# Agentic AI: State of play

Thematic paper

JULY 2026

*WEB4HUB: 'A SPACE FOR THE METAVERSE –VIRTUAL WORLD AND THE TRANSITION TO WEB 4.0'*

European Commission

PPMi Part of the Verian Group

WEB4HUB

**Authors**: Rūta Gabaliņa, Oleksandra Yevdokymova, Aida Zepcan.

**Acknowledgements**: Antoine-Alexandre Andre, Maryna Butym, Egidijus Barcevičius, Pierre Marro, Hubert Romaniec, Fabrizio Sestini, Dimitrios Thomas and others who contributed through interviews, workshops or reviews of the present paper.

## CONTENTS



## LIST OF FIGURES



## LIST OF TABLES

## LIST OF BOXES

# Introduction

This thematic paper examines the rapid development of **agentic artificial intelligence** (AI) and how it impacts various sectors and actors. Looking ahead, it explores the socio-economic forces driving and resulting from this transition, as well as the adaptation of policy and governance frameworks.

The paper is situated within the broader context of the **transition to Web 4.0**, which the European Commission defines as "*using advanced artificial and ambient intelligence, the internet of things, trusted blockchain transactions, virtual worlds and XR capabilities, digital and real objects and environments that are fully integrated and communicate with each other, enabling truly intuitive, immersive experiences, seamlessly blending the physical and digital worlds*"[1]. As elaborated in an earlier project paper[2], several technology clusters — including AI, internet of things, extended reality, advanced communication networks (5G/6G), blockchain, decentralised identity, privacy-enhancing technologies, cybersecurity, brain-computer interfaces and quantum technologies — serve as the core building blocks of Web 4.0.

Among these, AI — and specifically agentic AI — is emerging as a particularly **consequential driver of change**. Unlike technology clusters that largely supply capabilities like connectivity, agentic AI redefines *who* — or *what* — acts on the web. The shift from reactive generative models to autonomous systems capable of processing signals from their digital environment, planning across multiple steps, and transacting on behalf of users, reframes the internet from a repository of human-readable information into an infrastructure increasingly populated by non-human entities. This transition affects how economic value is created and captured, how public services are designed and delivered, how digital infrastructure is planned and provisioned, and how accountability is allocated when decisions are made at machine speed. It also cuts across other Web 4.0 building blocks — placing new demands on identity frameworks, communication protocols, decentralised architectures, security, and compute capacity.

This paper examines these dynamics along two complementary axes. First, it looks at **socio-economic drivers** accelerating adoption and deployment of AI agents. These drivers pull in different directions. On one side is the pursuit of societal and planetary good through innovative applications in healthcare, sustainability, public services and science. On the other side is the commercial logic reshaping enterprise software, which currently mobilises the bulk of capital and shapes the technology's trajectory. **These drivers are not mutually exclusive.** Commercial deployment can also generate broad societal benefit through wider access to services and infrastructure that public-interest applications build on. How these forces balance, and which kinds of agentic AI ultimately reach scale, is one of the central questions of the agentic transition. Another important question embedded in this discussion concerns **Europe's strategic position** — whether the EU's strengths in regulation, research, and human-centric values translate into a meaningful role in shaping the agentic web (see definitions in the following section), or whether structural gaps mean Europe inherits the agentic web rather than authoring it.

Second, the lens zooms in on the **implications** of agentic transition, including potential economic restructuring, the emergence of a new agentic dimension of the digital divide, and the energy and sustainability pressures created by persistent autonomous computation. Finally, the paper looks at **policy and governance** initiatives surrounding agentic developments — existing regulatory frameworks that apply (or fail to apply) to autonomous systems, the global standardisation efforts underway, and specific gaps that agentic AI exposes in accountability, transparency, and human oversight.

A defining feature of this transition is its **uncertainty**. The pace at which agent autonomy will continue advancing, the locus of value capture across the technology stack, the demands that persistent autonomous computation will place on energy and digital infrastructure, and the capacity of public institutions to steer the technology's development – none of these dimensions can be confidently predicted today. Rather than attempting to forecast a single trajectory, this paper therefore engages with this uncertainty directly, sketching out alternative pathways through which agentic AI could develop and what each would mean for Europe. The aim is not to identify which future is most likely, but to clarify the choices that will shape which one is realised.

This paper builds on a comprehensive research and stakeholder consultation exercise carried out under the project **"Web4Hub: A Space for the Metaverse – Virtual Worlds and the Transition to Web 4.0"**, implemented by PPMI and TNO on behalf of the European Commission. The analysis draws on an extensive review of scientific and policy literature, expert interviews, and insights from project activities, including the Global Multistakeholder High-Level Conference on Governance of Web 4.0 and Virtual Worlds held on 31–1 April 2025[3]. It also builds on the findings of earlier project deliverables, including the conference background and outcome documents and other thematic papers prepared as part of the project.

This paper is part of a **trilogy on agentic AI** and its impact on the future of the web. It provides an introduction and sets the context for the other two thematic papers: one on internet infrastructure, and the other on trust and architecture. By establishing a shared understanding of what agentic AI is, how it is developing, and the socio-economic and governance forces shaping its trajectory, this introductory paper provides the common analytical foundation on which the two thematic deep-dives build.

### Figure 1. The agentic paper trilogy

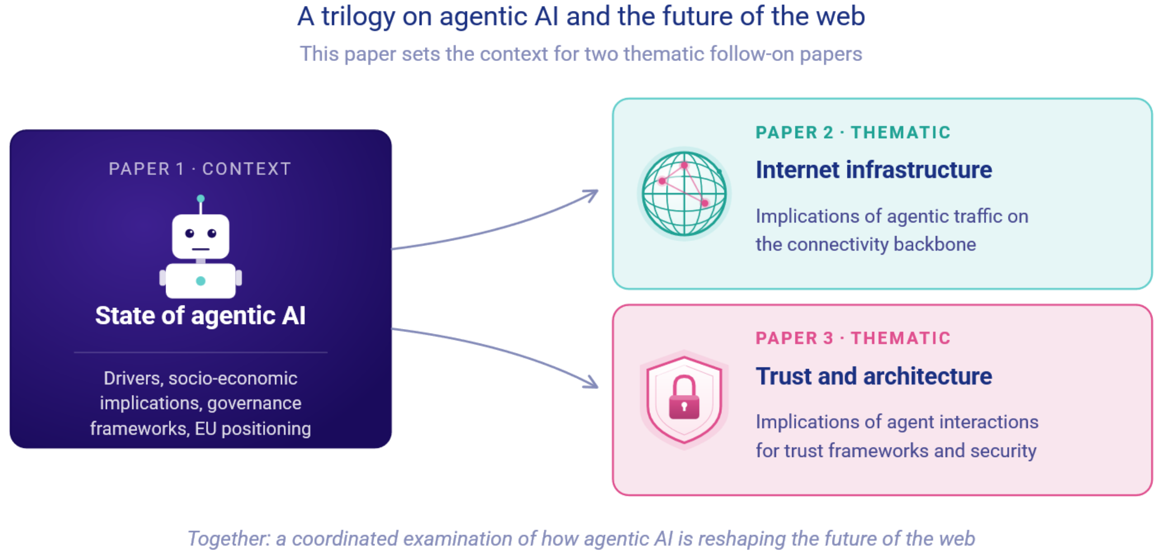


The paper is structured as follows: the remainder of this **introductory** chapter defines agentic AI, traces the trajectory of increasing agent autonomy and sets out the paper's objectives; **Chapter 1** examines the socio-economic dimensions of agentic AI; **Chapter 2** turns to the policy and governance landscape, analysing the existing gaps and the efforts to bridge them; the report concludes with **key findings** on the topic of agentic AI's evolution.

## Objectives of the paper

This paper aims to provide context for the thematic papers in the agentic AI trilogy and outline an evidence-based account of the current state of agentic AI and its trajectory. It does so with four cross-cutting objectives. First, to **clarify** what sets agentic AI apart from the generative systems that dominate public discourse, and what the shift toward autonomous, goal-directed computation means. Second, to **surface the tensions** shaping the technology's development: between societal ambition and commercial logic, between concentrated private investment and broad-based public benefit, and between the pace of deployment and the pace at which institutions can respond. Third, to **assess Europe's position**, identifying both the structural gaps that constrain the EU's competitive standing and the genuine assets on which a distinctive strategy could be built. And fourth, to **support decision-making under uncertainty**, with an aim to remain resilient to how the technology, the market and the geopolitical context may evolve over the coming decade.

Underlying these objectives is a question this paper takes seriously. Today, commercial logic is the dominant force directing AI investment, with capital flowing overwhelmingly to enterprise applications and the largest market opportunities. The pursuit of public good in domains attracts a small fraction of comparable resources. This paper does not treat that asymmetry as inevitable. It examines whether the EU, given its regulatory tradition, public-sector capacity, and stated commitment to human-centric digital values, will act to ensure that commercial logic is not the only force shaping how agentic AI develops, or whether it will accept the current default.

All in all, this paper addresses the questions shown in Figure 2 below.

**Figure 2. Questions the paper seeks to address**

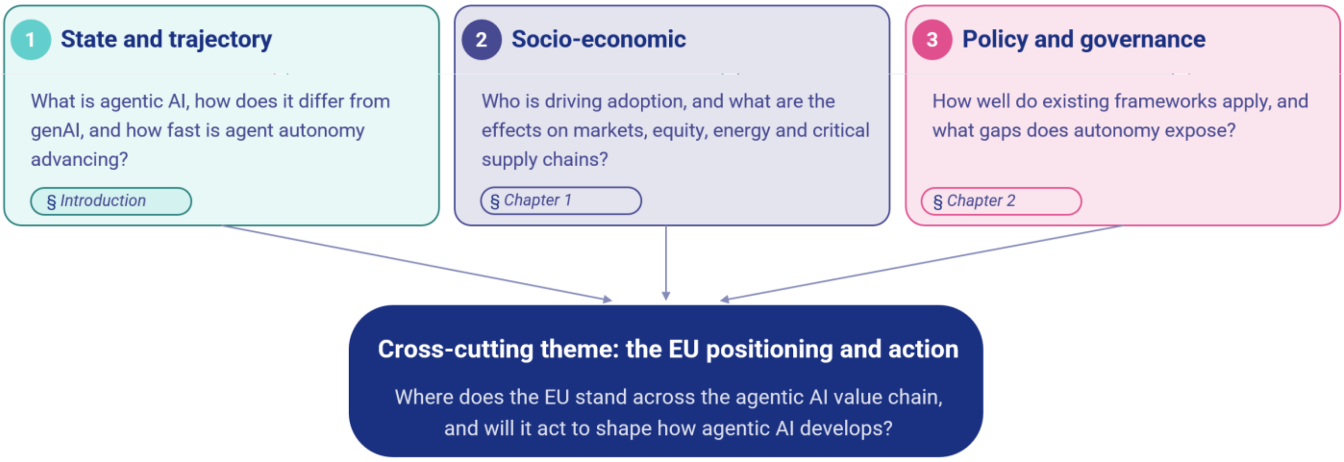


# What is agentic AI?

### *Differentiating between genAI and agentic AI*

The trajectory of the internet is currently undergoing a shift that transcends improvements in bandwidth speed or connectivity. The digital ecosystem is **transitioning from a human-centric web** designed for visual consumption and manual interaction to an infrastructure that must accommodate billions of **autonomous, non-human entities** that will need to seamlessly operate across data sources and organisations[4]. One of the key technologies enabling this is agentic AI. Figure 3 below provides a definition of agentic AI, alongside AI agents, generative AI and the agentic web. Importantly, while **AI agents** are single-entity systems, purpose-built for a specific, bounded task, **agentic AI is a system** comprising multiple, specialised agents that collaborate, communicate, and dynamically allocate subtasks toward a complex, high-level objective[5]. The transition from isolated, task-specific agents to complex agentic activity eventually shapes what is commonly called the **agentic web**, or web of agents. It is a dynamic, agent-driven environment where agents can perform complex tasks on users' behalf, which is considered a paradigm shift from a repository of human-readable information to an agent-centric, action-oriented framework[6,7].

## Figure 3. From reactive chatbots to agentic web

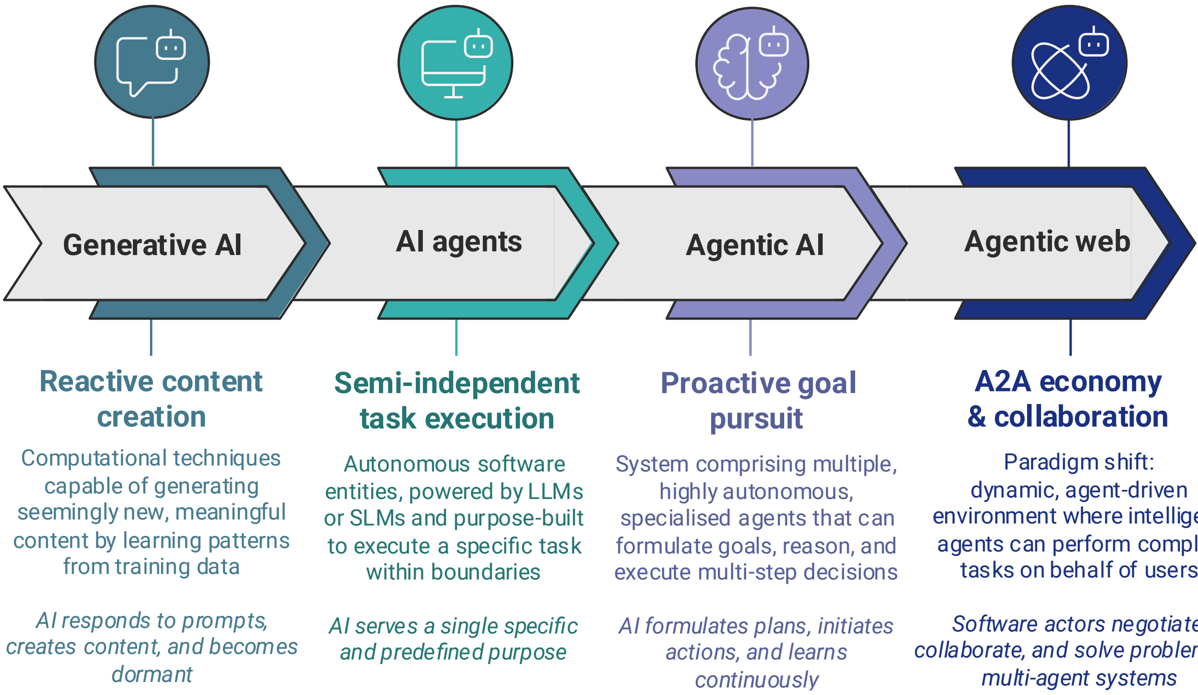


Source: authors' elaboration based on the literature review[8,9,10,11,12].

Another relevant concept is also the (agentic) **AI stack**, which includes the hardware, data, and software layers on which AI applications are built and operated[13]. The parts of the AI stack are analysed in further detail in a companion paper on "Agentic AI and internet infrastructure".

It is also important to give a clear distinction between agentic AI and generative AI (genAI) technologies that currently dominate public discourse.

- **Generative AI is fundamentally reactive,** functioning as a content creation engine, producing text, code, or imagery only in response to specific, isolated human prompts[14]. It typically executes a single inference pass to satisfy a request. Structurally, it executes a single "inference pass", meaning input goes in, the model calculates a probability, and output comes out. Once the output is delivered, the system returns to a dormant state.
- In contrast, **agentic AI is proactive**, defined by the systems' capacity for proactive agency. Agentic AI systems comprise multiple **AI agents**, autonomous software agents or entities, often powered by large language models (LLMs)[15], but now also increasingly often by small language models (SLMs) too. An important distinction is that while separate AI agents typically handle single tasks within their specific domain, agentic AI goes multiple steps further, achieving complex goals across environments. Agentic systems go beyond chatbots and pilots. They can independently formulate plans, execute multi-step workflows, and iterate to achieve high-level objectives with minimal to no human oversight[16]. They are also capable of perceiving their environment, reasoning, and planning in the long term[17,18], as well as adaptively refining their strategies based on feedback.

Figure 4 includes examples of current popular generative AI and AI agent models.

## Figure 4. Examples of generative AI and agentic AI models

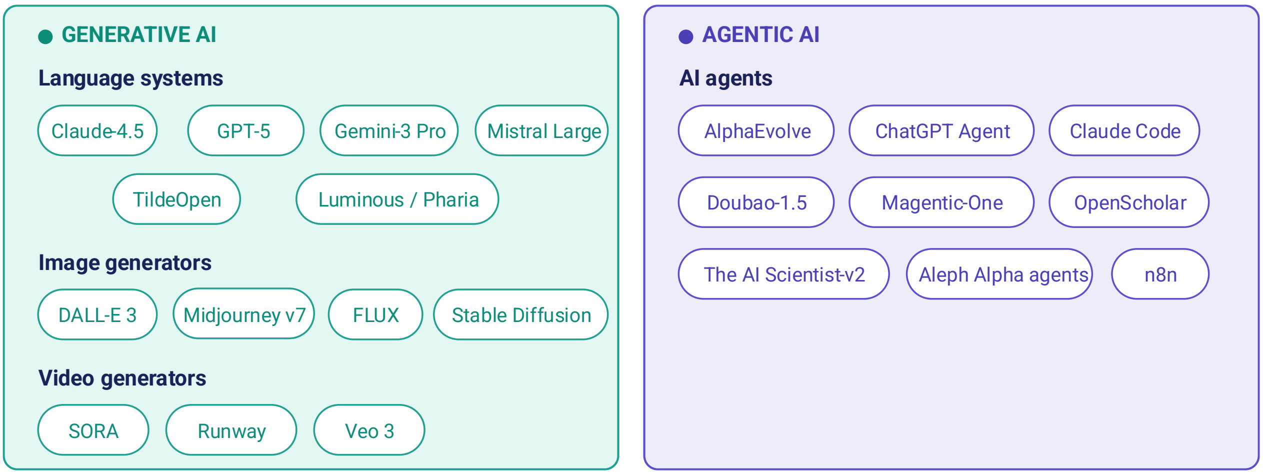


Source: Bengio, Y., et al. (2026).

Agentic AI is open-ended, non-deterministic and difficult to predict[19]. It is powered by a perception-reason-act-learn loop[20], which allows agents to continuously interpret data, evaluate scenarios, balance constraints, and execute autonomous actions[21]. Its architecture is explained in more detail under Annex 4.3.1. It is important to avoid **over-anthropomorphising AI agents**' capabilities[22]. While agentic systems demonstrate sophisticated pattern-matching and multi-step execution, their underlying mechanisms may be fundamentally different from human-like reasoning or intentionality, representing what some researchers characterise as advanced extrapolation rather than true, human-like, cognitive agency[23].

Agentic AI introduces additional layers of complexity that set it apart. The capacity for autonomous action necessitates **new technical components**, particularly communication protocols that enable agents to discover, authenticate, and interact with external services and other agents. Moreover, the **persistent, iterative nature** of agentic workflows generates substantially different bidirectional network traffic patterns compared to the request-response model of genAI. The implications of this for digital infrastructure warrant a separate discussion, to which the standalone thematic paper "Agentic AI and internet infrastructure" (published within this project) is dedicated.

At the same time, agentic and genAI do both rely heavily on LLMs and, increasingly, SLMs as their core reasoning engines and primary interface for understanding natural language[24]. Both agentic and genAI leverage transformer architectures and depend on massive-scale pretraining on vast corpora to generate contextually appropriate outputs. Both benefit from instruction fine-tuning and reinforcement learning from human feedback. They require sophisticated inference infrastructure and real-time access to large, high-quality datasets to ground their autonomous decisions in external knowledge. They both depend on computational infrastructure: cloud-based GPU clusters for large-scale pretraining and fine-tuning, high-performance computing for intensive experimentation, simulation, and multi-modal or ensemble architectures, and scalable data centres that provide the surrounding storage, networking and orchestration layers. Furthermore, they share challenges related to hallucination, bias propagation, and the interpretability of model outputs[25]. In fact, genAI often operates as a component within agentic systems: an autonomous agent may call on an LLM to compose an email, write code, or generate insights as a step towards a larger goal. This interdependence has an important implication for the scope of this paper, as explained in Box 1 below.

## Box 1. Methodological note: AI agent-specific evidence versus broader proxies

While the analysis presented in this paper focuses on agentic AI, evidence specific to agentic systems is not always available due to the novelty of technology. Given that agentic deployment at scale is still recent, with many systems only transitioning from prototype to production, agent-specific data on investment flows, energy consumption, labour-market effects, or adoption rates is frequently partial, proprietary, or not yet disaggregated from broader AI statistics. In these cases, drawing on evidence about AI or genAI provides the **most reliable available proxy**, which is highlighted where relevant in the paper.

As mentioned, Agentic AI is not a parallel category developing alongside genAI. It is, in substantial part, **built on top of it.** LLMs powering most contemporary agents are the same models that power consumer-facing genAI products, trained on the same corpora, served from the same compute infrastructure, and subject to the same reliability and safety constraints. Shifts in frontier training costs, data centre geography or regulation, therefore, reshape the conditions under which agentic systems operate, even where the evidence speaks of "AI" rather than "agents" specifically.

## Agentic leap forward: trends in the increasing autonomy of AI agents

While the concept of AI systems pursuing set goals is not inherently new, recent progress in LLM development has driven increased interest and enthusiasm about creating **sophisticated, general-purpose AI agents** among developers and investors[26]. In recent years or even months, agentic AI has rapidly transitioned from prototypes to wide deployment[27]. The Annual AI Governance Report published by the International Telecommunication Union (ITU) proclaimed 2025 "the year of agents" due to their rise in production and adoption[28]. While projections of how many AI agents will exist within the next 5-10 years are speculative and limited, it is highly likely that **billions of agents will be deployed** across consumer, business and public sector environments[29]. A 2026 study suggests that, with time, AI agents might become the main interface and gateway for users to access content and digital services[30].

As experts working with agentic AI emphasise, most of the AI agents being deployed today are not fully autonomous[31,32], meaning that their users primarily retain control as operators or collaborators (see Figure 5). The specific level of autonomy is not merely a matter of technical limitations of each system but also a deliberate choice of developers who define agents' capabilities.

### Figure 5. Levels of agents' autonomy

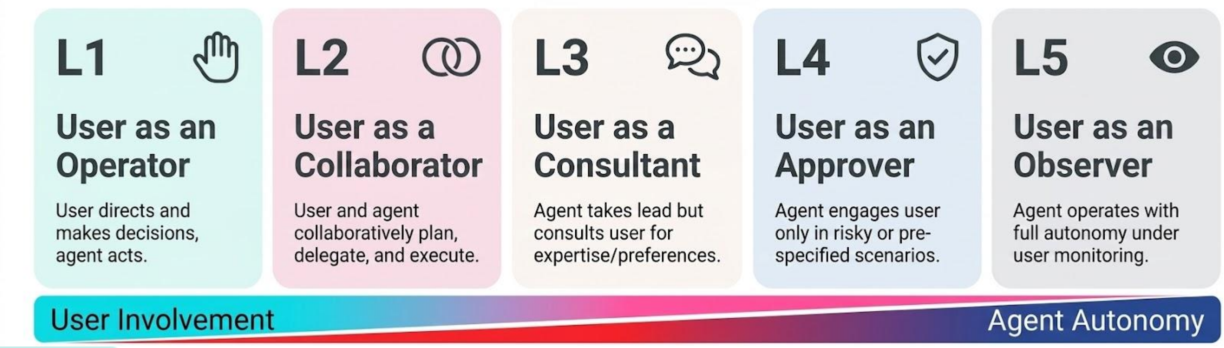


Source: based on Feng, K. J., McDonald, D. W., & Zhang, A. X. (2025).

There are contrasting perspectives on the **feasibility of fully autonomous agentic AI**. Some researchers argue that current agents continue struggling with multi-step tasks[33]. While recent studies have shown that agents now match human experts on engineering tasks that require 8 hours of work, human engineers still maintain a performance edge on tasks beyond that threshold[34]. Moreover, transformer-based LLMs have fundamental mathematical limitations that make them inherently unreliable for complex agentic tasks[35]. In reasoning tasks, AI also performs highly unevenly across languages[36]. In a paper on the possible AI trajectories towards 2030, OECD experts argue that if scaling AI models will stop delivering further capability gains or will run into substantial physical infrastructure, investment and resource limitations, then the abilities of independent agents to handle complex tasks in real-world environments will also remain far from human levels[37]. This means that even if developers in this scenario manage to advance agents' metacognition skills for computer-based tasks, their creativity and novel thinking will still be largely constrained.

Nonetheless, the current lack of agents' reliability and sophistication is often perceived by experts and executives primarily as a **temporary limitation** that will be resolved by the increasing performance of the underlying LLMs[38]. Some experts caution that relying exclusively on LLM scaling may not eliminate hallucinations but rather refine their obfuscation and reduce transparency if the chain of reasoning in natural language is removed for the sake of performance[39]. Others suggest that hallucinations can be managed through guardrails and continued technical improvement, as well as engineering solutions to filter unreliable outputs[40,41]. Importantly, given the probabilistic nature of AI agents, hybrid architectures should complement genAI with non-LLM methods for validation and control. While those constraints are being resolved for general-purpose agents, narrowly scoped agents stand to gain a significant advantage. Besides, developers are making progress in improving agents' reliability by breaking down their "thoughts" into step-by-step reasoning[42].

As discussed further in this paper, in both public and private sectors, research and science, agents are already being embedded across tasks and processes (see Section 1.1 for a more in-depth overview). According to the AI 2027 report[43], AI agents will be turning into increasingly **autonomous decision-makers** that can perform research and independently manage workflows. In the view of OECD experts, the scenario where agents' autonomy and metacognition accelerate can materialise if AI systems become able to identify algorithmic innovations autonomously and support human developers with innovating and experimenting[44]. Other consulted experts also foresee that, gradually, the cost and effort associated with building efficient and sustainable agentic solutions will decrease substantially to enable agents that perform complex, long running and open-ended tasks, overcoming today's constraints[45].

This represents a **qualitative leap from generative AI's initial capabilities**, moving from systems responding to individual prompts to systems that independently navigate across databases, applications, and decision points to complete entire workflows. While simple AI agents may perform singular tasks using tools like search engines, agentic AI goes further by coordinating multiple agents, managing their communication, and distributing tasks to accomplish larger, more complex objectives[46]. A crucial aspect is its ability to perform limited programming, and call other IT systems via APIs to gather information and adapt to its environment. Moreover, measurements show that the *length* of tasks AI agents can complete has been growing exponentially over the past 6 years, doubling every seven months: what used to be ten seconds in 2020 was around an hour in 2025[47].

Whether agentic AI will rapidly advance toward full autonomy or face prolonged bottlenecks remains contested. Optimists point to exponential growth in task complexity and the rapid move from prototypes to production. Sceptics, however, highlight such obstacles as the reliability gap in multi-step tasks, uneven reasoning performance across languages and problem types, and infrastructure and scaling ceilings that may constrain capability growth before human-level autonomy is reached. How quickly these bottlenecks are resolved will likely define the trajectory of agentic AI.

As the diversity of models is growing and the price of their development is decreasing, a particular value lies in the meaningful coordination and orchestration of distributed intelligent agents[48]. The previous iterations of the web focused on connecting humans to information (Web 1.0) or humans to humans (Web 2.0), built to mainly support human visual consumption and keyword search[49]. This new era, in contrast, envisions a machine-to-machine (M2M) or agent-to-agent (A2A) economy where software actors interact via natural language, run multiple processes in parallel, negotiate, collaborate, and transact to fulfil human requests. Such multi-agent systems (MASs) have been proving their ability to solve complex tasks in ways that single models cannot, like math Olympiad problems[50]. Annex 4.3.1 elaborates on MASs in further detail.

It is also important to take into consideration that recursive feedback loops in agentic AI and exponential growth in the number of interactions also demand **complex coordination and orchestration**, as well as constant alignment of stakeholders' preferences and values. The progressive shift toward the agentic web, while bringing higher value capture, foresees lower control, higher error propagation and higher requirements for trust infrastructure[51]. Besides, to achieve advanced multi-agent autonomy, significant progress is needed in integrating spatial intelligence, real-time analytics, and reinforcement learning[52].

Understanding the evolutionary trajectory of these technologies is essential, as each stage introduces distinct governance challenges. The shift from reactive tools to proactive and mature MASs alters the questions of accountability, liability, and oversight. The remainder of this paper examines the economic, societal and policy implications of this transition in detail.

# 1. Socio-economic considerations

This chapter examines the societal, economic, and sustainability aspects of agentic AI development. They are categorised into **drivers** (factors fuelling adoption) and **implications** (challenges and opportunities arising from deployment). While subsequent thematic papers prepared within this project address the effects on connectivity, trust and security frameworks, this chapter establishes why those infrastructural and architectural questions matter in the first place: because they are inseparable from questions of economic power, public and commercial benefit, equity, and governance.

### Figure 6. Structure of Chapter 1

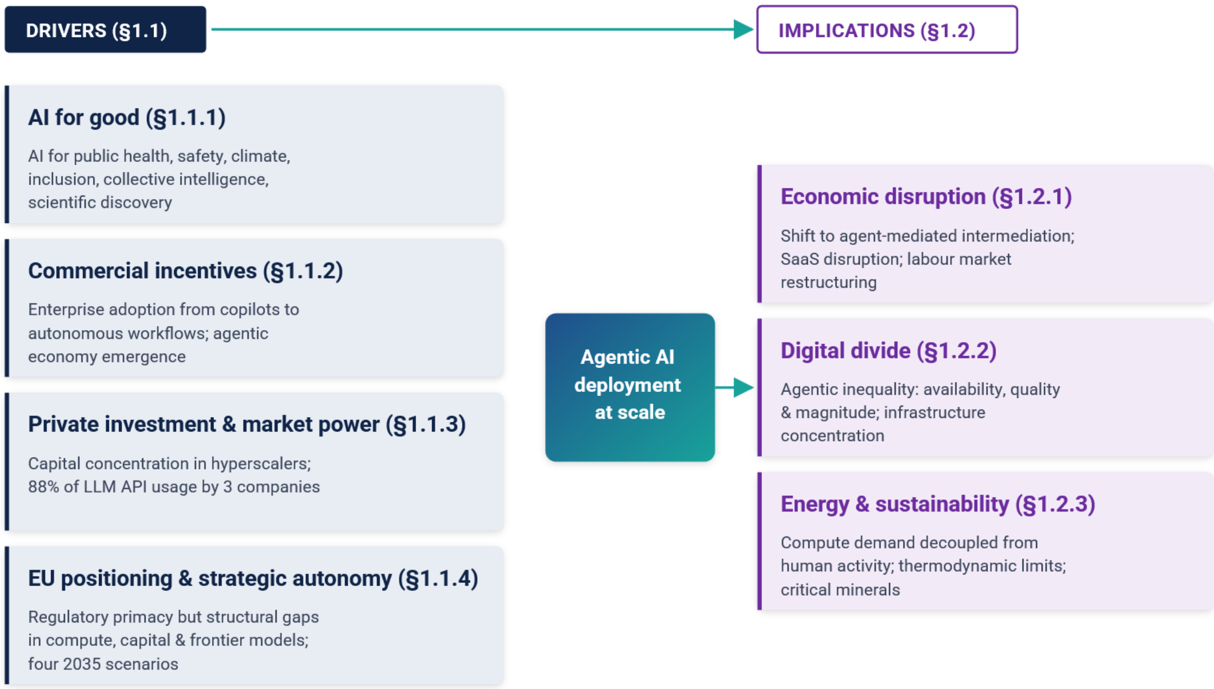


Agentic AI does not emerge in a vacuum. Its development trajectory is shaped by a specific configuration of societal ambitions, commercial incentives, capital flows, and geopolitical positioning. Section 1.1 traces these drivers, beginning with **two coexisting motivations** behind the agentic AI transition. These frequently overlap in practice, but are useful to distinguish:

- non-financial motivations, where agentic AI is pursued for societal good across healthcare, climate action, public services, and scientific discovery (see Section 1.1.1).
- financial motivations, where businesses adopt agentic AI to improve profitability, whether by raising revenue, cutting costs, or gaining competitive advantage (see Section 1.1.2).

These are not mutually exclusive, and in practice most deployments reflect both, but they pull development in different directions. Yet when examining the drivers of agentic AI development today, the answer is unambiguous: **private capital dominates**, with hyperscalers and frontier laboratories committing resources at scales that dwarf public investment and increasingly concentrating control over the infrastructure, models, and platforms on which agentic systems run. Current commercial incentives are oriented primarily towards financial returns, from automating enterprise workflows to boosting the productivity of firms of all sizes[53]. These investments often also generate wider societal benefits. However, those applications where societal purpose is primary (e.g. healthcare or climate science) still attract a comparatively smaller share of investment[54].

This raises a question for Europe. The EU's strengths in regulation, research, and human-centric values position it, in principle, to **lead on AI deployment for the public good**, as well as strategic, critical and highly regulated industries, where it holds a competitive advantage[55] Section 1.1.4 poses the question of whether this represents a genuine strategic opening or a consolation narrative for a continent already structurally outpaced in the race for frontier capabilities. This question is explored via an assessment of the EU's competitive standing across talent, compute, research, frontier models, and industry uptake. The section then presents **four scenarios** for how different combinations of AI progress and European strategic choices could reshape the continent's position by 2035.

Unlike genAI which produces outputs for human review, agentic systems act autonomously across interconnected environments, amplifying the consequences of both success and failure. Section 1.2 examines **three domains where these implications are most acute**: economic restructuring as agents become primary market intermediaries; the risk that unequal access to agentic capabilities deepens digital divides; and energy demands that strain decarbonisation efforts and critical supply chains. Each of these also has implications for how autonomous agents need to be steered and governed, as further explored in Chapter 2.

## 1.1. Socio-economic drivers

Whether agentic AI is steered toward broad societal benefit or narrower commercial gain is shaped less by the technology than by how these drivers interact. This section examines each in turn, as mapped in Figure 7, moving from the aspiration to serve societal and planetary benefit, to commercial demand accelerating adoption, to private investment concentrating behind it, and finally to what this configuration means for the EU.

**Figure 7. Socio-economic drivers of agentic AI**

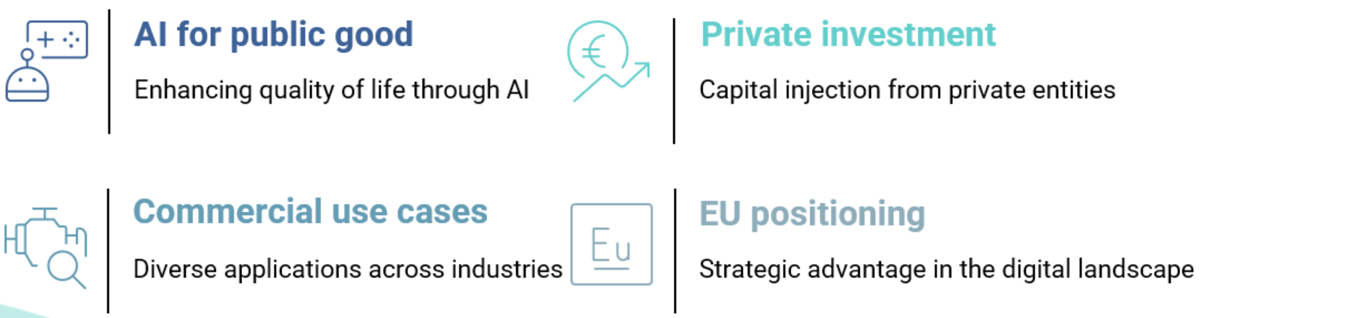


Source: authors' elaboration.

### 1.1.1. AI for good – leveraging technology for societal and planetary benefit

The transition to an agentic web presents a pivotal opportunity to deploy **"AI for good"** at scale that goes beyond productivity gains and operational efficiency. Today, there is an ambition in both policy and technology communities that agentic AI, like many other emerging technologies, can and should be **leveraged for the common benefit**, not only financial incentives (discussed in detail in Section 1.1.2). The emerging concept of "AI for good" looks beyond the profit-driven logic that often dominates technological innovation and calls for AI systems that are tailored, trained, tested and targeted to enhance human well-being and protect and preserve the planet[56]. Commercial and public-interest actors can pursue this aim together and in parallel.

