Policy Brief



November 2026

Transatlantic Lessons for the AI Value Chain

Designing policies to advance the public interest

By Colas Richard, Adrienne Goldstein, and Julia Tréhu

G | M | F

Washington, DC Berlin Brussels Paris Warsaw Belgrade Bucharest Ankara

November 2026



Table of Contents

Executive Summary. 3

Introduction. 5

Chapter 1: Building AI Data Centers on Foundations of Public Consent and Value. 9

Chapter 2: Leveraging the Transatlantic Partnership To Improve Data Rights and Outcomes. 14

Chapter 3: Integration and Innovation at the End of the AI Value Chain. 17

Conclusion. 21

November 2026

Executive Summary

AI is reshaping economies at unprecedented speed and scale. Governments in the United States and the EU have justified massive investments in the field in the name of competitiveness, technological leadership, and strategic rivalry with China. Public opinion, however, remains unconvinced. Underlying the skepticism is a sense that citizens have little control over who benefits from AI and who bears its costs. Governments have endorsed "public interest AI" in principle, but these commitments have not consistently translated into deployment or policy choices, contributing to growing mistrust of the technology. As AI moves from digital interfaces into physical infrastructure, public services and industry, public acceptance becomes a condition for its long-term economic and political viability.

This report examines how the AI value chain can be better aligned with the public interest. It draws on GMF's May 2026 Transatlantic Tech Exchange (TTX), during which 10 members of the European Parliament met with US federal and state policymakers, industry, investors, academics and civil society in Washington, DC and Boston. The report focuses on three nodes where public-interest tensions are sharpest and transatlantic cooperation most promising: data centers, data, and innovation policy.

On data centers, policymakers should:

- make transparency the baseline for local decisions and require projects to disclose estimated and actual impact on a common checklist of items. The process should be reinforced by shared learning through a transatlantic forum of cities and regions.

- require data centers to deliver lasting value to host communities through enforceable community benefit agreements, backed by baseline requirements throughout a facility's lifecycle.

- make AI infrastructure growth conditional on energy availability and efficiency. New data center demand should be tied to additional clean power and efficiency standards.

On data, policymakers should:

- encourage transatlantic interoperability, portability and standardization, including shared definitions of "anonymized" and "personal" data

- invest in transatlantic public datasets of nonpersonal data, building on initiatives such as the National AI Research Resource and the EU's Open Data Directive

- promote research on technical pathways for trusted data sharing, including federated learning and privacy-enhancing technologies

On innovation policy, policymakers should:

- use public funding, procurement, and approvals to fill gaps left by private capital, with support conditional on measurable public-interest outcomes
- strengthen existing innovation ecosystems while broadening participation to under-resourced regions
- align public and private capital with public-interest objectives, with better tracking of AI investment flows.

Introduction

The world is in the middle of one of the biggest infrastructure buildouts in history, and the public is unconvinced of the benefits. In the United States and Europe, AI proponents have invoked national interest, technological leadership, and competition with China to justify unprecedented investment. Still, over half of Americans are "more concerned than excited" about increased use of AI in daily life, and just 9% are "more excited than concerned". Views in the EU are only slightly more tepid. "More concerned than excited" exceeds its inverse by a margin of at least 10 percentage points in all nine countries included in a recent Pew Research Center survey.

Underlying this skepticism, particularly in the United States, is a sentiment of limited democratic control over the technology. Discussions about who benefits from AI, who bears its costs, and how public value should be incorporated into technological development often remain secondary to the urgency of the race itself.

Governments and multilateral institutions have recognized this tension. The OECD's 2019 AI Principles invoked "inclusive growth" and "well-being" as goals, and the Paris AI Action Summit in February 2025 chose "public interest AI" as one of five pillars for global governance. But principles have not consistently translated into those choices. As AI systems increasingly move beyond digital interfaces and become embedded in physical infrastructure, public services, workplaces, and industrial processes, questions about who benefits can no longer be treated as an afterthought. Moreover, closing this gap is not only a matter of societal positioning towards AI, but increasingly an economic and political imperative. Public acceptance is the condition for any long-term viability of AI development.

This report offers an alternative beyond the increasingly polarized choice between uncritical technological acceleration and blanket skepticism toward AI. It starts from the premise that AI is neither inherently beneficial nor inherently harmful. Different technological choices, governance arrangements, business models, and deployment pathways produce different outcomes. Applications can generate significant public benefits, from drug discovery to weather modeling to assisting people with severe vision impairments. But AI can also exacerbate social, economic, or environmental challenges. It can worsen healthcare inequalities and the digital divide among countries, and it has a larger carbon footprint than the global aviation industry. The central question becomes not whether AI should advance, but how its development and deployment can be aligned with broader societal goals.

