The Commission adopts a Delegated Act amending Regulation (EU) 2025/1999 to change the content of the 2029 survey.
  - **Q4 2026:** Eurostat also coordinates with Member State NSIs to immediately add the AI displacement subject at the national level.

- **Spring 2027:** JRC and Cedefop publish their annual exposure index based on the ISCO-08 four-digit level and quarterly exposure bulletin, respectively.

- **Success indicators**: The EU adopts the delegated acts and JRC and Cedefop update their statistics. The EU receives usage data for analysis by 2027.

CONSIDERATIONS

- The UK AI Economics Institute's joint statement with Anthropic, Google, OpenAI, and Microsoft (who later became members of a working group in June 2026) is an early precedent for formal lab–government cooperation. However, this institution is not yet fully operational, and its data access remains unproven.

- If EU data-protection rules constrain the collection of API usage data, as may have been the case with Google's ATLAS, regulators need to reduce legal uncertainty in analysing API data.

- While Eurostat already surveys enterprise AI adoption, it has not yet collected data on AI usage.

- The Commission should consider creating an EU AI Economics Institute to assess the economic impact of AI continuously. Either the JRC or the AI Office should host this research unit due to their research focus and network with AI companies, respectively.

02 Fund and steer conditional flexicurity adjustment

ACTION

The Commission should fund social programmes and incentivise flexibility reforms for Member States via the European Semester guidance.

- Commit funding from European Social Fund Plus (ESF+) to cushion workers from AI displacement by redeployment, employer-embedded retraining, and providing income bridges. Spending €2–2.3 billion on retraining for AI displacement from the €71.2 billion social earmark for the 2028–2034 period would be an order of magnitude above the European Globalisation Adjustment Fund for Displaced Workers (EGF)'s €35 million/year. This could finance ~130,000–150,000 participants at around €15,000 per head.

- Give Member States country-specific recommendations on high-earner flexibility now, with broader flexicurity reform explicitly conditional on the development of AI's impact.

- Complete, by September 2027, contingency design work (not a legislative proposal) for an EU-backed short-time work reinsurance for a fast-displacement scenario, modelled on the pandemic-era Support to mitigate Unemployment Risks in an Emergency (SURE).

IMPLEMENTATION

- **Potential instruments:**

  - European Semester: Deliver employment analysis and country-specific recommendations (CSRs) to Member States.

  - SURE could serve as precedent for a large-scale social programme by the EU (Council Regulation (EU) 2020/672 (Article 122 TFEU)). It was a ~€100 billion instrument that supported ~31.5 million people.

- **Potential first steps (next 12 months):**

  - **October 2026:** The Commission Work Programme 2027 should carry a non-legislative initiative on AI labour-market retraining.

  - **Q1–Q2 2027:** The Commission should encourage Member States, during the negotiation of the National and Regional Partnership Plans for 2028–34, to retain part of the 14% social earmark for retraining. Those should be paid out if monitoring indicators flag AI's impact.

- **By September 2027:** DG Economic and Financial Affairs and DG Employment, Social Affairs and Inclusion prepare plans for a SURE-style displacement facility, including trigger design, financing route (Article 122 precedent), and interaction with national short-time work schemes.
- **November 2027:** The Joint Employment Report draws on JRC and Cedefop series from Union level recommendation 1.

- **Success indicators:** The Commission commits €2 billion of ESF+ to AI-displacement adjustment measures over 2028–34; AI-displacement section present in ≥1 full Semester cycle (Autumn Package through CSRs); funding for AI automation scenarios retained in ≥5 adopted National and Regional Partnership Plans.

CONSIDERATIONS

- The reform should start with high earners, since the benefits are large and the political cost is low, while a broader flexicurity reform depends on monitoring results (see Union level recommendation 1).
- Retraining measures require careful design to be effective: the average US retraining programme raises employment by roughly 1.7 percentage points and earnings by $800/year per person offered the programme. Employer-embedded sector programmes deliver the largest gains, but they resist rapid scaling. Programme funding should be tied to sector-programme design features, fast-growing sectors with talent gaps, and a built-in evaluation mandate.
- Displacement adjustment should depend on aggregate monitoring indicators (see Union level recommendation 1) rather than on case-specific mobilisation or lengthy individual eligibility tests, to avoid failures of the EGF and the US Trade Adjustment Assistance mechanism.
- The EU should explore joint borrowing measures to provide security in a fast-automation scenario.

Recommendations at the national level

01	02
Feed into the AI labour market monitoring stack	Prepare high-earner-first flexicurity reforms
very high priority	very high priority

01

Feed into the AI labour market monitoring stack

ACTION

National statistical institutes (NSIs) add AI-displacement questions to national labour force surveys for the 2027 or 2028 survey wave, ahead of the common Eurostat module (Union level recommendation 1).

Each Member State can add national questions to its labour force survey under Regulation (EU) 2019/1700.

IMPLEMENTATION

- **Potential instruments:** Professional independence of national statistical institutes in developing statistics is legally protected (Article 2, Reg. 223/2009), and Reg. 2019/1700 mandates only some harmonised questions. Consequently, Member States can add questions.
- **Potential first steps (next 12 months):**
  - **2027:** NSIs add AI as a topic in their labour force surveys.
- **Success indicators:** 10 Member States roll out surveys by 2027 successfully. Eurostat compares this data across countries.

02 Prepare high-earner-first flexicurity reforms

ACTION

- Legislate reduced employment protection for high earners, measured against a national salary threshold set by Member State governments.
- Member States should lighten their dismissal rules for some professionals. They should pair these measures with stronger income support and active labour market spending, following for example the Danish model.
- Deliver retraining and benefit payments to displaced workers in weeks.

MEMBER STATE MODE

Employment and dismissal rules are mainly national competences. Member States therefore retain the principal responsibility for these reforms. A possible model is Germany's proposed relaxation of dismissal rules.

IMPLEMENTATION

- **Potential instruments:** Employment legislation such as the easing of employment protection for high earners should be balanced by active labour market policies to provide support for workers amid faster employment transitions. Portable, pre-funded severance arrangements on the Austrian model could make shorter notice periods more acceptable.
- **Potential first steps (next 12 months):**
  - **2026–27 legislative sessions:** Pass flexibility measures for the most dynamic and high-earning professionals.
  - **Q1–Q2 2027:** Member States prepare flexibility reforms for future implementation.
- **Success indicators:** 15 Member States enact the high-earner flexibility measures across Member States by late 2027; publication of contingency packages; shortening the length of time before payments arrive in cases of displacement.

CONSIDERATIONS

- The measures should be informed by the failures of the EGF: Case-by-case mobilisation was too slow and small relative to the shocks it aimed to address.
- This measure aims to reduce the costs of experimentation, where trying and failing fast is important. Workers employed by dynamic firms, for whom labour market mobility tends to be most important, are likely to be less vulnerable to firings. Targeted flexibility measures also reduce political opposition.
- Labour market flexibilisation for the broader workforce should be paired with income support and active labour market spending to deliver effective sector retraining programmes.

What would be required for a successful European frontier AI project?

The transformative AI strategy laid out in the previous sections does not include a recommendation for Europe to develop its own frontier AI models. Rather, it is a leverage-first strategy outlining how Europe can make itself irreplaceable in as many other layers of the AI value chain as possible and use that leverage to secure frontier model access and a stake in decisions over how this technology gets built. This is because, while per se highly desirable, truly catching up to the accelerating frontier would be politically and financially vastly more costly than any measures currently discussed in mainstream European AI policy.

However, in light of concerns about the recent access restrictions on some US frontier models, some European policymakers and experts have expressed renewed concern that betting on being able to negotiate guaranteed access to foreign frontier AI might be too risky. Commission President Ursula von der Leyen, for example, said when presenting the European Technological Sovereignty Package that Europe "cannot afford to depend on others for the technologies that keep our hospitals running, our energy grids stable, and our services secure."

Both a 'Leverage Approach' (as discussed in the previous sections) and a 'European Frontier AI Project Approach' are valid approaches for Europe to try to regain a position of strength in AI. Both require extraordinary political resolve and urgency given the several exponential trends discussed in the sections above. The European Frontier AI Project Approach has obvious political appeal over and above the Leverage Approach, but it also requires even more political and financial capital to even stand a serious chance of succeeding. A half-hearted attempt at a European frontier AI project consisting of bold political statements and ambitions but not backed up by the extraordinary resolve and resources required would lead Europe down the worst of all possible paths: still without domestic frontier AI capabilities, and without reliable ways to use other leverage to secure access to foreign frontier AI. Based on the past few years, this seems to be the default trajectory: The hopes of catching up to the frontier and thus achieving European sovereignty at the model layer have absorbed considerable political capital and attention, but not come close to being pursued with the resolve and resources needed to succeed.

Internalising this is crucial for European sovereignty and prosperity. Catching up to the frontier is in principle possible, but with US hyperscalers spending more than the equivalent of the Manhattan Project every month in today's money, doing so will be very expensive. As this section discusses, a serious attempt with a significant chance of succeeding would likely cost Europe approximately €790 billion over the first three years of such a project and would require a reordering of European industrial priorities as well as the commitment of most of Europe's leverage to a single, risky undertaking (see Cost summary).

The potential upside of safely developing such models in Europe, including the boost to European competitiveness and sovereignty, and the risk of losing access to frontier AI are both sufficiently high to justify such a bet, especially as European leverage to secure access to foreign frontier AI wanes over time. However, taking this bet is reasonable only *if* the right political conditions, described in this section, are in place. If those conditions are not in place in the near future, European decision-makers should pursue the default strategy presented in the other sections of this document, which, if executed with ambition and resolve, has a good chance of ensuring European competitiveness, sovereignty, and security as much as possible even in transformative AI scenarios. Efforts to develop domestic frontier models that are not backed up by the necessary level of ambition would leave Europe in the worst of all worlds.

This section lays out the conditions under which a European frontier project is a bet worth making, the best way for Europe to attempt this project if those conditions were met, and which of the main recommendations from the previous part of this strategy would change if Europe tried a European frontier project.

Starting conditions

Form a coalition of aligned countries

No government acting alone could fund a project intended to reach the AI frontier. Even under favourable conditions, only a committed bloc of countries would be able to mobilise the required resources and withstand the likely political pressure: Since a middle-power-led frontier project could run counter to US interests, it should anticipate attempts by the US government to slow down or even prevent the effort, either through AI-specific measures or escalation in adjacent domains, such as security cooperation and trade.

A coalition of liberal democracies, anchored in Europe, is the most promising starting point to address these challenges. Among likely participants, Europe itself retains the largest economic, industrial, and fiscal base and is the largest established political bloc. As such, it would be the natural foundation of a frontier project. By adding relevant assets, the participation of other countries would increase the likelihood of success: Canada and the UK could contribute technical talent, while the United Kingdom also has substantial state capacity. Japan and South Korea would add further economic weight and important positions in the semiconductor supply chain. Australia, Canada, and Norway could provide attractive locations for large-scale compute infrastructure.

However, there is a trade-off between coalition scale and political cohesion. Internal fragmentation would pose a major risk to coordinated action, and the withdrawal of any one member could introduce further instability into an already fragile arrangement. It would therefore be dangerous to define the coalition too broadly at the outset by including governments that are not firmly committed, only to lose them later. This risk would be compounded by strong US incentives to target individual participants bilaterally and induce them to defect. For example, if the US offered a close ally, such as the United Kingdom, large-scale data centre investment and a broad exemption from export controls, that ally might plausibly accept that offer. Close allies might also decline to join at the outset and instead pursue bilateral engagement with the US. Given the current state of bilateral relationships and the importance of US security guarantees, South Korea, Japan, and the United Kingdom are unlikely to be reliable participants in a European-led coalition, unless political conditions change (for example through further restrictions on US frontier models, which could lead more participants to join the coalition).