The "AI for good" perspective is broadly endorsed by the United Nations (UN), the EU, opinion leaders and experts in AI systems, among other actors. The European Commission, for example, pursues the goal of developing cutting-edge AI models and algorithms to tackle major societal challenges like **improving healthcare, optimising energy use, weather and climate forecasting**, mitigating the impact of disasters and crises, and facilitating the post-crisis reconstruction[57]. The EU's vision underlines that AI technology can bring transformative benefits to well-being, improve the quality and efficiency of public services, and be a powerful tool for preventing and combating discrimination and ensuring equal opportunities for all[58]. On the global level, the ITU established **AI for Good, an annual multi-stakeholder**

**forum** with a mission to unlock AI's potential to serve humanity through building skills, AI standards, and advancing partnerships while addressing the UN **Sustainable Development Goals** (SDGs)[59].

This vision of AI as a technology with a high potential for societal and planetary good drives innovation, experimentation, use case creation and deployment worldwide. It is also pertinent to **AI agents** specifically since they are emerging as **integrated collaborators** for improving not only business processes but also public services and everyday life[60]. Purposeful use of agentic AI for the public good can reshape social infrastructure, allowing autonomous coordination across sectors and regions to widen access to essential services. In such a context, resource allocation could better reflect shared societal priorities rather than market forces alone, human agency could be strengthened rather than displaced, and technological capabilities could be intentionally directed toward addressing common challenges. Furthermore, applications that address **concrete societal challenges** (e.g. energy, security, resource management) can help build trust towards agentic technologies and foster greater buy-in and safe adoption. The AI for Good forum highlights multiple use cases that range from AI-powered early disease detection, intelligent tutoring systems, smart electricity grids and connectivity networks to food waste mitigation, biodiversity monitoring and enhanced crime prediction for public safety[61].

For **individual users**, AI applications also yield significant benefits. Approximately one in three EU citizens uses genAI regularly, exceeding professional use across nearly all Member States[62]. The majority of internet users use AI for these primary reasons: therapy and companionship, organising daily life, finding purpose, enhancing learning and generating code or ideas[63]. Main topics of conversations include seeking specific information, editing texts, tutoring and teaching[64]. AI agents can proactively guide citizens through complex administrative processes like filing taxes or benefit applications, reducing friction and inequality. They can also optimise public transport routes or manage energy consumption in smart homes, improving daily comfort and convenience.

When it comes to specific processes, agents can effectively **synthesise disparate data streams**, from citizen feedback to environmental conditions, and iteratively refine solutions based on shifting community needs and planetary boundaries. They can also initiate and consistently implement **high-volume, rules-based tasks,** alleviating bureaucratic burden for both citizens and public servants as well as reducing working time[65]. Agentic AI's ability to maintain transparency and function within tightly defined frameworks coexists in balance with the capacity for autonomy, reasoning, and dynamic adaptation of behaviour to evolving societal contexts and conditions[66,67]. This enables systems that act on **human intentions, values, and collective objectives.** Besides, agents can **augment human agency** by interpreting complex goals, mediating between diverse interests, and coordinating collective action toward shared outcomes – at the pace that global challenges increasingly demand.

From passive, reactionary tools, AI applications evolve into autonomous multi-level decision-makers that can support entire organisations. The collective intelligence of humans and AI **amplify citizen science** by using agentic systems for task matching, data processing and idea generation[68]. Yet, for true **collective intelligence**, agentic frameworks need to be representative so that they can function and compromise in high-dimensional contexts, with factors to consider beyond economic costs, such as environmental benefits and social equity[69]. One interviewed expert underlined the particular value of the "world models"[70], which are neural networks that have an in-depth understanding of the real world's physical and spatial properties – an approach that a Turing Award winner, Yann LeCun, also endorsed[71].

Moreover, it is important for agentic systems to be resistant to risks revolving around data privacy, liability, security of software ecosystems, along with handling edge cases[72]. **Ethical frameworks** emerging from multistakeholder processes converge on a shared principle: to distribute benefits equitably rather than concentrating them among privileged demographics, AI systems should amplify human capabilities and agency rather than diminish them. As one of the interviewees within this project suggested, responsible use of AI can be supported through partnerships where universities and research centres develop algorithms and then, after proper testing, are deployed in the market[73]. Public procurement could also underpin support for trustworthy and value-driven products[74].

In interviews conducted as part of this project, experts emphasise the **need to carefully consider AI's limitations in sensitive applications**, particularly those involving vulnerable populations. One expert cited the example of AI psychological counselling for teenagers, noting these systems lack "lived experience", formal accreditation by recognised psychological associations, and transparent training beyond what is available to be read on the internet[75]. Such applications raise fundamental questions about deploying agentic systems in domains requiring human empathy and contextual judgment. ITU also calls upon giving voice in governance and oversight to people and regions that AI agents will affect, such as workers, civil society groups and governments in the Global South[76]. It is crucial for governance frameworks to reflect the diverse interests and contexts globally.

Across multiple high-level events within the last year, experts and inventors highlighted the role of AI in **high-impact scientific discoveries**, including at the 2025 AI in Science Summit[77]. Another example is the Paris Summit for Action on Artificial Intelligence, which endorsed scientific projects harnessing AI to tackle some of the world's most pressing challenges, like quality education, gender-based violence and saving biodiversity[78]. AI's role in vaccine development, drug discovery and climate modelling was also widely acknowledged during the 2025 World Economic Forum annual meeting, where the narrative generally shifted from naïve optimism to **cautious pragmatism**[79]. Finally, in 2024, the OECD updated its AI Principles, the first intergovernmental standard on AI. Among them, "Inclusive growth, sustainable development and well-being" focuses on responsible stewardship of trustworthy AI in pursuit of beneficial outcomes for people and the planet, e.g. advancing inclusion of underrepresented populations, reducing inequalities, and protecting natural environments[80].

*Demis Hassabis*, the co-founder and CEO of Google DeepMind and a 2024 Nobel laureate in Chemistry, also believes that AI – in particular, artificial general intelligence (AGI) – will be able to solve the "root-node" problems: curing terrible diseases, bringing healthier and longer lifespans, and finding new energy sources, and this could start happening towards 2030[81]. He insists, however, on cautious optimism, maximising AI's potential for good while careful risk management and effective global governance frameworks. *Yoshua Bengio*, a professor at the University of Montreal known as one of the "godfathers of AI", also emphasises the significant role AI already plays in researching applications in medical, environmental, social justice, and other "societal" fields. Yet, in line with ITU's vision, he also underlines that people whose lives will be most affected by these innovations (e.g. Global South, disadvantaged populations) should be involved to a much larger extent in the discussion and development of AI systems[82]. Section 2 delves more deeply into agent-related policy frameworks.

Box 2 below highlights examples of agentic AI solutions developed or deployed with an aim to solve pressing societal and planetary issues.

### Box 2. Illustrative applications of agentic AI for societal good

Agentic AI applications – those that truly leverage autonomous workloads – do not yet have a substantial catalogue of best practices from the public sector. The technology is rather new and is still developing, with many unknowns, including who builds agents, who governs them, and how we ensure even access for all citizens[83]. Nonetheless, there are **emerging examples** piloted across the world.

One recent research paper, for instance, explored the use of LLM-driven agents to simulate how different populations might respond to specific **policy interventions** before real-world implementation: e.g. for poverty reduction, climate action, and promoting inclusive institutions[84]. In this regard, agents can imitate the processes of deliberation, co-creation and emergence of community norms. Agentic-driven AI personas are also used in **humanitarian action**, for example, simulating conversations with refugees and conflict actors and training staff[85]. Scientific research has proposed numerous examples of multi-agent frameworks helping in predicting and preventing wildfires[86], planning land use[87], and other sustainability-oriented applications.

The Singapore government has developed a multi-agent system prototype that mimics data scientists collaborating on tasks, **delegating responsibilities** and **producing contextual insights** from CRM data[88]. A Centre for Responsible Business from the University of California, Berkeley, highlights real-life use cases for AI agents assisting dentists in diagnosing, verifying identities and credit scores of loan applicants, and sorting and processing police body-camera footage for underfunded legal aid offices[89]. The Ukrainian Ministry of Digital Transformation developed the world's first AI agent for government services, acting on behalf of a citizen requesting certain multi-step action via a chatbot[90]. This both makes services more accessible and helps in tackling corruption. Finally, a pilot study in Wales demonstrated how MASs improve the coordination of electricity, heat, storage and renewable generation in seaports, decarbonising energy-intensive industries[91].

Such applications can significantly improve the accessibility of **crucial public services** for disadvantaged groups and underserved communities. For instance, computer vision advancements boost the autonomy of AI-powered robots and their performance in assisting humans during emergencies, as well as in monitoring environments for disaster prevention[92]. A Whitepaper by the Berlin GovTech Centre forecasts that agents can make public services much more proactive and personalised, applying inclusive rules more consistently and prompting citizens toward beneficial next steps[93]. All these tools can amplify daily interactions with public administrations with fewer human resources involved, as a consulted expert confirmed[94].

Agentic AI's capabilities impact **local services**, too. For instance, CONSUL DEMOCRACY's open-source Civic Assistant, developed in a 2025-2027 European consortium, integrates agentic AI into participation platforms used by cities. It generates proposals, provides real-time translation, and facilitates voice input to boost engagement from marginalised groups, which leverages collective citizen input[95]. As IDC's FutureScape research on smart cities and communities predicts, by 2027, 65% of cities will deploy AI agents to orchestrate end-to-end workflows and reduce workloads, augmenting the capacity of public sector teams[96]. The processes it can affect vary from processing applications and issuing permits to reconciling budgets, detecting policy inconsistencies and planning infrastructure and zoning. For instance, in Pittsburgh, a multi-agent decision-making system coordinates traffic lights based on traffic conditions, resulting in improved traffic flow and a 25% reduction in average travel time[97].

### 1.1.2. Commercial use cases across sectors

Another important factor driving innovation, experimentation and deployment of agentic AI, is its **commercial value** and **market benefits**. Commercial deployment and societal benefit are not necessarily at odds. Many enterprise applications, like diagnostics, fraud prevention or energy optimisation, deliver efficiency and public value at once, and commercial scale is often what makes such benefits widely available. In the last several years, AI tools have witnessed a surge in adoption: in the EU, for example, the number of enterprises that introduced AI almost tripled between 2021 and 2025[98]. Google Cloud research based on a broad survey of technology leaders concludes that 98% of organisations are experimenting with, developing or actively using genAI in production[99]. Yet four in five enterprises report no material contribution to earnings from their genAI initiatives[100], and where revenue gains do materialise, they typically remain at 5% or below[101]. This disconnect reflects the predominance of horizontal tools that enhance individual workers' productivity but fail to transform business processes significantly, alongside deeper organisational challenges — from skills shortages and change management difficulties to the absence of clear leadership strategies for scaling AI beyond pilot stages.

Agentic AI offers a different proposition. Where general-purpose copilots assist individual users with discrete tasks, agentic systems can retain memory, recognise context, and **autonomously plan and coordinate actions across processes and business functions**[102]. This shift creates stronger potential for tangible enterprise value, though realising it will depend as much on organisational readiness as on the technology itself. As one expert interviewed for this paper noted, the efficiency gains on offer provide a strong incentive to push the boundaries and encourage investment[103].

Companies increasingly recognise agentic AI as distinct from prior AI waves precisely because it promises to automate and accelerate entire complex workflows and not just isolated tasks. Tech companies lead in the adoption of agentic systems, followed by communication and media companies, life sciences and manufacturing[104]. Gartner forecasts that by 2028, nearly a third of Fortune 500 firms will rely solely on a single AI-powered channel for customer service[105]. Some authors describe an expansive "agentic economy" evolving, where an open, interconnected, decentralised network allows AI agents to represent every person, business, and new digital services[106].

A variety of incentives backs agentic adoption by businesses: from efficiency gains and growth to market expansion, resilience, and risk management. Largely, agents' uptake is driven by pressures faced by organisations to boost **operational agility** and manage rising labour costs. Agentic AI offers SMEs capabilities previously available to large, multinational companies[107], reducing the need for human involvement and oversight, especially in routine, streamlined tasks with clear logic. This allows for much **faster and cheaper applications**, which is economically valuable and accelerates their development and deployment[108]. Leveraging this technology, manual workflows can be done up to 30% faster already in early adoption stages, reaching up to 60% with maturation[109]. The reduction of low-value work time also allows for more complex, value-adding tasks and sensitive issues. Agents promise significant economic growth by **lowering communication friction** and transaction costs, currently limiting market size and firm function[110].

The motivation for exploring agentic AI extends beyond efficiency gains. For instance, agentic solutions can boost organisations' **resilience** by monitoring disruptions, adapting process execution to unforeseen changes, and initiating safety measures, if necessary[111]. Furthermore, agents can help enterprises meet escalating customer expectations for **personalised**, always-on services by handling **unprecedented transaction volumes** 24/7. The buy-in among consumers is also growing as people become increasingly willing to let agentic AI handle grocery shopping or home maintenance[112]. A particular benefit stems from using a "swarm" of agents, where individual agents with sector-specific knowledge are integrated in one system to collaborate and leverage **collective intelligence** across many different processes[113].

AI agents are becoming a part of multiple business processes – from HR, finance and marketing to product development and equipment maintenance – across various sectors and industries. These applications primarily concentrate in domains where multi-step processes, real-time decision-making, and cross-system integration create clear value propositions. Agentic use cases are concentrated in enterprise software[114], but they nonetheless span a variety of other fields as well. Figure 8 highlights some examples that have already been deployed, while Annex 4.3.2 describes them in further detail.

## Figure 8. Illustrative sectoral applications of agentic AI in core activities

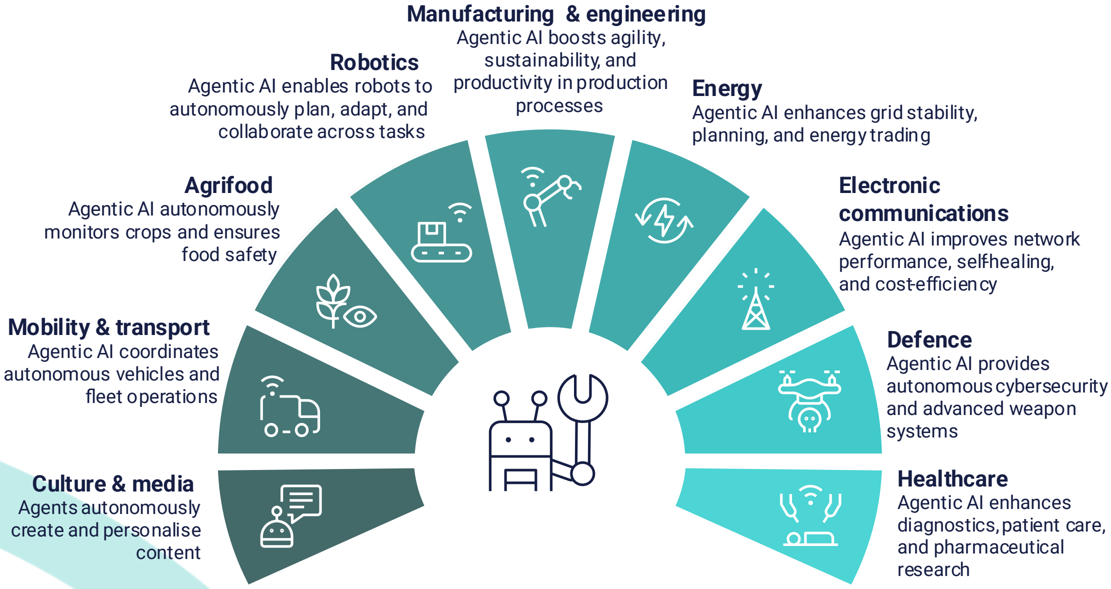


Source: authors' elaboration based on the reviewed literature. Note: the list of sectors reflects the strategic sectors of the EU Apply AI Strategy; public sector and climate applications are covered in section 1.1.1.

In 2026, the implementation of agentic AI is predicted to gain extra pace, encompassing more functionalities, cross-platform operation and hybrid modes of work between agents and employees[115]. Businesses plan substantial budget allocations into specifically agentic tools and expect to gain competitive advantages from them, as surveys show[116]. Many companies believe that investments will be primarily channelled in AI agents as applications instead of packaged apps enhanced by agents[117]. Annex 4.3.3 illustrates a possible timeline for such a trajectory.

Commercial applications also foster the emergence of new business models around the agent economy. For example, AWS and Anthropic introduced a marketplace of agents, where customers can find and buy third-party AI agent solutions from a catalogue[118]. An EU-based Sokosumi platform offers similar capabilities[119]. At the same time, intermediary low-code platforms emerge that allow users to build their own agents[120]. This indicates that the economy of agentic AI has moved away from theoretical research to a **commercial reality** – a trend also observed in agentic-related academic research that shifts focus from the "web of agents" as a theoretical construct to a practical necessity[121].

At the same time, Gartner predicts over 40% of agentic AI projects will be **cancelled** by the end of 2027 due to escalating costs, unclear business value or inadequate risk controls, as many experimental use cases are currently misapplied[122]. Besides, unstructured corporate data remains the primary bottleneck that limits the genuine impact of agentic solutions on the productivity levels in corporate workflows[123]. Realising the transformative potential of agentic AI in commercial settings will therefore require not only technical maturity, but also consistent investment strategies, robust data governance frameworks, and a clearer alignment between deployment ambitions and measurable value.

### 1.1.3. Private investment and market power

In 2024, **agentic AI received USD 1.1 billion (EUR 947.8 million) in equity investment globally**[124]. While this is a small share of the overall AI equity investment, amounting to USD 124 billion (EUR 106 billion) (see Figure 9 below), it is growing faster than any other technology and witnessing a corresponding demand for skilled talent. The share of agent-focused venture capital (VC) investment in Europe is also higher than in the U.S. and other regions, and the number of deals is steadily accelerating[125]. The AI agent market is predicted to have a compound annual growth rate of 45% over the next four years, quickly gaining traction[126]. In 2025, AI investments overall reached nearly 50% of total VC funding globally[127].

**Figure 9. Equity investment in AI companies globally, 2024**

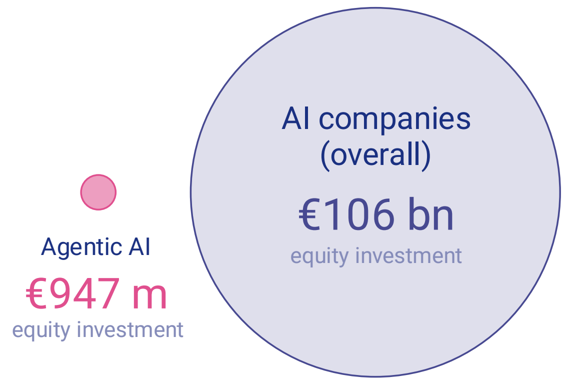


Source: authors' elaboration based on McKinsey & Company. (2025).

Currently, **private-sector investment and commercial interests are primary drivers of agentic AI development**. Although in the EU, public and quasi-public institutions play an admittedly important role in shaping the investment landscape, hyperscalers and frontier AI laboratories are committing resources at levels unavailable to other market players. In 2024, global private AI investment hit a record high with 26% growth[128]. Service-oriented sectors, like finance, telecommunications and software providers, lead the VC funding in the EU[129]. Over the past five years, companies have rapidly increased their share of AI supercomputer ownership, outnumbering public sector systems. In light of surging hardware costs and power needs, the public sector's share of supercomputers declined from 60% to 12% (see Figure 10)[130]. Goldman Sachs expects 2026 to undergo a significant surge in capital expenditure by AI hyperscalers, primarily shifting away from infrastructure to platforms providing database and development tools[131]. Nearly 90% of notable AI models in 2024 came from industry, up from 60% in 2023, while academia remains the top source of highly cited research[132]. The high marginal cost of agentic inference threatens to price out public-sector and non-profit applications[133,134]. If this overarching dominance of hyperscalers and verticals persists, **the core value in the agentic technologies will most likely be created and harnessed in the narrow ecosystem** of LLM developers, compute capacity providers and platforms.

## Figure 10. Share of aggregate supercomputer performance (16-bit FLOP/s)

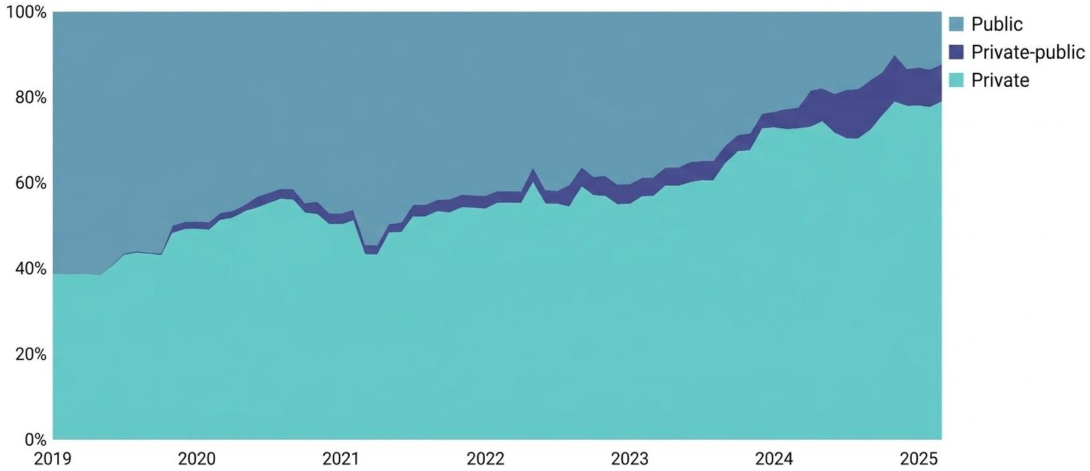


Source: EpochAI. (2025).

Box 3 highlights some of the **major commitments and ambitions** expressed by large technological corporations regarding AI and agentic AI for the coming years. This capital concentration reflects a consensus that **agentic AI infrastructure** is highly important for competitive advantage in the digital sphere over the next decade, encompassing data centres, high-performance computing clusters, and proprietary model-serving architecture. While general internet infrastructure is designed to connect users and services and move information across networks, agentic AI infrastructure is designed to connect and coordinate autonomous agents and support the traffic flows beneath them. For a more detailed discussion on internet infrastructure adaptation in light of agentic AI, see a companion paper published within this project. Already now, the providers of general-purpose AI systems are increasingly offering commercialised agents: from distinct ones to integrated product offerings[135]. At the same time, as an interviewed expert underlines, large corporations developing leading agentic systems do not necessarily need to have societal value at the core of their value proposition[136].

## Box 3. Agentic-AI investments and commitments from large companies

- Microsoft and NVIDIA launch an accelerator for agentic startups and scale-ups in the UK and Ireland[137].
- AWS plans to invest USD 100 million (EUR 86 million) to boost agentic AI deployment[138].
- In 2025, Alphabet announced a USD 75 billion (EUR 64.6 billion) investment in AI infrastructure. Later, the company presented its yet most powerful Tensor Processing Unit (TPU), Ironwood, powering proactive agentic AI for the "age of inference"[139].
- Siemens announced a EUR 1-billion investment in industrial AI capabilities, including "engineering agents"[140].
- Salesforce executed nine acquisitions in 2025 alone, specifically aiming to accelerate its Agentforce platform for businesses to build and customise autonomous agents[141].
- Chinese tech giants (Alibaba, Baidu, ByteDance) and startups are also actively launching agents based on in-house foundation models, albeit trailing behind American tools[142].

Driven by substantial computational costs of training, the need for specialised talent, and the demands of scaling infrastructure, **global AI investment is dominated by mega-deals above EUR 100 million, and these largest funding rounds in Europe are dominated by non-EU investors**[143]. The largest companies possess both the investment capacity and hardware access to develop vertically integrated platforms that combine foundational models, developer tools, and enterprise distribution channels, pivoting instantaneously toward breakthroughs. As a result, **just three companies – Anthropic, Google and OpenAI – account for 88%** of enterprise LLM API usage[144]. The cost structure of digital value is also shifting. Value capture is moving from the training of foundation models to the **inference phase**, or the continuous execution of reasoning loops required for agents to perform work[145]. In the underlying computing capacities, hyperscalers also hold a lion's share of initiative: in the near term, the largest tech companies will build at least 55 new hyperscale data centres[146].

Although right now, the agentic AI market in its nascent stage still has numerous vendors and platforms to choose from, experts warn about anticompetitive "bundling" of AI agents with other products or unfair practices towards competitors[147]. Frontrunners are already trying to hinder each other's agents: for example, Amazon demanded that Perplexity's Comet browser remove this marketplace from the websites it can access due to the "poor quality of service" Comet allegedly provides[148].

At the current moment, the expert community does not have certainty on how the agentic market will eventually be structured and whether it will be business-to-business or business-to-consumer applications that see greater commercial success[149]. What can already be observed is that despite fierce commercial competition, **large competitors are forming alliances**: Microsoft, Google, OpenAI, and Anthropic have joined forces to establish the Agentic AI Alliance, backed by the Linux Foundation[150]. Ostensibly, its mission is to develop open standards, as its members emphasise the current critical juncture between open, interoperable infrastructure and proprietary fragmentation[151]. Yet functionally, it also **enables incumbents to shape the technological trajectory** in ways that reinforce their market positions. In doing so, they also collectively raise the barriers for new entrants lacking similar capital depth and partnership leverage.

Apart from the Agentic AI Alliance, **other open initiatives have been established that support community-driven platforms**, **prioritise interoperability and collective development**. Prominent examples include AutoGPT for pioneering recursive self-prompting and goal-directed task execution, LangGraph for orchestrating complex multi-agent workflows, CrewAI for production-ready role-based task delegation, and SuperAGI for scalable multi-agent orchestration – all fostering interoperability and rapid iteration through thousands of global contributors. Besides, Nvidia unveiled a new family of open-source AI models with a mixture-of-experts architecture for creating custom and scalable MASs, with organisations' own data and values[152].

### 1.1.4. EU's positioning and strategic autonomy

The objective of this section is to analyse the **EU's positioning within the ecosystem of agentic AI and how it relates to the continent's sovereignty and competitiveness**. While the EU has established regulatory primacy through the AI Act and articulated ambitious industrial strategies, fundamental structural asymmetries in compute capacity, private capital formation, and frontier model development place the **EU at a competitive disadvantage relative to the U.S and China.** At the same time, according to McKinsey, "*adopting a holistic approach to capitalise fully on genAI's potential could boost European labour productivity by up to 3% annually through 2030*"[153].

Europe accounted for around 25% of the global AI agents' market as of 2024 (approximately EUR 1.1 billion), with Germany retaining the biggest share in revenue and France showing the highest growth rate[154]. The EU has **significant assets**: a large, high-quality AI talent pool, world-class research output, a strong position in producing critical nodes in the semiconductor supply chain, and a regulatory framework that may prove a differentiating advantage in ensuring trust and the safeguarding of individuals' rights. The precedent of social-media age restrictions points to another such strength: the EU could cultivate safe, homegrown agentic AI alternatives designed for users from an early age[155]. Moreover, the EU is particularly well-positioned to develop AI agents tailored to specific sectors, critical industries and public services, leveraging extensive domain-specific knowledge and data, an approach referred to as "verticalisation"[156] .

How these strengths and weaknesses interact will depend on both the EU's policy choices and the trajectory of agentic AI itself — whether the EU secures sovereign capacity over the infrastructure and standards underpinning agentic AI, and whether the resulting ecosystem develops along public-interest or commercially-driven lines.

Table 1 analyses **key dimensions of EU positioning for the next AI frontier**, such as talent pool, frontier model development and uptake. It is important to note that several of the data points cited are related to AI or genAI, not specifically agentic AI. Nonetheless, as explained in Box 1, this evidence is relevant on two grounds: as a proxy where agentic-specific data does not yet exist, and as a direct measure where competitiveness in AI and generative AI translates into competitiveness in agentic AI.

These dimensions inform the **scenarios** presented at the end of this section, which explore how different combinations of EU sovereignty and public interest versus commercial orientation could reshape the continent's infrastructure landscape and competitive standing by 2035.

### Table 1. EU positioning across key dimensions in (agentic) AI

| Dimension | EU position |
|---|---|
| **Talent pool** | • In terms of AI skills penetration[157], several EU countries are global leaders (e.g. Germany, France), with many other countries ranking at or above the OECD average (e.g. Spain, Netherlands, Italy, Sweden, Poland, Ireland, Denmark). While the U.S. remains the clear leader, top EU countries outperform or match other advanced economies such as Japan, Korea, and Canada.<br>• The EU produces more AI-specialised researchers than the U.S. every year.<br>• However, if underutilised, the AI talent may continue to migrate to other markets. For example, a 2023 analysis found that 14% of European-educated AI undergraduates moved to the U.S. for graduate studies.<br>• Initiatives such as the European strategy for AI in science (RAISE) and its "Choose Europe" campaign seek to reverse this trend by strengthening Europe's attractiveness for AI researchers. EU Visa Strategy and the proposed EU Talent Pool initiative also aim to attract and retain skilled talent by reducing administrative burdens.<br>• Agentic AI development draws on a subset of these capabilities — particularly in agent orchestration, tool integration and evaluation — so the broader AI talent base is a necessary but not sufficient indicator of agentic-specific capacity. |
| **Compute, chips and infrastructure** | • Around 74% of global GPU cluster capacity is based in the U.S., compared with about 15% in China and roughly 4.8% in the EU.<br>• The EU has launched a multi-tier compute strategy, investing EUR 10 billion to establish 19 interconnected AI Factories (including 3 in the global Top500) and 13 AI Factory Antennas supported by EUR 55 million in EU funding. The next generation foresees the creation of five AI gigafactories in Europe. Industry sees these capacities as a solid starting point but not yet sufficient to match the EU's economic and industrial ambitions.<br>• The EU hosts production of critical elements in the advanced semiconductor supply chain (ASML is the sole provider of EUV lithography systems, although 10% of its components and software originate in the U.S).<br>• Hyperscalers (AWS, Azure, Google) dominate Europe's cloud and AI stack. |
| **AI research** | • EU ranks second in scientific publications on AI with 18.2% (second to China with 20.1% and overcoming the U.S. with 14.8%) between 2020 and 2024; EU GenAI R&I activity growth: average annual growth of 32% (2019–2021).<br>• Unlike the U.S. and China, which dominate AI publications, no individual EU MS reaches a comparable volume. Instead, several EU countries (notably Germany, France, the Netherlands, Spain, and Italy) form a dense cluster with substantial AI research output.<br>• Global GenAI patents filed: >120,000 by 2024; however, the EU share of global GenAI patent filings: 2%[158]; the EU patenting activity is concentrated in only a few regions[159].<br>• Through Horizon Europe and Digital Europe, the Commission invests more than EUR 1 billion per year in AI, with the AI in Science Strategy aiming to double annual Horizon Europe AI investments to over EUR 3 billion.<br>• Agentic AI builds on a set of upstream research areas — including reinforcement learning, multi-agent systems, planning and alignment. Given the lack of availability of agentic AI disaggregated data, broader AI publication and patent figures are therefore used here as the closest available proxy. |
| **Frontier models** | • U.S. institutions produced 40 notable AI models in 2024; China produced 15; Europe produced only three (all from France: Mistral, Kyutai, LightOn).<br>• The gap between top U.S. and Chinese model performance narrowed to near-parity on major benchmarks (MMLU, HumanEval), while European models remain specialised, and only a few EU developers actively pursue frontier capabilities. Combined annual revenue of EU frontier developers is around 1% of that of leading non-EU peers.<br>• Focus instead is on the application layer, as foundation model development needs massive financial and computing resources European actors struggle to mobilise[160]. |

| Dimension | EU position |
|---|---|
| | • The AI in Science Strategy and its flagship RAISE initiative aim to pool EU funding, compute, data, and talent to advance frontier AI capabilities.<br>• Most AI agents today are built on frontier LLMs as their reasoning core, thus frontier-model development directly shapes where capable agents can be built. |
| **Industry uptake** | • According to Eurostat, AI use by EU enterprises grew by 50% between 2024 and 2025.<br>• GenAI uptake among businesses is similar in the EU (37%) and the U.S. (36%), according to the EIB investment survey. A McKinsey survey also shows similar AI use by organisations, with the U.S. slightly ahead, while other geographies, including Greater China, developing markets and Asia-Pacific, lag behind.<br>• A growing cohort of EU agentic AI startups — concentrated in hubs like Stockholm, Berlin, Paris, Munich, Amsterdam and Barcelona – are achieving unicorn-level valuations (e.g. Parloa, Lovable, Legora) and attracting substantial funding. The enterprise value of EU AI startups increased sixteen-fold over the past decade to around EUR 161 billion, with over 4,100 startups across priority industries. The number and scale of investment deals have been steadily accelerating since 2023.<br>• However, while several EU countries host significant numbers of AI startups, the EU still lags behind the U.S. and China, which lead the global AI startup landscape.<br>• About 10% of EU scaleups and nearly 30% of unicorns (2008–2021) relocate abroad — mostly to the U.S. — with moves especially common in digital sectors such as AI and biotechnology to access larger markets and investment.<br>• The Commission's Apply AI Strategy (October 2025) aims to accelerate AI adoption across strategic sectors and SMEs, complementing the broader AI Continent Action Plan and AI in Science Strategy.<br>• GenAI uptake establishes a baseline of organisational familiarity with AI workflows, but agentic deployment requires a further step, which in most organisational contexts has not yet begun or is at its very early stages[161]. |

Sources: stakeholder consultations; European Commission (2026)[162]; European Commission (2025a)[163]; European Commission (2025b)[164]; European Commission (2023)[165]; Epoch (2025a)[166]; Epoch (2025b)[167]; Abendroth Dias et al (2025)[168]; AI in Science Summit (2025); Stanford University (2025); the Paulson Institute (2023)[169] ; Negele et al (2025)[170]; OECD (2026)[171]; EIB (2026)[172]; EIB (2025)[173]; Eurostat (2025)[174]; Lemke and Schneider (2025)[175]; McKinsey (2024)[176]; StepUp StartUps (2026)[177]; AscendixTech (2025)[178]; Dealflow.eu & EU-Startups (2025)[179]; AI Policy Bulletin (2025)[180]; StepUp StartUps (2026)[181].

This analysis points to **three fronts where structural gaps require concentrated action**:

- scaling compute sovereignty and energy infrastructure to reduce dependence on foreign hyperscalers;
- mobilising private capital at a level that retains homegrown talent and prevents scaleup relocation;
- translating the EU's strong research base and regulatory frameworks into competitive agentic AI capabilities.

Closing these gaps does not, however, require the EU to replicate the full AI stack. Several experts consulted for this paper argued that a more selective approach — pragmatic rather than fully self-sufficient, balancing sovereignty with international partnerships — would make better use of Europe's actual strengths[182]. These include open-source R&I, public-interest applications in fields such as health and climate, identification and authorisation systems, security guardrails, and frameworks for testing model robustness and ensuring auditability and accuracy of the output[183,184].

The **four scenarios** shown in Figure 11 confront the EU with strategic choices about the development model it wants to pursue and the investment it is willing to make. Rather than asking whether agentic AI will advance — which is no longer a serious question — the scenarios are structured along two axes that define genuinely contested terrain: (1) the degree of EU tech sovereignty, including over the agentic AI ecosystem, encompassing the ability to act independently and develop key technologies, data, and infrastructure, while reducing reliance on non-EU providers (aligned with the Tech Sovereignty package[185]), as well as the openness of protocols and standards, and the scale of public and private investment; and (2) the tension between AI developed for the public good and commercially driven AI, reflecting whether agentic systems are designed around openness, interoperability and publicly accessible agents or whether they primarily serve commercial interests, concentrating capability among those who can afford premium products. More detailed analysis of the scenarios is presented in Annex 4.2.