This report looks at the lessons the transatlantic partners can learn from each other and concludes with recommendations for both parties for moving forward as the challenge of AI deployment looms.

Public Interest AI

For this report, the definition of "AI in the public interest" used is the one proposed by researchers Theresa Züger and Hadi Asghari, whose core assessment criteria are:

- public justification for the AI system
- an emphasis on equality
- a deliberation/ co-design process
- technical safeguards
- openness to validation
- sustainability (a later addition)

The term is used in policy and academic discussions, and this definition has been widely adopted because its breadth allows it to be applied to the full AI value chain, not just to algorithms and models.

Notably, public interest AI should not be confused with AI ethics guidelines, whose principles, when assessing AI systems, often include a combination of accuracy, robustness, explainability, transparency, fairness, privacy protection, accountability, and safety/security. AI in the public interest would necessarily incorporate such principles and practices, but such standards in and of themselves are insufficient to qualify an AI component or application as in the public interest.

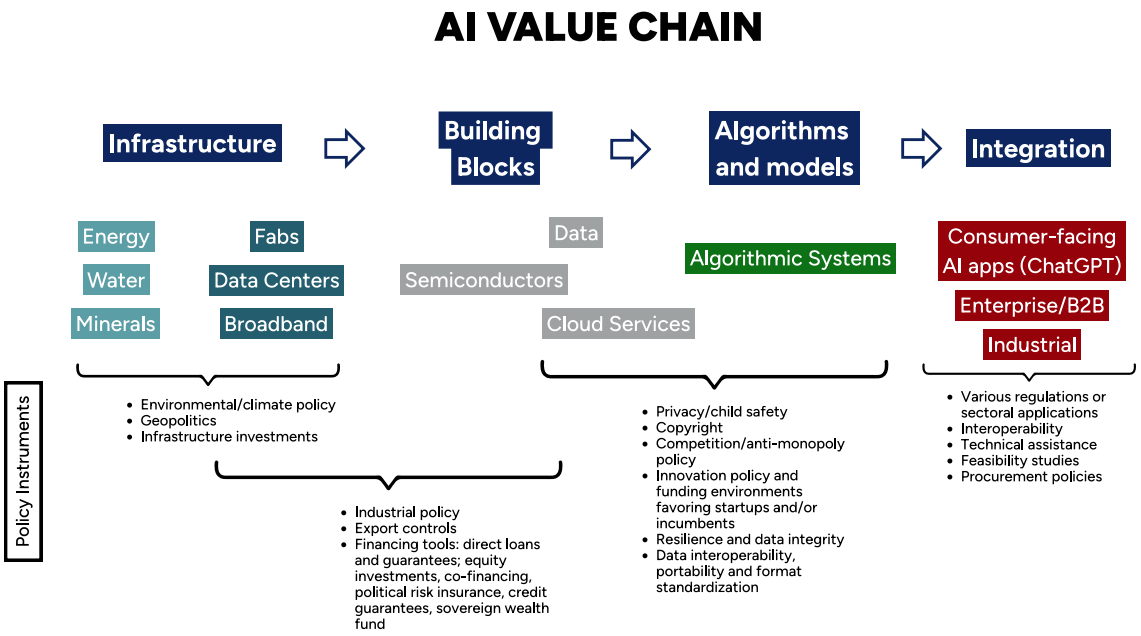
It is also necessary to guard against "public-interest washing", which permits claims of public interest to be used to superficially justify actions or proposals that fail to meet the aforementioned criteria.

Debates on "public interest AI" center on the definition, real-world implementation, models for accountability and oversight, the public infrastructure most useful for AI projects, and considerations of whether and in what cases AI-based solutions are an impactful and sustainable approach to solving societal and planetary problems.

This report's objective is not to create a checklist against which initiatives should be measured but to assess the state of play across subsections of the AI industry and propose policy recommendations.

The AI Value Chain

This report analyzes AI through a value chain framing, which the authors have previously defined as "the natural resources, industrial and network infrastructure, data, and public and private investments and inputs that create the machine learning systems known or marketed as AI". The value chain maps AI from start to finish and presents overlapping policy and innovation issues. The chain runs from infrastructure (natural resources, semiconductor fabs, data centers, and broadband) through building blocks (semiconductors, data, and cloud services) to algorithmic systems and, finally, integration into consumer, enterprise, and industrial applications.