The most likely founding alliance would therefore consist of value-aligned governments able to make substantial contributions: the European Union, Canada, and Australia. Participation should remain open to the United Kingdom, Japan, and South Korea if their strategic assessments change, while structured cooperation should be offered to close partners such as Norway, Switzerland, and New Zealand. The coalition would require a strong and binding commitment mechanism, as defection remains a risk even within this narrower group. One straightforward approach, although politically difficult, would be to require members to commit a majority of their funding, talent, and infrastructure access in forms that could not easily be withdrawn following a change in policy. Other interested governments could contribute funding in exchange for guaranteed access rights below those of full coalition members, but should not receive constituent voting rights, in order to preserve an effective decision-making structure.

Set up a new private company

A private company alone could not develop a frontier model on behalf of the coalition without sustained governmental support. Even in the US, the trend is moving towards stronger government involvement in the AI industry, despite a vital and well-capitalised private

ecosystem. A European-led coalition lacks the venture capital market to create a frontier company organically, and it does not have the time to build this market from scratch. If the trend towards securitisation continues in the US, by the time that Europe has a private frontier company up and running, the US will likely have taken equity stakes in its leading AI developers and interlinked them with the intelligence community. The coalition's private sector champion would thus be competing at least partly with the US government. Finally, any private company would have to be registered in one country and would be open to suspicion of state favouritism.

Instead, a frontier project needs the full backing of the coalition countries, including the support of their national security authorities to protect it from foreign interference, the full support of their trade authorities to secure the supply chains, and the full support of their treasuries to finance it.

Use previously successful institutional designs for developing frontier AI

Despite the need for government support, the vehicle responsible for executing a European frontier AI project should still be a private company. This company requires a highly empowered leadership group, selected on the basis of exceptional and demonstrated research judgement, experience building and scaling very large organisations, and an ability to communicate effectively with funders – which, in this case, would primarily mean a broad group of participating governments.

Below this founding team, the organisation would largely resemble today's most successful frontier AI developers. There is limited scope to experiment with alternative institutional forms: The private frontier developer model is the only structure that has so far demonstrated an ability to produce frontier AI systems. The location of the vehicle would be politically sensitive; ideally, its operations are distributed across several cities, both to accommodate the location preferences of scarce technical talent and to reduce concentration risk in any single jurisdiction. A distributed structure would also spread some of such a project's secondary economic benefits, including stimulation of local ecosystems and taxable income associated with salaries and infrastructure.

Existing teams from interested European and partner-country companies could be incorporated into this vehicle. European AI developers represent valuable concentrations of talent and initial expertise. However, a European frontier project could not be built around any of their existing corporate structures. European companies have not been able to reach the frontier so far, and so they are unlikely to recruit enough of the critical talent required for this. Top-tier researchers would be more likely to join a new organisation established with the explicit, government-backed goal of reaching the frontier. Acquiring approximately half of the relevant companies from coalition countries – including the teams, intellectual property, and valuable data – at close to current valuations could account for roughly €25 billion of such a project's total cost.

Since the vehicle would depend on very large-scale government support, and since governments would have legitimate interests in its approach to safety and security, the private company should be mirrored by a complementary government structure. This structure should remain institutionally separate from the vehicle's technical operations, to prevent bureaucratic processes from slowing execution and turning such a project into a conventional government programme. Each participating country should instead appoint a senior project lead at ministerial or state-secretary level who is capable of navigating the relevant national political institutions. These leads should be supported by high-calibre teams of frontier AI policy specialists, similar in structure and composition to the United Kingdom taskforce that subsequently became the world's first AI Security Institute. The national leads and their teams would act as the main interface between the frontier project and its constituent governments. They would manage communications, translate between technical and

policy language, ensure that the project respects non-negotiable political constraints, and ensure that governments understand when their interventions risk impeding execution. The quality of these government links would be central to such a project's political durability.

Put in place sufficient public budget funding

Once established, the vehicle would need to be sufficiently capitalised. As a rough approximation, a European frontier AI project could cost approximately €790 billion over three years. The figures in this section are stated in euros, but the principal inputs – including chips, construction, and frontier talent – are largely priced in dollars. The estimates therefore carry exchange-rate risk in addition to the underlying general uncertainty. The figures presented here and in the following sections should therefore be understood as order-of-magnitude estimates, benchmarked against available sources, such as disclosed spending by leading US developers and filings disclosing rental compute contracts. A consolidated ledger is provided at the end of this section.

Financing would need to come primarily from public budgets and dedicated debt instruments rather than from the private sector. European and allied governments retain substantial borrowing capacity, and a successful frontier project could eventually generate significant revenues. Until then, coalition governments would need to establish the legislative and executive mechanisms required to enable public investment. Treasuries, pension funds, sovereign wealth funds, and related institutions would provide direct funding where necessary. Over time, a European frontier project should seek to attract substantial private capital. Established industries across participating countries have significant financial resources and a strong interest in gaining exposure to AI and could be offered preferential access, model fine-tuning, or other commercial advantages in exchange for investment. The coalition should not, however, assume that private financing will be available at the required scale. The working assumption is that participating countries will bear most of the initial three-year cost.

The principal advantage of this model is that it creates a public ownership structure. Governments would absorb a European frontier project's early strategic and financial risk, while retaining ownership of the frontier company. Such companies can command very high valuations, and the sovereign equity generated by a successful European project would later become a significant public asset. In this sense, a European frontier project would be a highly concentrated and risky sovereign investment rather than pure fiscal expenditure. If it succeeded, the company could subsequently be privatised through European capital markets, subject to strategically acceptable limits on private ownership. Participating treasuries could thereby recover part of their initial investment through staged public offerings. Restrictions on the early withdrawal or sale of public equity would also serve as a commitment mechanism: A member that withdrew before completion would forfeit part or all of its initial contribution.

The overall cost of a European frontier project would be substantial but within the fiscal capacity of the proposed coalition. It would be comparable in scale to special debt packages previously undertaken by countries such as Germany, while being distributed across a much broader group. The principal constraint is thus not fiscal capacity, but a sustained political commitment.

Establish commitment mechanisms

A project of this scale faces two opposing failure modes: Participating governments may withdraw from a successful project if domestic political incentives change, and coalition members may continue to commit resources to a project that is visibly failing. The institutional design should address both risks, but the required safeguards point in opposite directions: Governments will want to retain the option of discarding a failed attempt, while

a credible commitment mechanism requires a degree of irreversibility. The appropriate balance is to pre-commit the main inputs while making such a project's technical milestones subject to review. Funding tranches, infrastructure access, and the equity lock described above should be committed for the full build period. A withdrawing member would forfeit its stake, and no single government's exit should be capable of depriving the project of the resources required to continue. Technical progress, by contrast, should be assessed independently at fixed intervals – for example after the first 18 months, and later more frequently – against the existing frontier. The coalition should agree in advance on criteria under which the attempt to reach the frontier would be discontinued and such a project converted into its fallback assets: European-only compute capacity and a strong concentration of technical talent. These conditions should determine whether the coalition spends *additional* resources beyond the initial commitment, rather than reopening that commitment. A project that can be deprived of funding in its second year following a national election in one of the participating countries would not be sufficiently credible to justify launching.

Compute

Building frontier models is extremely resource-intensive. European policymakers have sometimes assumed that spending these resources could be avoided, for example, through discovering a new, more efficient paradigm or advances in 'world models' that reach frontier capability at a fraction of the cost. Such approaches can be sensible bets to make that may or may not reach frontier performance on the most important tasks. One structural disadvantage they face is that, if such a fundamentally different approach proves viable, a compute-rich frontier developer in the US with access to a large concentration of talent may also be likely to discover or at least commercialise it.

A paradigm-shifting research programme can therefore form a valuable part of a broader portfolio. On its own, it would be insufficient for the task considered here: establishing a project with a proven path to the frontier that does not rely on contingent technical bets, and thus remains credible to fiscally conservative treasuries. This changes the calculus compared to different sovereignty approaches, where these technical approaches should play an important role (see Objective 2.2). Only one paradigm is currently known to lead to the frontier; alternative approaches could promise a more favourable balance of reliability, benefits, and risks, but remain as yet unproven at scale.

Build sufficient compute infrastructure and energy supply

Cost estimate

For a European frontier project to compete within the current technical paradigm, it requires a substantial amount of compute. On the one hand, this would likely be somewhat less than a leading US developer requires over the same period, since a European project would initially face lower inference demand due to a smaller customer base. Comparisons with aggregate hyperscaler capital expenditure are therefore misleading. At the same time, frequently cited estimates for the comparatively low training costs of Chinese open-weight models such as Kimi K2 and DeepSeek R1 are equally misleading. Even where the frontier model's final training run costs only several million dollars, the total compute requirement extends far beyond it, including the internal deployment of proprietary coding agents intended to accelerate research, the large number of research and development experiments needed to reach and advance the frontier, and pre-training and post-training at scale, both of which have become considerably more expensive since 2025.

The best available estimate is that a European frontier project would require approximately as much compute as the currently scheduled buildout of one major US frontier developer. It would need less inference compute than that developer, but potentially more R&D compute, while US developers are also likely to add further compute beyond their existing plans over the coming years. On this basis, a European frontier project would require roughly 16 GW of compute in IT load, of which roughly 12 GW would be available by the end of the three-year build (9 GW of built capacity, 3 GW of rented compute). This would come with a plausible cost of around €529 billion.

This is an estimate of the minimum amount of compute required by a European frontier project. Actual requirements could be substantially higher. Limited experience in data centre construction, the need to build more secure facilities, and the premium associated with entering an already contested compute market could, for example, increase the cost further. Note also that this estimate only concerns the compute requirements of a single frontier project, and is thus significantly below Europe's total compute target of 45 GW of total facility power that is recommended no matter what strategy Europe chooses (see Immediate Objective 3).

Buildout and energy

The logistical requirements for bringing this compute online are relatively well understood and have been detailed in a comprehensive analysis by the Carnegie Endowment. The coalition would need to select a portfolio of host countries with different comparative advantages and pursue accelerated infrastructure programmes in each. Some hosts would be chosen for their ability to begin construction quickly, allowing work to start as early as possible; others for their capacity to accommodate buildout at greater scale, giving such a project sufficient depth without encountering common bottlenecks. Some participating governments would seek to retain a share of the investment domestically, while others might require local instances of secure data centre capacity for national security purposes. Building the required compute would be an extremely ambitious infrastructure project, though at a scale that European governments have previously managed under crisis conditions. One relevant precedent is the emergency-law framework under which Germany constructed LNG terminals in less than a year during the 2022 energy crisis.

The same applies to the energy requirements of a European frontier project. Across the coalition, several options remain available for bringing data centres online quickly: renewable-heavy generation mixes in Australia, Scandinavia, and Iberia; nuclear-powered grid capacity in France, Finland, and potentially South Korea or Japan; and deindustrialised areas with high-capacity grid connections in Germany and possibly the United Kingdom. Distributing compute across these locations would reduce the risk that limited behind-the-meter generation capacity becomes an immediate constraint. A European frontier project would nevertheless place significant additional pressure on electricity grids and aggregate energy supply. Over the longer term, it would therefore need to offset at least part of this burden by developing new dedicated generation capacity.

Procure interim rental compute

Before the dedicated clusters become operational, a European frontier project would need to secure as much rental compute as possible. Two measures would be particularly important. First, participating governments would need to consolidate and requisition, with compensation, publicly owned supercomputers, redirecting them from their existing uses to the frontier project. Second, a European frontier project would need to procure any suitable capacity available on the GPU rental market. This compute is expensive to rent: Anthropic alone is reportedly paying $1.25 billion per month to use xAI's Colossus cluster. However, this approach would allow research to begin immediately. Both measures – requisitioning existing supercomputers and renting interim compute – would together cost approximately €105 billion.

Set up a coercion shield

Physically building out the data centres is the more tractable part of providing enough compute for a European frontier project. The challenging part is acquiring the chips without US intervention. While US AI export policy has been inconsistent, an administration willing to impose export controls on frontier models might similarly restrict exports of frontier chips to a project explicitly intended to develop frontier capabilities outside the US. Such restrictions would be relatively straightforward to implement: There is ample US demand for the available supply, the relevant legal authorities already exist, and the supply chain is concentrated around a single US chip company.