### Figure 11. Scenarios of the EU's position in the agentic AI

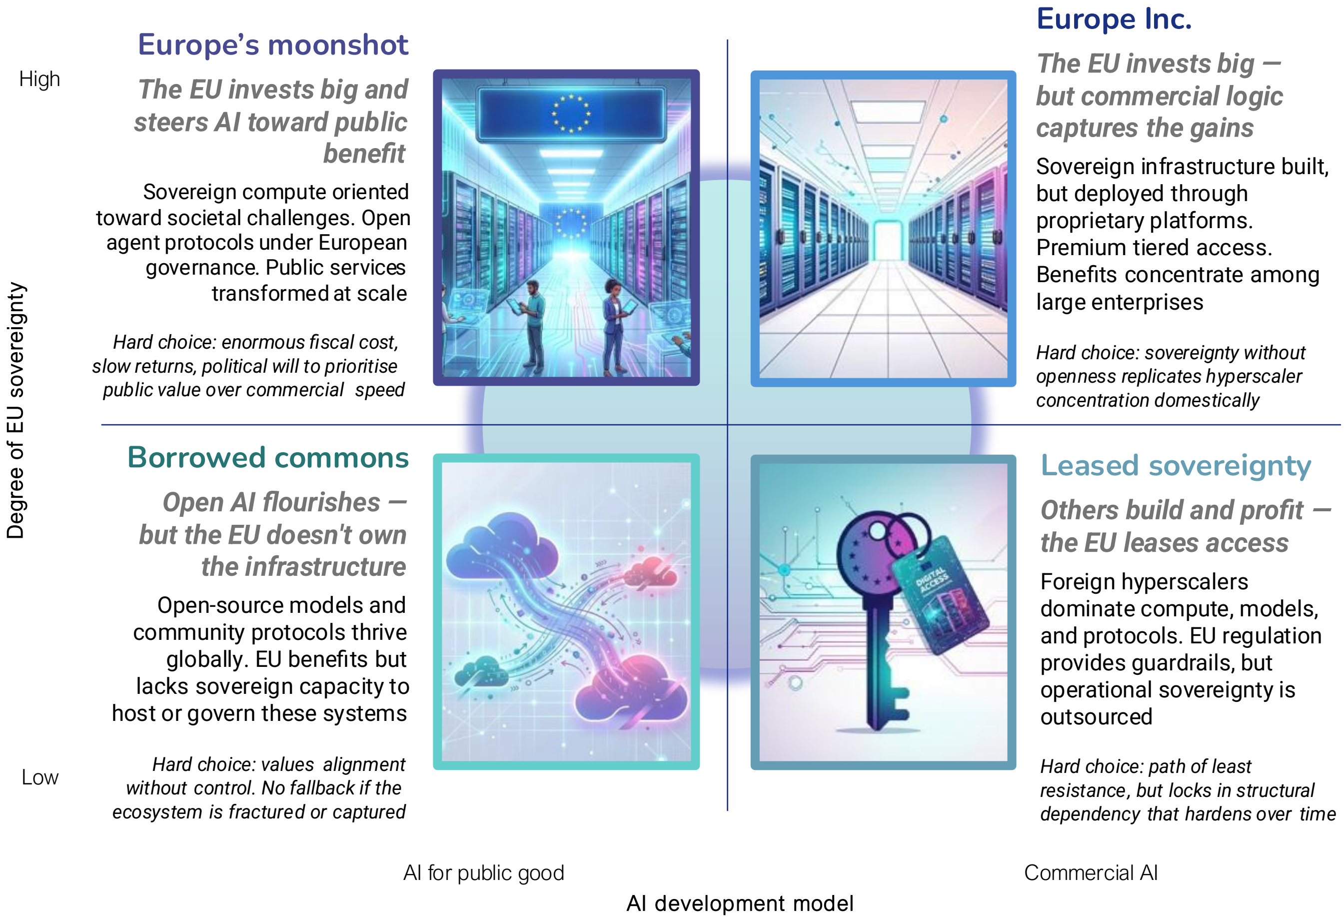


Source: authors' elaboration based on the analysis presented in this report; images generated using Adobe Firefly. Notes: for more information on how the scenarios were developed, see Annex 4.1. For more detailed scenario analysis, see Annex 4.2.

The two principal levers available to policymakers — **targeted investment and policy interventions** — cut across all three fronts, and their calibration will play an important role in determining which of the scenarios outlined in this section materialises.

As one of the two principal levers through which the EU can alter its competitive trajectory, **investment plays a critical role** in how the EU positions itself in agentic AI. The global scale of capital flow into AI illustrates both the urgency and the magnitude of the adjustment required if the EU were to aim to pursue one of the two scenarios characterised by aggressive AI capability building. Private investment in AI is heavily dominated by the U.S. In 2025, its funding of AI initiatives overall reached USD 159 billion (EUR 136 billion), compared to just USD 43 billion (EUR 36 billion) in the rest of the world[186]. Around 46% of the agentic AI market was concentrated in North America in 2024[187]. VC patterns reveal a similar asymmetry: between 2020 and 2025, total EU VC funding stood at USD 252 billion (EUR 216 billion), of which 18% was directed toward AI, compared with USD 1.33 trillion (EUR 1.14 trillion) in the U.S., where AI attracted 34% of all VC activity[188]. Moreover, OECD data shows that the average size of VC investment in the EU27 is lower than in China or the U.S., suggesting a degree of fragmentation in the EU investment landscape[189]. There is also a growing dependence on non-European investors to fund the continent's AI startups as the share of domestically sourced VC declines[190]. Late-stage funding and scale-up retention also are challenges, as many successful European AI companies exit through U.S. acquisition or relocation[191,192]. Experts also emphasise that the EU's persistent gap lies in **commercialisation**, with complex funding mechanisms and limited support for scaling, slowing the translation of innovation into market-ready products[193]. Simplifying funding access, leveraging public procurement, strengthening pathways between universities and industry and prioritising firms with demonstrated deployment and revenue capacity could strengthen the ability to build sustainable AI ecosystems.

### Figure 12. Venture capital investments in AI in USD millions by country

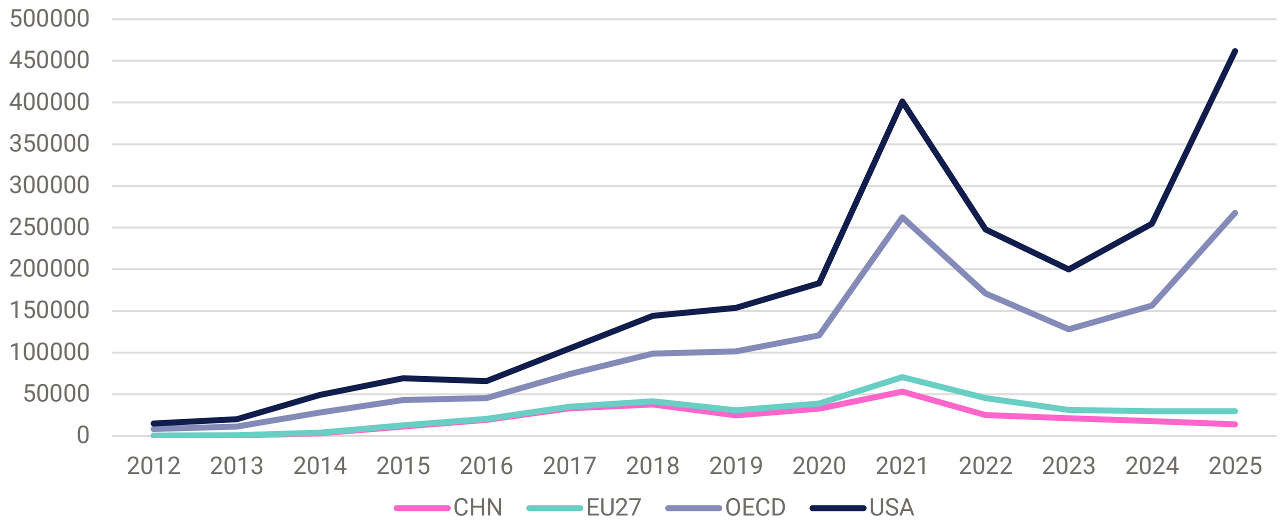


Source: OECD.AI (2026). Data from Preqin, www.oecd.ai

Annual global **investment in infrastructure** underpinning AI development is projected to grow by 10–15% per year, yet data centre capacity remains heavily concentrated in the U.S, which hosts over 40% of global capacity[194]. Meeting the anticipated global demand for AI compute would require roughly capital expenditure far exceeding any plausible level of public subsidy and necessitating large-scale private sector participation[195]. Between 2010 and 2024, more than USD 600 billion (EUR 515 billion) was channelled into AI-dedicated infrastructure globally, with the U.S. and China together capturing approximately 65% of all AI value-chain investment[196]. European countries, by contrast, accounted for roughly 10% of global AI infrastructure investment over the past fifteen years[197]. Furthermore, the U.S. controls approximately 74% of global AI supercomputing capacity, followed by China's 14%, while the EU holds roughly 4.8%[198]. This compute gap creates structural dependencies that analysts warn could undermine Europe's industrial competitiveness, critical infrastructure resilience and geopolitical leverage, creating compounding disadvantages across sectors and for geopolitical leverage as AI systems approach AGI[199]. As a positive example, the launch of JUPITER — Europe's first exascale supercomputer — marks a significant milestone in closing the compute gap, alongside the broader rollout of 19 interconnected AI Factories providing compute resources and technical support to businesses and researchers across the continent[200]. Ultimately, multiple consulted stakeholders see access to and exchange of high-quality, domain-specific, trustworthy and multilingual **datasets** as crucial for the EU to develop agents for specialised industrial, commercial and public-sector use[201].

The second lever available to policymakers is **policy and regulatory strategy**. A central question is whether the EU's regulatory frameworks can function as safeguards and as active instruments of digital sovereignty and global influence. Importantly, most agents deployed in consumer and business contexts are likely to be classified as low-risk and therefore subject to lighter obligations under the AI Act. To prevent a future in which agentic AI systems are overwhelmingly controlled by a small number of commercial actors, a growing body of research and expert commentary emphasises the strategic importance of investing in open-source models and publicly governed AI infrastructure as foundations[202,203,204]. As outlined in Chapter 2, the EU has put in place several policy initiatives to both boost its competitiveness in the AI market and promote the deployment of trustworthy and safe AI. Whether such initiatives will be sufficient depends in part on the reach and credibility of the EU's regulatory influence beyond its own borders.

The so-called **"Brussels Effect"** offers a potentially powerful mechanism for shaping international norms around trustworthy AI. A frequently cited example is the adoption of the OECD's AI principles, which drew from the EU's High-Level Expert Group on AI (AI HLEG) Ethics Guidelines for Trustworthy AI (see Box 4).

However, such leverage is not unconditional: it is strongest when stakeholders need access to the European single market or when soft guidance is concerned. Several experts suggest conditions under which the "Brussels Effect" may be weaker, including if Europe's regulatory environment comes to be seen as excessively complex, unpredictable, or inconsistently enforced[205,206] or if it concerns binding regulation instead of "softer" measures, such as guidelines or principles[207,208]. In fact, some experts highlight that neither the U.S. nor China has significantly adjusted its domestic approach in response to EU regulation[209]. The effectiveness of the policy lever, therefore, hinges not only on the ambition of individual initiatives but on the coherence and predictability of the broader regulatory environment and leverage in terms of investment and market access. As discussed in further detail in Chapter 2, the pace of technological development requires agility and results in frequent adjustments, posing challenges for both regulators and industry.

Europe's emphasis on human-centric AI could represent a genuine differentiation from large global players like the U.S. and China. However, it requires substantial investment to become and remain competitive[210] and, as explored earlier in this Chapter, the scale of investment in AI for good so far remains minuscule compared to funding for commercially driven AI[211,212]. Rather than pursuing technological independence, the EU's digital strategy may therefore be most effective when framed around cultivating shared technological sovereignty with allied nations, where mutual reinforcement of standards and capacities strengthens collective resilience against coercion and foreign influence[213].

### Box 4. EU's early leadership in AI norm-setting

The EU's AI HLEG published its Ethics Guidelines for Trustworthy AI in April 2019, which established the risk-based, human-centric framework. Soon after, the OECD adopted its AI Principles (the OECD Recommendation on AI), which closely mirrored the EU's approach[214,215]. The OECD Principles then went on to be endorsed by the G20 at the Osaka Summit in June 2019, effectively globalising the EU's normative framework.

This is a strong Brussels Effect example because the EU set the conceptual vocabulary and normative frame ("trustworthy AI," risk-based classification, human oversight) that then cascaded through multilateral governance institutions and influenced national AI strategies globally.

Across all four scenarios, the EU's trajectory hinges on a broader question of **AI sovereignty**. In line with the Tech Sovereignty package[216] it should be understood not as full ownership or control of infrastructure, but as the capacity to independently develop, deploy, and govern AI systems in accordance with domestic values, while reducing reliance on non-EU providers across the entire value chain[217]. The EU currently depends on other regions for around 80% of its core digital technologies[218]. Building the compute capacity to close that gap is itself difficult, constrained by bottlenecks in data centre construction, energy supply, and ownership[219]. These challenges are made more acute by structural dependencies on the U.S. for both security and critical digital inputs, including cloud computing[220,221,222].

In summary, the EU's competitive standing in agentic AI will be determined by its ability to move beyond primarily regulating technologies developed elsewhere and toward actively building its own components of the AI stack, infrastructure, and integrated ecosystems of excellence. This requires consolidating fragmented research efforts into strategic, European-scale initiatives that connect labs, talent, and deployment pipelines[223,224]. Interviews conducted in this project converge on the opinion that Europe's clear strengths are its specialists, open research activities and the built-in "safety and privacy first" approach. But acting collectively and decisively on this front is essential if Europe is to shape the trajectory of agentic AI in line with its values, rather than cede that capacity to others[225].

## 1.2. Socio-economic implications

The above section traced why agentic AI is being built. This section turns to what its deployment at scale sets in motion. Those consequences are harder to anticipate than for earlier AI, and not only because agents act with reduced human oversight. Development is concentrated within a small number of private actors whose proprietary methods and limited disclosure constrain public understanding of current capabilities and their attendant vulnerabilities[226].

The following sections examine **three interconnected domains** where these implications are highly acute: the potential for economic disruption driven by agentic automation; the uneven distribution of access and capabilities; and the energy and sustainability. Other important implications that should not be overlooked also include cybersecurity vulnerabilities inherent in autonomous systems and the challenges of agent identification and explainability. These are explored in further detail in a subsequent paper prepared within this project on the impact of agentic AI on trust and security.

### Figure 13. Socio-economic implications of agentic AI

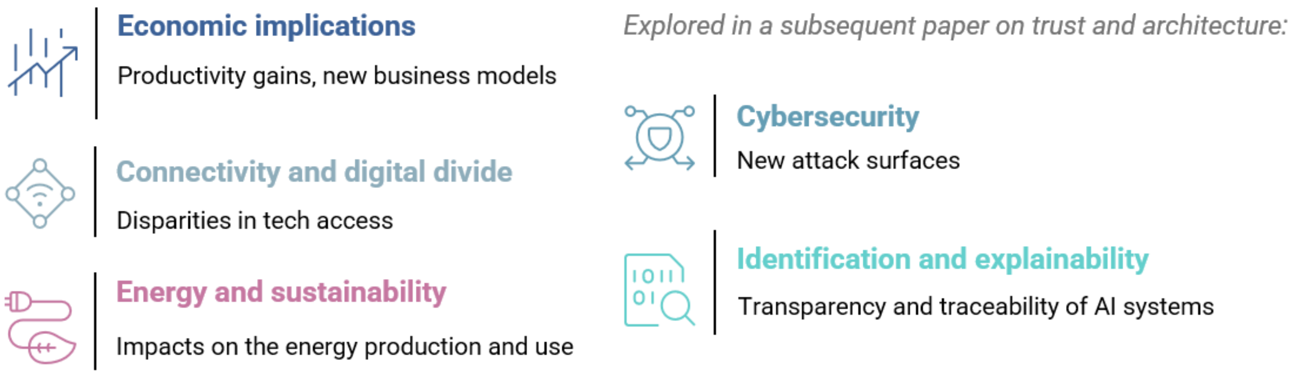


Source: authors' elaboration.

### 1.2.1. Economic implications

As agents acquire the capacity to negotiate, coordinate, and transact autonomously on behalf of users, economic exchange itself is being restructured, giving rise to an **"agentic economy",** in which autonomous systems displace humans as the primary intermediaries in routine economic activity[227]. A 2024 report estimates that the shift to agentic AI will unlock between EUR 2.2-3.8 trillion (USD 2.6-4.4 trillion) annually in global output by 2030, driven primarily by the decoupling of economic productivity from human labour supply[228,229]. For consumers and firms, this means wider choice, faster transactions, and more efficient allocation of resources across markets[230].

This section examines the **economic implications of that restructuring** in terms of digital markets and digital infrastructures, as well as spillovers to other sectors.

**Figure 14. The economic implications of agentic AI — value capture in digital markets, and spillovers to other sectors**


## DIGITAL MARKETS

AI integration can stimulate innovation by enabling platforms to develop new functionalities and **digital market business models**, intensifying competitive pressures[231]. Agentic AI is beginning to shift the locus of gatekeeping power in digital markets away from the providers of specific platform services, toward the companies that control the agent layer where users increasingly access those services. Open internet and network neutrality obligations, originally designed for internet service providers, later extended to "software gateways" such as operating systems, browsers and app stores. Experts argue that this "freedom of choice" must now extend further to agentic AI services, given their ability to interact directly with third-party services and make operational decisions on users' behalf[232].

The **attention economy** — built on capturing human engagement — **is being displaced by an agent-intermediated model**, in which autonomous systems discover, evaluate, and act on services on behalf of users without any direct human interaction. Across the globe, AI agents are strengthening businesses' positions: for example, China's local governments provide support to the one-person companies with a single founder using AI agents, as part of a broader promotion for an agent-powered industrial policy[233]. This disintermediation cuts both ways. As agents take over discovery and transactions, **control of valuable data**, such as customer relationships, shifts away from businesses that held it before. Without directly addressing transparency and data flow-through risks, it could disenfranchise European firms and consumers[234]. The most immediate evidence of this is visible in online search and web browsing, as explored in further detail in the subsequent paper prepared within this project, titled "Agentic AI, trust and architecture".

The Centre on Regulation in Europe warns that there are two compounding competition risks. More immediately, the Centre suggests that incumbents in markets on which AI agents rely actively constrain the growth of third-party agents that might otherwise challenge their position. The longer-term risk is that whoever controls the dominant agent layer becomes the **new gatekeeping layer of the digital economy**, including by algorithmically determining which services, content, and products users ultimately encounter, without those users ever making that choice consciously[235]. The EC High-Level Group for the Digital Markets Act has also highlighted that AI-driven network effects may create new gatekeepers or strengthen existing ones[236]. Large platforms are the ones that possess vast amounts of data required for effective AI systems with advanced capabilities. This **data advantage can lead to market concentration**, creating high entry barriers and stifling competition from smaller entrants[237].

To succeed in the agentic economy, businesses are being increasingly required to adapt and optimise their visibility within registries and search systems so they can be effectively discovered and selected by autonomous agents performing tasks. This has been termed **generative engine optimisation**[238], which favours structured, machine-readable data over traditional visual cues. As existing research shows, already now, AI-generated summaries tend to cite only a few widespread domains, affecting the visibility and quality of the content[239]. Across digital services, this is driving the emergence of a new layer of agent-facing infrastructure, which includes **capability registries** that enable agents to discover and verify what other agents or services can do[240], A2A communication protocols that allow autonomous systems to delegate tasks, share context, and coordinate workflows without human intermediation, and programmatic advertising infrastructure specifically redesigned around agent selection logic[241]. For a comprehensive discussion of agentic search, browsing and discovery, see the companion paper on trust and architecture.

Agentic AI is already disrupting traditional Software-as-a-Service (SaaS) models, creating structural economic pressures that flow directly into infrastructure requirements and investment patterns. This shift, sometimes termed "SaaSpocalypse"[242], has the potential to directly reshape infrastructure requirements, investment patterns, and architectural design principles for digital systems. Annex 4.3.4 further elaborates on the shift to agent-as-a-service.

Important structural shifts are also visible at the **infrastructure level**. Efficient execution of complex agentic workloads requires moving beyond conventional homogenous GPU deployments toward heterogeneous systems that combine diverse hardware[243]. In 2025, commercial hyperscalers fundamentally shifted infrastructure toward AI-optimised environments which prioritise high-value, continuous enterprise workloads such as large-scale model training and real-time inferencing over traditional scientific and civic use cases[244]. In the U.S, hyperscalers are expected to capture about 70% of the forecast data centre capacity to host advanced AI workloads[245], with already more than USD 300 billion (EUR 258 billion) committed by American hyperscalers in capital expenditure for AI infrastructure in 2025 alone[246]. This concentration enables rapid scaling and innovation, but it also concentrates control over agentic AI deployment in the hands of a small number of actors, as explored in detail in the companion paper on infrastructure.

### Value capture in the agentic economy

Overall, across the agentic AI stack, **economic value is not captured evenly**, and the contours of its distribution are already visible. At the base, upstream compute providers (chip manufacturers like NVIDIA and the hyperscalers operating data-centre capacity) capture enormous near-term revenues, remaining the clearest current beneficiaries of the agentic transition. Aggregate spending at this layer is expected to exceed USD 490 billion in 2026[247], capturing rents from structurally rising demand for inference and training capacity. Nonetheless, these players face intensifying competitive pressure as model efficiency improves and open-source alternatives proliferate.

As for **frontier model providers**, they capture value through API access and through their own emerging agentic products. They, too, face continuous margin pressure as inference costs fall sharply year-on-year – by 10 to 50 times since 2023[248] – and as open-weight alternatives continue to narrow performance gaps – from 8% to 1.7% in one year[249]. Yet strong open alternatives at every layer of the stack matter only if they are actually adopted, with adoption incentives differing across market segments: the public sector can be steered towards them through procurement policy, B2B adoption tends to follow existing reach, and consumer adoption is driven mainly by user familiarity[250].

The **agent and orchestration layer** is the most contested. Large incumbents are reinforcing their position by controlling data-centre capacity, low-latency networks, and the digital infrastructure on which cross-border agent workflows depend. But they are also integrating agents into their existing platforms and increasingly bundling agentic capabilities into existing enterprise contracts. As a wave of well-funded standalone providers attempts to cement their positions in the market, it is not yet clear whether agents will become commoditised features of larger ecosystems or sustain independent businesses. Mitigating this kind of lock-in starts with giving users genuine choice. This means supporting competition and diverse initiatives before converging on a common solution, since lock-ins have historically occurred precisely where competition was absent[251].

The **application and solution layer** (for example, agents purpose-built for healthcare, finance, legal, and other sectors) offers the most plausible avenue for durable value capture by new entrants and, notably, by European firms. Value here is secured through proprietary workflows, regulatory expertise, and access to domain-specific data, reinforced by what analysts term the "**context moat**": the institutional memory accumulated by agents trained on a firm's own processes[252]. Underlying all of these layers is the data layer, where **proprietary domain-specific** data is emerging as a critical bottleneck and a durable source of competitive advantage.

The distribution of value across the stack is not fixed and will depend significantly on how protocols, standards, and competition evolve. The **current trajectory points to disproportionate capture upstream at the compute and hyperscaler layers in the near term,** with the application, vertical agent and data layers offering the strongest foundations for long-run competitive positions. Such a shift is already reflected in the approximately 54.8% of agentic AI funding now flowing into developing vertical agents designed to manage specific workflows in a single domain[253].

## Figure 15. Value capture across the agentic stack in near- and long-term future

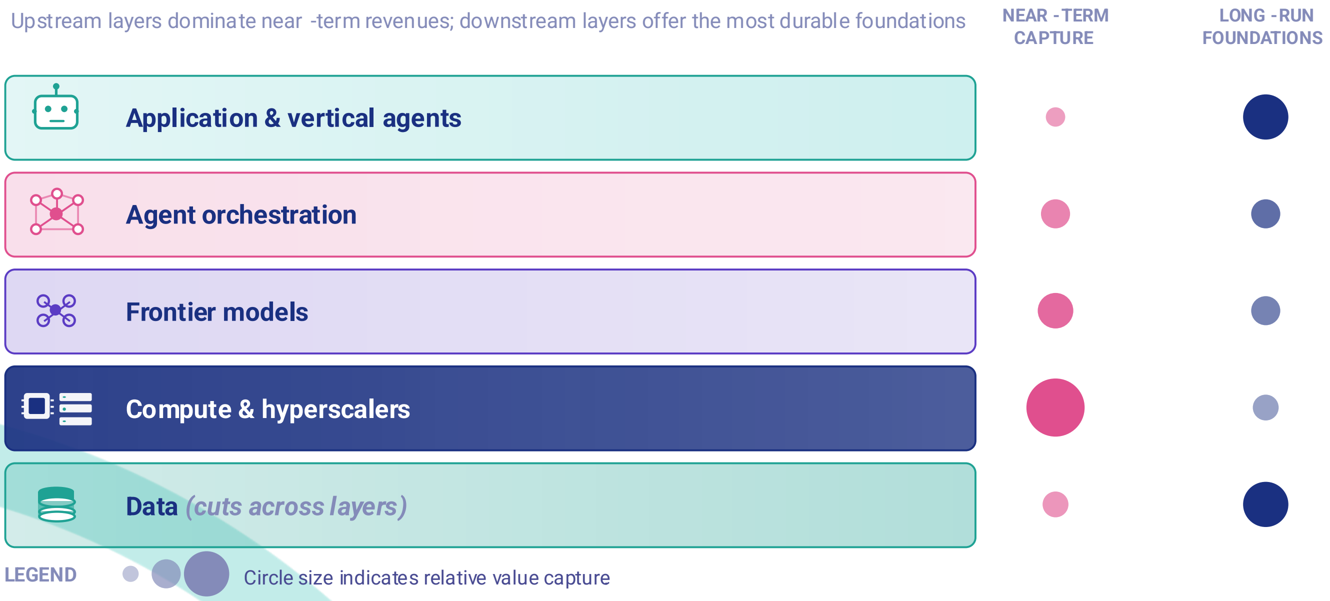


Source: authors' elaboration.

The current high cost of building bespoke, secure agentic infrastructure creates a formidable **barrier for smaller players**. Organisations often underestimate the substantial costs required to train and refine agentic systems[254]. While large enterprises can invest in private edge clouds and data modernisation to maintain control over their agents and workflows, SMEs often lack the same resources.

This **structural disadvantage compounds** with the data-related network effects as described above. Not only do SMEs face higher relative infrastructure costs, but they also accumulate far less of the proprietary, domain-specific operational data that increasingly determines agent performance. This can reinforce a two-tier dynamic in which the economic returns from agentic AI concentrate at the level of those with both the compute and the data to deploy it at a meaningful scale. The European Commission "Study on the Next Data Frontier"[255] confirms that access to high-quality, domain-specific, and multilingual datasets constitutes a critical bottleneck for European AI development, and that legal and technical barriers to addressing this bottleneck fall disproportionately on SMEs. There is a significant risk that SMEs become "tenant farmers" on the platforms of major tech giants, paying rents for the intelligence required to operate[256].

### Spillovers to other sectors

Agentic AI is also set to **transform other sectors of the economy**. Section 1.1.2 provides a comprehensive overview of how comparable restructuring is underway across supply chains, healthcare, manufacturing, finance, energy, and defence.

Across sectors, agentic AI is set to have profound implications for the **labour market**. Agents can chain together entire workflows with little or no human oversight, making them considerably cheaper to deploy than systems requiring human supervision at each step. Estimates of workforce exposure to AI vary widely depending on methodology. Studies measuring high-risk displacement place the current baseline at 23–34%[257,258], whilst broader assessments capturing any degree of task-level change reach 60% or above[259,260]. It is important to note a converging finding across all methodologies that the share of workers with zero meaningful exposure to AI is now small and declining. Cognizant's 2026 analysis found this group had already collapsed from 31% to only 7% of the workforce[261]. Allianz research from late 2025 concludes that agentic systems "may drive more substitution than previous technological waves", and under full agentic deployment scenarios, the high-risk exposure range is projected to rise further[262].

The **likely trajectory, however, is not sudden mass redundancy**. The Carnegie Endowment February 2026 analysis, based on evidence from more than 12,000 European firms, characterises this change as "incremental task substitution and workflow automation", whereby jobs are "hollowed out before eliminated", creating prolonged insecurity rather than immediate unemployment[263]. Their analysis found AI adoption increases productivity by approximately 4% with **no immediate employment losses**. Furthermore, an ECB analysis published in March 2026, finds that the most AI-intensive European firms are currently net hirers, with gains distributed unevenly across skill levels, firm sizes, and gender[264]. Moreover, women are almost twice as likely as men to hold roles with high AI exposure.

The resulting **disruption to skill demands and wages** is expected to occur more rapidly than earlier technological transitions, leaving policymakers less time to respond. Whilst some workers will benefit (e.g. through reduced administrative burden, access to previously unaffordable capabilities, improved coordination, greater transparency in evaluation, reduced exposure to physical hazards, and more flexible work arrangements), many others are likely to face displacement or wage pressure, particularly in low-to-medium-complexity cognitive roles[265,266].

Meanwhile, evidence from the adjacent literature on algorithmic management suggests that these effects will not be uniform: impacts on workers are strongly mediated by institutional frameworks, collective bargaining and the specific configuration of the technology deployed and can therefore vary considerably across sectors and geographic regions[267]. Thus, governance frameworks should explicitly manage the pace and distribution of agentic automation. The European Commission's Apply AI Strategy takes a step in this direction, committing to monitor AI's labour market impact and establishing an AI and Society Observatory to assess the changes these technologies may bring[268].

### *1.2.2. Connectivity and digital divide*

The deployment of agentic AI is not occurring on a neutral infrastructure landscape. It is being layered onto digital foundations already characterised by profound inequalities in connectivity, compute, energy, and skills. Agentic AI does not simply inherit these **inequalities — it can actively amplify them**. This section maps those divides using the three-layer framework established as part of the project[269,270], tracing how the access, capabilities, and outcomes dimensions of the digital divide are transformed by the shift to autonomous AI systems.

The background report prepared as part of this project[271], and outcome document of the Global Multistakeholder High-Level Conference on the Governance of Web 4.0 and Virtual Worlds[272], both conceptualise the digital divide through three layers: the **access divide** (whether connectivity and devices exist at all), the **capabilities divide** (whether people have the skills and literacy to engage meaningfully), and the **outcomes divide** (whether engagement translates into meaningful economic and social benefit), as visualised in the Figure 16 below.

### Figure 16. Three dimensions of the agentic divide

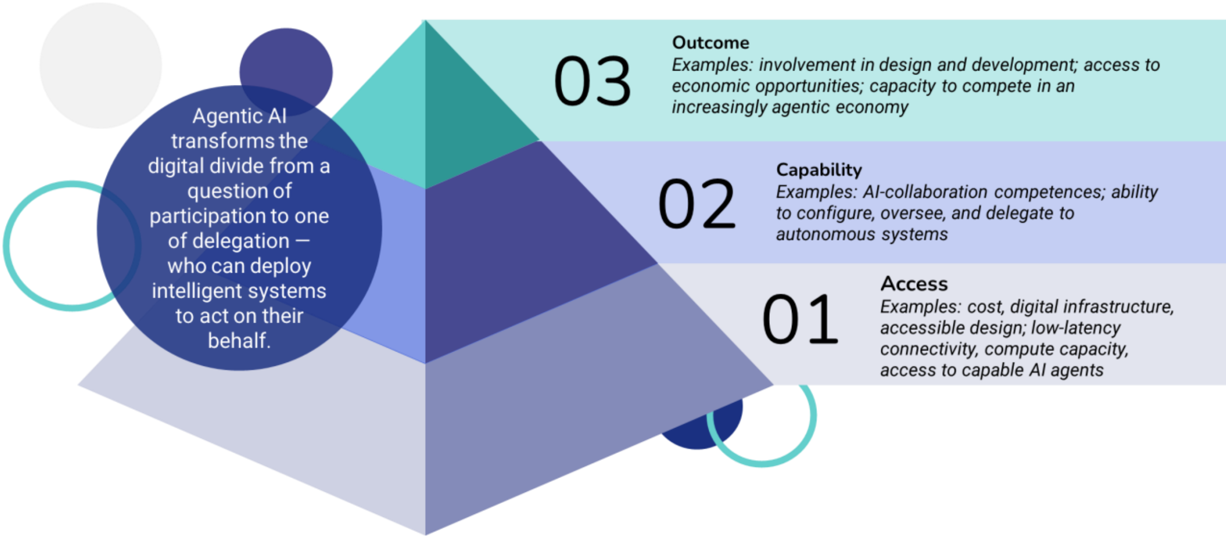


Source: dimensions – Wei et al. (2011); figures and examples – authors' own elaboration based on Barcevičius, E., et al. (2025).

**Agentic AI** does not simply add another layer to this framework — it **transforms its nature**. Where earlier digital divides concerned participation, that is, who could access, use, and benefit from digital services, agentic AI introduces a divide in delegation, or, who can deploy intelligent systems to act autonomously on their behalf, and who cannot[273]. This new form of stratification is not only about who has internet access, but about the active capacity to pursue goals using AI[274]. The inequality operates across three dimensions:

- **availability** (who has access to agents at all);
- **quality** (who has access to the fastest, most capable agents);
- **magnitude** (who can deploy coordinated swarms).

Critically, **these dimensions are self-reinforcing**. Actors with better agents can accumulate more operational data. In turn, this improves agents further and can widen the performance gap – a compounding dynamic that **connects directly to the economic concentration** discussed earlier in Section 1.1. Some experts also add a fourth dimension of divide: those who can trust their agents versus those who can't[275]. Actors with the resources to invest in low-latency private networks, advanced compute and high-quality, accurate training data will consistently outperform those relying on public internet access and commodity tools, with consequences that can extend to access to essential services such as healthcare, legal assistance, and financial products. This compounding logic also **intersects with the labour market aspects** discussed in the same section.

The remainder of this section traces how this delegation divide manifests across the access level through global and regional connectivity gaps, at the capabilities level through bias and asymmetric exposure, and at the outcomes level through the infrastructure ownership patterns that determine who captures value from the agentic transition.

The World Bank's 2025 Digital Progress and Trends report[276], stresses that AI's benefits depend on four foundational "Cs", namely connectivity, compute, context (data), and competency (skills). The **current distribution of these foundations is highly unequal**. Already, high-income countries account for 87% of notable AI models, 86% of AI startups, 91% of AI venture capital, and host 77% of global colocation data centre capacity, while representing only 17% of the world's population.

## Figure 17. Current distribution of AI foundations

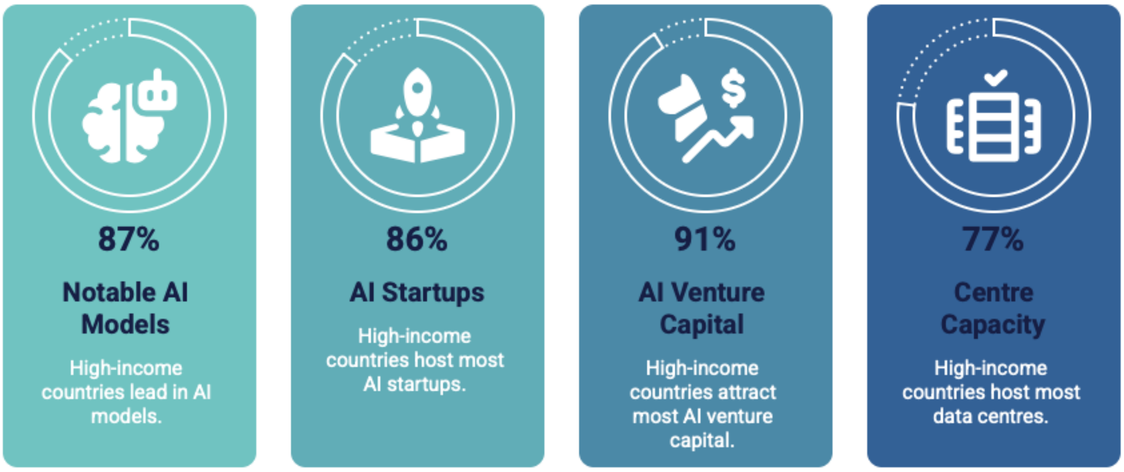


Source: authors' elaboration based on ITU (2024) and interview findings.