The diagram illustrates the cross-cutting policy issues that span these nodes, and which intersect levels spanning from local to international. Technology innovation and policy efforts often target individual pieces of the value chain, but there are opportunities to distribute more equitably the benefits and burdens of AI systems by using this broader lens.

This report draws on findings surfaced during GMF's Transatlantic Tech Exchange May 2026 study tour. During the week-long trip, 10 members of the European Parliament met with US federal and state policymakers, AI and data center industry experts, innovators, investors, academics, and civil society organizations in Washington, DC and Boston. The trip surfaced three areas—data centers, data policy, and innovation policy—in which there is the greatest opportunity to align the AI value chain with the public interest, and they form the basis of the chapters.

Two of the topics—data centers and data policy—are themselves nodes along the value chain while the third—innovation policy—is a driver of change in the AI implementation node. Each presents a unique public interest tension. Data centers pose a cost-benefit misalignment problem. Local communities bear the brunt of the harm they cause, while economic and security gains are geographically diffused. Creating datasets often comes at the expense of individuals' privacy or intellectual property, but data's value lies in its aggregation. Individuals have few avenues to protect their privacy or reap the benefits from the creation of datasets. And innovation policy faces a market imbalance, rewarding commercially viable AI but too often not direct innovation toward societal needs.

These topics are legislative priorities on both sides of the Atlantic, with policymakers focused on how best to compete with China on frontier AI while also ensuring that democratic, Western allies develop and deploy AI

to benefit their societies. There is an opportunity to drive the policy debate in these three areas compared to other nodes, such as semiconductors and algorithmic systems, on which there has long been significant policy attention.

Each chapter articulates the tensions at its node of the value chain, the policy gap that allows for such tensions, transatlantic learning and implications, and recommendations to better align the node with the public interest.

Transatlantic Approach

The study tour confirmed the continued value of and opportunity for a transatlantic approach to AI policy through shared learning and direct cooperation. There are three key reasons for the transatlantic partners' ability, despite recent tensions, to be collaborators on navigating this fundamental, technological shift.

First, regulatory approaches are converging. In the past five years, Brussels prioritized regulation through the EU AI Act and an industrial policy framework centered on sustainability reporting and regulation. After the Draghi report on EU competitiveness and innovation, concerns about both issues shifted the European Commission's policy focus to boosting EU technology innovation capacity through the Digital Omnibus on AI, the AI Factories initiatives, and the Cloud and AI Development Act (CADA). The United States, in contrast, fostered market-led deployment via hyperscale projects with Executive Order 14318 of July 2025. The White House later argued against strict copyright enforcement that makes AI development more difficult as that would "threaten national security". The administration also aimed to boost investment via new rules allowing faster expensing and depreciation for capital investments. But the United States is now seeing more involvement from state regulators, utility commissions, and other public authorities due to public concerns. Washington and Brussels can learn from each other through this convergence.

Second, the allies' economies are deeply entangled on AI. History shows that transformative technology booms trigger investment bubbles, overcapacity, and sharp corrections even when the technology succeeds. If Europe is less exposed to a home-grown AI bubble, it is vulnerable to US spillovers. American hyperscalers control over 70% of Europe's cloud market. A correction of the "Magnificent Seven"—the biggest US tech companies, which account for 35%-40% of the capitalization of the S&P 500—would hit euro-area households holding around €440 billion in US technology equities, and European insurers and pensions, holding respectively 5.1% and 4.4% of their assets in private credit. A US stock market decline resulting in slower American hyperscaler investment could impact European shareholders and awaken economic security and competitiveness concerns. The dependence is mainly asymmetric, but European companies are part of the value chain building out these AI capabilities, and Europe is a major market share for US technology companies. That means a severe European downturn would feed back into American firms' revenues and valuations.

Third, both sides face the same strategic challenge: How to build AI to sustain economic competitiveness, national security, and democratic resilience while enabling AI's responsible deployment and upholding the public interest. Local concerns must be balanced against the broader societal strategic considerations of the value of compute capacity, particularly as China expands its own AI and industrial capabilities.

Chapter 1
Building AI Data Centers on Foundations of Public Consent and Value

Public Interest Tension

There is perhaps no more visible and contested component of the AI value chain than a data center. As frontier AI capabilities and AI adoption have become increasingly dependent on compute, demand for the facilities that house and power advanced chips has surged, as have their physical footprints. The shift in public awareness can be observed from July 2025, after a series of major project announcements and a US executive order to accelerate permitting raised the profile of the buildout. But local opposition led to headlines, fueling broader cross-party and transatlantic skepticism. In September, 60% of Americans reported being uncomfortable with a data center operating in their area, and the results of 46 out of the 47 ballot initiatives related to data center projects, held during the US midterm elections, doomed construction. Europeans similarly showed broad support for stronger oversight of the centers' impact.