This is where Europe's semiconductor supply chain leverage, together with that of its coalition partners, becomes relevant. Under normal conditions, individual chokepoints provide only limited influence. But in the context of a European frontier project, they could serve as a credible anti-coercion tool. If the US sought to prevent the sale of advanced AI chips to such a project, the coalition could respond by restricting access to relevant inputs – such as EUV lithography machines, critical raw materials, and, if East Asian countries participated, memory chips – that US companies need for their own frontier development. Trade policy measures should be used for this purpose (see Immediate Objective 1, Union-level recommendations 1 and 2).

Two points would be essential. First, the mechanism should be strictly reciprocal and remain dormant unless US authorities intervene to block chip supplies to the European frontier project, and it should be suspended as soon as supply resumes. Premature or rhetorical deployment would exhaust the coalition's largest reserve of leverage without securing a corresponding benefit, while undermining the export control diplomacy pursued elsewhere in this strategy. Second, the mechanism would be more credible than attempts to leverage ASML alone because it would be deployed with broad support across the coalition and directed not only at the US government, but also at US chip companies, whose continued operations depend more directly on access to upstream semiconductor inputs. Properly designed, the coercion shield would create incentives for US chip companies to advocate internally against US restrictions. It would support these companies' interest in avoiding monopsonistic markets and upstream bottlenecks, make clear the risk of a severe disruption to the chip market, and establish continued sales to the European project as the least costly outcome.

This approach could work, but it would be highly risky. The coercion shield would commit some of the coalition's most important existing sources of leverage and economic participation to protecting the European frontier project, rather than using them to secure frontier model access directly or pursue other strategic objectives. In this respect, a European frontier project would constitute a genuinely concentrated strategic commitment: Europe and its partners would place much of their existing AI-related leverage behind the effort to establish a sovereign frontier capability. Without such a coercion shield, however, a European frontier project may not be able to acquire the chips required to sustain it.

Talent

A European frontier project needs two types of talent. First, the founders need to be truly exceptional leaders who strongly identify with the project. Second, the larger group of researchers need to be of the highest international quality and consist largely of former employees of frontier companies.

Attract leading frontier AI talent

In order to attract international top talent, a European frontier project would need to offer a challenging but credible mission and sufficient compensation. The latter is easier to achieve. Salary packages would have to match those offered by US frontier developers, including the expected value of equity, since a European project would present a substantially less attractive prospect of a future public listing than a US AI company. Securing sufficient senior talent could therefore cost close to the amounts Meta reportedly offered leading researchers, with individual senior packages reaching tens of millions of euros. A European frontier project would also need to provide standard enabling measures, including accelerated visa procedures.

Across such a project, total personnel expenditure could reach approximately €34 billion in the first three years, across three tiers of staff, equity buyouts, and operational costs. First, a European frontier project would require a founding group of approximately eight people, distributed across participating nationalities, with the calibre currently associated with the co-founders of a frontier company. Guaranteed packages in the hundreds of millions of euros per person may be necessary to convince such people to leave their current positions, implying a total cost of approximately €1.9 billion. Second, a European frontier project would need a senior research cadre of around 150 people recruited on terms comparable to Meta's reported offers, at €5–8 million per person per year. This would amount to roughly €3.6 billion over the first three years. Third, it would require a broader technical workforce of perhaps 3,000 staff. This is below the mid-thousands employed by leading frontier developers, as the European project would initially offer a smaller set of products. These staff would probably be more expensive for a European project than for an established frontier developer. Without the prospect of a comparable IPO or a large equity upside, a European frontier project would need to compensate in cash for value that US companies can offer in stock, at a likely total cost of €15–19 billion. In addition, such a project would need to compensate hired staff for unvested equity forfeited when they joined, requiring at least €10 billion in upfront expenditure, and operational costs over the first three years of approximately €1.5–2 billion.

The most effective way to establish a mission that researchers regard as credible would be to recruit and empower respected technical leaders. Researchers reasonably assume that a project's prospects depend heavily on the judgement of its leadership and on whether those leaders have sufficient authority to act on that judgement. The first recruitment priority should therefore be the founding group. There are likely to be European and allied nationals with some degree of patriotic motivation who are also seeking a new technical challenge. Some could be convinced to join if the coalition signalled a credible commitment to supporting their work and respecting their operational independence. Once recruited, the founders would become the project's representatives to the wider research community. They would communicate its technical vision, provide evidence of the commitments made by participating governments, and recruit the broader teams required to execute it.

Hiring former frontier company employees is not merely a means of recruiting the best researchers. It is the standard channel through which frontier knowledge nowadays diffuses across organisations. The movement of leading researchers between US frontier developers forms part of the sector's existing competitive equilibrium: Researchers carry techniques and organisational practices to their new employers, limiting the period for which such knowledge remains confined to a single company. By recruiting a sufficiently strong group of researchers, the European frontier AI project gains valuable frontier AI expertise. However, this channel may narrow over time. Frontier developers are becoming more internally siloed, while advances in internally deployed AI systems are reducing the need for information exchange between researchers. Therefore, the European project would have to hire leading researchers before this channel becomes less important for transferring frontier knowledge.

Secure access to frontier coding agents

AI agents already perform a substantial share of research and operational work at frontier companies: They write and execute code or complete administrative tasks, while human researchers are increasingly responsible for generating ideas and directing agent activity. The share of AI R&D tasks, including more open-ended work, conducted by AI agents is only set to increase as leading companies chase after recursive self-improvement. Securing access to advanced agents during the early stages of a European frontier project would therefore be a central condition for international top talent to join. Anthropic has already restricted the use of its most capable models for frontier development, and competing developers may adopt similar policies. Because a European project would not initially possess its own frontier model, it would also lack frontier-level coding agents, substantially reducing the productivity of its non-human workforce and slowing technical progress.

This constraint would diminish once a European frontier project developed its own coding agents and began to deploy them internally. It also illustrates why a sovereign frontier developer may be the only stable model for frontier AI access: A permanent fast follower would lack the internally deployed agents required to keep pace with compounding productivity gains at the frontier. Obtaining external coding-agent access would be difficult, but likely possible. A US developer might see commercial or reputational value in a large contract with the European project, or an attractive commercial structure might secure access during the first year, before the US government or companies concluded that the project had a credible chance of success. If a European frontier project secured sufficient compute, it could also offer to host US AI agents on its own infrastructure, increasing the financial incentive for US developers to provide access. Assuming that such access would command a substantial premium, a European frontier project should budget approximately €3 billion to procure US AI agents over its first 18 months. In addition to AI agents, such a project will have to budget for additional costs of approximately €3.5 billion for training data and RL environments.

Sustaining conditions

Even if a European frontier project started from the most favourable initial conditions and managed to develop a frontier model within the first two years after its launch, several further conditions would have to be in place to sustain this success over time.

Remain robust to domestic political pressures

First of all, a European frontier project would have to be robust against political pressures. The initial motivation for such a project could easily weaken over time, leading individual members to withdraw. AI itself would become an increasingly contested topic as labour market effects and misuse risks became more visible. Contributing governments might be reluctant to continue directing substantial resources towards frontier parity if the frontier itself was generating concern among their electorates.

A credible effort should mitigate these pressures from the outset. Relevant measures include commitments to protect households from project-related increases in energy prices; direct AI dividends for groups disproportionately affected by such a project's progress; electricity-price support for incumbent industries, so that compute buildout does not come at their expense; and payments to host regions to reduce the risk of local opposition to data centres. These measures could require approximately €80 billion. Many of these measures would also constitute independently defensible public investments of the kind already under consid-

eration by several European governments. Compared with the hundreds of billions of euros that European governments have spent insulating their economies from the effects of the COVID-19 pandemic and energy crises, the required package is relatively limited.

At the same time, a European frontier project could also receive greater societal acceptance than a similar effort in the US. For example, it would not be under the same commercial pressures to compete for consumer markets or sustain investor expectations through divisive applications, including for rapid labour automation or addictive consumer products. These structural advantages would not remove the need for sustained and effective public communication, but they could make a European frontier project less politically vulnerable than the current US model.

A European frontier project should also serve as the first model participant in the Union's enforcement regime. Full compliance with the systemic risk obligations for general-purpose AI models (Objective 3.1) would provide the least costly available demonstration that the regulatory framework can operate effectively at the frontier.

Aim for successful long-term commercialisation

If all of the above measures succeed, the European frontier project would eventually develop a frontier AI system. However, this would not guarantee sustained success. US AI developers are sophisticated product companies with established consumer distribution and portfolios of highly lucrative enterprise contracts. Frontier development is economically viable for them because the returns to technical leadership are substantial and can be reinvested in maintaining that lead. No European company currently has access to a comparable market position or product-development capability, and non-US companies have repeatedly struggled to scale at a similar level of ambition.

A European frontier project would initially become the main provider of frontier AI to coalition governments and to private-sector applications with high security requirements, especially if leading US systems were unavailable. Beyond these uses, the coalition might need to adopt measures to encourage domestic firms to use the project's models even where better US alternatives were available. Such measures would carry economic risks, but might be necessary to establish a reinforcing relationship between the coalition's critical industrial base and its frontier project. If such a European product were not genuinely competitive, this approach would require companies to adopt an inferior economic input. If it succeeded, however, privileged access to what would collectively constitute the world's largest market (if all potential coalition members joined) could help establish it as a durable commercial actor. Once it had developed a sufficiently capable and reliable product, it could also compete internationally. In global markets, it would have a distinct strategic position relative to US systems whose availability may be subject to political constraints and Chinese systems that are often regarded as insufficiently trustworthy.

Commercialisation would present a separate and potentially greater challenge: Building a successful frontier AI company is more difficult than building a frontier model. During its initial years, the coalition should expect to have to subsidise a European frontier project and its continued provision of frontier AI, treating it as similar to the costly investment in a strategically essential defence contractor, with the possibility – though far from a guarantee – of becoming a major commercial success. The upside would nevertheless remain substantial. A well-designed European frontier project could claim greater political legitimacy, both domestically and internationally, than US or Chinese AI developers.

Cost summary

ITEM	DESCRIPTION	APPROXIMATE CENTRAL ESTIMATE	APPROXIMATE RANGE	SOURCES
Sector buy-in	Buying some coalition AI companies outright, for their teams, intellectual property and datasets.	~€25 billion	~€10–40 BILLION	VALUATIONS (Mistral Series C, Mistral's reported 2026 mark, Cohere-Aleph Alpha, Black Forest Labs, ElevenLabs).
Compute	Accelerators and the data centres to house them: 16 GW of IT load built (roughly 18 GW of total facility power), of which roughly 9 GW of IT load would be running by 2029.	~€529 billion	~€320–640 BILLION	COST PER GIGAWATT (Plan Prométhée, Epoch AI), EUROPEAN BUILD COSTS (Turner & Townsend), FRONTIER COMPUTE LEVELS (Epoch AI on frontier laboratories).
Data centre operations	Electricity, maintenance, and taxes for running the data centres.	~€13 billion	~€11–43 BILLION	OPERATING COSTS (Plan Prométhée), ELECTRICITY PRICES (Eurostat), NEW GENERATION CONTRACTS (Bundesnetzagentur solar auctions, onshore wind auctions).
Interim compute	Compute rented from others to bridge the gap until the coalition's own clusters come online (around 1 GW in IT load in the first year and 3 GW in the second and third year each, for a total of 7 GW-years)	~€105 billion	~€58–158 BILLION	RENTAL CONTRACTS: Anthropic-xAI, IREN-Microsoft, Oracle-OpenAI. LEASING SCHEDULE: Plan Prométhée. MARKET RATES: SemiAnalysis.
Training data and RL environments	Expert human data and feedback, licensed content, training environments, and a one-off collection of books.	~€3.5 billion	~€2–6 BILLION	SPENDING PER LABORATORY: Labelbox's chief executive, TIME. VENDOR MARKET SIZE: Mercor's reported revenue. CONTENT LICENSING: licensing tracker. TRAINING ENVIRONMENTS: Epoch AI.
Personnel and operations	Pay for the founders, senior researchers, and ~3,000 further staff, buying out the equity they forfeit, plus recruiting and operational costs.	~€34 billion	~€15–40 BILLION	PAY LEVELS: Meta pay reporting, OpenAI on levels.fyi, Anthropic's audited UK accounts. UNVESTED EQUITY: Meta, Alphabet. ALTERNATIVE STAFFING MODEL: Plan Prométhée.
External coding agents	Access to frontier coding agents, bought from existing developers for at least the first 18 months.	~€3 billion	~€2–5 BILLION	NO GOOD REFERENCE CLASS ESTIMATES AVAILABLE. OVERALL ESTIMATE IS TAKEN FROM LEICHT (2026).