The connectivity baseline from which agentic AI is being deployed is deeply unequal. According to ITU, an estimated 2.6 billion people remain offline — nearly one third of the global population — with 1.8 billion of those in rural areas[277]. Internet use stands at 93% in high-income countries compared to just 27% in low-income countries, while 5G coverage reaches 84% of people in high-income economies but only 4% in low-income countries. A fixed-broadband subscription in low-income countries costs the equivalent of nearly a third of average monthly income. These conditions **matter acutely for agentic AI, which depends on low-latency, high-bandwidth connectivity to function reliably**. This makes the infrastructure gap not merely a background inequality but a direct constraint on who can deploy and benefit from agentic systems. AI uptake already stands at 27% in North America compared to just 9% in Sub-Saharan Africa, and adverse outcomes from AI arising from this uneven distribution are now ranked among the most significant long-term global risks in the WEF Global Risks Perception Survey 2025–2026, rising sharply from 30th in the two-year outlook to 5th over the ten-year horizon[278].

Within Europe, similar **disparities exist between urban and rural regions**, between member states with varying levels of digital infrastructure development, and between socioeconomic groups with differential access to quality connectivity. Satellite infrastructure, including low-Earth orbit (LEO) constellations, offers potential to bridge geographic gaps by providing connectivity to remote and underserved regions where fibre infrastructure is economically unfeasible[279]. However, the level of demand for LEOs[280] and the quality and latency characteristics of different connectivity modes can create tiered access. Actors with fibre connections and edge computing capabilities can operate fundamentally different classes of AI agents than those dependent on satellite or mobile networks with higher latency and lower bandwidth[281]. The thematic paper prepared within this project, "Agentic AI and internet infrastructure", offers a more detailed discussion on satellite and terrestrial infrastructure underpinning the agentic transition.

Roughly 97% to 99% of intercontinental internet traffic traverses submarine cable systems, and countries in the Global South depend disproportionately on these connections[282]. This is because most of the content, and increasingly, the computation that processes and delivers that content, is hosted outside the continent. Historically, submarine cables were built and owned by telecommunications consortia, or groups of national telcos who jointly funded cable construction. However, over the past decade, the **share of international subsea capacity used by four U.S. hyperscalers – Google, Meta, Microsoft, and Amazon – increased from 10% to 71%**. Google is today a part or sole owner of around 33 subsea cables; Meta approximately 15; Microsoft and Amazon a further nine between them[283]. As researchers have noted, major tech companies now possess the capacity to fund initiatives directly[284], that likely previously required a consortium for economic viability, raising significant concerns about the **concentration of digital infrastructure**. The ownership map of submarine cables thus increasingly mirrors ownership of data centres, cloud services, and AI models that run above them.

This infrastructure concentration means that as the agentic economy scales, the regions and actors that control connectivity, compute, and model access are already poised to disproportionately capture economic value, while others — despite being connected — **remain structurally dependent on services hosted, processed, and governed elsewhere**. The outcomes divide, in other words, is not primarily about whether people can use agents, but about whether the infrastructure through which agents operate is governed in their interest and whether the scaffolding, security, tool integration and other crucial layers constitute additional barriers to broader agentic use[285].

From the economic perspective, experts consulted for this project suggest that the emerging agent marketplaces are likely to develop tiered offerings, similar to existing software subscription models. More capable agents, with faster inference, greater autonomy, and broader data access, will thus become available only to those able to pay for premium infrastructure[286]. As more interactions are mediated through agents, actors without the infrastructure to run or reliably access these systems risk being de facto excluded from markets and decision-making, while reinforcing regional disadvantages[287]. Besides, the gap in agents' ability to follow context also reflects the geographic disparities, as organisations deploying off-the-shelf agents in the Global South encounter the challenge of tailoring them to local workflows, tax and payment systems, languages and other particularities[288].

Even where access exists**, biases embedded in training datasets can lead to the effective exclusion of users** who are mis-flagged, misidentified, or poorly served[289]. Facial recognition techniques, for instance, have been documented to carry significant disparities in accuracy based on users' race and gender[290,291]. As agentic systems take on high-stakes functions, like application screening, authorising transactions, or triaging healthcare queries, such biases can translate directly into unequal outcomes. An underexamined dimension of agentic inequality is **asymmetric exposure,** or where populations in lower-income regions, or with lower digital literacy, simultaneously lack access to capable agents and are disproportionately subject to agents deployed by large platforms in consequential decisions, for instance, credit scoring, price discrimination, and content moderation[292,293]. It is, thus, important to sustain effort in increasing **AI literacy** among citizens and clearly communicate how agents are used[294].

However, the outcomes described above are not inevitable. **Agentic AI can also be leveraged to reduce connectivity and service gaps** if deliberately directed. AI can analyse geographic, demographic, and traffic data to, for instance, prioritise broadband deployment, identify high-impact locations for backbone extensions or community networks, and reduce the cost per connected user[295]. Agents managing networks can detect congestion or failures, re-route traffic, and proactively schedule maintenance, improving reliability for underserved regions where outages are more frequent and repair capacity is limited[296] (see also "Agentic AI and internet infrastructure" companion paper). Additionally, the so-called "citizen agents" can help less advantaged people navigate interactions with the state and bridge the gap[297].

Whether these implications are addressed or entrenched depends on deliberate investment and governance choices. Infrastructure decisions are, as different experts put it, fundamentally political rather than technical, embedding choices about who participates, who bears costs, who governs, and who captures value in the digital economy of the decades ahead[298]. Additionally, **skills and capacity building** will play a fundamental role in the uptake[299]. In this light, Digital Skills Academies[300] can be a powerful means of strengthening the EU citizens' knowledge of agentic capabilities.

### *1.2.3. Energy and sustainability*

The energy demands of AI infrastructure represent one of the **most pressing risks of the current technological transition**, one with direct implications for decarbonisation commitments, resource sovereignty, and the competitiveness of regions that host AI infrastructure. Three major institutions converge on the same basic trajectory.

The **IMF** estimates that global electricity needs for AI could triple by 2030, reaching 1,500 TWh[301]. The **IEA**, looking from the supply side, shows data centres were already responsible for approximately 1.5% of global electricity consumption, or around 415 TWh, in 2024, and are projected to more than double to around 945 TWh by 2030, driven primarily by AI workloads[302]. In the EU, data centres already account for around 2% of electricity consumption[303]. The ITU's projects that in a scenario where billions of users each deploy multiple AI agents, **global energy demand from AI could reach on the order of several TWh per day**[304]. This is a significant share of the total world electricity consumption of approximately 80 TWh per day. To sustain this, global AI data centre capacity may need to reach 68 GW by 2027 and up to 327 GW by 2030, comparable to the entire electricity generation capacity of California[305]. In Europe alone, data centre capacity is projected to grow from 10 GW to 35 GW by 2030[306].

What makes these projections qualitatively different from earlier estimates of digital infrastructure growth is the **agentic compute multiplier**. Where a person might generate a handful of model requests per session, an autonomous agent acting on their behalf can generate hundreds within the same timeframe, at millisecond intervals and without pause[307]. The per-user demand multiplier is therefore not marginal but potentially orders of magnitude larger than baseline usage patterns suggest.

Understanding why these pressures are so difficult to engineer away requires looking beyond infrastructure planning to the underlying physics. The heat, the cooling burden, and the energy costs described above are **consequences of fundamental thermodynamic laws** that govern every computation performed, as described in further detail in Annex 4.3.5.

This shift matters not just for the overall scale of demand, but for where the energy burden falls. Training a frontier model is a discrete, energy-intensive event, but **inference**, the continuous act of running a deployed model, accumulates silently across millions of users and billions of agent interactions every day[308]. Because agentic AI operates persistently rather than episodically, inference becomes the dominant and structurally permanent component of AI energy consumption. This is compounded by a geographical constraint: inference clusters must remain close to population centres to minimise latency, meaning agentic AI concentrates energy demand precisely where grid capacity is already most constrained. The result is direct competition for power and cooling resources between data infrastructure and urban populations.

## Figure 18. Energy consumption for different AI tasks

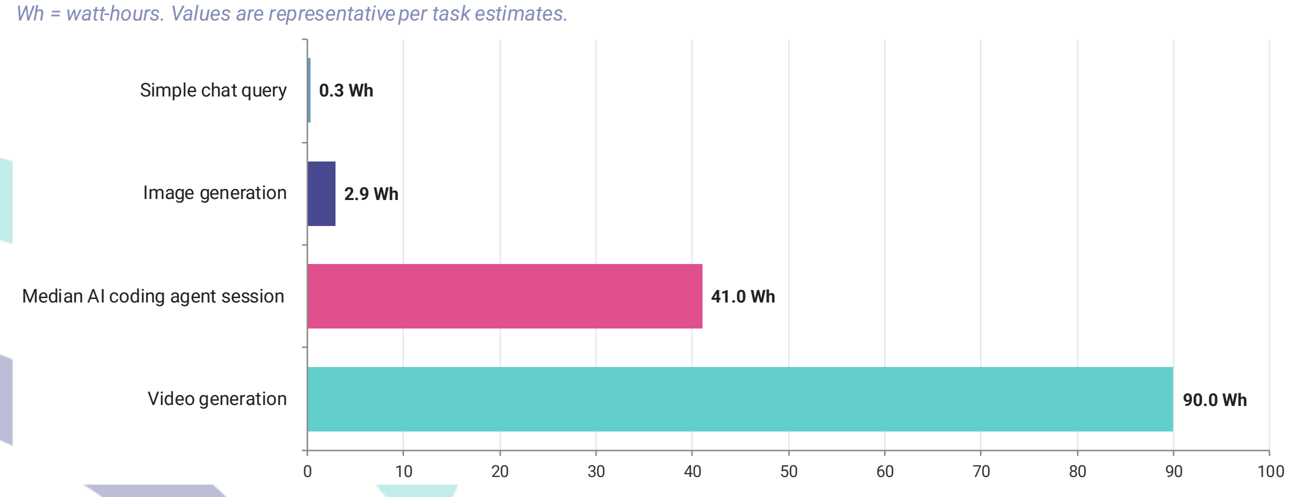


Source: authors' elaboration based on Couch, S. (2026); Chedraoui, K. (2025); Vincent, J. (2024).

The risks extend well beyond electricity consumption. The sheer scale of compute required threatens to overwhelm national decarbonisation efforts, particularly given that efficiency gains in hardware have historically been offset by rising demand. This dynamic, known as **Jevons' paradox**, suggests that more efficient AI chips may accelerate adoption and aggregate energy demand rather than relieve it[309].

Water consumption presents a parallel challenge. Large AI clusters require substantial cooling, creating direct resource competition with local communities in **water-stressed regions**[310,311]. Because low-latency requirements push infrastructure toward dense urban centres, these pressures concentrate precisely where sustainable energy and freshwater resources are already under strain[312]. EU evidence already points to large AI data centres intensifying water stress in drought-prone regions[313].

Beyond electricity and water, agentic AI intensifies demand for **critical minerals** required for hardware manufacturing, including lithium, gallium, germanium, and rare earth elements[314]. Supply chains for these materials are highly concentrated. The U.S currently imports 80% of its rare earth elements and is entirely dependent on imports for gallium, germanium, arsenic, and fluorite, all essential to AI chip fabrication[315]. China controls over 90% of global rare earth refining capacity[316]. In October 2025, it introduced stringent export controls on rare earths and magnets, directly threatening supplies of essential components for AI chips, server storage, and cooling systems[317]. This episode underscores the supply-chain fragility facing regions that host large concentrations of data centre infrastructure.

In this context, it is important to highlight the concept of "**environmentally sustainable digital sovereignty**", conceived in a recent publication by the Joint Research Centre. It is understood as the EU's ability to "exercise strategic independence in the digital domain while prioritising and aligning environmental sustainability, social responsibility, democratic values and international cooperation"[318]. Framing agentic AI within this paradigm makes clear that securing access to critical minerals and managing their environmental footprint, is not only a resilience issue for AI infrastructure, but a core test of the EU's broader claims to digital sovereignty.

Energy and grid capacity have already become binding constraints on AI expansion. **Data centres already consume approximately 21% of Ireland's electricity**[319], leading to a de facto prohibition on new connections/grid links in the Dublin area to prevent blackouts, which remained in place until late 2025[320]. The Netherlands similarly imposed a temporary ban on new hyperscale sites exceeding 10 hectares or 70 MW[321]. Grid upgrade lead times of more than five years are structurally misaligned with data centre planning cycles, and the wait for a new grid connection in the UK currently stands at 8 to 10 years[322]. These constraints are actively reshaping where AI infrastructure is built, and therefore, which regions can develop sovereign AI compute capacity at the pace the strategic situation requires. The companion paper on agentic AI and the internet infrastructure, examines in further detail the infrastructure adaptations these pressures are driving.

# 2. Policy frameworks and governance initiatives

This chapter elaborates on the **policy and governance landscape**, from the EU's regulatory stack to global initiatives, identifying both the foundations already in place and the significant gaps that remain, such as around agent-specific accountability, liability in multi-agent systems, and the pace mismatch between technological development and regulatory response.

As AI agents begin to negotiate, transact, and execute tasks across organisational boundaries, the digital ecosystem is witnessing a critical challenge to define guidelines and rules for safe use. Here, we provide an overview of the current political layer of regulation, where governments are attempting to codify the liability and safety of autonomous systems. While the technical layer of protocols and standards, where industry players are vying to establish the syntax of agent interaction, is no less important, we cover such initiatives in the subsequent paper on agentic AI, trust and architecture.

Figure 19 summarises key governance frameworks that are pertinent to AI agents and their capabilities. Annex 4.3.5 provides a comprehensive overview of the existing and emerging policy mechanisms at the EU and global levels. This section, in turn, focuses its discussion on the needs and gaps that are becoming apparent when it comes to steering AI agents.

**Figure 19. Key governance frameworks surrounding agentic AI**

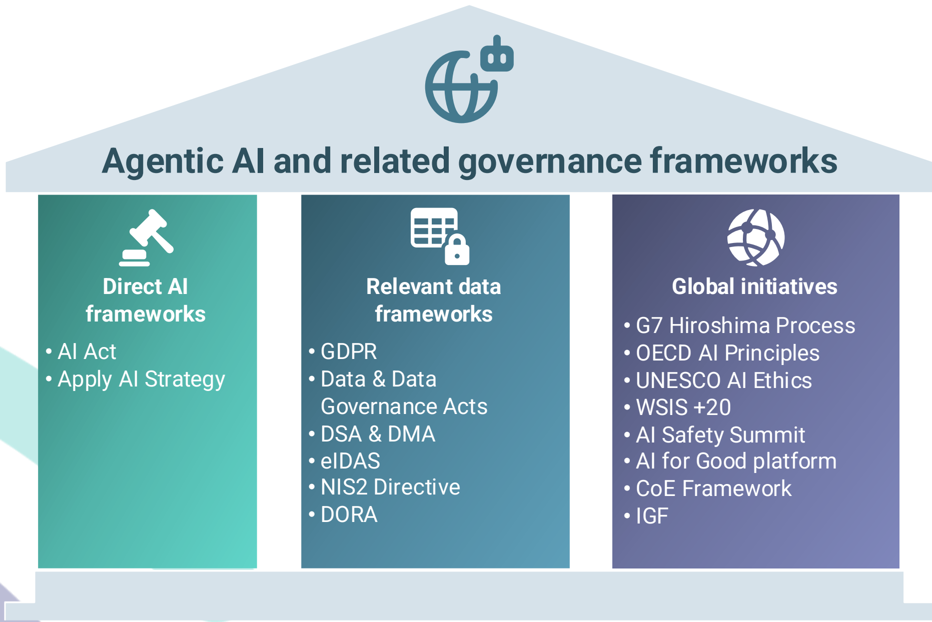


Source: authors' elaboration.

Overall, the EU's strong framework[323] around emerging technologies, support measures and experience, is often underlined by interviewed experts as a strong advantage. It helps build trust towards novel tools, test them safely and later export this **trustworthy technology with built-in safeguards**[324]. Nevertheless, several policy and governance gaps related to agentic AI persist, as shown in Figure 20 and elaborated in this section.

**Figure 20. Summary of policy and governance gaps related to agentic AI**

AGENTIC AI

**DEFINITIONS & EVIDENCE**

- **No legal definition**
  Major regulatory instruments lack agent-specific definitions, miscalibrating compliance obligations
- **Information asymmetry**
  Limited corporate disclosure and "build fast – fix later" widen the regulatory knowledge gap
- **Pace mismatch**
  Risks and capabilities evolve faster than governance cycles can adapt

**STANDARD-SETTING & CONCENTRATION**

- **Closed-off standardisation**
  Protocols shaped by incumbents in private consortia; high barriers to participation
- **Concentration of development**
  AI development in a few hands constrains democratic governance and control
- **Open-source sustainability**
  Without viable business models, open-source alternatives risk attrition to incumbents

**DECISION RIGHTS, OVERSIGHT & DATA**

- **Static role definitions**
  Provider/deployer/controller categories do not fit dynamic, delegating agent roles
- **Human oversight at machine speed**
  Meaningful human-in-the-loop control harder as autonomy and operating speed increase
- **Right to be forgotten**
  No procedural solution for erasure across distributed agentic memory architectures

**LIABILITY & ACCOUNTABILITY**

- **Diffused responsibility**
  Liability unclear across sub-agents and micro-decisions in multi-agent stacks
- **Identity beyond human users**
  Existing identity frameworks were not designed for autonomous, non-human subjects
- **Cross-jurisdictional reach**
  No framework addresses agents acting instantaneously across borders

## 2.1. Definition, evidence and standard-setting

Unlike internet governance, which evolved around a relatively coherent institutional core (ICANN, IETF, the IGF, the ITU), global AI governance has emerged **rapidly and in parallel across multiple forums, without a single coordinating architecture**. Moreover, while the initiatives listed in this section collectively address many dimensions of AI safety, ethics, and accountability, **none of them explicitly governs agentic AI as a distinct category**. While both the internet and AI governance have gone through an institutional experimentation phase before the roles are clarified between different bodies, the pace of AI evolution is faster, and the commercial stakes are extremely high from the outset, heightening the **risk that governance trails behind technological advancement**.

Moreover, the ITU highlights that the field of *agent governance* is currently "in its infancy" while at the same time, investment in building agents accelerates[325]. Such a **mismatch** makes both governments and industry poorly prepared to steer the future billions of autonomous digital actors, warns the organisation. While there is currently **no widely adopted legal definition of agentic AI** yet, the EU's Proposal for the **Cloud and AI Development Act** (CADA) was the first legal document to define AI agents: as "an AI system or a coordinated set of AI systems, that can perceive and act upon their environment, with a degree of autonomy, using tools as needed to achieve specific goals and adapt to changing inputs and contexts"[326].

In March 2026, a comparative survey analysed eleven regulatory instruments and industry frameworks, including the AI Act, OECD/G7 Principles, NIST, the UK ICO, and the European Commission's landscape report[327]. The study identified four definitional errors: anthropomorphising AI as if it reasons and plans, conflating any agentic AI with multi-agent systems, treating autonomy as a sliding scale rather than a qualitative architectural shift, and confusing database retrieval with genuine learning. The result is that **compliance obligations are often calibrated to the wrong features of the technology**, directing liability away from the developers who actually control the risks through deliberate architectural decisions.

At the same time, some experts caution against agent-specific policy formulations and emphasise regular updates to regulation[328]. Such a **technology-agnostic approach** could benefit both the EU and engineers, with more flexibility for future technology frontiers. Given current EU AI initiatives, including the AI Continent Action Plan, InvestAI, and AI Factories, remain distributed across multiple Directorates-General and Member States, the effectiveness of this approach will depend on adequate institutional coordination[329]. Consulted experts converge on the view that agentic AI does not need a governance body of its own either, and the multi-stakeholder model, including ICANN, remains fit for purpose for foundational infrastructure. There is no urgent governance gap that would necessitate new structures – rather, it is insufficiently active participation and resourcing within the existing ones[330].

The gap between practice and legislation is also acknowledged in a recent JRC brief, which states that empirical questions about legislation's effect on deployers can only be answered through implementation[331]. Policymakers have to make decisions about rapidly advancing AI without robust evidence, balancing premature or unnecessary mitigation against the danger of acting too late amid sudden capability leaps[332]. This is also complicated by an **information gap** caused by companies' limited disclosure of system details, as well as a "build fast – fix later" approach[333]. With research, development and deployment of AI concentrated in the hands of a few infrastructure owners, the possibilities for establishing democratic governance and control over these tools become especially constrained[334].

The building blocks most often identified as fundamental to address in agentic evolution are discovery, capability negotiation, agent-adapted messaging, exposed trust and security, fine-grained task-level authorisation, delegation, and auditing[335]. As several interviewees pointed out, some key standardisation efforts for agentic AI are already taking place among private companies and in their consortia, rather than in global, multilateral and multistakeholder governance mechanisms. This risks creating a hybrid system where open code and protocols are developed in **closed-off ecosystems** run by incumbents, threatening interoperability[336]. The risk is compounded by high barriers to entry in standard-setting bodies, where smaller actors and the Global South often lack the resources to participate effectively[337]. At the same time, without a sustainable business model or clear ownership, alternative open-source agentic AI initiatives risk becoming unmaintained, leaving the field open for well-funded incumbents to dictate the technology's direction[338].

## 2.2. Decision rights, oversight and data protection

The shift to agentic AI represents **not merely a technological upgrade but a transfer of decision rights**, challenging existing governance models[339]. A thematic paper, "Digital identity and identifiers in the era of Web 4.0 and virtual worlds"[340], prepared within this project, elaborates in more detail on the complexity of authorisation and delegation in agentic AI. Features of agents that add this complexity are their emergent, unpredictable behaviour, exposure to external integration through linked systems, and the diffused responsibility due to numerous micro-decisions[341]. Moreover, the existing EU legislation rests on static definitions of roles (such as provider, deployer or controller) and intended purposes in AI systems, which are inherently subject to a dynamic change in the case of AI agents[342].

The current legislative framework primarily **centres around human-mediated data governance and protection**, and while its requirements apply regardless of whether an AI system is agentic, it was not designed with the specific capabilities and behaviour of highly autonomous agents in mind — particularly those that interconnect and exchange sensitive information without the direct knowledge and control of end users[343]. One of the interviewees even argued for restraining the use of agents in critical sectors to self-hosted tools in control of EU-based users, to prevent fine-grained profiling and targeting[344]. Significant gaps also remain in terms of validating AI-generated outputs[345], ensuring explainability, tracking provenance, and mitigating bias[346]. Multiple experts also stress the importance of introducing clear requirements for ensuring **humans in the loop**. However, the EU legislative requirements for sufficient human oversight currently lack clarity, especially in their interpretation vis-à-vis agentic systems that act at machine speed[347,348].

Another point of tension emerges with the personal data AI agents cache. It remains unclear how "**the right to be forgotten**" can be enforced[349], especially if poorly managed, agentic memory can store excessive, irrelevant or obsolete personal data, as well as add new unauthorised recipients or change retention periods[350]. Agentic systems use complex memory architectures that span multiple components and jurisdictions. Responding to data subject access or erasure requests across these distributed memory systems presents challenges for which there is no established procedural solution under the GDPR. When agents' goals are set too broadly, and instructions are too thin, their actions can clash with purpose limitation and data minimisation principles.

Although the key EU legislation serves as a solid legal base, accountability is needed across toolchains, with continuous governance and focus on privacy and traceability[351]. A related gap sits even further upstream, at the root of development itself. The status of personal data once it has been absorbed into a trained model remains unclear, and so does the threshold under the AI Act and the GDPR at which synthetic data can be considered sufficiently anonymised[352].

## 2.3. Liability and accountability across multi-agent systems

The **liability and accountability** behind agents' actions constitute another widely discussed aspect of their use. In Web 4.0 environments, **authorisation must extend beyond human users** to a growing ecosystem of AI agents, digital twins, and automated services, as we discuss in further detail in the thematic paper "Digital identity and identifiers in the era of Web 4.0 and virtual worlds"[353]. Agentic AI fundamentally widens the scope of possible configurations and complicates traditional **distinctions between providers, deployers and users**[354]. A single human user may control multiple AI agents, own connected devices, and operate across different virtual worlds simultaneously. MASs and advanced AI collectives introduce a further layer of complexity, with supervisor agents managing strategic goals while delegating responsibilities to subordinate agents.

Yet, today's legislative landscape does not yet reflect these new realities. Identity for autonomous agents remains an open frontier, with early initiatives moving forward in several countries. On the EU level, while eIDAS 2.0 and the EUDI Wallet provide a robust foundation for human and entity identification, this framework was not designed to accommodate non-human subjects. Among Member States, there is a sustained public-sector effort to determine whether existing legal and technical frameworks suffice for agents acting on behalf of individuals, and what alterations are needed to capacitate the linking of agent credentials to individual identity numbers[355]. Some of the consulted experts believe that the experimentation with high-volume user journeys, especially in financial services, can help surface the architectural vulnerabilities and accelerate the development of solutions[356]. The subsequent paper, "Agentic AI, trust and architecture", as part of this trilogy, explores the identity question in further detail.

Apart from authentication and verification, consulted experts underline the importance of **maintaining a consistent and transparent log** of all agents' actions and algorithms. Agents need to be kept liable and visible, with transparency on what they do and who trained them and with clear allocated responsibility and redress[357]. Every agent operating in consequential domains, such as public decision-making and critical infrastructure, should maintain tamper-evident logs of its actions, decisions, and accessed data. At the same time, one interviewee argued against relying on one AI tool reviewing the outputs of another, as this presumes intentionality that AI does not possess[358].

Finally, assigning responsibility for security incidents becomes substantially more complex in MASs, where **actions of multiple agents can have an equally important weight in breaches or attacks**. Experts agree that this is a current grey area, as risks and vulnerabilities evolve much faster than governance, and the assignment of responsibility for decisions taken by complex autonomous systems needs clarity[359]. Agents can spawn sub-agents, delegate tasks laterally, and interact across organisational boundaries. The current regulatory framework, however, primarily assumes that only one system can be held accountable for a high-risk incident[360]. The diffusion of responsibility across model providers, platform operators, application developers, and deployers in multi-agent stacks so far has no clear resolution in regulatory frameworks. Equally, none of the existing governance frameworks addresses the fact that agents can operate **across jurisdictions** instantaneously, sharing data and triggering actions.

## 2.4. Looking ahead and closing the gaps

Taken together, these gaps point in a common direction: the governance challenge posed by agentic AI is not primarily about adding new rules to existing frameworks, but about critically examining existing frameworks, which are often built around assumptions that agentic systems systematically violate. Stable roles, bounded purposes, human-paced decisions and single-system accountability are not applicable for "swarms" of autonomous agents. Closing these gaps, however, does not imply duplicating existing instruments under an "agentic" label when the underlying provisions are technology-agnostic and reasonably applicable. Many of the challenges associated with AI agents are not entirely new – they represent an extension of broader issues already observed in AI governance[361].

The more productive path is, therefore, to identify the specific points at which agentic deployment exceeds the design assumptions of current frameworks and to address them through alternative means, such as **governance sandboxes, voluntary initiatives (e.g. codes of practice) and (pre-)standardisation work**. Consulted experts in the field overwhelmingly prefer greater legal clarity, practical guidance and implementation support over new, agentic-specific regulations[362]. Operating systems (OSs) illustrate the kind of audit this requires. Existing rules under the GDPR, the DSA and consumer protection law need to be examined to establish where responsibility for OS-level agent behaviour actually sits, and whether that responsibility should be limited to OSs that the DMA designates as gatekeepers or extend to all of them[363]. To address such mismatches, it is not a single intervention that is needed but **definitional clarity, open, multistakeholder standard-setting venues** and **institutional coordination** capable of keeping pace with technical change. Besides, evidence-generating mechanisms need to allow regulators to see inside deployed systems rather than rely on post-hoc disclosures.

The pressing pace of AI technology development can often lead to bespoke, one-off solutions, with standardisation efforts coming in later to address interoperability issues. Experts consulted also converged on a clear preference for **evolving existing protocols** rather than reinventing them. This would allow to build on foundations already laid by bodies such as the IETF and direct new work at **genuine gaps** instead of duplicating effort[364]. An equally important aspect is **interoperability and data portability**, which are fundamental to agents' ability to interact across platforms and services[365]. Timing reinforces the case for moving quickly on this technical track: once large providers consolidate their positions, developers naturally gravitate towards the most accessible solutions. Open technical solutions therefore need to be deployed early, to prove themselves scalable, low-latency and workable before the market settles on proprietary alternatives. Apart from standards, access conditions granted in licenses and contracts are also important determinants of agents' accessibility and compatibility[366].

Some of the experts stress that open standards ultimately prevail when they leverage **transparent, broadly participatory processes**, not through mandating adoption[367]. While some of the experts hold an opposing view, sceptical of the effectiveness of voluntary mechanisms, they, too, acknowledge that any obligatory rules have to stay flexible and forward-looking, applicable to the eventual A2A economy. Whichever side of this debate proves right, one point cuts across both: **users need to sit at the centre of co-creation**, otherwise solutions risk failing to scale. Rules alone, whether voluntary or binding, are not sufficient to build consumer-facing trust. Common terminology and potentially, standard contracts, are necessary to prevent dark patterns in which users effectively sign away their autonomy without realising it, echoing the past experiences with cookie consent[368].

When it comes to the EU' role in governance processes, its commitment to trustworthy and responsible (agentic) AI is already a substantial contribution to the international debate. The interviewed experts underlined that the Union's reputation allows it to play an active part in shaping safety guardrails, developing benchmarks and metrics to test the robustness of AI agents, and building trust in novel solutions[369]. All while working through existing multi-stakeholder processes rather than setting these terms alone. Nonetheless, it will be important to translate this commitment into a **competitive and viable model** with credible implementation results – not just a regulatory choice[370].

Alongside these regulatory and institutional responses, **initiatives and good practices** are being shaped across industries. In July 2025, the Coalition for Secure AI, with contributions from leading companies, presented "Principles for Secure-by-Design Agentic Systems", targeting agent developers, adopters, and security engineers[371]. Another example is TrustX, a partnership between several UK organisations in the health and care sector that joined efforts to verify, deploy, and test agentic AI safely[372]. Moreover, a complementary strand of work is emerging from the technical and research community, focused on the engineering practices and system-level capabilities that would make agentic AI safer *by construction* rather than by retrofit. A community of practice, "Safer Agentic AI", presented their vision of key drivers: positive practices and capabilities that, if widely promoted, can positively impact the development of safe agentic systems. These are summarised in Figure 21.

### Figure 21. Foundations of safe agentic AI systems

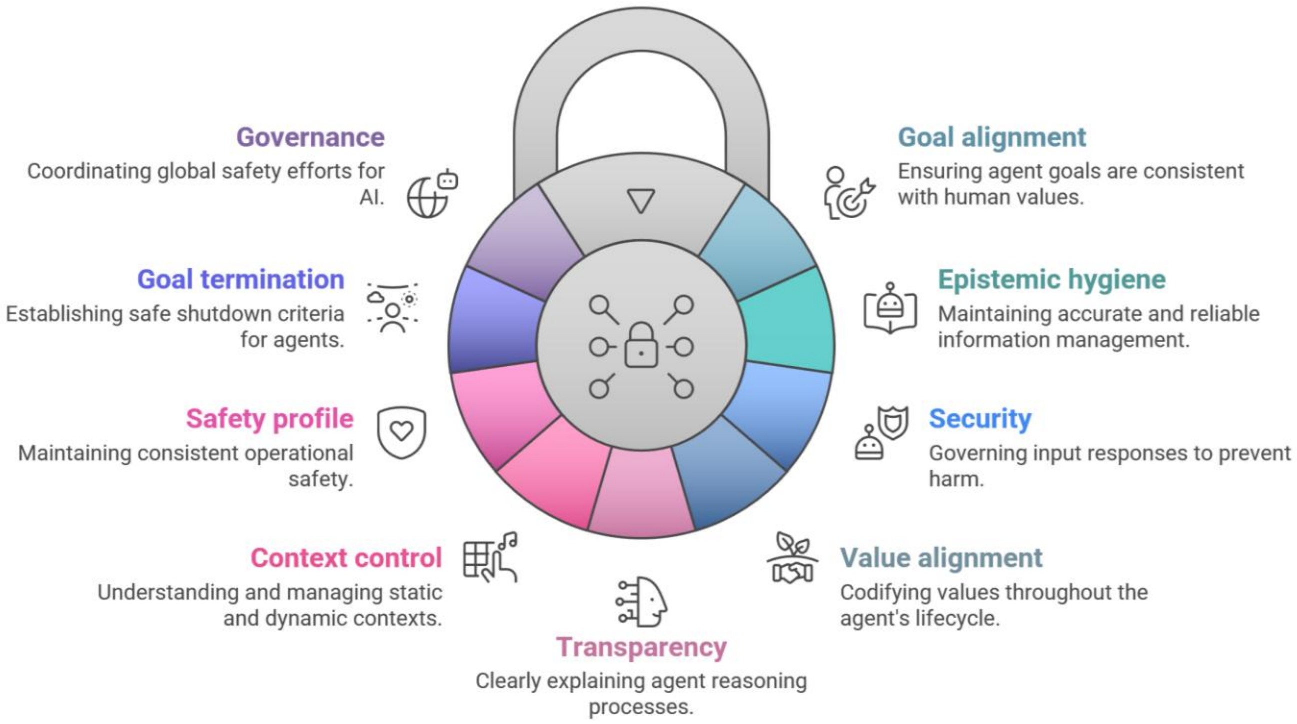


Source: authors' elaboration based on Safer Agentic AI. (2025).

# 3. Key findings

This section summarises the key findings from the analysis of the agentic web transition. The paper addressed this core research question: What are the socio-economic and governance implications of the proliferation of autonomous AI agents? It also examined the essence of agentic AI, the key forces driving its adoption, and how existing and emerging governance frameworks apply to the agentic web. Besides, the paper examined where the EU stands across the agentic AI stack, and how its position could evolve by 2035.

Across these dimensions, four threads run through the analysis. The first concerns the gap between what agentic AI can deliver for the public good and what the commercial logic shapes it to deliver. The second concerns the depth of the structural transformation it brings to labour markets, value capture and the broader digital economy. The third concerns the fit between the technology's autonomy and the governance frameworks it is encountering. The fourth concerns Europe's positioning to shape, rather than inherit, the trajectory of agentic AI. The findings below take each of these threads in turn.

While these key findings close this paper, they do not close the discussion of agentic AI within this project. The two thematic papers that follow in the trilogy take up specific dimensions in greater depth: one examines the implications of agentic traffic for the infrastructure backbone, and the other explores how agent interactions reshape trust frameworks and security architectures. Each builds on the analytical foundations set out here.

### *Finding 1: The benefits of agentic AI are real, but their distribution is a policy choice rather than a property of technology*

Agentic AI is generating **genuine value across multiple domains**. In healthcare, climate, and scientific research, agents are enabling discoveries and efficiencies that were previously out of reach. In commercial sectors, they are reshaping enterprise software, customer interaction, and operational workflows in ways that meaningfully increase productivity. In public administration, they hold the potential to improve service delivery and citizen interaction. **These benefits are not in dispute.** What is at stake is *who* captures them, *who* bears the costs, and *who* is left without a meaningful seat at the table where those questions are being answered.

The trajectory of agentic AI is shaped overwhelmingly by commercial priorities. These often deliver public value, but market incentives alone will not reliably steer the technology toward the areas of greatest societal need. Global investment in agentic AI for societal or scientific applications remains minuscule compared to funding flowing into commercial use cases[373]. Moreover, **the value of AI today is concentrated geographically and corporately** in a handful of firms that simultaneously dominate the cloud, model, and standard-setting layers. As agents become the primary intermediaries through which users access services, markets, and information, the infrastructure layer itself becomes a site of significant economic power. Businesses that cannot afford premium agentic infrastructure risk exclusion from markets increasingly mediated by autonomous systems, while the populations least equipped to deploy agents on their own behalf are simultaneously the most subject to decisions made by agents deployed by others.