The public tension stems less from the data centers themselves, which have quietly powered everyday digital services for decades, than from the rapid deployment of land-, capital- and resource-intensive AI-optimized data centers. They have given rise to a policy mismatch in which they are justified nationally but experienced locally. Proponents frame AI infrastructure as serving the public interest because it bolsters economic competitiveness, national security, and, in the United States, technological leadership or, in the EU, sovereignty. However, the centers' construction morphs public interest into local development considerations and calculations between concentrated community costs and the more diffuse nature of AI benefits. The rise of grassroot opposition, moratoria from New York State to Amsterdam, and the emergence of data centers as a voting factor in the recent US midterm elections suggest that the scale of the buildout has outpaced meaningful engagement with affected communities and left important public considerations unaddressed. A public interest approach to AI infrastructure must therefore consider national and local priorities, and put emphasis on the latter, where most of its tradeoffs are felt.

The Policy Gap

Aligning data center deployment with societal goals requires addressing environmental and socio-economic gaps in their buildout. Both factors call for a systems-level view that looks beyond the facility and across its full lifecycle, from development to operation to decommissioning.

Environmental impacts

Data centers, like other physical infrastructures, have significant land footprints. Their construction and operation can generate air, noise, light, dust, heat, and odor. Mitigation matters, but a public-interest approach

should act more quickly, through diligent siting and design choices that limit harm in the first place. Approval is then key to assessing a project's impacts on surrounding communities and ecosystems, including land conversion, habitat fragmentation, cumulative industrial development, and facilities' end-of-life considerations. Data centers built on former industrial land, such as Aligned's hyperscale campus in Ohio, for example, take a markedly different approach from greenfield megaprojects, such as Utah's 40,000-acres Stratos project, which requires land twice the size of Manhattan. Foresighted siting is also an opportunity to create public value, for example by locating facilities where waste heat can be recovered, as Dublin or Denmark have done.

The need for system thinking becomes even clearer when water and electricity are considered. Data centers consume water directly from cooling systems and indirectly through electricity generation. Operators have financial and reputational incentives to reduce consumption, but not all can afford the most efficient technologies, which can reduce a daily water footprint of 5 million gallons to that of a restaurant. Water is a hyper-local issue, and even modest consumption can strain local resources. Facility-level demand must therefore be assessed against watershed conditions and community needs. This requires reliable facility-level information, access to which TTX discussions highlighted as a common challenge in Europe and the United States. Failing to account for these conditions can trap communities in water deficit, derail projects, or require costly replenishment efforts.

For electricity, the issue are data centers' carbon intensity and the amount and type of energy they add to the grids. So far, data centers track and reinforce the share of non-decarbonized power in local grids. In 2026, 65% of US data centers power came from nonrenewable sources, higher than the 58% national grid average. Ongoing grid decarbonization and data center operators' renewable energy procurement could improve this balance, but rapidly growing demand complicates that trajectory. US data centers are set to more than double their electricity demand by 2030. A 70% increase in the EU is expected. To bypass grid-connection queues and constraints, operators turn to on-site, "behind-the-meter" power generation, which is overly reliant on gas turbines. Over time, this risks locking grids into fossil-fuel dependence and diverting from electrification elsewhere, complicating progress toward EU and US net-zero pledges.

Socio-economic opportunities

A data center can bring tangible benefits. Virginia's incentive program drove $80 billion in investment and $5.3 billion in revenue over five years compared to $3.2 billion in foregone sales tax. Research also associates facilities entry with higher local employment, wages, business formation, and household income. But aggregate benefits do not determine the net public value of a project because hungry-for-investment jurisdictions compete through zoning flexibility and tax breaks to secure investments, despite mixed evidence of such policy tools' influence on site selection. A complete public-interest assessment must therefore include what a community surrenders and anticipate benefits from a deal. As instances of scaled-back fiscal returns accumulate, incentives are becoming more contested. Overly successful incentives can generate their own externalities, such as the overbuilding and grid strain behind Dublin's 2021 moratorium. Consequently, jurisdictions in Nordic countries and eight US states have been repealing breaks for data center operators, or using conditional-use zoning in hotspots such as Loudoun County, Virginia.