				DIFFICULT TO MAKE A GOOD REFERENCE CLASS ESTIMATE, HENCE THE RANGE IS WIDE.
Political insulation	Payments shielding households, energy-intensive industry, and host regions from the project's effects, plus AI dividends.	~€80 billion	~€25–110 BILLION	EXAMPLES COULD INCLUDE: INDUSTRIAL ENERGY SUPPORT: European Commission. CRISIS-RESPONSE SCALE: Bruegel. PRICE-CAP OUTTURN: German government. HOST-REGION PAYMENTS: Danish Energy Agency, Irish government.
TOTAL		~€790 billion	~€445–1,040 BILLION	COMPARATOR: Plan Prométhée (~€620 BILLION). THE PROJECTION DOES NOT PRICE IN SECTOR BUY-IN, POLITICAL INSULATION, TRAINING DATA/RL ENVIRONMENTS, OR EXTERNAL CODING AGENTS. ITS COMPUTE BUILD COSTS ARE USED VERBATIM IN THE ESTIMATES ABOVE; ITS DATA CENTRE OPERATING COSTS ARE ADJUSTED TO EUROPEAN ELECTRICITY PRICES. ITS INTERIM COMPUTE VOLUMES ARE KEPT BUT RE-PRICED FROM OWNER'S COST TO MARKET RENTAL RATES.

TABLE 2 ORDER-OF-MAGNITUDE BOTEC FOR THE COST OF THE FIRST THREE YEARS OF A EUROPEAN FRONTIER AI PROJECT. All estimates are approximate and should be considered as order-of-magnitude estimates. Note that large public projects should also factor in a contingency allowance which may be needed, for example, because of overruns on the build, higher AI chip prices, delay, the need for higher-security data centres, and scope growth (Flyvbjerg, Holm & Buhl, 2002; Flyvbjerg & Bester, 2021; Flyvbjerg et al., 2022; UK Treasury, 2003; Mott MacDonald, 2002).

How would a European frontier AI project affect previous recommendations?

If Europe pursued such a project, most other recommendations in this strategy would remain unchanged: Europe should still, for example, build out compute, secure access to foreign frontier AI as long as its project's success is uncertain, ramp up societal resilience against AI-enabled harms, and support international coordination efforts. However, there are some areas where the recommendations in this strategy would change in emphasis or substance, or even reverse:

- **Compute buildout** (see Immediate Objective 3). The leverage approach envisions that, while some share of AI data centres on EU soil would support European researchers/start-ups and sensitive inference needs, the majority would serve US frontier companies – and function as physical leverage to secure access to their best models over time. If Europe pursues a frontier project, these priorities invert: Europe should use its most attractive sites to meet such a project's essential compute demands, pricing in enough headroom in case these turn out to be higher than expected, and only offer excess capacity to hyperscalers.[25] The goal of hosting 15% of global AI compute would remain intact, as hosting a greater share of this strategic infrastructure improves Europe's bargaining position whichever strategy it chooses.

25 Even if hyperscalers built the majority of AI data centres in Europe, a European frontier project could still rent from them, and it would likely do so to some extent anyway. However, this comes with a risk of high markups and potential restrictions by the US government.

- **Chip access** (see Immediate Objective 1, Objective 1.2)**.** Under a leverage approach, advanced AI chips are largely bundled with hyperscaler investment. This makes access negotiations easier and more narrowly scoped, as they would focus on securing enough chips for the minority share of European-only data centres. For a European frontier project, chip supply is a vital ingredient that could easily be restricted as the European effort gains credibility, since every chip used to boost European AI models is one that is not available to US companies. If Europe goes all in on such a project, advanced chip access therefore becomes a much higher priority and the leverage required to secure it should be in place and usable right from the outset.

- **Public procurement** (see Objective 1.1, Immediate Objective 2). A leverage approach would entail that public institutions buy the most capable model available, relying on portability safeguards in case access restrictions cannot be avoided. Under the conditions of a European frontier project, there is a stronger case for procuring from the European provider, even during the transitional period where the best European models lag behind the frontier.

The fast-follower alternative

The most frequently proposed alternative to a European frontier project is a 'fast follower': an AI company at some regular distance behind the frontier, something that Chinese and a few European developers have attempted with varying degrees of success. The relevant distinction is between fast following as economic policy and fast following as a solution to the sovereignty problem. As economic policy, fast followers are valuable. They should be decentralised, commercially driven, and encouraged rather than organised through a single state-backed project, and this strategy recommends them even under a leverage-first approach. They can reduce the dependence on foreign developers, provide secure access to models in domains where frontier performance is unnecessary, and are the basis for specialised applications built on sensitive or proprietary data. None of these benefits requires a large consolidated project, and they do not conflict with the proposal set out in this chapter.

As the primary strategy for protecting European sovereignty, however, fast following is less promising. It assumes that frontier developers will continue to release their best models, that it will be economically feasible to reproduce their performance at a stable delay of several months, and that the resulting model will retain strategic value. These assumptions may not hold. For example, if frontier AI systems turned out to be essential for Europe's economic competitiveness and national security, a fast follower would provide little protection against access restrictions by foreign governments or companies. Conversely, if frontier models remained widely available, the case for a domestic fast follower would also be weak. Europe would be better placed importing US models, avoiding the costs associated even with fast following, and concentrating resources on adoption and downstream value capture.

Fast following may also become progressively less feasible as frontier AI becomes more strategically important. If the most capable systems are available only to the US government and a small group of privileged users, the frontier will become less observable to outside developers. Followers would be unable to distil frontier systems and increasingly unable even to infer the structure and capabilities of state-of-the-art models. A widening gap between the restricted frontier and the commercially available models would also allow leading developers to use superior internal AI systems to accelerate their own research.

In short, the conditions under which Europe and its partners would realistically aim to build frontier models are those of exceptional necessity. These are precisely the conditions under which a fast-follower strategy would be insufficient and sovereign frontier capability would be required. Therefore, a frontier project should not aim at merely sub-frontier capability levels.

Deep dives

Compute deep dive

Analysis: Why Europe's compute-poverty endangers its prosperity, sovereignty, and security

In a world with transformative AI, access to computing power will be as crucial as access to energy. Some strategic resources are the backbones of economic prosperity. One example is energy: Essentially no nation in the world is both poor in energy and economically prosperous. This strong correlation between energy use and GDP per capita (see Figure 13) exists because energy underlies almost every process of value creation. Reliable access to energy resources is therefore a prerequisite to economic competitiveness and related values such as national sovereignty and security.

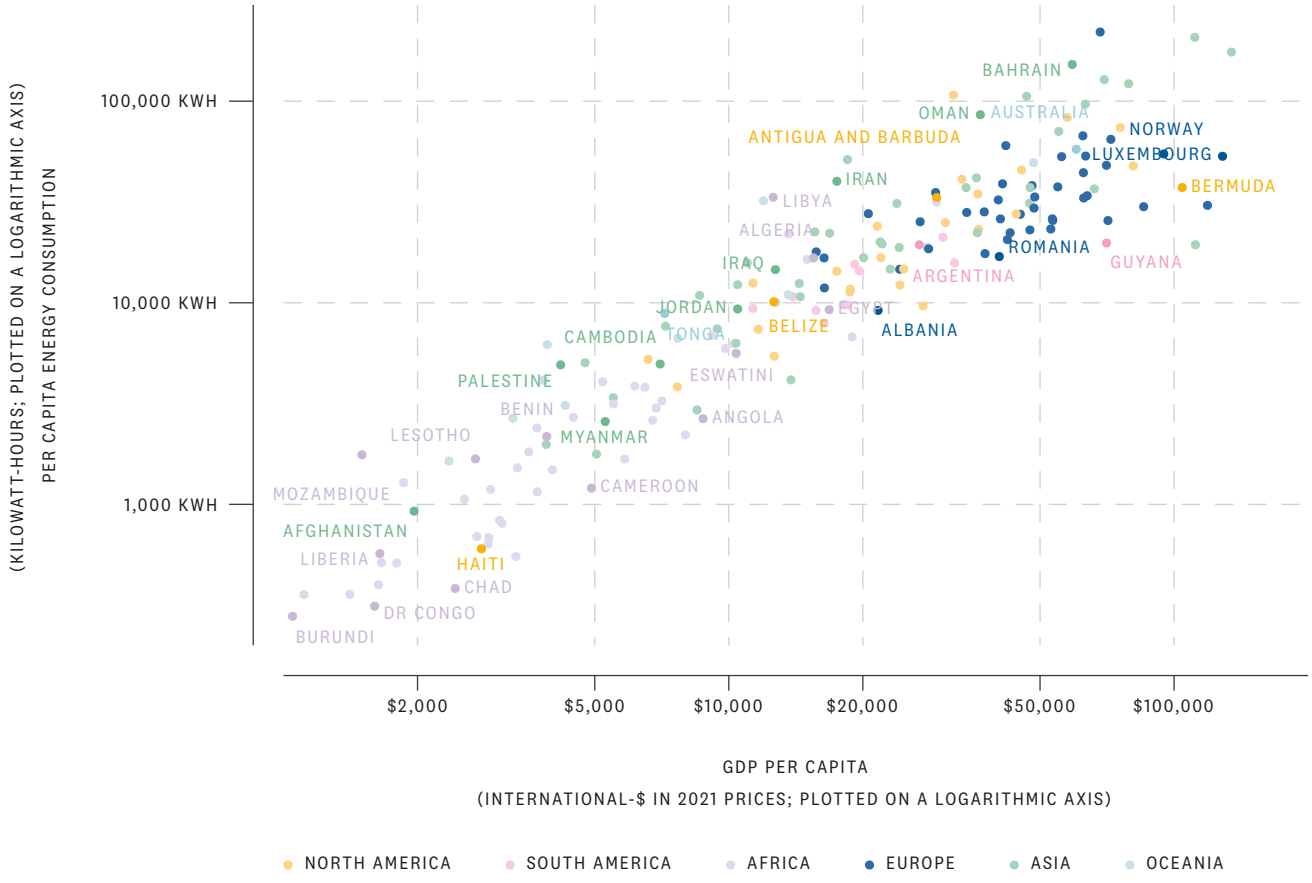

FIGURE 13

RELATIONSHIP BETWEEN ENERGY USE AND GDP PER CAPITA. Being able to access a strategic resource such as energy correlates strongly with average prosperity in a country. Source: Our World in Data (2026).

This strategy makes recommendations to help Europe prepare for a world in which AI becomes transformative within the next few years. It defines 'transformative AI' as *AI that changes the basic structures of society, including the economy, science, and geopolitics, materially faster than any technology in history*. If AI becomes transformative, reliable access to *computing power* (or 'compute' for short) will become as necessary for economic

competitiveness as energy is today, and for the same reason: As AI becomes transformative, it will underlie almost every instance of value creation in some way, and compute is the physical infrastructure on which AI depends. Countries without access to compute will be unable to use AI, and as a result, their companies, public services, and defence systems will become less competitive. Countries that are compute-poor – either because they have no data centres on their soil or because they lack reliable access to cloud services – will fail to be prosperous, sovereign, and secure.