If unchallenged, these asymmetries add a new and underestimated dimension to the **digital divide**, which existing connectivity-focused frameworks were not designed to address. Within standard-setting bodies, smaller actors and the Global South often lack the membership fees, resources, and specialised expertise required to participate effectively, meaning the rules governing how agents interact are written predominantly by those who can afford to be in the room.

None of these patterns is intrinsic to the technology. They reflect choices: about where capital flows, which use cases are funded, who participates in governance, how local communities are compensated, what conditions are attached to public investment, and how skills development is resourced. Distributional outcomes cannot be left to market dynamics that already concentrate gains; they must be designed into the policy architecture from the outset. This means deliberately resourcing public-interest deployment of agentic AI alongside commercial uptake, embedding equitable access provisions into procurement and infrastructure investment, and lowering the barriers to participation in the multilateral and multistakeholder venues. Besides, recent JRC work on environmentally sustainable digital sovereignty underscores public procurement as a lever for prioritising low-resource, low-footprint digital technologies[374]. Whether agentic AI narrows or deepens existing inequalities is, ultimately, a question of policy design rather than technological inevitability.

### *Finding 2: Agentic AI will transform digital markets and the broader economy*

Agentic AI is not simply another efficiency-enhancing technology layered onto existing economic structures. It is **reshaping how work is divided between humans and machines, how markets are organised, and how value is captured across the digital stack.** Where genAI augmented human tasks, agentic AI is already performing them. Where digital platforms intermediated between users and services, agents now intermediate between users and platforms. The result is a structural shift felt across labour markets, consumer markets, enterprise software, and the underlying infrastructure of the digital economy.

The labour-market implications are already visible and will intensify. Agents are increasingly taking on substantial portions of knowledge work across domains such as customer service, software engineering, legal research, finance, and administration. Productivity gains for those who work alongside capable agents will be real, but so will displacement pressures for those whose tasks are most readily automated. The resulting labour reconfiguration will require active policy responses to skills and training. The **mismatch is already considerable**: in the EU, 70% of enterprises that have explored AI adoption cite a lack of relevant expertise as their primary barrier[375]. The lack of expertise and talent mismatch affects the entire stack, from grids and data centres[376], to advanced analytics and cybersecurity[377]. University programmes, vocational training, and reskilling pathways have not yet adapted to a market in which the ability to operate agents is becoming a core competence rather than a specialist one[378].

The transformation of digital markets themselves is equally consequential. As agents become the primary interface through which users access services, information, and commerce, **the locus of value capture is shifting up the stack**. Compute providers and hyperscalers capture the largest share of near-term revenue, but the most durable foundations for long-run value creation are increasingly found in the application and data layers. Vertical agents purpose-built for healthcare, finance, legal, manufacturing, and other regulated sectors are emerging as the strongest candidates for sustained competitive position[379], with proprietary workflows and contextual data forming a new kind of protective moat. At the same time, growing dependence on dominant providers and model developers raises concerns about market concentration and the ability of European firms to scale. In this context, instruments like the DMA are important to ensure fair and non-discriminatory access for EU-based agentic solutions[380].

Whether European firms convert their depth in industrial and regulated-sector applications into a meaningful share of this value will depend on choices being made now. These include public investment in securing competitive access to compute and frontier models, the rules under which European industrial and scientific data become accessible for training and fine-tuning, and the speed at which education and reskilling systems produce the engineers, supervisors and auditors a deployed agentic stack will require.

### *Finding 3: Effective governance of agentic AI will depend on adaptability rather than speed of legislation*

Agentic AI introduces **governance challenges that differ in kind, not just in degree,** from those posed by earlier technological breakthroughs. The defining issue is autonomy: agents perceive, plan, decide and act on behalf of users, often through chains of sub-agents that operate across organisational and jurisdictional boundaries. This produces new challenges. Liability diffuses across model providers, platform operators, application developers and deployers. Accountability becomes harder to allocate when many micro-decisions are taken at machine speed. Authorisation must extend to non-human actors, and oversight must be meaningful in environments where humans cannot realistically follow each step. Existing frameworks were designed for systems with stable roles, bounded purposes and human-paced decisions. Agentic systems, especially multi-agent ones, violate those assumptions.

Nevertheless, many provisions in the current legal and regulatory stack are technology-agnostic and remain reasonably applicable. Layering parallel agent-specific rules on top of them risks fragmentation, slower implementation and reduced flexibility for whatever technology frontier comes next. The preferred response is **to identify the specific points at which agentic deployment exceeds the design assumptions of current frameworks and to address them through adaptable instruments**. Regulators need to be able to see inside deployed systems rather than rely on post-hoc disclosures. Standard-setting venues need to be open to voices beyond well-funded incumbents, and institutions need the capacity to keep pace with technical change.

The lessons of internet governance are instructive here. The internet evolved around a relatively coherent institutional core, with bodies such as ICANN, the IETF, the IGF and the ITU emerging over time to handle distinct functions through layered, multistakeholder arrangements. Global AI governance has emerged more quickly, in parallel across many forums, without an equivalent coordinating architecture. The pace of agentic AI development and its commercial stakes mean that the **institutional experimentation phase cannot extend indefinitely**. Europe is well placed to contribute to this process. Its regulatory tradition, experience with open standard-setting and commitment to human-centric digital values, are genuine assets. Translating that contribution into competitive and credible implementation, rather than rule-making alone, will determine whether Europe shapes the agentic web or follows the path set by others.

### *Finding 4: The EU's competitive standing in agentic AI is constrained by infrastructure gaps, and addressing them is inseparable from digital sovereignty*

The EU's competitive position in agentic AI is constrained not primarily by research quality or regulatory ambition, but by physical and financial infrastructure. Gaps in compute capacity, growth-stage capital , and frontier model development are deeply interdependent. Given that frontier training compute grows 4-5 times per year[381], the Union's limited compute capacity constrains the development of frontier models, and the lack of competitive electricity prices reduces the likelihood that EU-based data centres will host genAI applications and services[382]. Because agentic systems are built directly on top of frontier models and rely on the same compute and cloud substrate, **weaknesses in the broader AI and genAI position translate into a weaker position in agentic** AI as well, as the agent layer inherits the constraints of the layers beneath it. The absence of frontier models weakens the EU's ability to shape the agent layer through which economic value is increasingly captured. These issues are exacerbated by the insufficient and slowed commercialisation of EU-built agentic solutions[383]. Dependence on foreign hyperscalers for cloud infrastructure, model access, and submarine cable capacity means that even genuine strengths in research, talent, and regulatory credibility are partially undermined by the infrastructure foundations on which they must operate. Without sufficient domestically financed infrastructure and regulatory and fiscal reforms to **deepen capital markets**[384], Europe risks a position in which it consumes AI services developed, hosted, and governed elsewhere — the dynamic captured by the "Leased sovereignty" scenario in Annex 4.2.

At the same time, the EU possesses infrastructure-relevant **assets**. Supply of critical components (notably ASML and a strong market share in goods needed for AI semiconductor production[385]; see companion paper on "Agentic AI and internet infrastructure" for a further discussion on semiconductor value chain), Europe's first exascale supercomputer (JUPITER), 19 interconnected AI Factories, and the RAISE initiative, which pools sovereign compute for scientific institutions – all these represent meaningful foundations that gradually help address existing bottlenecks[386]. Edge infrastructure, where European actors hold a stronger position due to domain expertise and proximity to industrial deployment, represents a strategically important layer that does not require replicating the full hyperscaler compute stack. As the AI in Science Summit 2025 emphasised, the formula for competitiveness is frontier algorithms plus infrastructure. The EU's strong position in the **application layer, supporting services, trustworthy and domain-specific AI deployment** is a genuine asset too. The challenge is that these assets remain fragmented, underscaled relative to the pace of the transition, and insufficient on their own to close the gap in training compute, inference capacity, and protocol governance. Experts, therefore, believe that the EU needs to make **selective bets**, choosing a deliberate portfolio of high-risk yet high-return research directions for frontier AI, such as safety-critical sectors, industrial use cases, world models, and strategic, deep partnerships with leading enterprises[387].

Addressing these gaps is also inseparable from the broader question of **digital sovereignty**, understood as the capacity to intentionally develop, deploy, and govern AI systems in accordance with European values while reducing reliance on non-EU providers[388]. A Europe that accesses powerful agentic AI services through foreign-controlled infrastructure retains regulatory sovereignty on paper while ceding operational sovereignty in practice. The most credible path forward lies neither in attempting to replicate the full hyperscaler stack nor in accepting a permanent dependency, but in a selective, pragmatic approach with the right mix of public, private, and foreign ownership. This means prioritising sovereign capability at the layers that matter most and ensuring that Europe's human-centric values become a genuine competitive differentiator rather than a constraint on others. It is also about pursuing strategic partnerships where full self-sufficiency is neither feasible nor necessary, unlocking benefits for interoperability, scalability and technology diffusion[389]. More distributed regional capacity for low-latency inference and isolated facilities for developing high-security cases can also substantially benefit Europe's AI sovereignty[390]. Moreover, effective interoperability and transparent, non-discriminatory access rules help preserve the **open internet for users**, letting them freely choose their agents and the services they let them use[391]. As the four scenarios in Annex 4.2 illustrate, the right posture depends significantly on how much control the EU exercises over the elements of the agentic AI stack (see Sections 1.1.4 and Annex 4.2).

# 4. Annexes

## 4.1. Methodology

In the preparation of this trilogy of papers, the project team used a combination of **desk research and literature review, semi-structured stakeholder interviews, and scenario analysis** to examine the development of agentic AI, its implications for society, economy, governance, and the internet infrastructure and architecture.

This paper builds on the findings of earlier project deliverables and activities, including the **Global Multistakeholder High Level Conference on Governance for Web 4.0 and Virtual Worlds**, hosted by the European Commission and the Polish Presidency of the Council of the European Union on 31 March–1 April 2025, as well as its background and outcome documents[392] , and the stakeholder consultations carried out in the run-up to it. It also considers the findings of thematic papers developed earlier in the project covering topics, such as decentralised data and service architectures and digital identity.

The methods deployed during the preparation of this paper include:

- **Desk research and literature review**: the project team conducted extensive desk research drawing on the most recent academic literature, grey literature, and policy sources to analyse trends, define key concepts, and assess the current state of the art in technological, policy, and standardisation terms. Given the rapidly evolving nature of agentic AI, the research drew on a broad range of sources, including peer-reviewed journal articles and conference proceedings; reports and strategic publications from international organisations and research institutions; European Commission strategies, legislative texts, and standardisation roadmaps; white papers from key industry actors; expert discussions at relevant forums; and foresight studies addressing emerging and low-maturity technologies. The desk research process was conducted iteratively, with initial findings informing the refinement of focus areas and the structure of subsequent analysis.
- **Interviews and discussions**: the project team conducted 15 semi-structured interviews and discussions with stakeholders representing a diverse range of perspectives on agentic AI and its implications, including participants from academia, industry, civil society, international organisations, and the public sector. Stakeholders brought specialised knowledge in areas such as AI and data governance, internet architecture, network infrastructure, decentralised technologies, digital innovation policy, and emerging technology strategy. The interviews[393] were conducted between 2 February 2026 and 26 June 2026 and complemented and validated the findings of the desk research. The insights gathered contributed to identifying key trends shaping the development of agentic AI and its implications for internet infrastructure, assessing the associated risks and opportunities, and informing policy recommendations targeted at European decision-makers.

### Table 2. List of interviewed organisations

| Organisation | Date |
|---|---|
| Leiden University | 2 February 2026 |
| Constellation Research | 3 February 2026 |
| Dublin City University | 5 February 2026 |
| Internet Society | 10 February 2026 |
| Ericsson | 16 February 2026 |

| | |
|---|---|
| NLNet | 23 February 2026 |
| BGB Group | 23 February 2026 |
| Canadian Forces Morale and Welfare Services | 24 February 2026 |
| University of Sheffield | 13 May 2026 |
| Ayima | 26 May 2026 |
| Petrichor Labs | 8 June 2026 |
| IBM | 9 June 2026 |
| Black Forest Labs | 15 June 2026 |
| DeepL | 19 June 2026 |
| Infoblox | 26 June 2026 |

- **Horizon scanning**: structured horizon scanning was implemented across the WEB4HUB project, in which over 380 signals, including news, publications, policy initiatives, expert commentary, and use cases, were analysed across the project's thematic areas. The scan provided continuous evidence across the papers and was used to identify emerging trends throughout the analysis.
- **Expert roundtable**: on 2 July 2026, the project team organised a roundtable gathering European Commission representatives and experts in agentic AI, decentralised technologies and internet governance. The roundtable aimed to exchange perspectives on the future of Web 4.0 and identify strategic priorities for Europe in the emerging Web 4.0 and agentic AI landscape. A total of 29 participants joined the roundtable. **Scenario[394] analysis**: four structured scenarios were developed to explore how uncertainties surrounding agentic AI and Europe's strategic posture could shape future outcomes. The scenarios are constructed along two critical axes[395]:
  - **Degree of EU sovereignty:** ranges from European-built infrastructure, substantial public-private investment, and domestically governed standards, to dependence on foreign hyperscalers, externally set protocols, and limited domestic investment.
  - **AI for good vs. commercially-driven AI:** the tension between the agentic AI ecosystem developing primarily for proprietary commercial products and the focus on openness, interoperability and agents that broadly serve society.

  The resulting four scenarios (*Europe's moonshot, Europe Inc., Leased sovereignty,* and *Borrowed commons)* represent internally consistent combinations of AI performance, infrastructure development, investment timing, and governance choices. They are not forecasts, but analytical devices used examine trade-offs and assess how different policy decisions may influence Europe's economic resilience, technological autonomy, and capacity to generate public value. In the methodology, scenarios function as tools for structured exploration of uncertainty rather than as preferred or normative outcomes. The scenario construction built on the evidence base developed through horizon scanning, literature review and desk research, as well as interviews covering socio-economic drivers, commercial applications, investment, EU strategic positioning, infrastructure requirements, and governance. It was further informed by external foresight work, including the OECD's exploration of possible AI trajectories through 2030[396], the project's preceding deliverables, and probabilistic forecasting generated through The Forecasting Machine[397], an AI-assisted strategic foresight platform.

The preparation of this paper was supported by the use of generative AI. This was used for tasks such as generating ideas, fine-tuning text, preparing figures, etc. All contributions generated by generative AI have been thoroughly reviewed, edited and integrated by the authors, in combination with their own research and insights. The authors take full responsibility for the final content and conclusions presented in this paper.

## 4.2. Scenario analysis

This annex presents four scenarios that confront the EU with the hard choices it faces regarding the development model for agentic AI and the investment it is willing to make. The scenarios are structured along two axes:

- **Axis 1 — Degree of EU sovereignty**: How much control does the EU exercise over the protocols, standards, infrastructure and investment that underpin the agentic AI ecosystem? This ranges from high sovereignty (European-controlled infrastructure, substantial public-private investment, domestically governed standards) to low sovereignty (dependence on foreign hyperscalers, externally set protocols, limited domestic investment).
- **Axis 2 — AI for good vs. commercially-driven AI**: Does the agentic AI ecosystem develop primarily around openness, interoperability and publicly accessible agents that serve broad societal benefit, or does it consolidate around proprietary commercial products where capability is concentrated among those who can afford them? This axis captures the tension between public-interest AI (open protocols, public agents, equitable access) and commercially driven AI (closed ecosystems, premium services, market-determined access).

### Figure 22. Summary of the scenarios

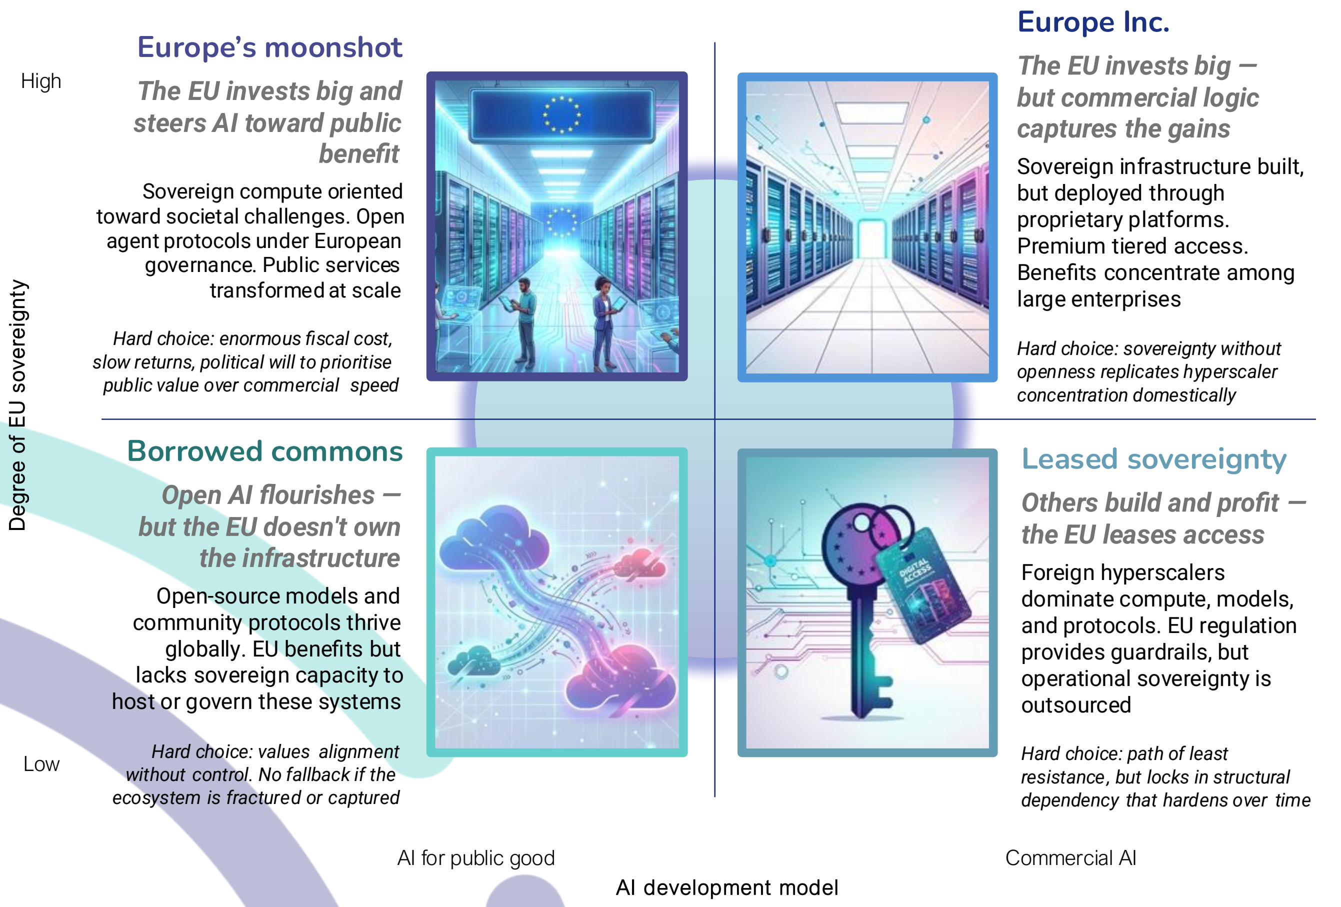


Table 3 highlights the key characteristics of each scenario.

### Table 3. Key characteristics of the four scenarios for agentic AI

| Dimension | Europe's moonshot | Europe Inc. | Leased sovereignty | Borrowed commons |
|---|---|---|---|---|
| **EU sovereignty** | High — sovereign infrastructure, | High — sovereign infrastructure, large-scale investment | Low — dependence on foreign hyperscalers, externally set protocols, | Low — limited domestic infrastructure, reliance |

| Dimension | Europe's moonshot | Europe Inc. | Leased sovereignty | Borrowed commons |
|---|---|---|---|---|
| | substantial public-private investment | | limited domestic investment | on foreign platforms, incremental investment |
| **AI development model** | Public interest — AI Factories and gigafactories oriented toward societal challenges; open protocols for agent interoperability; publicly accessible agents | Commercial — European champions build sovereign capacity but deploy it through proprietary platforms, premium tiered access, market-driven logic | Commercial — hyperscaler-dominated infrastructure; proprietary agent platforms; capability concentrated among those who can afford it | Public interest — open-source models, community-driven protocols, equitable access aspirations, but limited sovereign capacity to sustain them |
| **Agentic AI role** | Agent-to-agent workflows in public services (healthcare, transport, administration) and industry; European-governed agent identity frameworks | Autonomous workflows concentrated in large enterprises and well-resourced sectors; agents optimised for commercial productivity | Tightly coupled to hyperscaler ecosystems; EU actors as consumers of foreign-hosted agent services | Domain-specific agents built on open-source models; incremental deployment; limited scale due to infrastructure constraints |
| **Key strategic question** | Is the EU willing to sustain the fiscal and political cost of building sovereign infrastructure oriented toward public value rather than commercial returns? | Can sovereign investment avoid reproducing domestically the same concentration that the EU criticises in foreign hyperscalers? | Is the EU prepared to accept that the path of least resistance locks in permanent structural dependency on foreign-controlled infrastructure? | Can an open, public-interest AI ecosystem survive without the sovereign infrastructure to host, govern and guarantee its continuity? |
| **AI for social good** | Strong public-sector transformation; healthcare, transport, administration upgraded at scale; equitable access by design | Commercially viable public-sector applications deployed, but access stratified by ability to pay; digital divide risks within Europe | High-quality services delivered but dependent on foreign-controlled infrastructure; EU regulation provides guardrails but not operational control | Open tools improve services incrementally; high public trust but limited scale; dependent on global open-source community for continued development |

### 4.2.1. Scenario 1: "Europe's Moonshot"

The EU makes a generational investment in sovereign compute and connectivity infrastructure and deliberately steers the agentic ecosystem toward societal benefit. AI Factories and gigafactories are oriented toward public challenges — healthcare, climate, administration — alongside industrial applications. European-governed agent identity and trust frameworks set the standard.

The hard choice Europe is facing in this scenario is that this demands fiscal commitment at a scale the EU has rarely sustained in technology, political will to prioritise public value over commercial speed, and acceptance that returns will be slow and uncertain. If the investment is not matched by organisational capacity and deployment readiness, sovereign infrastructure risks becoming expensive and underutilised.

**What would need to happen:**

- The EU mobilises large-scale public-private investment before 2030, with AI Factories and gigafactories operational and oriented toward public-sector and scientific workloads, not only commercial ones.
- European-led agent identity and interoperability standards gain international traction, creating a trust framework that attracts global adoption.
- Energy infrastructure expands fast enough to support sovereign AI compute without undermining decarbonisation commitments.
- Political consensus holds across Member States for sustained investment with backloaded returns, resisting pressure to redirect funds toward more immediately visible priorities.
- Talent outflow reverses: competitive research environments, credible research-to-deployment pipelines, and startup ecosystems keep European AI specialists on the continent.

- Public-sector institutions develop the organisational capacity to deploy agentic systems meaningfully, not just procure them.

### 4.2.2. Scenario 2: "Europe Inc."

The EU invests heavily in sovereign infrastructure — and succeeds. European neoclouds, sovereign cloud providers and industrial champions build substantial domestic compute and connectivity capacity. But the ecosystem develops along commercial lines: proprietary agent platforms, premium tiered access, and market-driven deployment dominate. Sovereignty is achieved but its benefits concentrate among large enterprises and well-resourced actors.

The hard choice facing the EU in this scenario is that the EU must decide whether sovereignty without openness is worth the investment. Building a European version of the same power concentration that it criticises in US hyperscalers may protect against foreign dependency, but deepen inequality within Europe.

**What would need to happen:**

- Large-scale sovereign investment proceeds as in 'Europe's moonshot', but without deliberate steering toward public benefit or open interoperability. This may also be coupled with efforts to make the EU particularly attractive for attracting AI-based businesses.
- European commercial actors — neoclouds, industrial platforms, telecoms — capture the bulk of the new infrastructure and deploy it through proprietary offerings.
- Agent interoperability standards fragment along commercial lines, with dominant European platforms setting de facto standards that favour their own ecosystems.
- Public-sector and SME access to sovereign infrastructure remains secondary to enterprise contracts, reproducing domestically the tiered access that characterises hyperscaler markets.
- The regulatory framework focuses on sovereignty and competitiveness metrics (market share, infrastructure ownership), as well as support for commercial innovation, such as via start-ups and scale-ups, while less focus is given to ensuring equitable access or public-interest deployment.

### 4.2.3. Scenario 3: "Leased sovereignty"

Low EU sovereignty + Commercial AI. The EU does not build sovereign infrastructure capacity at the scale required, and the agentic ecosystem develops along commercial lines dominated by foreign hyperscalers. European citizens and businesses access powerful AI services, but those services are hosted, processed and governed by non-European entities. EU regulation provides meaningful guardrails on safety, transparency and fundamental rights — but operational sovereignty over the digital systems on which the economy and public services depend is effectively outsourced.

This is the path of least resistance, a continuation of current trends requiring no politically difficult investment decisions. But it locks in structural dependency that will be progressively harder to reverse as agentic systems become embedded in critical functions.

**What would need to happen:**

- No geopolitical trigger proves severe enough to force the investment step-change required for Scenarios 1 or 2. Consensus-based EU decision-making operates as usual; incremental investment through existing instruments continues.
- The venture capital gap persists. Promising European AI startups continue seeking American financing, often relocating or being acquired.
- Hyperscalers physically expand in Europe because it is a large market, but infrastructure remains foreign-governed.
- Energy costs remain structurally high for European actors while hyperscalers negotiate preferential access.
- Agentic capabilities become tightly coupled to proprietary model ecosystems and vertically integrated cloud platforms controlled outside Europe.

- The "Brussels Effect" in AI regulation proves weaker than hoped: leading AI 'superpowers' do not follow EU's example, and large foreign companies comply minimally and lobby against stricter regulation in the EU.
- A geopolitical crisis — supply chain disruption, trade dispute, unilateral export controls — could suddenly restrict European access to critical AI infrastructure, with cascading effects on dependent public services and economic functions.

### 4.2.4. Scenario 4: "Borrowed commons"

The global agentic ecosystem develops along public-interest lines — open-source models flourish, community-driven protocols gain traction, and equitable access to AI agents improves worldwide. The EU benefits from and contributes to this open ecosystem, and European values of transparency, interoperability and public benefit are reflected in the tools available. But the EU lacks the sovereign infrastructure to host, govern or guarantee the continuity of these systems domestically. European actors use open tools built and maintained elsewhere.

However, openness without sovereignty means that the EU cannot guarantee that the open ecosystem will remain open, that critical services built on it will remain available, or that governance decisions made by global communities will continue to reflect European priorities.

**What would need to happen:**

- Open-source AI development accelerates globally, with smaller specialised models, federated architectures and community-driven protocols reducing the dependency on hyperscalers for many agentic applications.
- The EU contributes meaningfully to open-source research and standards but does not invest in sovereign infrastructure at scale. The incremental investment approach continues.
- European researchers and institutions pioneer efficient architectures — smaller models, neuromorphic computing, novel approaches — that require less raw compute, playing to Europe's research strengths rather than its capital constraints.
- Protocol development proceeds through global standards bodies toward convergence rather than consolidation by a few dominant stakeholders.
- Public trust in AI is higher in Europe than in economies where commercial over-deployment has damaged confidence. Safe, efficient, domain-specific European AI emerges as a quality differentiator.
- If the global open-source ecosystem fractures, is commercially captured, or is disrupted by geopolitical events, the EU has no fallback sovereign capacity.

## 4.3. Supporting information and explainers

### 4.3.1. Agentic architectures and multi-agent systems

The architecture of agentic AI typically includes a memory module for long-term and short-term context, a planning module (often an LLM or a small language model, SLM), and an action module comprising APIs and system integrations (tools) that enable agents to interact and manipulate environments[398]. In simple terms, agents have a so-called "scaffolding" software built around an LLM that acts as an interface between the model and the external world, generating prompts (instructions) for the model and converting its outputs into interactions with external systems – both real and virtual ones[399]. Figure 23 below provides a schematic illustration of this scaffolding.

## Figure 23. Depiction of an LLM-Plus-Scaffolding style agent

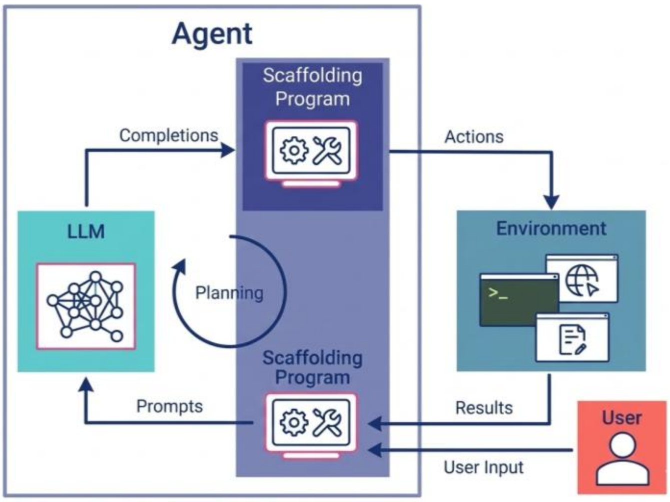


Source: Adapted from Kinniment et al. (2023).

In complex multi-agent systems, AI agents closely interact, negotiate, collaborate, and transact to fulfil human requests. Due to their parallelism, they can perform multiple tasks separately and merge the results afterwards, cross verifying each other's outputs. Experts anticipate that in Web 4.0, "supervisor" agents will be managing strategic goals and delegating responsibilities to subordinate agents, with MASs operating as "collectives" or "hiveminds"[400,401,402].

MASs can be configured in different ways, depending on a user's needs and resources. Figure 24 represents the five agentic architectures conceptualised by the Google Research Team:

- Single-agent SAS: a solitary agent executing all reasoning and acting steps sequentially.
- Independent MAS: multiple agents working on sub-tasks in parallel without communicating, aggregating results only at the end.
- Decentralised MAS: a peer-to-peer mesh where agents communicate directly and reach consensus.
- Centralised MAS: a central orchestrator delegates tasks and synthesises outputs.
- Hybrid MAS: a combination of hierarchical oversight and peer-to-peer coordination.

## Figure 24. Types of agentic architectures

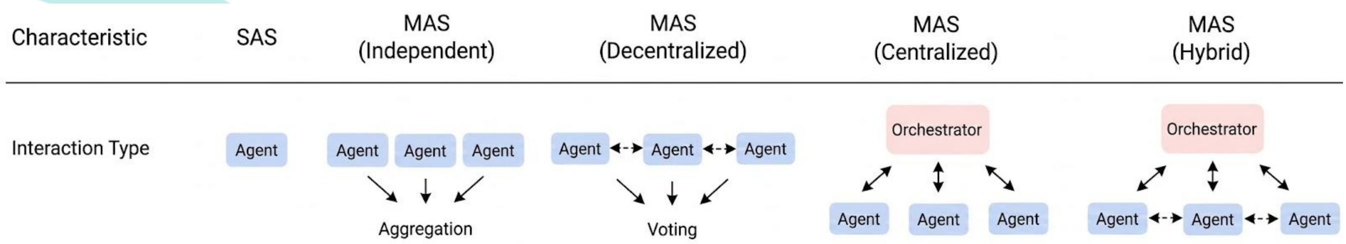


Source: Kim, Y. and Liu, X. (2026).

### 4.3.2. Sectoral applications of agentic AI

AI agents are becoming a part of multiple business processes – from HR, finance and marketing to product development and equipment maintenance – across various sectors and industries. Box 5 describes a variety of examples that have already been deployed.

## Box 5. Illustrative sectoral applications of agentic AI

Agentic AI has a wide range of applications in **serving customers** and interacting with them: it can learn from different data sources and provide context-aware, better-tailored responses, or act on behalf of a customer service agent, updating CRM or ERP systems[403]. This reduces the cost per interaction and waiting time for customers. Gartner's prediction is that by 2029, agentic AI will autonomously resolve 80% of common customer service issues without human intervention[404].

In **retail**, agentic commerce is a trend on the rise, moving towards an entirely agent-driven checkout process. Vendors, marketplaces and fintech companies introduce tools for autonomous purchases that remember users' preferences, details and order history, making the processes much faster and highly personalised[405]. For Shopify, AI-driven traffic rose sevenfold in 2025[406]. Morgan Stanley forecasts that "agentic shoppers" could drive between USD 190-385 billion (EUR 163.5-331.4 billion) in U.S. e-commerce spending by 2030, representing up to 20% of total market share[407].

In **banking**, agents can drive complex decision-making, monitor transactions, dynamically adjust prices, and streamline compliance checks[408]. Agentic commerce also drives changes. As major service providers experiment with transacting on behalf of customers, new concerns emerge around authorisation, consent attestation and fraud prevention.

**Supply chain and inventory management** is one of the spheres where companies that adopt AI report the largest cost savings and revenue gains[409]. Agents can also optimise complex processes, embedded in operational workflows like procurement, contract negotiation and dynamic sourcing[410]. One of the real-life examples is how DHL implements agents for appointment scheduling, transport status calls, and warehouse coordination[411].

In **manufacturing**, autonomous and adaptive agents can provide additional agility and creativity, advance sustainability and productivity[412]. They can optimise production schedules, equipment maintenance, and quality control, freeing human workers to focus on complex, higher-value problem-solving and innovation. Beverage producer PepsiCo collaborates with NVIDIA and Siemens, using agentic AI to simulate, test, and refine system changes and identify issues before modifying any production processes[413].

In the **energy sector**, planning agents can propose and visualise operational plans for transmission grids based on real-time observations and forecasts[414]. Agentic AI can also improve grid stability by monitoring performance, detecting anomalies and disruptions, and optimising the load[415]. Besides, agents can do energy trading on people's behalf[416].

For **telecommunications**, this technology can augment or replace existing manual processes to improve network performance and reduce costs. It can also power network self-healing and reconfiguration during disruptions[417]. Telenor and Ericsson, for example, developed a system that learns to optimise capacity and corresponding power consumption of the Radio Access Network[418].

The **defence sector** is witnessing the emergence of autonomous cybersecurity agents capable of self-learning and adapting to evolving threats, engaging in vulnerability patching[419]. For instance, Thales' AI Deployable Agent is built to protect aircraft systems from cyberattacks, detecting and responding to cyberthreats in real time[420]. Besides, a defence startup, Scout AI, just recently demonstrated autonomous weapons powered by AI agents that identify targets, assess threats and coordinate lethal strikes and multi-step attack sequences[421].

The benefits of agentic AI for **healthcare** are primarily concentrated in enhanced diagnostics, decision support, patient care, treatment planning, and robotic surgery[422]. Besides, they can significantly reduce administrative overhead for hospitals[423]. Increasingly adaptive AI-driven robots with advanced cognitive processing and computer vision can also lead to improved, more dynamic interactions with patients[424]. In **pharmaceutical** research, AI agents have been shown to reduce time-to-candidate identification by 40%[425]. In rehabilitation, AI already powers exoskeletons to provide autonomous self-balancing support for people with mobility issues[426].

In the **cultural, creative and media sectors**, agentic AI systems are being deployed to personalise content experiences at scale, going beyond static recommendation engines. Agents can continuously observe user micro-behaviours and autonomously adjust discovery feeds, thumbnails, and stream quality in real time[427]. In cultural heritage, AI agents are used for automated semantic enrichment and 3D digitisation of cultural sites and artefacts.

In **precision agriculture**, agentic AI enables multi-agent farming systems where specialised agents — covering irrigation, nutrient management, pest control, and crop health — operate in real-time coordination, autonomously deciding when and how to water, fertilise, or treat crops. Studies demonstrated that a federated agentic AI framework for crop disease classification and weed detection can achieve 96.4% accuracy using distributed sensing agents across multiple farm nodes[428]. On food safety, agents are further deployed to automate farm-to-fork traceability[429].