The same scrutiny should apply to social impact. Jobs are the most expected local benefit, yet a data center directly employs just 25 to a few hundred permanent workers. The largest employment boosts are in temporary construction jobs, and skilled hires may not be local. American unions back this tradeoff, but an increasing number of jurisdictions requires local hiring or training commitments as compensation. Taxes are a second benefit, providing revenue for local communities. This is key especially for burdened communities, where current data center projects are concentrated and raise equity concerns. New projects in smaller cities in "Tier 2/3 markets" strain local institutions, which have fewer resources to negotiate with developers and manage impact. Moreover, hosting a data center does not guarantee wider benefits. Evidence shows that rural communities reap disproportionately fewer economic spillovers than data center clusters and cities, even when these communities host the facilities. Community benefit agreements (CBAs) increasingly seek to correct these imbalances by making explicit the expected public returns of a project and the distribution of spillovers.

Finally, longer-term economic risks also require attention. A public-interest approach must keep households and businesses away from grid cost-shifting while rewarding facilities that provide system benefits. Regulatory and contractual safeguards, such as the US Ratepayer Protection Act and Virginia's GS-5 electricity tariff, are emerging to make data center operators internalize the economic risks. In Europe, CADA aims to include data centers in grid flexibility programs. In addition, relying heavily on a single facility's tax base exposes local administrations to concentration risk. In states such as Wisconsin, local democratic safeguards against this risk have eroded under the buildout competition, leaving operators with extensive leverage to demand concessions and towns facing major revenue losses if a project stalls or closes. Other actors, such as utilities, face investment decisions on an unprecedented scale as they balance uncertain demand forecasts with opportunities to modernize. In the meantime, risk-transfer rises as developers limit their exposure by structuring projects through separate operating, insurance, and financing entities. A public-interest approach must ensure that these risks are distributed and that developers bear an appropriate share of responsibility for stranded assets from a burst AI bubble or from slower demand growth as edge computing shifts part of the workload to end devices.

The Broader Needs

Communities need time and information to weigh a project's costs and benefits. Moratoria provide this time, but hyperscalers and developers may push their projects to other investment-hungry communities in the interim. Time should therefore be used well, to gather data and bring developers, utilities, regulators, and residents together so that communities can move beyond simply attracting or resisting data centers and negotiate projects that offer public value. During TTX, the Massachusetts Green High-Performance Computing Center (MGHPCC) offered an example of this. While its public-led model is not transferable to every commercial project, it shows how hydro-powered computing infrastructure can be structured around industry partnerships, research, workforce development, and broader public value.

The Transatlantic Dimension

The transatlantic opportunity lies primarily in shared learning when local conditions and challenges are comparable, and in cooperation when common standards can steer innovation and investment. Transatlantic AI infrastructure policies are converging, with the EU shifting from regulation to capacity building and the United

States from a compute rush to increasing governance. Both now face the same citizen concerns about energy prices, strain on local resources, and the distribution of benefits. Meanwhile, interconnected AI technology stacks require cooperation on shared infrastructure standards and dependencies to manage common risks and to seize the opportunity of a joint buildout. The two sides have complementary strengths. The United States leads in chips and cloud while Europe brings chipmaking equipment, grid components, and expertise in heat networks and renewable integration. Common goals and standards would efficiently allow transatlantic companies to boost mutual trade and competitiveness.

Recommendations

Transparency and shared learning should be the baseline for local decisions. Most public-interest tradeoffs are hyper-local and require community-led decisions. Transparency is a prerequisite for localities to democratically weigh those tradeoffs and build public acceptance. Stronger reporting standards are needed in the EU, where disclosure remains restricted and incomplete, and in the United States, where non-disclosure agreements (NDAs) have historically kept information hidden. Pennsylvania has already ended NDAs, and other states and corporations have taken a welcome step to drop them. Europe's forthcoming data center rating scheme gives more information on a facility's design with aggregated efficiency measurements, but the effort still lacks the granular data necessary for local governments to compare those measurements to hyper-local resources conditions. Projects should therefore disclose estimated impacts before approval and disclose actual impacts during operation, both against a common checklist of indicators. A forum of transatlantic city and county networks would help consolidate this checklist and allow communities to learn earlier in the buildout process from existing hubs such as Northern Virginia, California, Dublin, and Amsterdam.