Compute is already scarce today. Compute demand is already increasing rapidly, even with current systems that are much less transformative than the systems that are the focus of this strategy. Around the world, AI companies are scrambling to get as much compute online as possible. At the same time, demand is growing so fast that the world is likely headed for a 'compute crunch', where demand outpaces supply and customers compete for scarce access to frontier AI. According to the research nonprofit Epoch AI, current trends suggest that such a compute crunch is "near, if not already here", especially for long-context, agentic AI workloads. Different lines of evidence point in this direction:

- **AI demand appears to be growing much faster than supply.** Epoch AI finds that global inference supply – the ability to serve AI models at scale – is growing ~3.4x per year, driven by gains in both the amount and the efficiency of chips. However, several proxies suggest that demand for AI at a fixed model size and price is growing closer to 10x per year. This is driven to a large extent by the highly compute-intensive, agentic AI systems that are unlocking particularly valuable applications (e.g. in AI-assisted coding, already a multibillion-dollar market). If these trends continue, rising prices for frontier AI – and potentially export restrictions by the governments of leading AI nations – could force European users to switch to cheaper, less capable 'sub-frontier' models.

- **Market prices reflect acute scarcity.** The Nvidia H100 is a legacy AI chip, which came out in 2022 and is two generations behind the Rubin chips likely to be rolled out this year. Instead of depreciating steadily, one-year H100 rental prices actually increased by approximately 40% between October 2025 and March 2026 (Figure 14). According to research company SemiAnalysis, "on-demand GPU rental capacity is sold out across all GPU types", which makes the search for AI chips "like trying to book airplane tickets on the last flight out", characterised by "high prices, and almost no availability".

- **Companies are sourcing power from increasingly improbable assets.** In the US, grid constraints are forcing companies to seek creative solutions to power data centres in the face of rapidly increasing demand. For example, some are using retired commercial aircraft and ship engines to supply on-site ('behind-the-meter') power for data centres. One company has proposed repurposing nuclear reactors from old Navy warships to meet customer demand for AI services.

- **Frontier companies in both the US and China have explicitly stated that they are compute-bound, not demand-bound.** According to Anthropic's CEO Dario Amodei, the company faced an acute compute shortage when, in the first quarter of 2026, its revenue and usage were growing at an annualised rate of 80x, far above the roughly 10x annual growth rate it had planned for. Amin Vahdat, Google's infrastructure VP, reports that Google has to double its AI serving capacity roughly every six months to keep pace with demand. Chinese AI companies are even more compute-constrained. DeepSeek's CEO Liang Wenfeng stated in 2024 that the company is more constrained by US chip export controls than by a lack of capital. Similarly, Alibaba's CEO Eddie Wu said that "the pace at which we can add new servers is insufficient to keep up with the growth in customer orders", noting that all their GPUs are running at full capacity.

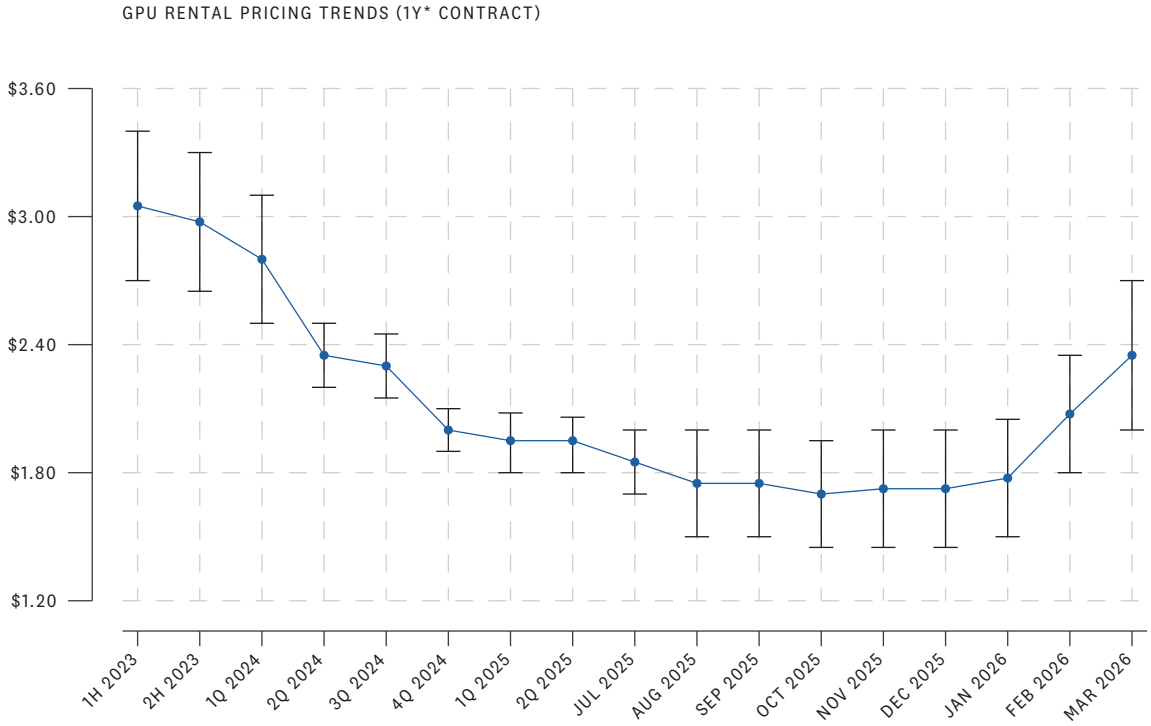

FIGURE 14

RENTAL PRICES OF LEGACY CHIPS OVER TIME. Between October 2025 and March 2026, one-year rental prices for an Nvidia H100, a four-year-old chip generation, increased by approximately 40%. Source: SemiAnalysis (2026).

Compute will likely remain scarce for the foreseeable future. Chipmaking companies plan to invest over $100 billion in 2026 alone to expand the global supply of AI compute. Despite these efforts, several trends suggest that compute will remain a scarce resource over the coming years:

- **Compute demand for both training and inference will likely continue to grow rapidly.** As described above, several proxies suggest that inference demand at a fixed model size and price is likely expanding ~10x per year, much faster than supply. Similarly, the compute required to develop a frontier model in the first place is also increasing exponentially, growing 4–5x per year. According to Epoch AI, this trend is likely to continue until 2030, meaning that by then a single frontier training run could draw 4–16 GW of power. Even the lower bound of this range exceeds the entire AI compute supply in Europe – including countries such as the UK and Norway – by the end of 2026 (~2 GW).

- **Cheaper inference appears to be tightening rather than relieving scarcity.** Inference prices – the cost of serving a trained AI model to customers – are steeply declining for any given level of capability. For example, thanks to greater algorithmic and hardware efficiency, the cost of achieving GPT-4-level performance on PhD-level science questions has plummeted by 40x per year. These efficiency gains can relieve compute scarcity for routine, short-context workloads, where cheaper smaller models substitute for last year's frontier AI models. But at the frontier, demand so far appears to be growing faster than supply. Two factors play a role here. The first is a classic rebound effect: AI models are a general-purpose tool that can be used for a wide variety of purposes. As the tool becomes cheaper, people deploy it more widely. The second factor is that while AI capabilities at any given level become cheaper over time, AI companies are also releasing more capable – and usually more compute-intensive – models at a rapid pace. For example, the reasoning models at the heart of agentic AI systems like Claude Code have many more useful real-world applications than their predecessors, which leads to greater demand for these models despite the fact that they require more computational resources and are thus more expensive.

- **Supply chain bottlenecks are likely to remain persistent.** Cutting-edge AI chips are the product of the world's most complex supply chain, which is already running into severe technical and logistical limits to further scaling. For example, high-bandwidth memory – a key component of modern AI chips – is fully sold out for 2026, a bottleneck that industry and independent experts think could persist through 2027 and beyond. Similarly, TSMC's 2-nanometre chip production process, required for the latest generation of AI chips, is reportedly sold out until 2028. From 2028 and beyond, some experts expect extreme ultraviolet lithography (EUV) machines to become the biggest bottleneck – chipmaking machines so complex that only one company in the world, the Dutch ASML, is able to produce them. This combination of supply chain constraints and rapidly increasing demand means that compute is unlikely to become an abundant commodity anytime soon.

  Given compute scarcity, universal access to frontier AI capability is not a given. In a world with transformative AI, this will turn into a decisive geopolitical factor: As AI becomes an essential driver of economic, political, and military power, compute access will directly determine a region's place in the global distribution of power.

  Without physical control, access to compute is precarious. Countries and regions whose AI workloads mostly run on foreign infrastructure – such as the EU, which only hosts less than 5% of global AI compute – could be cut off from data centres and the frontier models that run on them at any time. This could happen for a variety of reasons, ranging from coercion, to misuse concerns (e.g. terrorists using AI to design explosives), to the simple fact that, in a compute-scarce world, host countries might prioritise their own customers before supplying compute to the rest of the world.

  Access restrictions around Anthropic's Mythos illustrate this. Mythos is an AI model with strong offensive cyber capabilities and thus high misuse potential. Anthropic launched it in April 2026 through the tightly controlled Project Glasswing, initially restricting it to selected partners; no European actor was included. Only after sustained negotiations did Anthropic admit the EU cybersecurity agency ENISA. Less than two weeks later, the US government ordered Anthropic to suspend all access to Mythos and its consumer counterpart, Fable 5, for any foreign national, prompting Anthropic to disable both models worldwide. If the EU hosted a sizeable fraction of physical AI infrastructure on its soil, including AI data centres used by US companies, it would be in a better position to negotiate frontier model access for EU companies and governments, ensuring that they are able to use state-of-the-art capabilities in areas such as cyberdefence and to stay economically competitive (see Objective 1.1 for a more detailed explanation of this 'compute-for-access' logic).

  Within the EU, pooled infrastructure is a reasonable policy: Some countries are more natural homes for large data centres than others, including countries with cheap or reliable access to green energy, such as the Nordic countries, France, or Spain; or countries with decommissioned industrial sites that have existing, GW-scale grid connections, such as Germany. It makes sense to concentrate compute in these regions and distribute access fairly across the Union.

  By contrast, dependence on *non-European* countries for compute supply is a risky bet, especially in the current geopolitical climate. The more critical a sector is for the functioning of Europe's economy and public institutions, the more worrying such a dependence becomes. In sectors such as finance, energy, or healthcare, Member States need reliable access to compute, rather than relying merely on the trust and goodwill of foreign actors. The most straightforward way to achieve this is to ensure that a large fraction of European AI demand is served through data centres on EU soil, which are ultimately under European countries' physical control.

  Of an already scarce resource, the EU has very little. Researchers estimate that Europe as a whole – including non-EU countries such as the UK and Norway – will have around 2 GW of AI compute by the end of the year, constituting a less than 5% share of 45 GW globally by

that time.[26] This share is likely to stay roughly flat over the next five years if current trends continue, with approximately 21 GW of AI compute in Europe and ca. 370 GW globally in 2031. In the default scenario, and despite the AI Continent Action Plan's stated goal to triple the EU's data centre capacity over the next years, the EU's share of global AI compute would therefore remain significantly below its 18% share of global GDP.

The EU should aim for a compute share in proportion to its share of the world economy. To fight its acute compute scarcity and control a meaningful share of the world's AI infrastructure, EU Member States should collectively aim to host at least 15% of global AI compute by the end of the decade – roughly proportional to the EU's 18% share of global GDP. In line with the compute forecast of the Europe 2031 scenario paper, this strategy assumes a global stock of AI compute of 300 GW by the end of 2030. This is roughly in line with estimates recently published by leading AI hardware forecasters, as Table 3 below shows.[27]

SOURCE	ESTIMATE OF GLOBAL AI COMPUTE	TIME
Europe 2031	300 GW	2030 (END OF YEAR)
RAND	327 GW	2029 (END OF YEAR)
AI 2040	286 GW[28]	2030 (END OF YEAR)
SemiAnalysis	239–380 GW[29]	2030

TABLE 3 A COMPARISON OF RECENT EXPERT FORECASTS OF GLOBAL AI COMPUTE BY THE END OF THE DECADE.