Modern **autonomous driving** deployments rely on multi-agent architectures in which specialised agents for perception, trajectory prediction, path planning, and safety monitoring operate in parallel, coordinating driving decisions in real time[430]. Commercial robotaxi services exemplify agentic AI at scale in mobility[431]. At the fleet management level, companies such have developed multi-agent AI platforms that move fleet operations beyond reactive monitoring into autonomous action, including predictive maintenance, intelligent dispatch, and route optimisation

In industrial and logistics settings, agentic AI is shifting **robotics** from pre-programmed automation to continuously adaptive systems. Robots dynamically plan, reason, and self-correct in real time without pre-written scripts. Multi-agent robot swarms are also being deployed for autonomous inspection of critical infrastructure, replacing manual teams with sensor-fused, self-coordinating agents[432]. In surgical robotics, agentic AI systems go beyond teleoperation, analysing tissue conditions in real time to anticipate and autonomously execute precision procedures[433].

### 4.3.3. Agentic systems in enterprise applications

Experts in agentic deployment in enterprises expect that by the early 2030s, agents will to a significant degree replace entire applications: for example, in processes like supply chain management. At the same time, reliance on traditional interface and UI design will be increasingly replace by agent-driven interfaces[434]. Table 4 illustrates possible timelines and distributions of various types of agent integration in enterprises.

**Table 4. The agentic evolution of enterprise applications**

| | | Timing and distribution of adoption by level | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Level** | **Description** | **2025** | **2027** | **2029** | **2031** | **2033** | **2035** | **2037** |
| **Agents as apps** | Agents replace entire applications (e.g. CRM Agent/Ag-fleet). Traditional interfaces are rarely used. | 0% | 0% | 5% | 20% | 35% | 50% | 65% |
| **Agent led** | Agents replace entire functional areas within the application. Agentic interfaces dominate (text prompt and voice). Traditional interfaces are used infrequently. | 0% | 5% | 20% | 40% | 50% | 35% | 25% |
| **Ag-enhanced** | Applications significantly enhanced and augmented by Ag-driven capabilities. Ag-interfaces become more dominant, reliance on traditional interfaces fades. Slowing development of As/Ad as traditional interfaces become less prominent. | 3% | 30% | 40% | 25% | 10% | 10% | 5% |
| **As/Ad-enhanced, Ag-supplemented** | Most functionality is As/Ad-enhanced, embedded Ag grows. Mostly traditional interfaces, some replaced with Ag-interface. | 17% | 35% | 20% | 10% | 5% | 5% | 5% |
| **As/Ad-enhanced** | Growing As/Ad-enhanced functionality. Traditional interfaces remain. | 50% | 20% | 10% | 5% | 0% | 0% | 0% |
| **Traditional** | Traditional functionality and interface. | 30% | 10% | 5% | 0% | 0% | 0% | 0% |

Source: IDC. (2025). The Agentic Evolution of Enterprise Applications. Available at: https://www.idc.com/resource-center/blog/the-agentic-evolution-of-enterprise-applications/. Notes: Legend: As = Assistants, Ad = Advisors, Ag = Agents.

### 4.3.4. From Software-as-a-Service to Agent-as-a-Service

Agentic AI is already disrupting traditional SaaS business models, creating what industry observers have termed the SaaSpocalypse"[435]. This shift has the potential to directly reshape infrastructure requirements, investment patterns, and architectural design principles for digital systems.

**Figure 25. Software companies' share prices rebased in USD terms**

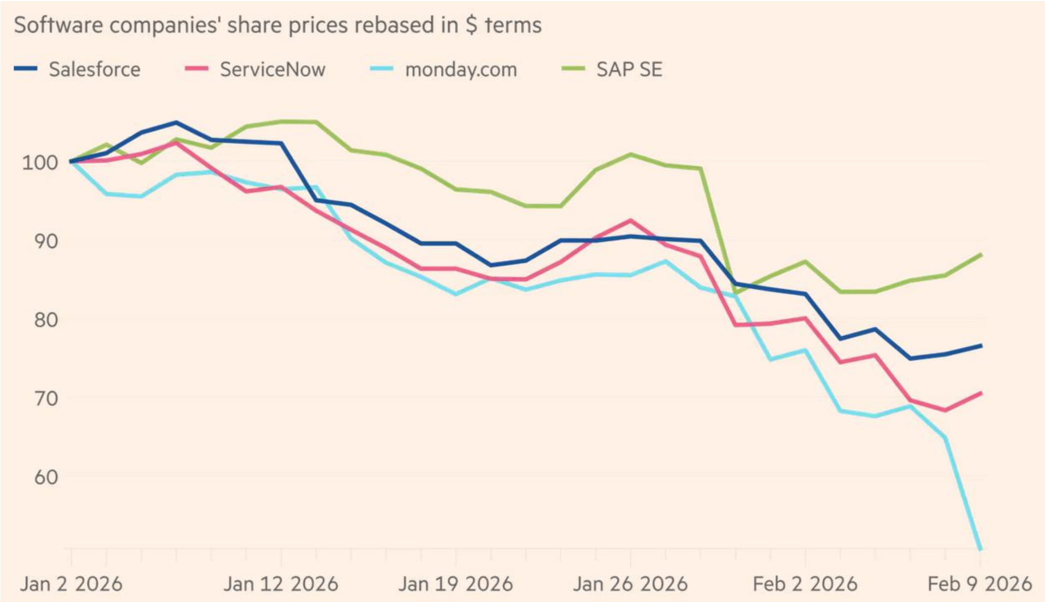


Source: Palladino, C. (2026).

Traditional SaaS pricing rests on per-seat subscriptions: infrastructure costs are relatively fixed, revenue scales with human users, and usage patterns are predictable. AI agents render much of this model obsolete.

A single agent can perform work previously requiring dozens of employees while generating continuous, machine-speed computational loads — fundamentally different from the intermittent, human-paced interactions that SaaS infrastructure was designed to support. A 2026 market forecast projects that this structural mismatch could precipitate a collapse in recurring-revenue-backed SaaS valuations by 2028, as the link between user count and business value breaks down. Companies such as ServiceNow[436], which automates IT workflows, and Zendesk[437], which manages customer service operations, are identified in the forecast as acutely exposed. Both platforms were built on the assumption that revenue scales with headcount. As agents begin executing those same workflows autonomously, the human users underpinning the subscription model — and the revenue logic the valuation was built on — risk disappearing with them.

Agents interact programmatically via APIs, bypassing visual interfaces entirely. This forces a shift from serving HTML and managing browser sessions to handling high-volume API traffic with strict service-level requirements on availability, latency, and throughput.

Moreover, agents require persistent state management at scale. As an expert interviewed as part of this project pointed out[438], while human SaaS sessions last minutes or hours, agents maintain context across days or weeks of autonomous operation. Infrastructure needs to support these long-duration connections, alongside durable storage and consistency across distributed agent workflows.

### 4.3.5. Thermodynamic foundations of intelligence and computation

Even the most sophisticated microchips generate heat primarily due to acts of erasure. When a computer processes information, it constantly discards 1s and 0s to make room for new states. This establishes a deep connection between physics and computational information, indicating that **information is indeed physical**[439].

This relationship is codified by Landauer's Principle, which defines the **minimum amount of entropy** needed to erase one bit of information. The principle states that the erasure of one bit releases a minimum amount of heat. At room temperature, erasing one bit of information generates approximately three times $10^{-21}$ joules of heat, or about a hundred trillion trillionth of the energy needed to boil a kettle[440]. Across the billions of transistors switching billions of times per second in a modern processor, the aggregate effect is enormous. In the United States alone, it is estimated that waste heat from computers drains GDP by over USD 30 billion (EUR 25 billion) a year[441].

In 1973, IBM physicist Charles H. Bennett extended this insight by demonstrating that computation can, in principle, be made logically reversible[442]. I.e., **if no bits are erased, no heat needs to be produced**. In practice, today's computers operate far from this theoretical limit, or rather, they are inefficient compared to these limits, dissipating energy at rates exceeding the Landauer limit by six to nine orders of magnitude (millions to billions of times more energy)[443]. Nonetheless, the principle grounds the search for more efficient architectures in fundamental physics rather than incremental engineering optimisation.

The **upper limit on thermodynamic efficiency** (the efficiency of any heat engine) was established by French physicist Sadi Carnot in 1824, in his book "*Reflections on the motive power of fire, and on machines fitted to develop that power*"[444], a principle that remains crucial for modern thermodynamics. Carnot's Theorem states that the efficiency (h) of a thermodynamic cycle depends solely on the temperature difference between the heat source ($T_{H.}$– the data centre equipment) and the cold reservoir ($T_C$ – the ambient environment or a cooling medium). Real-world cooling systems always operate at lower efficiencies than the Carnot limit due to irreversibilities such as friction, pressure drops, and non-ideal heat transfer processes. Engineers design systems to minimise this gap.

While originally applied to steam engines, Carnot's law is inextricably linked to **modern data centres, which function effectively as a reverse heat engine**. They use high-grade electrical work to transfer heat from their cold reservoir (the servers) to the hot reservoir (the outside atmosphere)[445]. The efficiency of this system is traditionally measured by Power Usage Effectiveness, defined as the ratio of total facility energy to IT equipment energy, where a score of 1.0 would represent perfect efficiency with no overhead losses[446].

As the power consumption of GPU begins to exceed 1,000 watts[447], producing intense, localised heat that air cannot dissipate fast enough, the efficiency of this removal process becomes a hard constraint on computational scaling. **By 2035, experts predict GPUs could consume more than 15,000 watts per unit**[448]. At these thermal densities, air cooling becomes physically unmanageable: the temperature gradients required to move heat into air demand enormous airflow volumes and disproportionate additional energy. This phenomenon can also be described as "**hitting an energy wall**", or where computing is hitting a thermal ceiling where chips cannot be cooled fast enough to process the data required, particularly for future leaps in capability[449].

## 4.4. Policy frameworks and initiatives

This annex provides a comprehensive overview of policy and governance frameworks that are directly or indirectly impacting the development of AI agents and related technologies.

### *4.4.1. AI, data and privacy regulations*

The years between 2015 and 2019 witnessed the most notable growth in AI-related policy initiatives, indicating policymakers' closer attention[450]. Lately, the focus has shifted from new legislation to developing, implementing, or refining the existing frameworks.

Among the most recent pieces of legislation on the EU level, the AI Strategy and the AI Act are the most relevant to autonomous AI systems and applications. The **Apply AI Strategy** explicitly mentions *agentic AI* as part of the EU's plans to strengthen the deployment of advanced AI technologies in strategic sectors and public services. In particular, the strategy states that the Commission will support the development of sovereign frontier models, deploy AI agents for healthcare professionals and also create a technical and policy toolkit to support generative and agentic AI solutions in public administrations[451]. These measures are presented as part of a broader effort to accelerate the adoption of scalable European AI solutions and improve the quality of services provided to citizens[452].

While **the AI Act** does not explicitly refer to agentic AI or AI agents, it covers autonomous and adaptive AI[453]. AI agents built on general-purpose AI models may therefore be subject to existing obligations under the Act, including rules applicable to GPAI systems and, depending on their intended purpose, potentially the high-risk regime[454]. The Act's human oversight requirements are particularly relevant in this context, as higher levels of autonomy may make it more difficult to ensure effective monitoring, intervention, and allocation of responsibility across the AI value chain[455]. To support implementation of the EU AI Act, European Standards Organisations are developing harmonised technical standards, the AI Office has led the drafting of a General Purpose AI Code of Practice adopted in August 2025, and the Commission has published a series of guidelines, delegated acts, and implementing acts clarifying key provisions including the definition of an AI system, prohibited practices, and obligations for general-purpose AI model providers[456].

Alongside, a number of other policy frameworks indirectly affect agentic AI in the EU. **General Data Protection Regulation (GDPR)** establishes a unified data protection framework across EU member states, protecting privacy rights, including processing, profiling, and automated decision-making[457]. GDPR relates to agentic AI by requiring lawful bases for data processing, transparency under Article 22, human oversight to prevent solely autonomous decisions and accountability via data minimisation.

With regard to governing the data necessary for training and refining agentic systems, the **Data Act** intends to grant users rights to access, control, and share data from connected products, enabling fair data economies with rules for data portability and cloud access, including dependencies[458]. It facilitates agentic AI by providing data access for training while imposing real-time controls as third-party ICT under government oversight. In addition, the **Data Governance Act** aims to enhance trust in voluntary data sharing via regulated intermediaries and public sector reuse, including data portability and access rules[459]. It aids agentic AI through secure data pooling, complementing GDPR, and enabling users to act on behalf of digital identities.

**Digital Services and Markets Acts** (DSA and DMA) also contribute to the overall governance framework. The DSA sets rules for online services like marketplaces to ensure safety and rights, covering data portability and cloud access with dependencies[460]. It regulates agentic AI on platforms via risk assessments, moderation transparency, and obligations for identity-based user actions. As to the DMA, it imposes obligations on gatekeeper platforms for data access, interoperability, and fair competition, including intermediary access rules[461]. For agentic AI, it ensures real-time data provision and third-party integration while addressing digital identity actions.

The **eIDAS Regulation** provides an EU framework for electronic signatures, trust services, and identification, including digital services, data portability, and cloud access rules[462]. By design, the regulation and the EU Digital Identity Wallets (EUDIWs) that it establishes cover natural and legal persons rather than non-human entities. This is a deliberate policy choice reflecting current priorities. However, as agentic AI systems increasingly require secure authentication for autonomous transactions, the question of how to handle machine identities and non-human agents represents an emerging area for future policy consideration. In the security realm, the **NIS2 Directive** strengthens cybersecurity for networks and systems, with incident reporting and supply chain resilience obligations[463]. Its requirements around security standards and incident reporting are relevant to agentic AI insofar as autonomous agents and the APIs they rely on form part of the networks and systems that fall within NIS2's scope. Additionally, the **Digital Operational Resilience Act** establishes ICT risk management, incident reporting, and resilience testing requirements for EU financial entities and their third-party providers. It is relevant to agentic AI insofar as autonomous ICT tools used in finance would fall under its real-time controls and third-party risk management obligations.

Finally, important initiatives on the EU level are also supporting the development of the **underlying infrastructure layer**. The subsequent thematic paper, "Agentic AI and internet infrastructure", prepared within this project, examines them in greater detail, illustrating how the connectivity backbone is being governed and developed.

### 4.4.2. Global initiatives in AI governance

On the global level, several governance mechanisms for AI have taken shape in recent years, encompassing voluntary principles, normative instruments, multilateral partnerships, safety-focused summit processes (see Figure 26), and the first binding international treaty, though most of these address broader topics and are not focused solely on agentic AI.

**Figure 26. Examples of important global governance initiatives in AI**

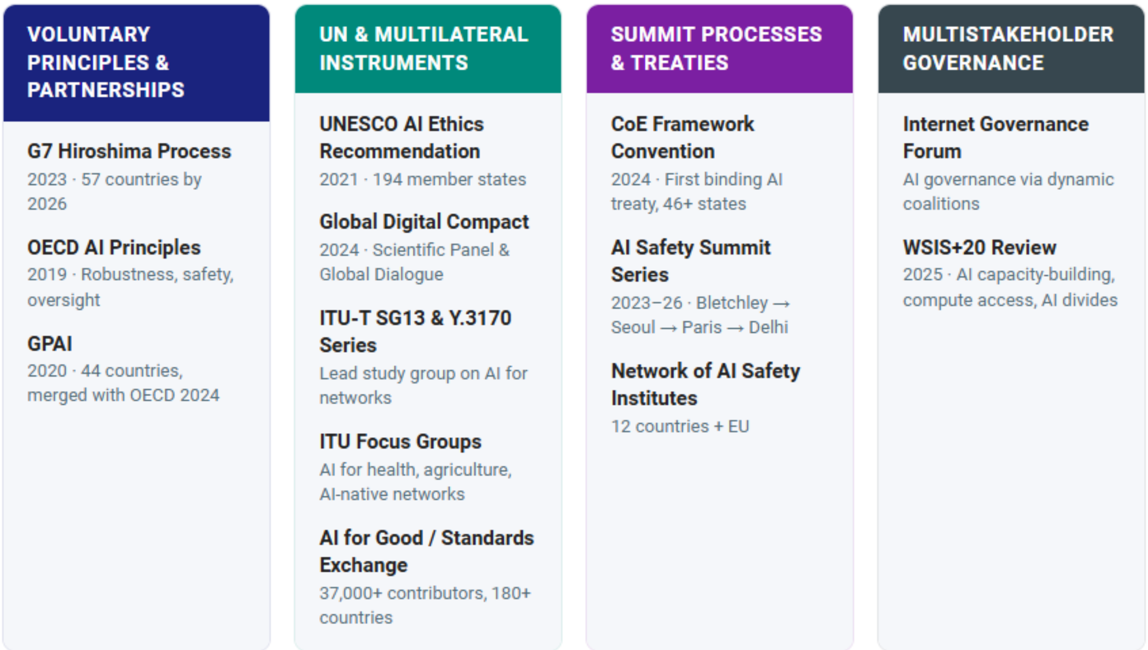


Some of the most pertinent examples include:

- The **G7 Hiroshima Process** principles, adopted in 2023, provide voluntary guidance for organisations developing advanced AI systems like foundation models to ensure they are safe, secure, and trustworthy[464]. They relate to agentic AI through risk-based approaches for autonomous deployment, emphasising misuse mitigation and transparency without explicit "agentic" terminology. The G7 Hiroshima Process has since broadened through a Friends Group reaching 57 countries by early 2026[465].
- The **OECD AI principles**, covered in more detail in Section 1.1.1, also indirectly cover agentic AI via robustness, safety, and human oversight pillars[466]. The Global Partnership on AI, launched in 2020 and merged with the OECD's AI work in July 2024, now brings together 44 countries to advance responsible AI through expert working groups[467].
- The **UNESCO Recommendation on the Ethics of AI** (2021) was the first global normative instrument on AI ethics, adopted by all 194 member states[468]. The **Global Digital Compact** (2024) committed UN member states to establishing an International Scientific Panel on AI and a Global Dialogue on AI Governance[469,470].
- The **World Summit on the Information Society (WSIS) +20 review**, concluded in 2025, represents the most comprehensive reassessment of the global digital cooperation framework since the original WSIS outcomes of 2003 and 2005. The outcome document[471] encourages international partnerships on AI capacity-building, including education and training programmes, increased access to open AI models, training data, and compute, and support for SME participation in the digital economy[472]. It calls on the UN system to map and strengthen AI capacity-building for

developing countries through mechanisms including an AI fellowship, an Independent International Scientific Panel on AI, and a Global Dialogue on AI Governance.

- The **AI Safety Summit series** — from Bletchley Park (2023) through Seoul (2024), Paris (2025), and New Delhi (2026) — has catalysed frontier AI safety commitments from governments and companies alike and launched an International Network of AI Safety Institutes now spanning twelve countries and the EU[473,474].
- The **Council of Europe Framework Convention on AI and Human Rights** (2024) is the first legally binding international treaty on AI[475]. The Framework Convention was drafted by 46 member states of the Council of Europe, with additional participation from 5 observer states (Canada, Japan, Mexico, the Holy See, and the United States), the European Union, and 6 non-member states (Australia, Argentina, Costa Rica, Israel, Peru, and Uruguay).
- The **ITU-T Study Group 13**, the lead study group on AI for future networks, has produced a full series of AI-related recommendations, while domain-specific focus groups have addressed AI governance in areas from health to agriculture[476,477].
- The **ITU's AI for Good platform** convenes over 37,000 contributors from more than 180 countries and has launched an International AI Standards Exchange jointly with ISO and IEC[478].
  - The **Internet Governance Forum** (IGF) provides an additional multistakeholder platform for discussions related to AI risks and governance using a multistakeholder format.

### 4.4.3. Standardisation and related initiatives

A large and growing ecosystem of standards is relevant to AI governance across international, regional, and national bodies. While a comprehensive mapping is beyond the scope of this section, several emerging frameworks merit specific mention.

When it comes to standardisation processes, **ISO/IEC 42001** (2023) is the first international standard for AI management systems, specifying requirements for ethical AI governance, risk treatment, and continual improvement[479]. It applies to agentic AI by establishing controls for autonomous decision-making, resource management, and transparency in AI operations. It sits within the broader work programme of ISO/IEC JTC 1/SC 42, the principal international standards subcommittee for AI.

Another important framework is **ITU-T Y.3172 ML Framework.** It defines an architectural framework for machine learning systems, covering data, training, and deployment in telecommunications networks[480]. It relates to agentic AI through scalable ML autonomy in networks, enabling self-managing agents via standardised interfaces, although not explicitly termed "agentic".

Related to this is the **IETF's Artificial Intelligence Framework for Network Management** working group has begun standardisation work on applying AI to network management tasks involving complex reasoning and event processing[481,482], which has direct implications for the deployment of autonomous network-managing agents.

Other notable efforts include the IEEE Standards Association's portfolio of AI-related standards, including its **Global Initiative on Ethics of Autonomous and Intelligent Systems**, which promotes fairness, accountability, and transparency principles increasingly relevant to autonomous agent design[483,484], while the European Telecommunications Standards Institute's (ETSI) **Securing AI Technical Committee** addresses threat mitigation for AI systems, including risks arising from autonomous agent behaviour[485]. At the European level, **CEN-CENELEC JTC 21** is developing harmonised standards that will grant adopting organisations a presumption of conformity with the EU AI Act, though publication has been delayed to the summer of 2026[486]. Additionally, the **W3C's semantic web standards**, which enable machine-readable data sharing across the web, provide foundational interoperability infrastructure upon which agentic AI systems can operate across distributed environments[487]. In addition to this, different types of protocols are emerging for AI agents and their integration, covered in more detail in the forthcoming paper on agentic AI and trust prepared within this project.

Among **industry-driven protocol development initiatives**, prominent examples are the Frontier Model Forum and the Agentic AI Foundation. The former, launched by Anthropic, Google, Microsoft, and

OpenAI, is an industry-supported non-profit created to identify best practices in safety and security, support standards development, fund safety research, and facilitate information-sharing with governments, academia, and civil society[488]. The Agentic AI Foundation is a Linux Foundation–hosted consortium that brings together companies, researchers, and developers to create open, interoperable standards, protocols, and open-source infrastructure[489].

Finally, it is worth noting that the National Institute of Standards and Technology, an agency of the U.S. Department of Commerce, has been active in AI governance through its AI Risk Management Framework and AI test, evaluation, validation, and verification programme[490,491]. In February 2026, its Centre for AI Standards and Innovation launched the **"AI Agent Standards Initiative"**. This initiative solicits industry input on standards for agentic AI, focusing on security risks from autonomous actions. It aims to cement U.S. dominance at the AI frontier amid China's growth, covering governance for self-acting agents without prescriptive rules[492].

## 4.5. References

[1] COM(2023) 442/final. Communication From the Commission to The European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions. An EU initiative on Web 4.0 and virtual worlds: a head start in the next technological transition. Strasbourg, 11 July 2023. Available at: https://digital-strategy.ec.europa.eu/en/library/eu-initiative-virtual-worlds-head-start-next-technological-transition

[2] Barcevičius, et al. (2025). Background document: Input for the Global Multistakeholder High-Level Conference on Governance for Web 4.0 and Virtual Worlds, 31 March – 1 April 2025. European Commission, PPMI, TNO, & WEB4HUB. Available at: https://ec.europa.eu/newsroom/dae/redirection/document/113701

[3] More information available at: https://digital-strategy.ec.europa.eu/en/policies/event-web-4-governance

[4] Rubio-Licht, N. (2025). Isolating AI agents may be ‘strangling’ your enterprise. The Daily Upside. Available at: https://www.thedailyupside.com/cio/enterprise-ai/isolating-ai-agents-may-be-strangling-your-enterprise/

[5] Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges. *arXiv preprint arXiv:2505.10468*.

[6] Petrova, T., et al. (2025). From Semantic Web and MAS to Agentic AI: A Unified Narrative of the Web of Agents. *arXiv preprint arXiv:2507.10644*.

[7] Yang, Y., et al. (2025). Agentic web: Weaving the next web with AI agents. *arXiv preprint arXiv:2507.21206*.

[8] Feuerriegel et al. (2023), *Business & Information Systems Engineering*, published in the peer-reviewed Springer journal

[9] Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges. *arXiv preprint arXiv:2505.10468*.

[10] Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges. *arXiv preprint arXiv:2505.10468*.

[11] Petrova, T., et al. (2025). From Semantic Web and MAS to Agentic AI: A Unified Narrative of the Web of Agents. *arXiv preprint arXiv:2507.10644*.

[12] Yang, Y., et al. (2025). Agentic web: Weaving the next web with AI agents. *arXiv preprint arXiv:2507.21206*.

[13] Flinders, M. & Smalley, I. (n.d.). What is AI Infrastructure? IBM. Available at: https://www.ibm.com/think/topics/ai-infrastructure

[14] Majumder, A. (2025). Rise and impact of AI agents in the digital landscape. American Journal of Intelligent Systems, 14(1), 10-17. Available at: https://www.researchgate.net/profile/Anirban-Majumder-16/publication/389875872_Rise_Impact_of_AI_Agents_in_Digital_Landscape/links/67d5a44e478c5a3feda23da8/Rise-Impact-of-AI-Agents-in-Digital-Landscape.pdf

[15] Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges. *arXiv preprint arXiv:2505.10468*.

[16] Murugesan, S. (2025). The rise of agentic AI: implications, concerns, and the path forward. IEEE Intelligent Systems, 40(2), 8-14. Available at: https://ieeexplore.ieee.org/abstract/document/10962241/

[17] KPMG. (2025). Agentic AI: The future of autonomous intelligence. KPMG India. Available at: https://assets.kpmg.com/content/dam/kpmgsites/in/pdf/2025/10/agentic-ai-the-future-of-autonomous-intelligence.pdf

[18] Yang, Y., et al. (2025). Agentic web: Weaving the next web with ai agents. arXiv. Available at: https://arxiv.org/pdf/2507.21206?

[19] LF Decentralized Trust. (2025, September 4). *Trusted AI Agents: Architecting Identity and Granular Access for the Agentic Web* [Video]. YouTube. Available at: https://www.youtube.com/live/SJ8rFKJ8NHw?si=vgucs7PI95WIu8oe

[20] Pol, Dr. (2025). Generative AI, AI Agents, and Agentic AI : An Overview of Current AI Technologies. International Journal for Research in Applied Science and Engineering Technology. 13. 2650-2652. 10.22214/ijraset.2025.75710.

[21] KPMG. (2025). Agentic AI: The future of autonomous intelligence. KPMG India. Available at: https://assets.kpmg.com/content/dam/kpmgsites/in/pdf/2025/10/agentic-ai-the-future-of-autonomous-intelligence.pdf

[22] Interview IDs IT_2, IT_4, IT_7. (2026).

[23] Ji, J. J. (2024). Demystify ChatGPT: anthropomorphism around generative AI. GRACE: Global Review of AI Community Ethics, 2(1). Available at: https://ojs.stanford.edu/ojs/index.php/grace/article/view/3222

[24] Red Hat. (2025). Agentic AI vs. generative AI. Available at: https://www.redhat.com/en/topics/ai/agentic-ai-vs-generative-ai

[25] Pol, Dr. (2025). Generative AI, AI Agents, and Agentic AI: An Overview of Current AI Technologies. International Journal for Research in Applied Science and Engineering Technology. 13. 2650-2652. 10.22214/ijraset.2025.75710.

[26] Toner, H., et al. (2024). Through the Chat Window and Into the Real World: Preparing for AI Agents. *Center for Security and Emerging Technology. https://doi. org/10.51593/20240034*.

[27] World Economic Forum. (2025). AI Agents in Action: Foundations for Evaluation and Governance. Available at: https://reports.weforum.org/docs/WEF_AI_Agents_in_Action_Foundations_for_Evaluation_and_Governance_2025.pdf

[28] ITU. (2025). The Annual AI Governance Report 2025: Steering the Future of AI. Available at: https://www.itu.int/epublications/en/publication/the-annual-ai-governance-report-2025-steering-the-future-of-ai/en

[29] Rothschild, D. M., et al. (2025). The Agentic Economy. arXiv preprint arXiv:2505.15799.

[30] Arcep (2026). Generative AI: Challenges for the Future of the Open Internet. Available at: https://www.arcep.fr/uploads/tx_gspublication/report-generative-AI-challenges-open-internet-january2026.pdf

[31] LF Decentralized Trust. (2025, September 4). *Trusted AI Agents: Architecting Identity and Granular Access for the Agentic Web* [Video]. YouTube. Available at: https://www.youtube.com/live/SJ8rFKJ8NHw?si=vgucs7PI95WIu8oe

[32] Feng, K. J., McDonald, D. W., & Zhang, A. X. (2025). Levels of Autonomy for AI Agents. *arXiv preprint arXiv:2506.12469*.

[33] Bengio, Y., et al. (2025). International AI Safety Report. Available at: https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025

[34] Wijk, H. et al. (2025). RE-Bench: Evaluating Frontier AI R&D Capabilities of Language Model Agents against Human Experts, arXiv [cs.LG]; Available at: http://arxiv.org/abs/2411.15114

[35] Levy, S. (2026). The Math on AI Agents Doesn’t Add Up. Available at: https://www.wired.com/story/ai-agents-math-doesnt-add-up/

[36] Xuan, W. et al. (2025), “MMLU-ProX: A Multilingual Benchmark for Advanced Large Language Model Evaluation”, arXiv. Available at: https://doi.org/10.48550/arXiv.2503.10497.

[37] Hobbs, H. et al. (2026). Exploring possible AI trajectories through 2030. *OECD Artificial Intelligence Papers*, No. 55, OECD Publishing, Paris, https://doi.org/10.1787/cb41117a-en.

[38] Toner, H., et al. (2024). Through the Chat Window and Into the Real World: Preparing for AI Agents. *Center for Security and Emerging Technology. https://doi. org/10.51593/20240034*.
[39] Keynote address by Yoshua Bengio, Professor of Computer Science at Université de Montreal: Applying AI in Research, AI in Science Summit 2025.
[40] Bengio, Y., et al. (2025). International AI Safety Report. Available at: https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025
[41] Levy, S. (2026). The Math on AI Agents Doesn't Add Up. Available at: https://www.wired.com/story/ai-agents-math-doesnt-add-up/
[42] Hobbs, H. et al. (2026). Exploring possible AI trajectories through 2030. *OECD Artificial Intelligence Papers*, No. 55, OECD Publishing, Paris, https://doi.org/10.1787/cb41117a-en.
[43] Kokotajlo, D., et al. (2025). AI 2027. Available at: https://ai-2027.com
[44] Hobbs, H. et al. (2026). Exploring possible AI trajectories through 2030. *OECD Artificial Intelligence Papers*, No. 55, OECD Publishing, Paris, https://doi.org/10.1787/cb41117a-en.
[45] Findings of stakeholder consultations carried out within the project.
[46] Goldstein, A. (2025). Agentic AI. European Data Protection Supervisor (EDPR). Available at https://www.edps.europa.eu/data-protection/technology-monitoring/techsonar/agentic-ai_en
[47] METR. (2025). Measuring AI Ability to Complete Long Tasks. Available at : https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/
[48] The Collective Intelligence Project. (2025). There will soon be millions of AI agents working on our behalf for work, commerce, and democracy. Available at: https://blog.cip.org/p/there-will-soon-be-millions-of-ai
[49] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[50] Ibid.
[51] LF Decentralized Trust. (2025, September 4). *Trusted AI Agents: Architecting Identity and Granular Access for the Agentic Web* [Video]. YouTube. Available at: https://www.youtube.com/live/SJ8rFKJ8NHw?si=vgucs7PI95WIu8oe
[52] WEF (2025). Technology convergence report: Insight report. Available at: https://reports.weforum.org/docs/WEF_Technology_Convergence_Report_2025.pdf
[53] Interview IDs IT_1, IT_3, IT_4, IT_8, IT_9.
[54] High level panel: Shaping the Future of Science with AI: Europe's Strategic Moment. AI in Science Summit 2025.
[55] Stakeholder consultation organised during the project.
[56] Walther, C. (2025). Why Prosocial AI Matters For The Future Of Planetary Health. Forbes. Available at: https://www.forbes.com/sites/corneliawalther/2025/10/15/why-prosocial-ai-matters-for-the-future-of-planetary-health/
[57] European Commission AI Office. (n.d.). AI for Public Good. Available at: https://digital-strategy.ec.europa.eu/en/factpages/ai-public-good
[58] European Commission. (2025). The AI Continent Action Plan. Available at: https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan
[59] AI for Good. (n.d.). About us. Available at: https://aiforgood.itu.int/about-us/
[60] World Economic Forum. (2025). AI Agents in Action: Foundations for Evaluation and Governance. Available at: https://reports.weforum.org/docs/WEF_AI_Agents_in_Action_Foundations_for_Evaluation_and_Governance_2025.pdf
[61] ITU. (2024). AI for Good. Impact Report. Available at: https://s41721.pcdn.co/wp-content/uploads/2025/03/eBat-2402309_AI-for-Good-Impact-Report-E-v6.pdf
[62] Eurofound. (2026). Who uses Generative AI? Patterns and inequalities across the EU. Available at: https://www.eurofound.europa.eu/en/publications/all/who-uses-generative-ai-patterns-and-inequalities-across-the-eu
[63] DataReportal. (2026). Digital 2026: Global Overview Report. Available at: https://datareportal.com/reports/digital-2026-global-overview-report
[64] Ibid.
[65] Shark, A. (2025). How Agentic AI Will Transform Local Government: A Forewarning. StateTech Magazine. Available at: https://statetechmagazine.com/article/2025/08/how-agentic-ai-will-transform-local-government-forewarning
[66] Klover.ai. (2025). Governance 2.0: The Rise of AI Agents in Policy Making. Available at: https://www.klover.ai/governance-2-0-the-rise-of-ai-agents-in-policy-making/
[67] Shark, A. (2025). How Agentic AI Will Transform Local Government: A Forewarning. StateTech Magazine. Available at: https://statetechmagazine.com/article/2025/08/how-agentic-ai-will-transform-local-government-forewarning
[68] Cui, H., & Yasseri, T. (2024). AI-enhanced collective intelligence. *Patterns*, *5*(11).
[69] The Collective Intelligence Project. (2025). There will soon be millions of AI agents working on our behalf for work, commerce, and democracy. Available at: https://blog.cip.org/p/there-will-soon-be-millions-of-ai
[70] Interview ID IT_7. (2026).
[71] Chen, C. (2026). Yann LeCun's new venture is a contrarian bet against large language models. MIT Technology Review. Available at: https://www.technologyreview.com/2026/01/22/1131661/yann-lecuns-new-venture-ami-labs/
[72] Pepper, T. (2025). Securing Agentic AI: Resilience for nations, enterprises and society. Government Transformation Magazine. Available at: https://www.government-transformation.com/data/securing-agentic-ai-resilience-for-nations-enterprises-and-society
[73] Interview ID IT_3. (2026).
[74] Findings of stakeholder consultations carried out within the project.
[75] Interview ID IT_2. (2026).
[76] ITU. (2025). The Annual AI Governance Report 2025: Steering the Future of AI. Available at: https://www.itu.int/epublications/en/publication/the-annual-ai-governance-report-2025-steering-the-future-of-ai/en
[77] See more at: https://ais25.eu
[78] Paris Peace Forum. (2025). AI Action Summit: 50 innovative projects selected by the Paris Peace Forum following its Call for AI Projects. Available at: https://parispeaceforum.org/press-releases/sommet-action-ia-50-projets-innovants-selectionnes/
[79] Digital Watch Observatory. (2025). World Economic Forum 2025 at Davos. Available at: https://dig.watch/event/wef-2025-davos
[80] OECD. (n.d.). Inclusive growth, sustainable development and well-being (OECD AI Principle). Available at: https://oecd.ai/en/dashboards/ai-principles/P5
[81] Levy, S. (2025). Demis Hassabis Embraces the Future of Work in the Age of AI | WIRED. Available at: https://www.wired.com/story/google-deepminds-ceo-demis-hassabis-thinks-ai-will-make-humans-less-selfish/
[82] Democracy Now! (2025). Artificial Intelligence "Godfathers" Call for Regulation as Rights Groups Warn AI Encodes Oppression. Available at: https://www.democracynow.org/2023/6/1/ai_bengio_petty_tegmark