Require data centers to deliver lasting value to the communities that host them. National, regional, and city governments should require enforceable CBAs and similar public-value deals built on baseline requirements to prevent another race to the bottom. Strong agreements should include contributions proportionate to project size and follow New York State's recommended benchmarks of a $1 million investment in host communities for every megawatt of utility demand, commitments to grid flexibility and renewable energy integration, and provisions for eventual decommissioning. Agreements should also be public, independently reviewed, and backed by penalties or clawbacks if commitments are not met. In Massachusetts, state agencies cannot issue permits to covered projects unless the developer submits a CBA. Additionally, the state's environmental justice office has sole authority to determine if community benefit requirements are met. In the EU, CADA negotiations should make measurable local commitments and compensation mechanisms a condition for priority permitting and grid access. Most ambitiously, compute credits or shared capacity for universities, researchers, and startups can be negotiated to create more tangible AI benefits for local ecosystems.

Make AI infrastructure growth conditional on energy availability and efficiency. Data centers create costs beyond their immediate communities, affecting national electricity prices, grid reliability, and regional emissions. Yet neither the EU's ambition to triple data center capacity nor US efforts to accelerate permitting adequately addresses the amount of energy AI can or should consume. Natural and economic resources are finite, and data centers compete with other industries and citizens for power. In 2008, projections suggested an exascale supercomputer would need around 150 megawatts. Rather than accept that figure, the US Department

of Energy set a 20-megawatt cap based on the price it was willing to pay for electricity, and the constraint drove a decade of innovation until the target was met in 2022. Ireland learned a similar lesson. After freezing new data center connections around Dublin in 2021, the country reopened the grid four years later with a new term: Facilities must match 80% of their demand with additional renewables. CADA could combine both lessons, with the EU setting a binding energy budget for data centers that expands only as they bring additional clean power, and enforcing that through standards each new facility must meet. The Federal Energy Regulatory Commission and state utility commissions can develop a US counterpart. Aligned across the Atlantic, these standards would create a predictable market for clean firm power, storage, efficient chips, cooling and grid equipment, giving investors a clear signal and opening room for transatlantic trade and co-investment in climate tech.

Chapter 2
Leveraging the Transatlantic Partnership To Improve Data Rights and Outcomes

Public Interest Tension

Whereas the infrastructure node of the AI value chain is marked by clear, tangible impacts, from electricity use and tax revenue to noise and job creation, the data node is marked by opacity. AI developers can create datasets without the knowledge of those who own or created the content. Many datasets in the United States and Europe are inaccessible to the public. Yet data decides what a model knows and whom it serves. Limitations in the quality and utility of training data are the second most common reason AI projects fail, and, in practical AI applications, 80% of machine learning engineers' activities involve data cleansing and preparation.

Data is the raw material that is the basis of AI in the public interest, but there is a public interest tension between quality data and individual rights. Developing accurate and representative AI systems requires access to diverse, high-quality datasets, but expanding data collection and transatlantic sharing can raise concerns about privacy, consent, and intellectual property rights. Transatlantic cooperation could help reconcile these competing interests by establishing trusted frameworks for sharing and accessing data. Developers could improve AI systems' representativeness and accuracy while protecting individual rights and respecting intellectual property and copyrights.

The Policy Gap

The United States and Europe have pathways for training AI models to be a "legitimate interest" for the fair use of copyrighted materials, but significant privacy and intellectual property concerns exist.

Privacy and intellectual property

On both sides of the Atlantic, AI companies have used personal data without legal basis, undermining the public's trust and raising broader concerns about the use of personal data without guardrails. European regulators have repeatedly found that scraped personal data lacked a clear legal basis. In 2024, France, Italy and the Netherlands fined Clearview AI €20 million-€30.5 million for harvesting facial images. In the United States, Anthropic agreed to a $1.5 billion settlement after part of its training library was downloaded from pirate sites In September 2026, a court of appeals rejected fair use for a tool trained on Thomson Reuters' Westlaw summaries of law and case holdings.

Representation and accuracy

AI models require data quality and diversity to perform for different populations, and in different languages and real-world contexts. A 2019 National Institute of Standards and Technology (NIST) study found that the majority of facial recognition algorithms studied exhibited demographic differences in accuracy, with higher false-positive rates for certain racial and ethnic groups. AI models trained from large chest X-ray datasets sourced from academic healthcare facilities have been shown to under-detect disease in specific patient populations, including female, Black, Hispanic, and low-income. The challenge of representative datasets is also evident in Europe, where the dominance of English-language content in AI training datasets limits the performance of models in other languages. For example, the European Commission found that Latvian, Irish, and Maltese account for just 0.09%, 0.07%, and 0.03%, respectively, of the Common Crawl dataset, a widely used source of training data.

The Broader Needs

Transatlantic cooperation creates an opportunity to pool or share existing data to build better, more representative models. To facilitate this, governments should create frameworks for greater dataset compatibility and techniques to share data that are legal and protect privacy. Governments should also pilot collaboration on public datasets to benefit AI development, showcase the value of such partnerships, and create the infrastructure for further exchanges, including those in the private and academic sectors.