This strategy's estimate of 300 GW of global AI compute (total facility power) by 2030 already prices in a slowdown of annual AI compute growth, mainly due to supply chain bottlenecks, from over 3x per year historically to around 1.25x per year by 2031. Similarly, it prices in that AI hardware will become more efficient over time. In a world where AI systems have transformative impact, however, greater hardware efficiency would likely have two main effects: (1) As a given level of AI capability becomes cheaper, customers will demand more of it. (2) AI companies can turn a given level of investment into greater AI capabilities. Both

26 All GW estimates of European and global AI compute in this deep dive refer to total facility power.

27 Some AI experts have put forward more aggressive estimates, such as Leopold Aschenbrenner in his article 'Situational Awareness'. His back-of-the-envelope calculation projects that by 2030, the world's AI data centre capacity will be equivalent to total US power generation – around 4,250 TWh according to Aschenbrenner, likely referring to 2023, equivalent to an average load of 485 GW. Apart from AI forecasting experts, consultancies have also estimated AI compute growth, with McKinsey expecting 90 GW by 2030 and Goldman Sachs 24 GW by 2030 (as presented by RAND). The 24 GW estimate by Goldman has likely been overtaken by reality already, with Epoch AI estimating approximately 30 GW of AI compute at the end of 2025. As RAND notes, conservative estimates from Goldman and McKinsey only assume a "brief period of exponential growth" in AI. This is contrary to the premises of this strategy, which assumes that AI capabilities and the required inputs will continue to grow until AI has transformative effects.

28 The AI 2040 authors state their estimate of global AI compute in H100 equivalents (H100e), a different measure of AI data centre capacity than power draw, measured in GW. For converting the H100e estimate into GW, this strategy relies on AI 2040's assumptions about chip efficiency growth.

29 It is unclear whether the SemiAnalysis estimate refers to IT load or total facility power. If it is the former, then for comparability it should be multiplied by a 1.1x PUE factor to yield 263–418 GW of total facility power.

factors incentivise the AI industry to grow the global stock of AI compute far beyond current levels, even if the growth rate slows down over time. If AI compute growth slows down less than expected, a global supply >400 GW by 2030 is not inconceivable.

Assuming a total AI compute stock of 300 GW by 2030, a steady ramp-up of compute in the EU could look like this:

YEAR	EU AI COMPUTE	GLOBAL AI COMPUTE	EU SHARE
2026	1.8 GW	45 GW	4%
2027	4.9 GW	82 GW	6%
2028	12 GW	147 GW	8.2%
2029	24 GW	207 GW	11.6%
2030	45 GW	300 GW	15%

TABLE 4 A POTENTIAL RAMP-UP OF AI COMPUTE IN THE EU TO REACH THE 15% TARGET BY 2030.

Reaching a goal of 45 GW by 2030 would require one of the most ambitious infrastructural programmes in the history of post-war Europe. Operating AI data centres with 45 GW capacity at full utilisation would increase the EU's electricity consumption by 15.8% (394/2,492 TWh). A recent study found that this is feasible, subject to enough political will: In a "crisis mobilisation" scenario, the EU's major power markets alone – France, Germany, Iberia, and the Nordics (ex-Norway) – could connect 35.6 GW of AI compute to their grids by 2030.[30] The remaining 9.4 GW could plausibly be provided by other EU countries, and potentially by relying on on-site power generation as a bridge solution while the European grid is being expanded.

The EU and its Member States should aim for this ambitious but feasible goal because transformative AI – understood here as AI systems that would materially affect the EU's long-term role in the distribution of economic and geopolitical power – demands industrial-scale efforts that far exceed the pace of infrastructure expansion in normal times. The required effort could be reduced somewhat if the EU plans part of the compute expansion together with allied democracies such as Norway, Iceland, Canada, or Australia. Even in that case, a highly ambitious effort will still be required on EU soil. This will require difficult trade-offs and heavily rethinking established norms and procedures, but without it, EU countries are at the risk of losing their sovereignty and economic standing in a world where a region's share of compute is a central determinant of its power and influence. The EU and its Member States should act quickly and with resolve to guard against this scenario, because the time to lock in geopolitical leverage through domestic compute on EU soil is now, rather than in the mid-2030s when AI companies could start to build data centres in space.

30 In a crisis mobilisation scenario, policymakers would unlock additional AI compute capacity mainly via flexible grid connections for AI data centres and shifting electricity demand away from use cases such as hydrogen production towards AI data centres. In this scenario, the markets named above plus the UK and Norway could together connect 47.6 GW by 2030.

This strategy's assessment is that, if the EU does not find creative ways to reach 15% of global AI compute by 2030, it will have much less leverage to secure access to frontier technology than in a world where it succeeds in this generational buildout. The situation is especially risky in light of a likely future compute crunch: The less physical leverage they have over computing infrastructure, the more Member States will have to fight for compute access by making potentially costly concessions elsewhere. Below, this deep dive outlines a difficult, but feasible, path towards realising an ambitious compute target. This strategy recommends this path not because it is politically easy to implement, but because Europe otherwise risks losing access to the strategic resource on which transformative AI depends.

The EU and its Member States can achieve 45 GW of AI compute by 2030

Who: The large majority of AI compute buildout in the EU should be done by the private sector. 45 GW of AI compute (total facility power) requires an enormous amount of investment. At an estimated capital expenditure of approximately $34 billion for a 1 GW AI data centre, the total cost of building these data centres would amount to approximately €1.3 trillion.[31] Only the private sector is able to mobilise capital at this order of magnitude. Goldman Sachs estimates that $7.6 trillion will be spent on AI infrastructure over the next six years.[32] This year, the US hyperscalers alone are on track to spend over $700 billion on capital expenditure, the majority of which is for AI data centres. Individual investors and consortia have similarly announced plans to channel large amounts of capital into AI infrastructure, for example the $500 billion Stargate Project, Saudi Arabia's $100 billion Project Transcendence, and Brookfield's $100 billion AI Infrastructure Program. The question is not whether this capital exists, but whether a significant fraction of it chooses the EU.

Moreover, Europe itself controls significant private capital that could be leveraged for large-scale AI compute. Asset managers in Europe as a whole oversee around €23 trillion of institutional capital;[33] redirecting capital equivalent to roughly 1.9% of this base annually over three years would be enough to finance the 45 GW ambition. That said, the majority of the buildout would have to be financed by US AI companies, as only the US hyperscalers and frontier developers currently generate enough demand to justify a double-digit GW buildout (that is, unless Europe attempted its own frontier project – see What would be required for a successful European frontier AI project?). As proposed above, US data centres on EU soil would serve as physical leverage to secure access to leading US models. This is why the alternative to 45 GW of AI compute majority-built by US companies is *not* 45 GW of European-only compute, but simply much less AI compute in Europe overall.

To further illustrate the scale, note that the proposed investment is less than what some other countries outside the US and China are spending on AI compute, relative to their size: For example, the Republic of Korea has announced private-led investment that is roughly 3x as much relative to its GDP.[34]

---

31 Epoch AI estimates that 1 GW of IT load costs slightly below $38 billion for an AI data centre. Assuming a PUE (power usage effectiveness) of 1.1, this translates into $34.4 billion per 1 GW of total facility power. At current exchange rates, the total investment of $1.55 trillion would correspond to approximately €1.3 trillion.

32 At $38 billion per GW, this implies a total AI data centre capacity of 200 GW by 2031. This is likely an underestimate; expecting 300–400 GW of global AI compute by 2031 is more plausible.

33 This figure includes non-EU countries such as the UK and Switzerland, which owned almost one half of Europe's assets under management in 2024.

34 Korea's 550 trillion won investment by 2029 is a 19.7% share of its 2026 nominal GDP of 2,788 trillion won, or approximately 6.6% per year when averaged over 2027–29. The proposed EU investment of $1.55 trillion by 2030 would be a 6.7% share of its 2026 nominal GDP of $23.03 trillion, or approximately 2.2% per year when averaged over 2027–29.

While financing a 45 GW compute buildout should be mostly private sector-led, a public contribution from the EU and Member States of €100 billion to the €1.3 trillion overall investment would help to mobilise private funds and increase the share of European-only AI compute, owned and operated by European organisations on EU soil. These data centres could be used, for example, for sensitive workloads, public good use cases, or boosting European AI researchers and startups. For the most sensitive use cases, such as AI workloads in national security or defence, Member States should consider direct equity investment into AI data centres alongside private sector partners to ensure that the most critical AI infrastructure is publicly co-owned. Moreover, if the compute needs of a European developer of fast-follower models (see Objective 1.1) or emerging frontier company increased, it could readily tap into this European-only compute pool, rather than compete for scarce cloud compute with foreign frontier companies that locked in long-term contracts in advance. If, however, Europe owned more compute than domestic companies and researchers needed, it could rent out its compute to foreign customers at standard rates, which are likely to remain lucrative as compute remains scarce in a world with transformative AI. Raising €100 billion for AI compute through public funding is ambitious, but possible if different buckets are combined, with (for example) shares from the EU's proposed Digital Leadership (€48.5 billion) and cohesion (€405 billion) budgets or the InvestEU Fund (€30 billion). A potential IPCEI for Frontier AI Compute would help to unlock corresponding amounts from Member States.

Where: Suitable sites for building large AI data centres quickly are limited, as a number of conditions all have to be in place: dozens or even hundreds of hectares of land, high-voltage grid connections (or options for on-site power generation), cooling infrastructure, fibre connections, and a favourable regulatory environment.

Most EU countries will have some sites that fulfil these conditions, and building 45 GW of compute will generally require a collective effort. However, some countries are much more favourable places for building data centres than others, for example, due to cheap or abundant green energy, a cold climate that eases cooling demands, or industrial brownfields with existing grid connections that can be repurposed for AI data centres. Therefore, there is a strong case for prioritising regions with the most favourable conditions, turning them into European compute centres whose benefits and burdens are shared across all EU countries. The table below offers a preliminary overview of what these regions could be:

COUNTRY	SUITABLE REGIONS (EXAMPLES)	ADVANTAGES
Denmark	Esbjerg (West), Varde, Ølgod	offshore wind surplus, cool climate
Finland	Lappeenranta, Mäntsälä, Kajaani, Myllykoski	abundant nuclear and hydropower, cool climate, strong tenant base
France	Hauts-de-France, Grand Est	nuclear surplus, industrial sites
Germany	Rhineland, Lusatia, Northern Germany	industrial sites, high grid reliability, wind energy in the north
Poland	Silesia	industrial sites
Portugal	Sines	planned GW-scale projects with sea water cooling, abundant solar/wind
Romania	Cernavodă, Doicești	nuclear surplus, industrial sites

Spain	Aragón, Castilla-La Mancha	abundant solar/wind, experience through existing hyperscaler projects
Sweden	Norrland, Falun, Strängnäs	abundant hydropower, cool climate

TABLE 5 ATTRACTIVE AI DATA CENTRE REGIONS IN THE EU. SOURCE: OWN FIGURE.

In order to leverage those assets, the EU should seek a mechanism whereby all countries contribute to ambitious data centre buildout in the most favourable regions and receive fair access to the resulting compute in return.

How: Only the private sector can realistically fund a compute buildout on the order of 45 GW, with public funding only playing a limited enabling role. Therefore, the role of the EU and its Member States lies mainly in two areas:

1. Making the EU one of the world's most attractive regions for private companies to build AI data centres.

2. Securing the political success conditions for this buildout, both vis-à-vis (i) the US as the most powerful actor in the global chip supply chain and (ii) local communities and the broader population as the groups directly or indirectly affected by new data centres.

For concrete recommendations on how to achieve this, see Immediate Objective 3 in Part B of this strategy.