[83] Ilves, L., Kilian, M., Peixoto, T. and Velsberg, O. (2025). An Introduction to Agentic AI for Governments. Berlin Global GovTech Centre. Available at: https://www.globalgovtechcentre.org/an-introduction-to-agentic-ai-for-governments
[84] Liu, J. A., et al. (2025). Agentic AI for sustainable development: Leveraging large language model-enhanced agent-based modeling for complex policy strategies. Emerging Media, 3(3), 401-413. Available at: https://journals.sagepub.com/doi/pdf/10.1177/27523543251365678
[85] United Nations University. (2025). The role of AI Agents in humanitarian action: balancing ethics and innovation. Available at: https://unu.edu/cpr/blog-post/role-ai-agents-humanitarian-action-balancing-ethics-and-innovation
[86] Caron, N., et al. (2025). AI for Wildfire Management: From Prediction to Detection, Simulation, and Impact Analysis—Bridging Lab Metrics and Real-World Validation. *AI*, *6*(10), 253.
[87] Zhou, Y., et al. (2025). Multi-agent systems in climate-resilient land-use planning: a review. *International Journal of Digital Earth*, *18*(1), 2487051.
[88] Government Technology Agency Singapore. (n.d.). Applications for GovTech and the Singapore Government. Available at: https://playbooks.aip.gov.sg/agentic-ai-primer/09_applications_for_singapore_govt/
[89] Center for Responsible Business. (2025). Agentic AI for the Public Good: Expanding Access to Critical Services. UC Berkley Haas. Available at: https://haas.berkeley.edu/responsible-business/blog/posts/agentic-ai-for-the-public-good-expanding-access-to-critical-services/
[90] Public Sector Tech Watch. (2025). Diia.AI: Ukraine's National AI Agent for Government Services. Available at: https://interoperable-europe.ec.europa.eu/collection/public-sector-tech-watch/diiaai-ukraines-national-ai-agent-government-services
[91] Amin, A., et al. (2025). Multi-agent systems for optimising smart energy clusters: A case study on cost and emission reduction in industrial seaport facilities. *Energy Reports*, *14*, 5061-5071.
[92] United Nations University. (2025). Unlocking AI's potential to serve humanity. Available at: https://unu.edu/cpr/project/unlocking-ais-potential-serve-humanity
[93] Ilves, L., et al. (2025). An Introduction to Agentic AI for Governments. Berlin Global GovTech Centre. Available at: https://www.globalgovtechcentre.org/an-introduction-to-agentic-ai-for-governments
[94] Interview ID IT_1. (2026).
[95] CONSUL DEMOCRACY. (n.d.). CONSUL DEMOCRACY. Available at: https://consuldemocracy.org/
[96] IDC. (2025). How Agentic and GenAI Will Advance Smarter, More Resilient Cities. Available at: https://www.idc.com/resource-center/blog/how-agentic-and-genai-will-advance-smarter-more-resilient-cities/
[97] Opdyke, H. (2019). Surtrac Allows Traffic To Move at the Speed of Technology. Carnegie Mellon University. Available at: https://www.cmu.edu/news/stories/archives/2019/october/traffic-moves-at-speed-of-technology.html
[98] European Commission. (n.d.). Digital Decade DESI visualisation tool. Available at: https://digital-decade-desi.digital-strategy.ec.europa.eu/datasets/desi/charts/compare-countries-progress?period=desi_2024&indicator=desi_ai&breakdown=ent_all_xfin&unit=pc_ent&country=EU
[99] Google Cloud. (2025). 2025 State of AI Infrastructure Report. Available at: https://cloud.google.com/resources/content/state-of-ai-infrastructure?hl=en
[100] McKinsey. (2025). The state of AI: How organizations are rewiring to capture value. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value
[101] Stanford University. (2026). The 2025 AI Index Report. Available at: https://hai.stanford.edu/ai-index/2025-ai-index-report
[102] McKinsey. (2025). Seizing the agentic AI advantage. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage
[103] Interview ID IT_3. (2026).
[104] Databricks. (2026). State of AI Agents. Available at: https://www.databricks.com/resources/ebook/state-of-ai-agents
[105] Gartner, Inc. (2024). Gartner predicts that 30% of Fortune 500 companies will offer service through only a single, AI-enabled channel by 2028 [Press release]. https://www.gartner.com/en/newsroom/press-releases/2024-12-11-gartner-predicts-that-30-percent-of-fortune-500-companies-will-offer-service-through-only-a-single-ai-enabled-channel-by-20
[106] Rothschild, D. M., et al. (2025). The Agentic Economy. arXiv preprint arXiv:2505.15799.
[107] AImonks. (2024). AI Agents Are the Key to Small and Mid-Sized Business Success in 2025 Medium. Available at: https://medium.com/aimonks/ai-agents-are-the-key-to-small-and-mid-sized-business-success-in-2025-82d3ab3321d4
[108] Y. Bengio, et al. (2025). International AI Safety Report. Available at: https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025
[109] Awad, N., Serry, M. & Vasquez, J. (2025). How Agentic AI is Transforming Enterprise Platforms. Boston Consulting Group. Available at: https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms
[110] Yang, Y., et al. (2025). Agentic web: Weaving the next web with ai agents. arXiv. Available at: https://arxiv.org/pdf/2507.21206?
[111] McKinsey. (2025). Seizing the agentic AI advantage. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage
[112] PYMNTS Intelligence. (2025). From Assistive to Agentic AI: Consumers Wade Into Autonomous Commerce. Available at: https://www.pymnts.com/study_posts/from-assistive-to-agentic-ai-consumers-wade-into-autonomous-commerce
[113] Alfaris, A. (2025). The rise of the cognitive enterprise: Why agentic AI platforms are the next great business revolution. World Economic Forum. Available at: https://www.weforum.org/stories/2025/06/cognitive-enterprise-agentic-business-revolution/
[114] StepUp StartUps. (2026). European agentic AI landscape. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[115] GlobalData. (2025). 2026 promises to be an intriguing year for the digital workplace. Available at: https://www.globaldata.com/newsletter/details/2026-promises-to-be-an-intriguing-year-for-the-digital-workplace_397617/
[116] PwC. (2025). PwC's AI Agent Survey. Available at: https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html
[117] IDC. (2025). The Agentic Evolution of Enterprise Applications. Available at: https://www.idc.com/resource-center/blog/the-agentic-evolution-of-enterprise-applications/
[118] AWS. (2025). Introducing AI agents and tools in AWS Marketplace. Available at: https://aws.amazon.com/about-aws/whats-new/2025/07/ai-agents-tools-aws-marketplace/
[119] See more at: https://www.sokosumi.com/
[120] The Future Society. (2025). *Ahead of the curve: Governing AI agents under the EU AI Act*. Available at: https://thefuturesociety.org/wp-content/uploads/2023/04/Report-Ahead-of-the-Curve-Governing-AI-Agents-Under-the-EU-AI-Act-4-June-2025.pdf
[121] Petrova, T., et al. (2025). From Semantic Web and MAS to Agentic AI: A Unified Narrative of the Web of Agents. *arXiv preprint arXiv:2507.10644*.

[122] Gartner. (2025). Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027. Available at: https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027
[123] Negele, M., et al. (2025). Europe and the geopolitics of AGI: The need for a preparedness plan. RAND Corporation. Available at: https://www.rand.org/pubs/research_reports/RRA4636-1.html
[124] McKinsey & Company. (2025). The top trends in tech. Available at: https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-top-trends-in-tech
[125] StepUp StartUps. (2026). European agentic AI landscape. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[126] BCG. (n.d.). AI Agents: What They Are and Their Business Impact. Available at: https://www.bcg.com/capabilities/artificial-intelligence/ai-agents
[127] StepUp StartUps Consortium. (2025). Funding the AI economy: Strengthening Europe's investment capacity. Publications Office of the European Union. Available at://digital-strategy.ec.europa.eu/en/library/funding-ai-economy-strengthening-europes-investment-capacity (based on Pitchbook data)
[128] Stanford University. (2026). The 2025 AI Index Report. Available at: https://hai.stanford.edu/ai-index/2025-ai-index-report
[129] Ibid.
[130] EpochAI. (2025). Trends in AI supercomputers. Available at: https://epoch.ai/blog/trends-in-ai-supercomputers
[131] Goldman Sachs. (2025). Why AI companies may invest more than $500 billion in 2026. Available at: https://www.goldmansachs.com/insights/articles/why-ai-companies-may-invest-more-than-500-billion-in-2026
[132] Stanford University Human-Centered Artificial Intelligence. (2025). AI index report 2025. Available at: https://hai.stanford.edu/ai-index/2025-ai-index-report
[133] Larson, J., & Smith, A. (2025). The Age of Agentic AI. Booz Allen. Available at: https://www.boozallen.com/insights/velocity/the-age-of-agentic-ai.html.
[134] Lawton, C. (2025). Building national foundation models. AWS Public Sector Blog. Available at: https://aws.amazon.com/blogs/publicsector/building-national-foundation-models/
[135] The Future Society. (2025). Ahead of the curve: Governing AI agents under the EU AI Act. Available at: https://thefuturesociety.org/wp-content/uploads/2023/04/Report-Ahead-of-the-Curve-Governing-AI-Agents-Under-the-EU-AI-Act-4-June-2025.pdf
[136] Interview ID IT_3. (2026).
[137] Digital Watch Observatory. (2025). Microsoft launches agentic launchpad to support AI start-ups. Available at: https://dig.watch/updates/microsoft-launches-agentic-launchpad-to-support-ai-start-ups
[138] Amazon. (2025). AWS announces new innovations for building AI agents at AWS Summit New York 2025. Available at: https://www.aboutamazon.com/news/aws/aws-summit-agentic-ai-innovations-2025
[139] Google. (2025). Ironwood: The first Google TPU for the age of inference. Available at: https://blog.google/innovation-and-ai/infrastructure-and-cloud/google-cloud/ironwood-tpu-age-of-inference/
[140] Siemens. (2026). Siemens brings AI to the physical world with Eigen Engineering Agent. Available at: https://press.siemens.com/global/en/pressrelease/siemens-brings-ai-physical-world-eigen-engineering-agent
[141] Salesforce Ben. (2025). Salesforce acquires AI marketing giant Qualified. Available at: https://www.salesforceben.com/salesforce-acquires-ai-marketing-giant-qualified/
[142] Lo, K. (2025). China's AI agents take on OpenAI in global automation race. Rest of World. Available at: https://restofworld.org/2025/china-ai-agent-openai/
[143] StepUp StartUps Consortium. (2025). Funding the AI economy: Strengthening Europe's investment capacity. Publications Office of the European Union. Available at://digital-strategy.ec.europa.eu/en/library/funding-ai-economy-strengthening-europes-investment-capacity (based on Pitchbook data)
[144] Tully, T., et al. (2025). 2025: The state of generative AI in the enterprise. Menlo Ventures. Available at: https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/
[145] Asgar, Z., Nguyen, M., & Katti, S. (2025). Efficient and Scalable Agentic AI with Heterogeneous Systems. arXiv. Available at: https://arxiv.org/pdf/2507.19635?
[146] Gandhi, H., & Chandran, R. (2025). We mapped the world's hottest data centers. Rest of World. Available at: https://restofworld.org/2025/data-center-heat-map/
[147] Suarez, C., et al. (2025). Agentic AI: Future Issues at the Intersection of Technology, Innovation, and Competition Policy. *Competition Policy International's TechREG CHRONICLE*. Available at: https://www.steptoe.com/en/news-publications/agentic-ai-future-issues-at-the-intersection-of-technology-innovation-and-competition-policy.html
[148] PYMNTS. (2025). Amazon and Perplexity Clash Over Use of Third-Party AI Agents. Available at: https://www.pymnts.com/artificial-intelligence-2/2025/amazon-and-perplexity-clash-over-deployment-of-third-party-ai-shopping-agents
[149] The Future Society. (2025). *Ahead of the curve: Governing AI agents under the EU AI Act*. Available at: https://thefuturesociety.org/wp-content/uploads/2023/04/Report-Ahead-of-the-Curve-Governing-AI-Agents-Under-the-EU-AI-Act-4-June-2025.pdf
[150] Bellan, R. (2025). OpenAI, Anthropic, and Block join new Linux Foundation effort to standardize the AI agent era. TechCrunch. Available at: https://techcrunch.com/2025/12/09/openai-anthropic-and-block-join-new-linux-foundation-effort-to-standardize-the-ai-agent-era/
[151] Block. (2025). Block, Anthropic, and OpenAI launch the Agentic AI Foundation. Available at: https://block.xyz/inside/block-anthropic-and-openai-launch-the-agentic-ai-foundation
[152] GlobalData. (2025). Nvidia unveils Nemotron 3 open models for agentic AI. Available at: https://www.globaldata.com/newsletter/details/nvidia-unveils-nemotron-3-open-models-for-agentic-ai_397368/
[153] Sukharevsky, A., et al. (2024). Time to place our bets: Europe's AI opportunity. McKinsey & Company. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/time-to-place-our-bets-europes-ai-opportunity
[154] StepUp StartUps. (2026). European agentic AI landscape. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[155] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[156] Findings of stakeholder consultations carried out within the project.
[157] Based on self-reported LinkedIn data, relative to the OECD average.
[158] For comparison, South Korea share of global GenAI patent filings: 7%; US share of global GenAI patent filings: 6%. More information available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC142598
[159] Over the past eight years, Inner London, Île-de-France, Noord-Brabant, and Oberbayern have accounted for more than a quarter of all AI patent filings, while Bucharest, Estonia, Budapest, and Brussels show strong relative specialisation (according to OECD RegPat database,

version 2024). More information available at: https://op.europa.eu/en/publication-detail/-/publication/e13d1844-e462-11ef-be2a-01aa75ed71a1/language-en
[160] Total AI investment in the EU27 in 2023 is estimated at EUR 220–294 billion (central estimate approx. EUR 257 billion); for comparison United States invested around USD 90 billion in AI R&D alone in 2023 and 200 billion in data and equipment. More information available at: https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/09/advancing-the-measurement-of-investments-in-artificial-intelligence_7f58ff65/13e0da2f-en.pdf
[161] Official statistical sources tend to lack disaggregated agentic-specific data. McKinsey's The state of AI in 2025: Agents, innovation, and transformation (2025), based on 1,993 respondents across 105 countries, finds that 62% of respondent organisations are at least experimenting with AI agents, but in no single business function does the share of organisations reporting agents as scaled or fully scaled exceed roughly 10%. The gap between broad experimentation and narrow scaling reflects the additional integration, governance and workflow-redesign requirements that agentic deployment imposes beyond standard generative AI use.
[162] European Commission. (2026). AI Office publishes frontier AI expert findings on EU competitiveness, sovereignty and security. Available at: https://digital-strategy.ec.europa.eu/en/library/ai-office-publishes-frontier-ai-expert-findings-eu-competitiveness-sovereignty-and-security
[163] European Commission, Directorate-General for Research and Innovation. (2025). A European model for artificial intelligence. Publications Office of the European Union. https://doi.org/10.2777/8034640
[164] European Commission, Directorate-General for Communications Networks, Content and Technology. (2025). *Study on the next data frontier – Generative AI, regulatory compliance and international dimensions: Executive summary* (CARSA, ICF, Timelex, & VTT). Publications Office of the European Union. https://doi.org/10.2759/3909870
[165] European Commission, Directorate-General for Research and Innovation. (2023). Trends in the use of AI in science: A bibliometric analysis. Publications Office of the European Union. https://doi.org/10.2777/418191
[166] Pilz, K. F., Rahman, R., Sanders, J., & Heim, L. (2025). *Data on GPU clusters* [dataset]. Epoch AI.
[167] Pilz, K. F., et al. (2025). *The US hosts the majority of GPU cluster performance, followed by China*. Epoch AI.
[168] Abendroth Dias, K., et al. (2025). Generative AI outlook report: Exploring the intersection of technology, society and policy (E. Navajas Cawood, M. Vespe, A. Kotsev, & R. Van Bavel, Eds.). Publications Office of the European Union. https://doi.org/10.2760/1109679
[169] The Paulson Institute. (2023). *The Global AI Talent Tracker 2.0*. MacroPolo.
[170] Negele, M., et al. (2025). Europe and the geopolitics of AGI: The need for a preparedness plan. RAND Corporation.
[171] OECD.AI (2026). OECD AI Policy Observatory. Visualisations powered by JSI using data from OpenAlex, LinkedIn, www.oecd.ai
[172] EIB (2026). *Drivers of relocation by innovative EU startups and scaleups*. Publications Office of the European Union.
[173] EIB (2025). European firms show resilience, invest in green transition, and match US companies in adopting AI technologies, new EIB survey shows.
[174] Eurostat. (2025). 20% of EU enterprises use AI technologies [News article]. Available at: https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2
[175] Lemke, N., & Schneider, C. (2025). *The European Union's AI Factories: Lessons for public investment in AI infrastructure in Europe* [Policy brief]. Interface. Available at: https://www.interface-eu.org/publications/ai-factories
[176] Singla, A., et al. (2024, May 30). The state of AI in early 2024: Gen AI adoption spikes and starts to generate value. McKinsey & Company. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024
[177] StepUp StartUps Consortium. (2026). Agentic AI: Leveraging European AI talent and regulatory assets to scale adoption. Publications Office of the European Union.
[178] AscendixTech (2025). How many AI companies are there in the world?, Ascendix Technologies.
[179] Dealflow.eu & EU-Startups (2025). Startups driving Europe's AI transformation: Supporting the Apply AI strategy — AI startup and investment activity across 10 key industrial sectors. Innovation Radar Bridge, European Commission.
[180] AI Policy Bulletin. (2025). Bargaining Chips: Could the EU Leverage ASML to Influence U.S. AI Policy? Available at : https://www.aipolicybulletin.org/articles/bargaining-chips-could-the-eu-leverage-asml-to-influence-u-s-ai-policy
[181] StepUp StartUps. (2026). European agentic AI landscape. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[182] Findings of stakeholder consultations carried out within the project.
[183] Interview IDs IT_1, IT_6, IT_8, IT_9. (2026).
[184] Findings of stakeholder consultations carried out within the project.
[185] European Commission. (2026). Strengthening Europe's Tech Sovereignty | Shaping Europe's digital future. Available at: https://digital-strategy.ec.europa.eu/en/policies/eu-tech-sovereignty
[186] Teare, G. (2025). 6 charts that show the big AI funding trends of 2025. Crunchbase News. Available at: https://news.crunchbase.com/ai/big-funding-trends-charts-eoy-2025/
[187] Precedence Research. (2025). Agentic AI market size to reach USD 199.05 billion by 2034 driven by autonomous decision-making and North America's leadership. GlobeNewswire. Available at: https://www.globenewswire.com/news-release/2025/09/04/3144393/0/en/Agentic-AI-Market-Size-to-Reach-USD-199-05-Billion-by-2034-Driven-by-Autonomous-Decision-Making-and-North-Americas-Leadership.html
[188] StepUp StartUps Consortium. (2025). Funding the AI economy: Strengthening Europe's investment capacity. Publications Office of the European Union. Available at://digital-strategy.ec.europa.eu/en/library/funding-ai-economy-strengthening-europes-investment-capacity (based on Pitchbook data)
[189] OECD.AI (2026). Available at: www.oecd.ai
[190] StepUp StartUps Consortium. (2025). Funding the AI economy: Strengthening Europe's investment capacity. Publications Office of the European Union. Available at://digital-strategy.ec.europa.eu/en/library/funding-ai-economy-strengthening-europes-investment-capacity (based on Pitchbook data)
[191] OECD.AI (2026). Available at: www.oecd.ai
[192] Presentation by Cecile Huet (AI Office) in "Workshop 2: Community & Society" of the AI in Science Summit 2025.
[193] Findings of stakeholder consultations carried out within the project.
[194] World Economic Forum. (2026). Rethinking AI Sovereignty: Pathways to Competitiveness through Strategic Investments. Available at: https://www.weforum.org/publications/rethinking-ai-sovereignty/
[195] Bain & Company. (2025). How can we meet AI's insatiable demand for compute power? Available at: https://www.bain.com/insights/how-can-we-meet-ais-insatiable-demand-for-compute-power-technology-report-2025/
[196] World Economic Forum. (2026). Rethinking AI Sovereignty: Pathways to Competitiveness through Strategic Investments. Available at: https://www.weforum.org/publications/rethinking-ai-sovereignty/
[197] Ibid.

[198] Haag, A. (2025). The State of AI Competition in Advanced Economies. U.S Federal Reserve Notes. Available at: https://www.federalreserve.gov/econres/notes/feds-notes/the-state-of-ai-competition-in-advanced-economies-20251006.html
[199] Shamir, A & Negele, M. (2025). RAND Europe's response to the EU Cloud and AI Development Act call for evidence. RAND Europe. Available at: https://www.rand.org/content/dam/rand/pubs/testimonies/CTA4100/CTA4156-1/RAND_CTA4156-1.pdf
[200] Oxford Insights. (2025). Government AI Readiness Index 2025. Available at: https://oxfordinsights.com/ai-readiness/government-ai-readiness-index-2025/
[201] Findings of stakeholder consultations carried out within the project.
[202] Mishra, A & Kinning, L. (2025). Reclaiming AI's Promise: A Roadmap for Public Interest AI. Equitech Futures.
[203] Zagorski, D. (2025). Open letter: Harnessing open source AI to advance digital sovereignty. Available at: https://wikimedia.brussels/open-letter-harnessing-open-source-ai-to-advance-digital-sovereignty/
[204] Forum on Information and Democracy. (2024). AI as a public good: Ensuring democratic control of AI in the information space. Available at: https://informationdemocracy.org/wp-content/uploads/2024/03/ID-AI-as-a-Public-Good-Feb-2024.pdf
[205] Negele, M., et al. (2025). Europe and the geopolitics of AGI: The need for a preparedness plan. RAND Corporation. Available at: https://www.rand.org/pubs/research_reports/RRA4636-1.html
[206] Schneider, I. (2025). Reclaiming digital sovereignty: The EU's role in the geopolitics of digital governance (Policy Paper No. 13, Geopolitics of the Digital Era). IE University, Center for the Governance of Change. Available at: https://static.ie.edu/CGC/CGC_ReclaimingDigitalSovereignty_PolicyPaper_01_2025_v2.pdf
[207] Almada, M., & Radu, A. (2024). The Brussels side-effect: How the AI act can reduce the global reach of EU policy. German Law Journal, 25(4), 646-663.
[208] Crum, B. (2025). Brussels effect or experimentalism? The EU AI Act and global standard-setting. Internet Policy Review, 14(3), 1-21.
[209] Hoch, S., & Mügge, D. (2025). EU influence in global AI governance and its limits (RegulAIte Working Paper No. 01/2025). Amsterdam: RegulAIte.
[210] High level panel: Shaping the Future of Science with AI: Europe's Strategic Moment. AI in Science Summit 2025.
[211] Keynote address by Yoshua Bengio, Professor of Computer Science at Université de Montreal: Applying AI in Research, AI in Science Summit 2025.
[212] High level panel: Shaping the Future of Science with AI: Europe's Strategic Moment. AI in Science Summit 2025.
[213] Schneider, I. (2025). Reclaiming digital sovereignty: The EU's role in the geopolitics of digital governance (Policy Paper No. 13, Geopolitics of the Digital Era). IE University, Center for the Governance of Change. Available at: https://static.ie.edu/CGC/CGC_ReclaimingDigitalSovereignty_PolicyPaper_01_2025_v2.pdf
[214] Feldstein, S. (2024). Evaluating Europe's push to enact AI regulations: how will this influence global norms?. Democratization, 31(5), 1049-1066.
[215] Smuha, N. A. (2019). The EU approach to ethics guidelines for trustworthy artificial intelligence. Computer Law Review International, 20(4), 97-106.
[216] European Commission. (2026). Strengthening Europe's Tech Sovereignty | Shaping Europe's digital future. Available at: https://digital-strategy.ec.europa.eu/en/policies/eu-tech-sovereignty
[217] World Economic Forum. (2026). Rethinking AI Sovereignty: Pathways to Competitiveness through Strategic Investments. Available at: https://www.weforum.org/publications/rethinking-ai-sovereignty/
[218] Chartomatsidis, C., 2025, How the DIGITAL Building Blocks can help bring EuroStack's vision of European digital sovereignty to life, European Commission.
[219] Negele, M., et al. (2025). Europe and the geopolitics of AGI: The need for a preparedness plan. RAND Corporation. Available at: https://www.rand.org/pubs/research_reports/RRA4636-1.html
[220] Ibid.
[221] Abendroth Dias, K., et al. (2025). Generative AI outlook report: Exploring the intersection of technology, society and policy (E. Navajas Cawood, M. Vespe, A. Kotsev, & R. Van Bavel, Eds.). Publications Office of the European Union. https://doi.org/10.2760/1109679
[222] Gineikyte-Kanclere, V., Eggert, M., & Skiotyte, G. (2025). European software and cyber dependencies. Policy Department for Transformation, Innovation and Health, Directorate-General for Economy, Transformation and Industry, European Parliament. PE 778.576. Available at: https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778576/ECTI_STU(2025)778576_EN.pdf
[223] Presentation by Cecile Huet (AI Office) in "Workshop 2: Community & Society" of the AI in Science Summit 2025.
[224] Conclusions from the AI in Science Summit 2025.
[225] Ibid.
[226] Y. Bengio, S. Mindermann, et al. (2025). International AI Safety Report. Available at: https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025
[227] Rothschild, D. M., et al. (2025). The Agentic Economy. arXiv. Available at: https://arxiv.org/pdf/2505.15799
[228] Kamal, A., et al. (2024). Agentic AI: The new frontier in GenAI. PwC Middle East. Available at: https://www.pwc.com/m1/en/publications/agentic-ai-the-new-frontier-in-genai.html
[229] Occhipinti, J.-A., et al. (2024). The recessionary pressures of generative AI: A threat to wellbeing. arXiv. Available at: https://arxiv.org/abs/2403.17405
[230] Rothschild, D. M., et al.. (2025). The Agentic Economy. arXiv. Available at: https://arxiv.org/pdf/2505.15799
[231] Cennamo, C., Kretschmer, T., Constantinides, P., Alaimo, C., & Santaló, J. (2023). Digital Platforms Regulation: An Innovation-Centric View of the EU's Digital Markets Act. Journal of European Competition Law & Practice, 14(1), 44-56. Available at: https://academic.oup.com/jeclap/article-abstract/14/1/44/6761966
[232] Arcep (2026). *Generative AI: Challenges for the Future of the Open Internet*. Available at: https://www.arcep.fr/uploads/tx_gspublication/report-generative-AI-challenges-open-internet-january2026.pdf
[233] Chandran, R. (2026). AI agents will exacerbate global economic inequality. Rest of World. Available at: https://restofworld.org/2026/ai-agent-inequality/
[234] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[235] Centre on Regulation in Europe (CERRE). (2025). Is the DMA ready for agentic AI? CERRE. Available at: https://cerre.eu/wp-content/uploads/2025/07/Is-the-DMA-Ready-for-Agentic-AI_Final.pdf
[236] European Commission High-Level Group for the Digital Markets Act. (2024). Public statement on artificial intelligence. Available at: https://digital-markets-act.ec.europa.eu/high-level-group-digital-markets-act-public-statement-artificial-intelligence-2024-05-22_en
[237] Geradin, D. (2023). The Digital Markets Act and E.U. Competition Policy: A Critical Perspective. Journal of European Competition Law & Practice, 14(2), 78-89. Available at: https://link.springer.com/article/10.1007/s40926-022-00213-4

[238] Kim, J. (2025). Advertising in the Age of Agentic AI: Call for Research. Journal of Interactive Advertising, 25(3), 215–221. Available at: https://doi.org/10.1080/15252019.2025.2557107
[239] Arcep (2026). *Generative AI: Challenges for the Future of the Open Internet*. Available at: https://www.arcep.fr/uploads/tx_gspublication/report-generative-AI-challenges-open-internet-january2026.pdf
[240] Singh, A., et al. (2025). Evolution of AI Agent Registry Solutions: Centralized, Enterprise, and Distributed Approaches. arXiv preprint arXiv:2508.03095. Available at: https://arxiv.org/pdf/2508.03095
[241] Feger, A. (2026). GenAI will take over programmatic advertising in 2026 and agentic AI isn't far behind. EMARKETER. Available at: https://www.emarketer.com/content/genai-will-take-over-programmatic-advertising-2026-agentic-ai-isn-t-far-behind
[242] Palladino, C. (2026). How the SaaSpocalypse will hit private equity. Financial Times. Available at: https://www.ft.com/content/9c55f944-456d-4fec-ac05-e57e2e0444b9
[243] Asgar, Z., Nguyen, M., & Katti, S. (2025). Efficient and Scalable Agentic AI with Heterogeneous Systems. arXiv. Available at: https://arxiv.org/pdf/2507.19635?
[244] Merizzi, N., Thomas, C., & Burns, E. (2025). The AI infrastructure reckoning: Optimizing compute strategy in the age of inference economics. In Tech trends 2026 (Chapter 3). Deloitte Insights. Available at: https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/ai-infrastructure-compute-strategy.html
[245] Srivathsan, B., Sorel, M., & Sachdeva, P. (2024). AI power: Expanding data center capacity to meet growing demand. McKinsey & Company. Available at: https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand
[246] Acharya, V. (2025). The strategic imperative of AI consulting in the era of hyperscaler-driven AI infrastructure expansion. Medium. Available at: https://medium.com/@explorify.tech/the-strategic-imperative-of-ai-consulting-in-the-era-of-hyperscaler-driven-ai-infrastructure-b0fb9de5c3d8
[247] Neiger, C. (2025). AI Infrastructure Spending to Hit $490 Billion in 2026 -- Here's Who Wins and Who Loses. Yahoo!Finance. Available at: https://finance.yahoo.com/news/ai-infrastructure-spending-hit-490-091500744.html
[248] TokenMix. (2026). AI API Pricing History: GPT-4 $60 to GPT-5.4 $15 (50x Drop). Available at: https://tokenmix.ai/blog/ai-pricing-trends-history
[249] Stanford University. (2025). Artificial Intelligence Index Report 2025. Available at: https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf
[250] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[251] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[252] D'Ornaro, R. (2025). Second Law of Value in the Agentic Era: Context is the New Moat. Available at: https://www.decodingdiscontinuity.com/p/second-law-of-value-agentic-context-new-moat
[253] New Market Pitch. (2026). Agentic AI Startup Funding 2025-2026. Available at: https://newmarketpitch.com/blogs/news/agentic-ai-funding-analysis
[254] Gross, G. (2026). AI agent evaluations: The hidden cost of deployment. CIO. Available at: https://www.cio.com/article/4123029/ai-agent-evaluations-the-hidden-cost-of-deployment.html
[255] European Commission, Directorate-General for Communications Networks, Content and Technology. (2025). Study on the next data frontier: Generative AI, regulatory compliance and international dimensions — Final report. Publications Office of the European Union. Available at: https://op.europa.eu/en/publication-detail/-/publication/f79578fd-cccf-11f0-8da2-01aa75ed71a1/language-en
[256] Cambridge Judge Business School. (2025). Navigating AI partnerships between big and small tech firms. Available at: https://www.jbs.cam.ac.uk/2025/navigating-ai-partnerships-between-big-and-small-tech-firms/
[257] Lagarde, C. (2025). The transformative power of AI [Speech]. ECB Conference on "The Transformative Power of AI: Economic Implications and Challenges", Frankfurt, Germany. European Central Bank. Available at: https://www.ecb.europa.eu/press/key/date/2025/html/ecb.sp250401_1~d6c9d8df11.en.html
[258] International Labour Organization. (2025). One in four jobs at risk of being transformed by GenAI, new ILO–NASK Global Index shows. Available at: https://www.ilo.org/resource/news/one-four-jobs-risk-being-transformed-genai-new-ilo%E2%80%93nask-global-index-shows
[259] Cedefop. (2025). Skills empower workers in the AI revolution: First findings from Cedefop's AI skills survey. Publications Office of the European Union. Available at: https://www.cedefop.europa.eu/en/publications/9201
[260] Georgieva, K. (2024). AI will transform the global economy: Let's make sure it benefits humanity. IMF Blog. Available at: https://www.imf.org/en/blogs/articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity
[261] Cognizant. (2026). New work, new world 2026: How AI is reshaping work. Available at: https://www.cognizant.com/us/en/aem-i/ai-and-the-future-of-work-report
[262] Pelayo Romero, P., & Utermöhl, K. (2025). Agentic AI: The self-driving economy? Allianz SE. Available at: https://www.allianz.com/en/economic_research/insights/publications/specials_fmo/250916-agentic-ai.html
[263] Coakley, A. (2026). How Europe can survive the AI labor transition. Carnegie Endowment for International Peace. Available at: https://carnegieendowment.org/europe/strategic-europe/2026/02/how-europe-can-survive-the-ai-labor-transition
[264] Lebastard, L., & Sondermann, D. (2026). Artificial intelligence: friend or foe for hiring in Europe today? ECB Blog. Available at: https://www.ecb.europa.eu/press/blog/date/2026/html/ecb.blog20260304~d9e34fc95f.de.html
[265] Toner, H., et al. (2024). Through the chat window and into the real world: Preparing for AI agents. Center for Security and Emerging Technology. Available at: https://cset.georgetown.edu/publication/through-the-chat-window-and-into-the-real-world-preparing-for-ai-agents/
[266] Mariscal de Gante Martín, A., & Villani, D. (2026). Algorithmic management and working conditions in Europe: Evidence from the AIM-WORK Survey (JRC Working Papers Series on Labour, Education and Technology No. 2026/05, JRC147505). European Commission. Available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC147505
[267] Ibid.
[268] European Commission. (2025). Apply AI Strategy. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[269] Barcevičius, E., et al. (2025). Background document: Input for the Global Multistakeholder High-Level Conference on Governance for Web 4.0 and Virtual Worlds, 31 March – 1 April 2025. European Commission, PPMI, TNO, & WEB4HUB. Available at: https://ec.europa.eu/newsroom/dae/redirection/document/113701
[270] PPMI & TNO. (2025). *Outcome of the Global Multistakeholder High-Level Conference on Governance of Web 4.0 and Virtual Worlds*. Available at: https://ppmi.lt/news-insights/outcome-global-multistakeholder-high-level-conference-governance-web-40-and-virtual-worlds
[271] Barcevičius, E., et al. (2025). Background document: Input for the Global Multistakeholder High-Level Conference on Governance for Web 4.0 and Virtual Worlds, 31 March – 1 April 2025. European Commission, PPMI, TNO, & WEB4HUB. Available at: https://ec.europa.eu/newsroom/dae/redirection/document/113701