The Transatlantic Dimension

Transatlantic actors can model new initiatives on current data-sharing initiatives. Some data partnerships exist already. The EU–US Transatlantic Open Data Partnership combines economic data from the US Bureau of Economic Analysis and the European Commission's DG CONNECT and Eurostat, and the OECD operates a global Bilateral Trade in Goods by End-Use database. A lack of data interoperability and portability, however, can impede efforts for further data partnerships.

Privacy makes moving data across borders difficult. EU Commission data privacy rules require equivalent regimes from partners before data can be transferred across a border. As the United States lacks federal privacy legislation, Washington and Brussels have negotiated a series of data privacy agreements that have been subject to numerous legal challenges from European citizens and the European Parliament. This makes future exchanges of personal data under the existing agreement unpredictable. As a result, industrial and anonymous data are the more likely avenue for transatlantic cooperation.

Creating datasets involving personal, if anonymous, information could have enormous public benefits. The health sector offers the clearest examples, such as pooling data on rare diseases. Some global infrastructure already exists here, too. The International Rare Diseases Research Consortium features databases with registries of active rare diseases. In areas such as these, in which cross-border data centralization raises legal and privacy concerns, privacy-enhancing technology can promote data-sharing. These technical solutions are also necessary because anonymizing datasets through de-identification and sampling is insufficient to ensure meaningful

privacy. A study found that 99.98% of Americans would be correctly re-identified in any dataset using 15 demographic attributes. The study's authors suggest that more complex anonymization standards are needed to satisfy data protection laws and protect privacy.

Recommendations

Encourage transatlantic interoperability, portability, and standardization. The United States and Europe, seeking to involve other global partners when possible, should harmonize data formats, metadata schemes, and technical protocols to facilitate cross-border data-sharing. In the EU, much of the work to coordinate interoperability and portability is being done centrally through the bloc's data strategy, while in the United States much is done through antitrust enforcement. As such, US antitrust litigators should consider modeling sector-specific interoperability and portability remedies after European established practices.

The United States and Europe should also continue to develop standard definitions of key AI terms. This work commenced under the Trade and Technology Council (TTC), and the EU-U.S. Terminology and Taxonomy for Artificial Intelligence provided 65 key AI terms with shared definitions in 2023. Further collaboration between the NIST and the EU's AI Office can clarify the definition of terms such as "anonymized" and "personal" data.

Invest in transatlantic public datasets of nonpersonal data. The TTX trip found that US-state and local government initiatives to make nonsensitive government data publicly available were successful ways of broadening data access. Transatlantic policymakers should look to build on the success of initiatives such as the National AI Research Resource, Massachusetts AI Hub's Data Commons Collaborative, and the EU's Open Data Directive and global initiatives such as the International Rare Diseases Research Consortium. The TTX trip found that US states and local governments could successfully promote competition by making datasets with nonsensitive government data publicly accessible.

Promote research on technical pathways for trusted data sharing. The United States, the EU, academia, and technologists should improve technical pathways for data-sharing, including federated learning methods and privacy enhancing technology. Federated learning software allows models to train on local data without centralizing the raw data. A hospital can train an AI model on patient records, for example, but will send only the updated model—not patient records—to the central server. Other privacy-enhancing technologies, such as encryption and multi-party computation, can build technological pathways for trusted data-sharing. Meanwhile, synthetic datasets, which are created using deep learning techniques and mimic the attributes of real data but are not collected from real-world data have promising uses in fields with heightened data privacy needs. But these datasets come with risks of data quality degradation and even model collapse. Further research on these techniques will improve their reliability and create additional opportunities for cross-border data-sharing.

Chapter 3
Integration and Innovation at the End of the AI Value Chain

Public Interest Tension

Integration is the final node along the value chain. It is where AI's building blocks come together into outputs such as consumer-facing AI applications, enterprise and B2B software, and industrial applications. Policy at this stage matters because it shapes how citizens and consumers interface with AI, and it has feedback effects on developments earlier in the chain.

AI is a unique case study for designing innovation policy in the public interest. AI development is largely concentrated among a small number of private-sector firms in the United States, with a market-oriented logic incentivizing the creation of products for commercialization. At the same time, AI's societal impacts are global, albeit unevenly distributed, making the "public" in public interest expansive and hard to define. Contrary to techno-solutionist narratives, AI applications have frequently undermined the public interest in fields such as education, the provision of public services, and medical record-keeping. "Innovation" as a buzzword can cloud these tradeoffs, conflating progress in AI with potential societal benefits. Governments should therefore use innovation policy to correct specific gaps in the direction, participation, and distribution of AI innovation, without replacing bottom-up innovation with top-down technological choices. Innovation policy refers to the rules, policies, and strategies that shape AI's development, distribution, and integration across an economy and a society.