Robotics deep dive

Robotics will likely become an important economic and strategic driver as AI becomes transformative, but Europe will play a leading role only if it substantially strengthens its position. Europe has historically been strong in manufacturing and deploying industrial robots, particularly precise, high-quality machines that performed narrow, repetitive tasks in heavily engineered environments. However, if current trends in AI continue, the value of manufacturing traditional industrial robots may decrease relative to the production of AI-integrated robots, which can adapt to changes, understand natural language, reason about unfamiliar tasks, and transfer what they learn across environments and robot embodiments. Europe retains deep hardware expertise and manufacturing capacity, but those advantages will not necessarily translate into leadership in this new paradigm: Europe lags in AI integration, does not have assured access to the frontier AI models underpinning the most advanced robotic systems, has limited manufacturing capacity in other embodiments, and is not well positioned to manufacture at scale. Remaining a leading robotics producer will therefore require Europe to reorient its robotics capacity towards AI-integrated systems rather than relying on traditional industrial robotics to provide an enduring edge.

A note on definitions
In this strategy, 'frontier AI models' refers to the most capable, general-purpose AI models, such as the models behind popular AI applications like Claude or ChatGPT. 'AI-integrated robots' refers to robotic systems that use models trained for functions such as perception, language understanding, planning, or control, and which generally enable greater adaptation to changes in tasks or environments. 'Frontier robotics models' refers to the most capable, relatively general-purpose models capable of operating a robot – systems that can accept natural-language instructions, reason about unfamiliar tasks, perform multi-step tasks, or

transfer across environments and robot embodiments. Frontier robotics models can include advanced vision-language-action models specifically trained for robotic control, like Gemini Robotics 2. This can also include frontier AI models, which have demonstrated the ability to control robots in some tasks. Robots can take various shapes. 'Form factor' refers to the broad category of robot, including categories like fixed-base arms, autonomous mobile robots, mobile manipulators, humanoids, and drones, while 'embodiment' refers to the particular body or hardware on which a model operates (e.g. a specific company's humanoid). This strategy uses 'traditional industrial robots' to principally refer to fixed-base industrial arms used in industrial settings.

Europe has historically been strong in robotics

Historically, the robotics industry was built around precision and repeatability rather than adaptability. Industrial arms, which have been used in industrial processes like automotive manufacturing since the 1960s, were deployed before there was any form of 'intelligent' operating system for robots. Instead, the movements of industrial arms were programmed to perform the same motions repeatedly. Because the sensing and corrective capacities of these robots were limited, both production lines and machines required precision: production lines were heavily engineered to ensure that products moving down a production line appeared before an industrial robot in the same position each time, and industrial robots were optimised for 'repeatability' – the ability to perform the same motions over and over with as little variance as possible. Accordingly, the most competitive manufacturers of industrial robots were typically those who could produce robots with the greatest precision. Doing so requires high-quality hardware and specialised manufacturing processes – often themselves involving industrial robots. Advances in computer vision and motion-control software have made industrial robots better at following programmed trajectories accurately and making small corrections in response to variation. Even so, most remain limited to narrow, repeatable tasks.

Traditional industrial robots are costly to build and deploy, restricting their use primarily to highly standardised manufacturing. Because of the quality demands, industrial robots have typically been expensive investments. In addition to the cost of the robot itself, deploying traditional industrial robots has typically involved re-engineering production lines to operate with the precision required by industrial robots. Doing so is sufficiently capital-intensive that industrial robots were disproportionately deployed by large manufacturers, like automotive companies. The constraints of industrial robots – the lack of significant sensing or corrective capacity and the cost of deploying them – have meant that industrial robots are deployed in only a subset of industrial processes: industrial applications with high-throughput, low-variance production processes, especially in large manufacturers.

Europe and Japan dominate the traditional industrial robotics paradigm. Seventeen of the top 20 industrial robot manufacturers are Japanese or European, with seven headquartered in Europe (see Table 6). European robot manufacturers account for roughly one-third of estimated global industrial robot revenue. The industry is largely incumbent-dominated, with many of the largest companies having entered the robotics industry in the 1970s or 80s. For example, ABB (headquartered in Switzerland) and KUKA (headquartered in Germany but owned by China's Midea Group), the two largest industrial robot manufacturers headquartered in Europe, both began manufacturing industrial robots in the early to mid 1970s. This is not to say that all European industrial robotics firms are this old – Universal Robots, for example, released its first robot in 2008 – but the majority of Europe's robotics revenue is concentrated in incumbent firms. Europe is also a major deployer of industrial robots: Western Europe has a higher industrial robot density – a measure of industrial robot deployment relative to the number of manufacturing workers in a region – than North America or Asia.

RANK	COMPANY NAME	COUNTRY OF MANUFACTURING HEADQUARTERS	COUNTRY OF OWNERSHIP	ROBOTICS REVENUE (USD)	REVENUE SHARE	CUMULATIVE REVENUE SHARE
1	FANUC	Japan	Japan	$2,486M	16.7%	16.7%
2	ABB	Switzerland	Switzerland*	$2,331M	15.7%	32.5%
3	Yaskawa Electric (Motoman)	Japan	Japan	$1,622M	10.9%	43.4%
4	KUKA	Germany	China	$1,085M	7.3%	50.7%
5	Mitsubishi Electric (Factory Automation)	Japan	Japan	$873M	5.9%	56.6%
6	Kawasaki Heavy Industries (Robotics)	Japan	Japan	$610M	4.1%	60.7%
7	Stäubli Robotics	France	Switzerland	$515M	3.5%	64.2%
8	Brooks Automation	US	US	$372M	2.5%	66.7%
9	Dürr	Germany	Germany	$360M	2.4%	69.1%
10	Universal Robots	Denmark	US	$308M	2.1%	71.2%
11	Rorze	Japan	Japan	$285M	1.9%	73.1%
12	Nidec	Japan	Japan	$263M	1.8%	74.9%
13	Estun Automation	China	China	$255M	1.7%	76.6%
14	OTC Daihen	Japan	Japan	$215M	1.4%	78.0%
15	Franka Emika / Agile Robots	Germany	Germany	$215M	1.4%	79.5%
16	HD Hyundai Robotics	South Korea	South Korea	$189M	1.3%	80.7%
17	Yamaha Motor (Robotics)	Japan	Japan	$183M	1.2%	82.0%
18	Nachi-Fujikoshi	Japan	Japan	$166M	1.1%	83.1%
19	Neura Robotics	Germany	Germany	$157M	1.1%	84.1%
20	Epson Robots	Japan	Japan	$154M	1.0%	85.2%

TABLE 6

Highest-earning industrial robot manufacturers. European-headquartered manufacturers are highlighted in blue. Source: FAI 2026. *SoftBank Group (Japan) has agreed to acquire ABB Robotics; closing is expected in the second half of 2026, subject to approvals.

AI integration into robotics reduces Europe's historical advantages in industrial robotics

AI integration could have a transformative effect on robotics. AI integration into industrial robots has traditionally been limited. Some industrial robots are equipped with simple machine learning-based systems, like small computer vision models configured to recognise a defined set of objects. These are typically heavily tailored to the specific task. Frontier robotics models, by contrast, are large AI models – often referred to as vision-language-action models (VLAs) – trained on much larger and more heterogeneous datasets. For example, Physical Intelligence's π_0 model was pretrained on 10,000 hours of dexterous-manipulation data spanning 68 tasks and seven robot configurations, in addition to large public datasets that pool more than one million robot trajectories. The training process for frontier robotics models is increasingly intensive as the models increase in size. OpenVLA – a seven-billion-parameter robot model trained on 970,000 real-world trajectories – required 64 Nvidia A100 GPUs running for 15 days; Generalist AI's models now exceed 10 billion parameters, and Generalist reports that its most recent model trained continuously for eight months. These robot-specific models often make use of existing large vision-language models (VLMs) for natural language, image processing, and reasoning ability. For example, Figure pairs a VLM with simulation and teleoperation data to train a robot able to process language commands and analyse scenes, and Gemini Robotics uses its Gemini AI models to enable robots to reason and problem-solve before passing off control for motor commands to a VLA. Additionally, preliminary evidence suggests that sufficiently advanced frontier AI models, like Claude, could be used for direct robotic control.

Transformative AI could dramatically accelerate robotics R&D. In addition to direct applications for robotic control, transformative AI could be used as a tool to speed up the development of robotics software and hardware. For example, frontier AI models could improve the success of simulation-based training, which is currently not sufficiently realistic to broadly enable successful transfer of simulation-trained skills to the real world without augmenting with real-world training data. Doing so would reduce the need for hard-to-scale real-world data, significantly accelerating progress in training advanced robotics models.

AI integration into robots could give robots new capabilities that enable them to succeed in a wider range of tasks. Many of the most advanced capabilities described below come from company demonstrations, and so should be treated as preliminary examples of advancing capabilities.

- **Flexibility**. AI enables robots to adapt to changes in a task, without requiring production processes to be standardised and invariable. For example, an AI-integrated robot using vision to pick up objects can navigate to them even as they move and appear in different positions, rather than assuming specific coordinates, and figure out how to grasp objects of different shapes. Simple versions of these already exist in industrial robots. For example, ABB's Robotic Item Picker system uses a suction gripper to pick up stationary objects of different sizes and shapes. Frontier robotic models go further, demonstrating the ability to fold clothes they have never seen before, do task sequences they were not trained on based just on verbal instructions, and adapt to new tools and environments. Frontier models are also able to transfer skills learned in one embodiment to a different embodiment (for example, one with longer arms or different hands) – this cross-embodiment transfer enables deployment in a wider range of embodiments without requiring further training. This enables the deployment of robots in cases where they would previously have been uneconomical, including low-volume or high-variance production processes.

- **Hard-to-program manipulation.** Some motions are difficult to reduce to a fixed trajectory because the object changes shape or moves unpredictably. Tasks involving these motions, including manipulating fabric, routing wire harnesses, and handling food, have historically

been difficult to automate with industrial robots. Frontier demonstrations now include folding clothing, handling fruit, doing origami and other deformable-object tasks.

- **Lower dependence on mechanical precision.** Improved perception and correction reduce the required mechanical repeatability in certain contexts. For example, if a robotic arm with embodied intelligence can repeatedly evaluate and replan its trajectory, it does not need to be as mechanically precise in its movements. This does not make hardware quality obsolete – it would be difficult for intelligence to substitute for core features like payload tolerance – but it may reduce the relative value of mechanical precision.

- **Ease of deployment.** Flexible robots reduce the need for reorganising a production line to ensure the robot sees low-variance processes. Instead, more intelligent robots could handle higher-variance processes without having to reconfigure the production processes. For example, frontier demonstrations show robots sorting packages of varying sizes, shapes, and placements in a warehouse, and successfully placing items in their intended containers even as the containers move around. Additionally, more intelligent robots are easier to teach new tasks: frontier robotics models are able to learn to do a task from in-context information (i.e. without having to update their model weights), and can do so from a single example, even just watching a human do a task (see, for example, Generalist AI or Skild AI).

- **Reasoning and intelligence.** Many physical tasks are difficult to reduce to pre-planned motions because they require understanding a goal, planning subtasks, adapting to changing conditions, solving problems, or language- or vision-based judgement. Sufficient intelligence enables robots to complete more complex tasks than they would have otherwise. Frontier robotics models exhibit some reasoning abilities, like stringing together multiple subtasks to perform long-horizon tasks, understanding and acting on natural language commands, and even searching the internet for context. Further advances, including greater integration of frontier AI models into robot control systems, could increase the complexity of automatable tasks. For example, sufficient intelligence would enable tasks like reading and understanding a natural-language assembly manual, doing construction work in new terrain, repairing machinery with unknown problems, exploring novel environments, and communicating with other robots or humans to complete a complex task.

The above capabilities could allow robots to perform many more tasks across the economy. If AI becomes transformative, it is possible that sufficiently capable systems will eventually automate most, if not all, physical work now performed by humans. The potential economic value of developing frontier robots could therefore be very large.