[272] PPMI & TNO. (2025). *Outcome of the Global Multistakeholder High-Level Conference on Governance of Web 4.0 and Virtual Worlds*. Available at: https://ppmi.lt/news-insights/outcome-global-multistakeholder-high-level-conference-governance-web-40-and-virtual-worlds
[273] Sharp, M., et al. (2025). Agentic Inequality. arXiv. Available at: https://arxiv.org/pdf/2510.16853
[274] Ibid.
[275] Chandran, R. (2026). AI agents will exacerbate global economic inequality. Rest of World. Available at: https://restofworld.org/2026/ai-agent-inequality/
[276] World Bank. (2025). Digital progress and trends report 2025: AI foundations. Available at: https://www.worldbank.org/en/publication/dptr2025-ai-foundations
[277] Ibid.
[278] WEF (2026). The global risks report 2026 (21st ed.). Available at: https://reports.weforum.org/docs/WEF_Global_Risks_Report_2026.pdf
[279] Tony Blair Institute for Global Change. (2024). Bridging the digital divide in Africa: The promising role of LEO satellites. Available at: https://institute.global/insights/tech-and-digitalisation/bridging-the-digital-divide-in-africa-the-promising-role-of-leo-satellites
[280] Dou, S., et al. (2025). SpaceCache+: Towards pervasive content delivery via low-Earth orbit mega-constellations. IEEE Transactions on Services Computing. Available at: https://ieeexplore.ieee.org/abstract/document/11095828
[281] TencentCloud. (2025). What are the advantages of edge computing in AI agent deployment? Available at: https://www.tencentcloud.com/techpedia/126538
[282] Interview ID IT_4. (2026).
[283] Brown, J. (2026). I traced who owns the undersea cables that carry 95% of global internet traffic — the map is a colonial one. Silicon Canals. Available at: https://siliconcanals.com/sc-d-i-traced-who-owns-the-undersea-cables-that-carry-95-of-global-internet-traffic-the-map-is-a-colonial-one/
[284] Khanal, S., Zhang, H., & Taeihagh, A. (2025). Why and how is the power of Big Tech increasing in the policy process? The case of generative AI. Policy and Society, 44(1), 52-69. Available at: https://academic.oup.com/policyandsociety/article/44/1/52/7636223?login=false
[285] Chandran, R. (2026). AI agents will exacerbate global economic inequality. Rest of World. Available at: https://restofworld.org/2026/ai-agent-inequality/
[286] Interview ID IT_4. (2026).
[287] Kende, M. (2024). The AI revolution will reinforce regional disadvantages unless steps are taken to close the data divide. Analysys Mason. Available at: https://www.analysysmason.com/about-us/news/newsletter/ai-data-divide/
[288] Chandran, R. (2026). AI agents will exacerbate global economic inequality. Rest of World. Available at: https://restofworld.org/2026/ai-agent-inequality/
[289] Ibid.
[290] Raposo, V. L. (2024). When facial recognition does not 'recognise': erroneous identifications and resulting liabilities. AI & SOCIETY, 39(4), 1857-1869.
[291] US General Services Administration (2022). Executive Order 13985—Equity Action Plan (January 20, 2022). Available at: https://www.gsa.gov/cdnstatic/GSA_Equity_Action_Plan_2022.pdf
[292] European Union Agency for Fundamental Rights. (2022). *Bias in algorithms: Artificial intelligence and discrimination*. Publications Office of the European Union. Available at: https://fra.europa.eu/sites/default/files/fra_uploads/fra-2022-bias-in-algorithms_en.pdf
[293] Stanford University Human-Centered Artificial Intelligence. (2021). *How flawed data aggravates inequality in credit*. Available at: https://hai.stanford.edu/news/how-flawed-data-aggravates-inequality-credit
[294] Findings of stakeholder consultations carried out within the project.
[295] Ocheltree, K. (n.d.). Bridging the digital divide with artificial intelligence. LTEN. Available at: https://www.l-ten.org/Web/Web/News---Insights/focus-articles/Bridging-the-Digital-Divide-with-Artificial-Intelligence.aspx
[296] MIT Technology Review Insights. (2025). Designing digital resilience in the agentic AI era. MIT Technology Review. Available at: https://www.technologyreview.com/2025/11/20/1127941/designing-digital-resilience-in-the-agentic-ai-era/
[297] Chandran, R. (2026). AI agents will exacerbate global economic inequality. Rest of World. Available at: https://restofworld.org/2026/ai-agent-inequality/
[298] Interview IDs, IT_1 & IT_2. (2026).
[299] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[300] See more at: https://digital-strategy.ec.europa.eu/en/policies/digital-skills-academies
[301] Bogmans, C., et al. (2025). AI needs more abundant power supplies to keep driving economic growth. IMF Blog. Available at: https://www.imf.org/en/blogs/articles/2025/05/13/ai-needs-more-abundant-power-supplies-to-keep-driving-economic-growth
[302] International Energy Agency. (2025). Energy and AI (World Energy Outlook Special Report). IEA. Available at: https://www.iea.org/reports/energy-and-ai
[303] Colom, A. et al. (2026). Environmentally Sustainable Digital Sovereignty. JRC. Available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC147383
[304] International Telecommunication Union. (2025). The annual AI governance report 2025: Steering the future of AI. ITU. Available at: https://www.itu.int/epublications/publication/the-annual-ai-governance-report-2025-steering-the-future-of-ai
[305] ibid.
[306] Granskog, A., et al. (2024). The role of power in unlocking the European AI revolution. McKinsey & Company. Available at: https://www.mckinsey.com/industries/electric-power-and-natural-gas/our-insights/the-role-of-power-in-unlocking-the-european-ai-revolution
[307] Interview ID IT_8. (2026)
[308] Naveen, S., Kounte, M. R., & Ahmed, M. R. (2021). Low latency deep learning inference model for distributed intelligent IoT edge clusters. IEEE Access, 9, 160607-160621. Available at: https://ieeexplore.ieee.org/iel7/6287639/6514899/09628063.pdf
[309] Alcott, B. (2005). Jevons' paradox. Ecological economics, 54(1), 9-21. Available at: https://www.sciencedirect.com/science/article/pii/S0921800905001084
[310] Yañez-Barnuevo, M. (2025). Data centers and water consumption. Environmental and Energy Study Institute. Available at: https://www.eesi.org/articles/view/data-centers-and-water-consumption
[311] Hope, B. (2025). Why AI is bringing the water wars. The Closer. Available at: https://www.thecloser.fm/why-ai-is-bringing-the-water-wars/
[312] Future Bridge NetZero. (n.d.). AI data centres are overloading Europe's power grid: What comes next. Available at: https://netzero-events.com/ai-data-centres-energy-crunch-europe/
[313] Colom, A. et al. (2026). Environmentally Sustainable Digital Sovereignty. JRC. Available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC147383

[314] Kande, M. (2025). Two keys for governments to seize the potential of the Intelligence Age. World Governments Summit. Available at: https://www.worldgovernmentssummit.org/observer/articles/detail/two-keys-for-governments-to-seize-the-potential-of-the-intelligence-age
[315] Ibragimova, N. (2025). AI Runs on Rocks: Why Central Asia Belongs in America's Critical Minerals Strategy. The Hague Research Institute for Eastern Europe, the South Caucasus & Central Asia. Available at: https://hagueresearch.org/ai-runs-on-rocks-why-central-asia-belongs-in-americas-critical-minerals-strategy/
[316] Kim, T.-Y., et al. (2025). With new export controls on critical minerals, supply concentration risks become reality. International Energy Agency. Available at: https://www.iea.org/commentaries/with-new-export-controls-on-critical-minerals-supply-concentration-risks-become-reality
[317] Baskaran, G. (2025). China's new rare earth and magnet restrictions threaten U.S. defense supply chains. Center for Strategic and International Studies. Available at: https://www.csis.org/analysis/chinas-new-rare-earth-and-magnet-restrictions-threaten-us-defense-supply-chains
[318] Colom, A. et al. (2026). Environmentally Sustainable Digital Sovereignty. JRC. Available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC147383
[319] Ambrose, J. (2024). Ireland's datacentres overtake electricity use of all homes combined, figures show. The Guardian. Available at: https://www.theguardian.com/world/article/2024/jul/23/ireland-datacentres-overtake-electricity-use-of-all-homes-combined-figures-show
[320] Bloomberg. (2025). Ireland ends its moratorium on new power links to data centers. Politico Pro. Available at: https://subscriber.politicopro.com/article/eenews/2025/12/15/ireland-ends-its-moratorium-on-new-power-links-to-data-centers-00688969
[321] IDCA. (2022). The Netherlands temporarily bans the construction of new hyperscale data centers. Available at: https://www.idc-a.org/news/The-Netherlands-temporarily-bans-the-construction-of-new-hyperscale-data-centers
[322] World Economic Forum. (2026). Rethinking AI Sovereignty: Pathways to Competitiveness through Strategic Investments. Available at: https://www.weforum.org/publications/rethinking-ai-sovereignty/
[323] Leone de Castris, A. (2025). AI governance around the world: European Union. The Alan Turing Institute. https://doi.org/10.5281/zenodo.17054419
[324] Interview IDs IT_1, IT_3. (2026)
[325] ITU. (2025). The Annual AI Governance Report 2025: Steering the Future of AI. Available at: https://www.itu.int/epublications/en/publication/the-annual-ai-governance-report-2025-steering-the-future-of-ai/en
[326] European Commission. (2026). Strengthening Europe's Tech Sovereignty | Shaping Europe's digital future. Available at: https://digital-strategy.ec.europa.eu/en/policies/eu-tech-sovereignty
[327] Osmond, M., & Jego, T. (2026). Mind The Gap: How The Technical Mechanism Of Agentic AI Outpace Global Legal Frameworks. *arXiv preprint arXiv:2603.27075*.
[328] Interview IDs IT_6, IT_9. (2026)
[329] Negele, M., Juijn, D., Shamir, A., Janků, D., Özcan, B., Soder, L., Velasco, L., Reddel, M., Bakker, M., Pacchiardi, L., & Andriushchenko, M. (2025). *Europe and the geopolitics of AGI: The need for a preparedness plan*. RAND Corporation. Available at: https://www.rand.org/pubs/research_reports/RRA4636-1.html
[330] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[331] Rodriguez Müller, P. et al. (2026). The European approach to artificial intelligence policymaking. European Commission. Available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC146313
[332] AI Security Institute. (2025). International AI Safety Report 2025. Available at: https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025
[333] Interview ID IT_3. (2026)
[334] Azizi, O. (2025). AI and the New Digital Feudalism: Users as Serfs in Cognitive Estates. Available at: https://www.researchgate.net/publication/393999111_AI_and_the_New_Digital_Feudalism_Users_as_Serfs_in_Cognitive_Estates
[335] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[336] Interview IDs IT_1, IT_4, IT_9 (2016).
[337] Barcevičius, E., et al. (2025). Background document: Input for the Global Multistakeholder High-Level Conference on Governance for Web 4.0 and Virtual Worlds, 31 March – 1 April 2025. European Commission, PPMI, TNO, & WEB4HUB. Available at: https://ec.europa.eu/newsroom/dae/redirection/document/113701
[338] Interview conducted as part of the stakeholder consultation in the run up to the "Global Multistakeholder High Level Conference on the governance of Web 4.0 and Virtual Worlds" on 31 March and 1 April 2025. More information available at: https://digital-strategy.ec.europa.eu/en/policies/event-web-4-governance
[339] Isenberg, R., & Rahilly, L. (2026). Trust in the age of agents [Podcast]. The McKinsey Podcast. McKinsey & Company. Available at: https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/trust-in-the-age-of-agents
[340] Barcevičius, E., Gabaliņa, R., Butym, M., Wójtowicz, A., & Zepcan, A. (2026). Digital identity and identifiers in the era of Web 4.0 and virtual worlds: Thematic paper. Publications Office of the European Union. https://doi.org/10.2759/6450018.
[341] StepUp StartUps. (2026). European agentic AI landscape. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[342] eyreACT. (2026). AI Agents and EU AI Act: What's Changing? Available at: https://eyreact.com/ai-agents-eu-ai-act/
[343] Interview IDs IT_1, IT_4. (2026)
[344] Interview ID IT_7. (2026)
[345] To be explored in more depth in a future thematic paper of the project.
[346] Virtual Worlds Partnership. (2025). Strategic research & innovation agenda: Virtual worlds — Solving real world problems. Available at: https://www.virtualworldsassociation.eu/actions/strategic-research-innovation-agenda-virtual-worlds-eu
[347] Donnelly, R. (2026). EU AI Act: Compliance Flexibility & Strategy. Enzai. Available at: https://www.enz.ai/blog/the-eu-ai-act-bends-for-agentic-ai-it-need-not-break
[348] StepUp StartUps. (2026). European agentic AI landscape. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[349] Interview ID IT_7. (2026)
[350] AEPD. (2026). Agentic Artificial Intelligence from the Perspective of Data Protection. Available at: https://www.aepd.es/en/guides/agentic-artificial-intelligence.pdf
[351] StepUp StartUps. (2026). European agentic AI landscape. Available at: https://digital-strategy.ec.europa.eu/en/library/agentic-ai-leveraging-european-ai-talent-and-regulatory-assets-scale-adoption
[352] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.

[353] Barcevičius, E., Gabaliņa, R., Butym, M., Wójtowicz, A., & Zepcan, A. (2026). Digital identity and identifiers in the era of Web 4.0 and virtual worlds: Thematic paper. Publications Office of the European Union. https://doi.org/10.2759/6450018.
[354] Findings of stakeholder consultations carried out within the project.
[355] Ibid.
[356] Ibid.
357 ITU. (2025). The Annual AI Governance Report 2025: Steering the Future of AI. Available at: https://www.itu.int/epublications/en/publication/the-annual-ai-governance-report-2025-steering-the-future-of-ai/en
[358] Interview ID IT_2. (2026)
[359] Interview IDs IT_1, IT_3, IT_9. (2026)
[360] Ashman, N. et al. (2026). EU Regulations Are Not Ready for Multi-Agent AI Incidents. TechPolicy.Press. Available at: https://www.techpolicy.press/eu-regulations-are-not-ready-for-multiagent-ai-incidents/
[361] Findings of stakeholder consultations carried out within the project.
[362] Findings of stakeholder consultations carried out within the project.
[363] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[364] Id.
[365] Findings of stakeholder consultations carried out within the project.
[366] Arcep (2026). Generative AI: Challenges for the Future of the Open Internet. Available at: https://www.arcep.fr/uploads/tx_gspublication/report-generative-AI-challenges-open-internet-january2026.pdf
[367] WEB4HUB expert roundtable. (2026, 2 July). Opportunities for Europe in the Web 4.0 and Agentic AI Era.
[368] Ibid.
[369] Interview IDs IT_1, IT_2, IT_8 (2016).
[370] Rodriguez Müller, P. et al. (2026). The European approach to artificial intelligence policymaking. European Commission. Available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC146313
[371] Coalition for Secure AI. (2025). Announcing the CoSAI Principles for Secure-by-Design Agentic Systems. Available at: https://www.coalitionforsecureai.org/announcing-the-cosai-principles-for-secure-by-design-agentic-systems/
[372] Responsible AI Institute. (2025). New Agentic AI Health Initiative Launched in the UK. Available at: https://www.responsible.ai/news/trustx-press-release/
[373] High level panel: Shaping the Future of Science with AI: Europe's Strategic Moment. AI in Science Summit 2025.
[374] Colom, A. et al. (2026). Environmentally Sustainable Digital Sovereignty. JRC. Available at: https://publications.jrc.ec.europa.eu/repository/handle/JRC147383
375 Eurostat. (n.d.). Artificial intelligence by size class of enterprise. Available at: https://ec.europa.eu/eurostat/databrowser/bookmark/9dd0df02-604e-4600-8169-76985f5fab92?lang=en&createdAt=2026-02-06T11:28:19Z
[376] Yee, L., et al. (2025). McKinsey Technology Trends Outlook 2025. McKinsey & Company. Available at: https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-top-trends-in-tech
[377] World Economic Forum. (2026). Rethinking AI Sovereignty: Pathways to Competitiveness through Strategic Investments. Available at: https://www.weforum.org/publications/rethinking-ai-sovereignty/
[378] Interview IDs IT_3, IT_8, IT_9.
[379] Findings of stakeholder consultations carried out within the project.
[380] Findings of stakeholder consultations carried out within the project.
[381] European Commission. (2026). AI Office publishes frontier AI expert findings on EU competitiveness, sovereignty and security. Available at: https://digital-strategy.ec.europa.eu/en/library/ai-office-publishes-frontier-ai-expert-findings-eu-competitiveness-sovereignty-and-security
[382] Sukhaversky, A. et al. (2024). Time to place our bets: Europe's AI opportunity. McKinsey Global Institute. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/time-to-place-our-bets-europes-ai-opportunity
[383] Findings of stakeholder consultations carried out within the project.
[384] European Commission. (2026). AI Office publishes frontier AI expert findings on EU competitiveness, sovereignty and security. Available at: https://digital-strategy.ec.europa.eu/en/library/ai-office-publishes-frontier-ai-expert-findings-eu-competitiveness-sovereignty-and-security
[385] Ibid.
[386] Conclusions from the AI in Science Summit 2025.
[387] European Commission. (2026). AI Office publishes frontier AI expert findings on EU competitiveness, sovereignty and security. Available at: https://digital-strategy.ec.europa.eu/en/library/ai-office-publishes-frontier-ai-expert-findings-eu-competitiveness-sovereignty-and-security
[388] European Commission. (2026). Strengthening Europe's Tech Sovereignty | Shaping Europe's digital future. Available at: https://digital-strategy.ec.europa.eu/en/policies/eu-tech-sovereignty
[389] World Economic Forum. (2026). Rethinking AI Sovereignty: Pathways to Competitiveness through Strategic Investments. Available at: https://www.weforum.org/publications/rethinking-ai-sovereignty/
[390] Negele, M., Juijn, D., Shamir, A., Janků, D., Özcan, B., Soder, L., Velasco, L., Reddel, M., Bakker, M., Pacchiardi, L., & Andriushchenko, M. (2025). Europe and the geopolitics of AGI: The need for a preparedness plan. RAND Corporation. Available at: https://www.rand.org/pubs/research_reports/RRA4636-1.html
[391] Arcep (2026). *Generative AI: Challenges for the Future of the Open Internet*. Available at: https://www.arcep.fr/uploads/tx_gspublication/report-generative-AI-challenges-open-internet-january2026.pdf
[392] Barcevičius, E., et al. (2025). Background document: Input for the Global Multistakeholder High-Level Conference on Governance for Web 4.0 and Virtual Worlds, 31 March – 1 April 2025. European Commission, PPMI, TNO, & WEB4HUB. Available at: https://ec.europa.eu/newsroom/dae/redirection/document/113701
[393] Beyond the interviews listed above, the paper incorporates insights from confidential stakeholder engagements, which cannot be attributed to named individuals or organisations.
[394] A plausible, internally consistent description of how key uncertainties and structural drivers may interact. Scenarios are used to explore alternative pathways and test policy robustness, not to predict future developments.
[395] A scenario-building technique that structures uncertainty around two critical and independent variables, generating four distinct but plausible combinations for analysis.
[396] Hobbs, H., et al. (2026). Exploring possible AI trajectories through 2030 (OECD Artificial Intelligence Papers, No. 55). OECD Publishing. https://doi.org/10.1787/cb41117a-en
[397] Public beta version. More information at: https://theforecastingmachine.com/
[398] Virodhula, D. V. (2025). AI Agents: The Future of Autonomous Intelligence. Available at: https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=5214616

[399] Toner, H., et al. (2024). Through the Chat Window and Into the Real World: Preparing for AI Agents. *Center for Security and Emerging Technology. https://doi. org/10.51593/20240034*.
[400] Andre, D. (2025). Hierarchical AI agents: Redefining Task Management in Artificial Intelligence. All About AI. Available at: https://www.allaboutai.com/ai-agents/hierarchical-agents/
[401] Kokotajlo, D., et al. (2025). AI 2027. Available at: https://ai-2027.com
[402] Logan, D., & Zaidi, E. (2024). Generating Value Together: From Fundamentals to AI Readiness to Collective Intelligence [Video]. Gartner Data & Analytics Summit Opening Keynote. Available at: https://www.youtube.com/watch?v=uVoPJdltu14
[403] Drenik, G. (2025). Agentic AI Is Transforming Enterprises—Here's Where To Start. Forbes. Available at: https://www.forbes.com/sites/garydrenik/2025/04/10/agentic-ai-is-transforming-enterprises-heres-where-to-start/
[404] Gartner. (2025). Gartner Predicts Agentic AI Will Autonomously Resolve 80% of Common Customer Service Issues Without Human Intervention by 2029. Available at: https://www.gartner.com/en/newsroom/press-releases/2025-03-05-gartner-predicts-agentic-ai-will-autonomously-resolve-80-percent-of-common-customer-service-issues-without-human-intervention-by-20290
[405] Rubin, B. F. (2025). What is agentic commerce? Your guide to AI-assisted retail. Mastercard. Available at: https://www.mastercard.com/global/en/news-and-trends/stories/2025/agentic-commerce-explainer.html
[406] PYMNTS. (2025). Shopify President: Agentic AI Drives 11x Spike in Orders. Available at: https://www.pymnts.com/news/ecommerce/2025/shopify-president-agentic-ai-drives-11x-spike-in-orders
[407] Morgan Stanley. (2025). Here come the shopping bots. Available at: www.morganstanley.com
[408] Garvey, K. et al. (2024). How Agentic AI will transform financial services. World Economic Forum. Available at: https://www.weforum.org/stories/2024/12/agentic-ai-financial-services-autonomy-efficiency-and-inclusion/
[409] Stanford University. (2026). The 2025 AI Index Report. Available at: https://hai.stanford.edu/ai-index/2025-ai-index-report
[410] IBM. (2025). Scaling supply chain resilience: Agentic AI for autonomous operations. Available at: https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/supply-chain-ai-automation-oracle
[411] DHL Group. (2025). DHL boosts operational efficiency and customer communications with HappyRobot's AI Agents. Available at: https://group.dhl.com/en/media-relations/press-releases/2025/dhl-boosts-operational-efficiency-and-customer-communications-with-happyrobots-ai-agents.html
[412] Fujitsu. (2025). AI Agents and the Pathway to Evolving Intelligent Manufacturing. Available at: https://global.fujitsu/en-global/technology/key-technologies/news/ta-ai-agents-intelligent-manufacturing-20250625
[413] Siemens. (2026). Siemens unveils technologies to accelerate the industrial AI revolution at CES 2026. Available at : https://press.siemens.com/global/en/pressrelease/siemens-unveils-technologies-accelerate-industrial-ai-revolution-ces-2026
[414] AI on Demand. (n.d.). AIPlan4Grid. Available at: https://www.ai4europe.eu/business-and-industry/case-studies/aiplan4grid
[415] Anglen, J. (n.d.). AI Agents Revolutionizing Grid Stability 2025. Available at: https://www.rapidinnovation.io/post/ai-agents-for-grid-stability
[416] See, for example, https://gridsingularity.com/
[417] United Nations University. (2025). Unlocking AI's potential to serve humanity. Available at: https://unu.edu/cpr/project/unlocking-ais-potential-serve-humanity
[418] Ericsson. (2025). Telenor and Ericsson successfully create Agentic AI. Available at: https://www.ericsson.com/en/press-releases/3/2025/telenor-and-ericsson-successfully-co-create-agentic-ai-for-radio-controls-and-expand-innovation-partnership
[419] Reding, D.,F., et al. (2023). Science & Technology trends 2023-2043. NATO Science & Technology Organization. https://www.nato.int/nato_static_fl2014/assets/pdf/2023/3/pdf/stt23-vol2.pdf
[420] More information available at: https://www.thalesgroup.com/en/news-centre/press-releases/protecting-aircraft-artificial-intelligence-thales-and-partners-selected
[421] The Tech Buzz. (2026). Scout AI Demos Lethal Autonomous Weapons Powered by AI Agents. Available at: https://www.techbuzz.ai/articles/scout-ai-demos-lethal-autonomous-weapons-powered-by-ai-agents
[422] Karunanayake, N. (2025). Next-generation agentic AI for transforming healthcare. *Informatics and Health*, *2*(2), 73-83.
[423] Laviola, E. (2025). What Is Agentic AI, and How Can It Be Used in Healthcare? HealthTech. Available at: https://healthtechmagazine.net/article/2025/05/what-is-agentic-ai-in-healthcare-perfcon
[424] United Nations University. (2025). Unlocking AI's potential to serve humanity. Available at: https://unu.edu/cpr/project/unlocking-ais-potential-serve-humanity
[425] Virodhula, D. V. (2025). AI Agents: The Future of Autonomous Intelligence. Available at: https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=5214616
[426] United Nations University. (2025). Unlocking AI's potential to serve humanity. Available at: https://unu.edu/cpr/project/unlocking-ais-potential-serve-humanity
[427] Futureweek. (2026). Week In Review: Cannes Lions 2026, Omnicom Media and Netflix Partner, Gap Announces AI Transformation, and Adobe Enters a Series of Partnerships https://futureweek.com/week-in-review-cannes-lions-2026-omnicom-media-and-netflix-partner-gap-announces-ai-transformation-and-adobe-enters-a-series-of-partnerships/
[428] Srinivasu, P. N., Pavate, A., JayaLakshmi, G., Shafi, J., Choi, J., & Ijaz, M. F. (2025). Agentic AI for smart and sustainable precision agriculture. *Frontiers in Plant Science*, *16*, 1706428.
[429] Pargesoft. (n.d.). Agentic Food & Agriculture Solutions. Available at: https://www.pargesoft.co.uk/en/industries/food-agriculture/
[430] Digiqt Blog. (2026). 5 AI Agents in Autonomous Driving. Available at: https://digiqt.com/blog/ai-agents-in-autonomous-driving/
[431] Lambert, F. (2026). Waymo begins fully autonomous ops with 6th-gen Driver, targets 1M weekly rides | Electrek. Available at: https://electrek.co/2026/02/12/waymo-begins-fully-autonomous-ops-with-6th-gen-driver-targets-1m-weekly-rides/
[432] Robotics Herald. (2025). First Agentic AI Robotics. Available at: https://roboticsherald.com/first-agentic-ai-robotics/
[433] Karunanayake, N. (2025). *Next-generation agentic AI for transforming healthcare. Informatics and Health, 2 (2), 73-83.*
[434] IDC. (2025). The Agentic Evolution of Enterprise Applications. Available at: https://www.idc.com/resource-center/blog/the-agentic-evolution-of-enterprise-applications/
[435] Palladino, C. (2026). How the SaaSpocalypse will hit private equity. Financial Times. Available at: https://www.ft.com/content/9c55f944-456d-4fec-ac05-e57e2e0444b9
[436] ServiceNow. Available at: https://www.servicenow.com/
[437] Zendesk. Available at: https://www.zendesk.com/
[438] Interview ID IT_4. (2026).
[439] Davies, P. (2019). The demon in the machine. Penguin Books.
[440] Ibid.
[441] Ibid.

[442] Bennett, C. H. (1973). Logical reversibility of computation. IBM journal of Research and Development, 17(6), 525-532. Available at: https://www.cs.princeton.edu/courses/archive/fall04/cos576/papers/bennett73.html
[443] Enriquez, R. P. H. (2026). The Fluctuation-Dissipation Compilation Theorem: A Unified Framework for Thermodynamically Optimal Computing. Available at: https://www.researchsquare.com/article/rs-8653565/v1
[444] Carnot, S. (1824). Reflections on the motive power of fire, and on machines fitted to develop that power. Paris: Bachelier, 108(1824), 1824. Available at: https://www.asme.org/getmedia/e29ff2e8-c1b0-419b-8a15-7648a258aa65/carnot-brochure-final-2021.pdf
[445] Schmidt, R., & Iyengar, M. (2009). Thermodynamics of information technology data centers. IBM Journal of Research and Development, 53(3), 9-1. Available at: https://ieeexplore.ieee.org/abstract/document/5429022/
[446] Brady, G. A., Kapur, N., Summers, J. L., & Thompson, H. M. (2013). A case study and critical assessment in calculating power usage effectiveness for a data centre. Energy Conversion and Management, 76, 155-161. Available at: https://www.sciencedirect.com/science/article/pii/S0196890413004068
[447] Shilov, A. (2024). Nvidia's next-gen AI GPUs could draw an astounding 1000 watts each, a 40 percent increase — Dell spills the beans on B100 and B200 in its earnings call. Tom's Hardware. Available at: https://www.tomshardware.com/tech-industry/artificial-intelligence/nvidias-b100-and-b200-processors-could-draw-an-astounding-1000-watts-per-gpu-dell-spills-the-beans-in-earnings-call
[448] Shilov, A. (2025). Future AI processors said to consume up to 15,360 watts of power — massive power draw will demand exotic immersion and embedded cooling tech. Tom's Hardware. Available at: https://www.tomshardware.com/pc-components/cooling/future-ai-processors-said-to-consume-up-to-15-360w-massive-power-draw-will-demand-exotic-immersion-and-embedded-cooling-tech
[449] Bolgar, C. (2025). AI chips are getting hotter. A microfluidics breakthrough goes straight to the silicon to cool up to three times better. Microsoft Source. Available at: https://news.microsoft.com/source/features/innovation/microfluidics-liquid-cooling-ai-chips/
[450] Marques Santos, A., Molica, F., & Torrecilla Salinas, C. (2025). EU-funded investment in artificial intelligence and regional specialization. Regional Science Policy & Practice, 17(7), 100190. Elsevier. https://doi.org/10.1016/j.rspp.2025.100190; based on OECD AI policy observatory.
[451] European Commission. (2025). Apply AI Strategy (COM(2025) 723 final). Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025DC0723
[452] Ibid.
[453] Pearcy, S. (2025). Governing agentic AI. HiddenLayer. Available at: https://www.hiddenlayer.com/insight/governing-agentic-ai
[454] Oueslati, A., & Staes-Polet, R. (2025). Ahead of the curve: Governing AI agents under the EU AI Act. The Future Society. Available at: https://thefuturesociety.org/wp-content/uploads/2023/04/Report-Ahead-of-the-Curve-Governing-AI-Agents-Under-the-EU-AI-Act-4-June-2025.pdf
[455] Nightingale, P., & Kirkhope, A. (2025). Agents for good? Reconciling agentic AI with existing AI governance frameworks. techUK. Available at: https://www.techuk.org/resource/agents-for-good-reconciling-agentic-ai-with-existing-ai-governance-frameworks.html
[456] Leone de Castris, A. (2025). AI governance around the world: European Union. The Alan Turing Institute. https://doi.org/10.5281/zenodo.17054419
[457] GDPR info. Available at: https://gdpr-info.eu
[458] European Commission. (n.d.). Data Act. Available at: https://digital-strategy.ec.europa.eu/en/policies/data-act
[459] European Commission. (2022). Data Governance Act (Regulation (EU) 2022/868). EUR-Lex. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32022R0868
[460] European Commission. (n.d.). Digital Services Act. Available at: https://digital-strategy.ec.europa.eu/en/policies/digital-services-act
[461] European Commission. (n.d.). Digital Markets Act. Available at: https://digital-markets-act.ec.europa.eu/index_en
[462] European Commission. (n.d.). eIDAS Regulation. Available at: https://digital-strategy.ec.europa.eu/en/policies/eidas-regulation
[463] European Commission. (n.d.). NIS2 Directive. Available at: https://digital-strategy.ec.europa.eu/en/policies/nis2-directive
[464] G7. (2023). Hiroshima Process international guiding principles for organizations developing advanced AI systems. Ministry of Foreign Affairs of Japan. Available at: https://www.mofa.go.jp/files/100573471.pdf
[465] Ministry of Internal Affairs and Communications, Japan. (2025). Latest update on the Hiroshima AI Process Friends Group and Partners' Community. Available at: https://www.soumu.go.jp/main_sosiki/joho_tsusin/eng/pressrelease/2025/10/29_1.html
[466] OECD. (n.d.). OECD AI Principles. Available at: https://www.oecd.org/en/topics/sub-issues/ai-principles.html
[467] OECD. (n.d.). About GPAI. Available at: https://oecd.ai/en/about/about-gpai
[468] UNESCO. (n.d.). Recommendation on the ethics of artificial intelligence. Available at: https://www.unesco.org/en/artificial-intelligence/recommendation-ethics
[469] Slotin, J. (2024). The Global Digital Compact: A win for digital cooperation. Global Partnership for Sustainable Development Data. Available at: https://www.data4sdgs.org/blog/global-digital-compact-win-digital-cooperation
[470] United Nations Regional Information Centre. (2026). General Assembly appoints artificial intelligence panel. Available at: https://unric.org/en/general-assembly-appoints-artificial-intelligence-panel/
[471] UN General Assembly (2025). Outcome document of the high-level meeting of the General Assembly on the overall review of the implementtation of the outcomes of the World Summit on the Information Society.
[472] While the document does not reference agentic AI explicitly, its emphasis on bridging AI divides in access to compute, models, and skills addresses several of the foundational inequalities that agentic deployment is likely to amplify.
[473] Bengio, Y., et al. (2025). International AI safety report 2025. Available at: https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025
[474] Allen, G. C., & Adamson, G. (2024). The AI Safety Institute International Network: Next steps and recommendations. Center for Strategic and International Studies. Available at: https://www.csis.org/analysis/ai-safety-institute-international-network-next-steps-and-recommendations
[475] Council of Europe. (n.d.). The Framework Convention on Artificial Intelligence. Available at: https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence
[476] Carugi, M. (2025). ITU-T SG13 AI activities and focus groups [Conference presentation]. ETSI AI Conference 2025, Sophia Antipolis, France. Available at: https://docbox.etsi.org/Workshop/2025/02_AICONFERENCE/SESSION07/ITU-T%20SG13_CARUGI_MARCO.pdf
[477] ITU. (n.d.). Joint Coordination Activity on Machine Learning (JCA-ML). Available at: https://www.itu.int/en/ITU-T/jca/ml/Pages/default.aspx
[478] ITU. (n.d.). AI Standards Exchange. AI for Good. Available at: https://aiforgood.itu.int/ai-standards-exchange/
[479] Microsoft. (n.d.). ISO/IEC 42001:2023 artificial intelligence management system. Microsoft Learn. Available at: https://learn.microsoft.com/en-us/compliance/regulatory/offering-iso-42001
[480] ITU. (2019). Architectural framework for machine learning in future networks including IMT-2020 (Recommendation ITU-T Y.3172). Available at: https://www.itu.int/itu-t/recommendations/rec.aspx?rec=13894

[481] Barcevičius, E., et al. (2025). Background document: Input for the Global Multistakeholder High-Level Conference on Governance for Web 4.0 and Virtual Worlds, 31 March – 1 April 2025. European Commission, PPMI, TNO, & WEB4HUB. Available at: https://ec.europa.eu/newsroom/dae/redirection/document/113701
[482] Martinez-Julia, P., Homma, S., & Lopez, D.R. (2024). Artificial Intelligence framework for network management. https://www.ietf.org/archive/id/draft-pedro-nmrg-ai-framework-05.html
[483] IEEE (no date). Ethics in Action: Advancing ethical AI and technology. IEEE. Available at: https://ethicsinaction.ieee.org/
[484] Leikas, J., Koivisto, R., & Gotcheva, N. (2019). Ethical framework for designing autonomous intelligent systems. Journal of Open Innovation: Technology, Market, and Complexity, 5(1), 18. Available at: https://www.mdpi.com/2199-8531/5/1/18/pdf
[485] ETSI. (2025). ETSI TR 104 221 V1.1.1 – Securing Artificial Intelligence (SAI): Problem Statement. ETSI. Available at: https://www.etsi.org/deliver/etsi_tr/104200_104299/104221/01.01.01_60/tr_104221v010101p.pdf.
[486] Leone de Castris, A. (2025). AI governance around the world: European Union. The Alan Turing Institute. https://doi.org/10.5281/zenodo.17054419
[487] Barcevičius, E., et al. (2025). Background document: Input for the Global Multistakeholder High-Level Conference on Governance for Web 4.0 and Virtual Worlds, 31 March – 1 April 2025. European Commission, PPMI, TNO, & WEB4HUB. Available at: https://ec.europa.eu/newsroom/dae/redirection/document/113701
[488] Frontier Model Forum. (n.d.). Frontier Model Forum. Available at: https://www.frontiermodelforum.org/
[489] Established by Anthropic, Google, Microsoft, and OpenAI. More information available at: https://aaif.io/
[490] NIST (n.d.). AI Risk Management Framework (AI RMF). NIST. Available at: https://airc.nist.gov/AI_RMF_Knowledge_Base/AI_RMF
[491] NIST (2024). AI Test, Evaluation, Validation, and Verification (TEVV). NIST. Available at: https://www.nist.gov/ai-test-evaluation-validation-and-verification-tevv.
[492] Dunn, J. E. (2026). US dominance of agentic AI at the heart of new NIST initiative. CSO Online. Available at: https://www.csoonline.com/article/4134743/us-dominance-of-agentic-ai-at-the-heart-of-new-nist-initiative.html