European leadership in robotics is far from guaranteed

There are several reasons why Europe's historical strengths in industrial robotics may not translate into comparable gains in an AI-driven robotics paradigm:

European robotics manufacturers are mostly not building custom AI-integrated robots. European industrial-robot incumbents use AI mainly for narrow applications. For example, ABB's Robotic Item Picker, as discussed above, applies AI-based vision for defined pick-and-place problems. Certain European companies are partnering with non-European companies to allow their robots advanced AI integration: for example, ABB is integrating Nvidia simulation and foundation-model tools into its platform to allow users to incorporate more advanced AI systems, and Universal Robots' AI Accelerator incorporates Nvidia hardware and software into its robots for developer use. These efforts may make existing products more capable. But there appear to be few, if any, major European industrial-robot incumbents that are developing a proprietary, general-purpose robot model able to follow natural-language instructions and operate across varied tasks, environments, and embodiments. The leading incumbents are therefore positioned mainly as hardware and integration providers, rather

than developers of core models. While new European robotics companies have emerged in recent years that are oriented towards frontier AI integration, such as Agile Robots or Neura Robotics, the ecosystem remains small and few European companies have publicly demonstrated robots that use frontier AI models for perception, planning, or control (and fewer still are developing the underlying models themselves). Additionally, established manufacturers like ABB and KUKA are mostly not developing frontier VLAs. Europe's frontier-robotics sector is therefore emerging for the most part outside the incumbents that underpin its existing robotics strength. These startups are also much smaller and less well capitalised than leading US competitors, reflecting Europe's broader difficulty in financing and scaling new technology companies (see Objective 2.1).

Europe's hardware lead is concentrated in traditional industrial robots. Europe is most dominant in fixed-base industrial robots. While industrial robots will almost certainly continue to be an important form factor, they are only one of many poised to benefit from AI integration into robotics, which also include autonomous mobile robots, mobile manipulators, humanoids, and drones, among others. Europe's manufacturing presence in these other sectors is smaller: European shipments accounted for about 11% of autonomous mobile robots for intralogistics in 2024, and a negligible share of humanoids in 2025. While the hardware supply chains and manufacturing capacity of industrial arms are well-suited for some level of transfer to other form factors, Europe is not, by default, poised to manufacture most of the form factors that could benefit from the returns to AI integration.

Europe's continued dominance in traditional industrial robotics is also not a given. While traditional industrial robots will likely remain valuable, China is rapidly encroaching on incumbent market share. China's share of global industrial robot production has increased from 9% in 2016 to over 30% in 2024, and China has significantly reduced its import reliance for industrial robots. China has also made significant inroads as an upstream hardware supplier. For example, Chinese manufacturer Leaderdrive began producing an increasingly large share of Universal Robots' strainwave reducers in the late 2010s, accounting for around 30% of Universal Robots' reducers in 2019 (though recent figures are unclear). Additionally, several European-headquartered industrial robotics companies have lost European ownership: KUKA is owned by China's Midea, Universal Robots by US-based Teradyne, and ABB has agreed to sell its robotics division to Japan's SoftBank. Traditional industrial robotics is not a guaranteed fallback if Europe fails to lead in AI-enabled robotics.

Europe's dominance in high-precision, high-quality hardware might be less important than the ability to manufacture at scale. As the scope of tasks that robotic systems are able to automate broadens beyond narrow, high-precision industrial tasks, the ability to swiftly increase production may matter more than the ability to manufacture exceptionally high-quality robots – where Europe's dominance has historically been. While Europe remains strong in the final assembly of traditional robots, it is more reliant on foreign suppliers for raw and processed materials than China is. More broadly, China produced 35% of global manufacturing output, compared with 13% for the euro area, and more than the next ten largest manufacturing countries combined. China's manufacturing base has also expanded considerably faster: in the first quarter of 2026, its output was 39% above its 2020 level and 5.6% higher than a year earlier, while European output was 12% above its 2020 level and 1.2% lower than a year earlier. In practice, China's dense hardware ecosystems and broader manufacturing capacity appear to enable Chinese robotics companies to develop and manufacture robots faster and more cheaply.

Europe lags in frontier models. Some leading robotics companies, such as Physical Intelligence in the US, focus on developing frontier robotics models rather than building their own embodiments. Physical Intelligence, along with other leading robot-model companies such as Nvidia, Figure and Skild AI, are mostly American. While Europe does have a few developers, such as Agile Robots, building robotics foundation models, no European company has emerged as a clear peer to the leading US companies in the demonstrated generality and capabilities of its models. Europe also lags behind on general-purpose AI models, which increasingly provide robots with perception, language understanding, and reasoning. The 2026 Stanford AI Index attributes 59 notable AI models released in 2025 to the United States,

35 to China, and two to Europe. Even if a different architecture replaced LLMs as the main paradigm behind general-purpose models in the future, this would not necessarily favour Europe: while it could create opportunities for new entrants, developing successful robotics models under any paradigm would still require frontier AI talent, compute, data, and capital. These resources are currently concentrated in foreign AI companies, and a greater share would end up in Europe only through an ambitious attempt to attract them.

Europe lacks reliable access to frontier models built elsewhere. While European companies could use open-source models, the most capable frontier models are closed-source, including an increasing share of frontier robotics models. For example, many frontier robotics model developers like Figure, 1X Technologies, Generalist AI, and Skild AI do not release their models. Access to frontier models from foreign developers is not guaranteed (see Objective 1.1). Ongoing access to open-source models is not guaranteed, either; for example, Physical Intelligence – a frontier research lab which open-sourced its first generations of models – has not released its two most recent models. As robots take on increasingly complex tasks, differences in reasoning and planning ability will increasingly determine what they can do reliably and how economical they are to deploy; if European robotics companies must rely on less capable general-purpose AI or robotics models, they could become less competitive than companies with access to frontier models. Gaps in access to the best models could also compound over time: Frontier general-purpose AI models could accelerate robotics R&D (e.g. through improved sim-to-real transfer), while frontier robotics models could enable more capable deployments, generating real-world interaction data that can be used to improve subsequent models.

Europe's industrial data will not necessarily provide strong advantages in frontier robotics. European factories generate extensive production telemetry, quality records and maintenance data. In an effort to take advantage of the data generated by Europe's industrial base, the EU has launched several initiatives to facilitate the sharing of manufacturing data, including data generated by industrial robots and other machinery. Training frontier robotics models, however, often needs synchronised camera or tactile observations, robot states, actions, task labels, and outcomes. Legacy arms were not necessarily designed to record or share all that data. For example, many industrial arms do not collect synchronised video and joint-level motion data. Data is also more valuable for training if it has greater variance, leading frontier robotics labs like Figure to create custom data collection platforms to build more diverse training sets. By contrast, industrial robots often collect data from repetitive motions in environments with little variation, like industrial production lines. Additionally, to be deployed, robots usually must have very high success rates in their tasks, making it difficult to use their data to train on the difference between successful and unsuccessful operation. While data collected from deployed robots may still be valuable, the high repetition from industrial robots and the low failure rate may limit the usefulness of the data for training generalised physical skills.

Going forward: adapting Europe's robotics capacity

Europe is not starting from zero. European companies remain globally competitive in industrial robotics, with significant hardware production capacity and deep expertise in high-quality industrial robot manufacturing processes. Additionally, industrial robotics capacity is, to some extent, general purpose, with many areas of hardware and manufacturing overlapping. For example, humanoids, like industrial robots, need joint-level actuators and operable end-effectors, and mobile robots use similar types of motors to industrial arms. Under any future, robots will require reliable embodiments, and Europe has many of the supply chains and production processes to make the relevant components. If Europe were to reorient towards the frontier of robotics – AI integration into varied autonomous form factors – it has deep resources on which to draw. Doing so could give Europe a shot at retaining its strength in robotics.

Europe should use its industries as a testbed for further industrial automation. Western Europe already has a more automated industrial sector (as measured by the ratio of industrial robots to manufacturing workers) than Asia or North America. This gives Europe a dense and diverse base of factories and production processes that can serve as testing grounds for early robotics deployment. The EU and Member States could co-fund production pilots that pair European robotics and AI developers with manufacturers and support the deployment of advanced robotic systems. Operating in real factories would give European robotics companies access to diverse task and failure data, and allow them to demonstrate their aptitude for greater commercialisation.

Data-sharing efforts should focus on the most valuable sources of data for robotics. As discussed above, Europe's large industrial base may not automatically produce a data advantage: its data is often narrow and repetitive and may not include the necessary task annotations. Developing advanced robotics models, by contrast, requires data that captures varied tasks, environments, and outcomes, which might involve specialised data collection (e.g. collection that explicitly labels whether a robot was successful in a task). European robotics initiatives should therefore prioritise creating systems to share the types of data that are most valuable for advanced model development, while protecting commercially sensitive information.

Europe should fund robotics R&D and commercialisation. Robotics is a capital-intensive industry, and Europe's broader difficulties in financing and scaling new technology companies (see Objective 2.1) will impede Europe's success in frontier robotics. The EU and Member States should consider expanding funding for frontier robotics research and commercialisation efforts. This could include support for research groups and startups, as well as incumbent robotics companies or hardware manufacturers attempting to transition into frontier robotics. Prioritising robotics companies' access to compute could also be a useful lever for accelerating European development of frontier robotics models.

Europe could benefit from well-designed partnerships with foreign companies. European robot manufacturers are unlikely to develop their own frontier models in the near term, so partnerships with foreign model developers could accelerate the integration of advanced AI into European hardware. Such partnerships should be designed in a way that protects European intellectual property, for example by using privacy-enhancing technologies when sharing sensitive proprietary data.

Standardised regulations should facilitate robot deployments across the EU. Differing regulatory standards by jurisdiction could slow down real-world testing and deployment of advanced robotic systems. The Commission should ensure that regulations guiding the production and deployment of advanced robots – including those, like humanoids, with high levels of autonomy – are clear and harmonised across the EU.

Legal notice

Published by
This strategy is a project of individual experts who contributed in a personal capacity; institutional affiliations are given for identification purposes only. For questions related to this publication, please contact the KIRA Center:

Zentrum für KI-Risiken & -Auswirkungen gUG (haftungsbeschränkt) (KIRA Center)
Prenzlauer Allee 186, 10405 Berlin, Germany
Registered at Amtsgericht Charlottenburg (Berlin), HRB 251628 B
Managing Director (Geschäftsführer): Daniel Privitera

Responsible under German press law (V.i.S.d.P.): Daniel Privitera, address as above.

Version
Version 1.1, 15 September 2026: corrected edition; typographical and consistency fixes, no changes to recommendations. The strategy may be updated to correct errors; the current version is available at https://transformative-ai.eu.

How to cite
Schnitzer, M., Privitera, D., Bucknall, B., Dreksler, N., Fox, P., Leicht, A., McGlynn, C., Rignell, M., Stehr, F., Aghion, P., Winter, C., Bakker, M., Ischinger, W., Vestager, M., Mądry, A., Fuest, C., Kwiatkowska, M., Varadkar, L., Osika, A., Acemoglu, D., Bengio, Y., Angeli, M., Blom, O. E., Doan, R., Martens, B., Mindermann, S., Mökander, J., van Putten, J.-W., Arcesati, R., Benner, T., Burga, T., Dhaliwal, J., Freund, J., Garicano, P., Gillett, E., Goutbeek, J., Greenblatt, R., Growiec, J., Hacker, P., Heintz, F., Hobbhahn, M., Juijn, D., Levie, A., Moës, N., Praas, R., Prins, J., Radu, R., Riekeles, G., Schaal, J., Schenk, T., Sevilla, J., Shamir, A., Sheikh, H., Teutloff, O., Touzet, C., Uszkoreit, H., Uuk, R., Velliet, M., von der Wippel, P., Welling, M., Willemyns, J. (2026): A *Transformative AI Strategy for Europe*. transformative-ai.eu.

Credits
Design: Paul Jochum, David Reitenbach

A Transformative AI Strategy for Europe
Read online: https://transformative-ai.eu
Published in September 2026 by the KIRA Center.