The Policy Gap

The policy gap lies in ensuring that the conditions under which innovation develops better reflect public-interest objectives. The objective is not government's determination of the technologies or firms that succeed but the creation of the conditions for diverse innovation ecosystems to emerge, compete, and respond to societal needs. Thus, regulation can establish overarching guardrails, while selected policy tools address gaps that markets leave unresolved.

AI and public interest innovation

Legislation and regulation provide a first layer. Many AI uses already fall under existing privacy, liability, nondiscrimination, consumer protection, and sectoral rules. The EU AI Act intersects with and relies on sectoral regulations governing product safety, cybersecurity, medical devices, automotives, machines, and toys. Emerging US legislation is reflected in the Sectoral AI Governance Act, introduced in 2026. Such regulations are and should remain guardrails establishing overarching conditions under which innovation occurs, with a more limited role in determining which societal problems should be prioritized.

In sectors such as employment, housing, finance, and healthcare, in the United States and the EU, AI slots into a dense existing web of consumer protection, antitrust, or other measures designed to protect the public interest and ensure that products or services brought to market respect certain norms and standards. Existing frameworks such as the Americans with Disabilities Act, UN Economic Commission for Europe regulations and international ISO standards apply to AI. But AI-specific standards can go further, and time matters before AI deployment practices become entrenched. Because standards define outcomes rather than prescribe technologies, they shape the conditions for innovation without picking winners. Shared, interoperable standards also lower entry barriers for smaller firms and give the United States and the EU a natural area for convergence.

Public procurement and funding are better policy tools to address this gap more efficiently. Fully 70% of research and development (R&D) occurs at firms in OECD countries, and tax incentives for R&D can provide a boost but fall short in channeling innovation toward specific societal needs. Public procurement can therefore stimulate demand or lower barriers where societal value is significant but commercial incentives are weaker, including interpretability, transparency, independent evaluations, privacy- and rights-preserving technologies, and tools for communities with limited purchasing power. Existing studies show that public procurement can drive innovation and R&D.

The notion of mission-driven innovation represents another framework for guiding innovation in the public interest. When it comes to AI, such an approach focuses on pre-distributive structures, principles, and conditionality for public funding. The Massachusetts AI Hub, visited by TTX participants, illustrates an enabling approach through public funding. It combines $100 million in public and private investment with university partnerships, access to compute, talent development, and targeted AI application programs to address local inequalities and create the conditions for Massachusetts' AI innovation ecosystem to grow.

Innovation ecosystems

Innovation does not occur in a vacuum. It depends on the interactions among entrepreneurs, universities, risk-capital entities, corporations, and government, as emphasized in innovation policy literature. The implication for public-interest innovation policy is that it must facilitate conditions under which these actors can connect and compete, not attempt to reproduce successful technology clusters from above. Frequently cited examples of this beyond Silicon Valley include Boston and Cambridge for biotechnology and Copenhagen for quantum computing. Portugal's Network for Participative Municipalities and the Netherlands' Association of Dutch Municipalities are examples of networks that include civil society.

Successful ecosystems depend on local culture and context, experimentation, and competition. Attempts to force specialization or transplant a model can crowd out alternatives or misread local strengths. A public-interest approach should therefore identify local assets and missing links, such as access to research, infrastructure, capital, talent, or potential customers, and enable bottom-up development accordingly. Incubators and universities are relevant intermediary institutions to scale these connections.

Public-private partnerships (PPPs) are often cited as a solution for integrating private-sector expertise into public-sector programs, especially for AI. But evidence of their effectiveness is mixed, raising questions about the future of AI-focused PPPs for the public interest. Examples of successful PPPs include the 16 nationwide

Advanced Manufacturing Institutes in the United States and Finland's Research and Innovation Council. More recently, Germany's SPRIND innovation agency and the EU's AI Factories have attempted to incorporate the private sector. According to another study that examined PPPs, three key factors to success are "a commitment to a strong partnership beyond the terms of the contract, built-in mechanisms to share perspectives about the project, and effective ways to rebound from failures to deliver".

Financing and capital

Finally, the financing environments for AI differ significantly between the United States and the EU. An "AI-driven investment gap" is frequently cited. In the Draghi report, the gap is explicitly mentioned as a key factor that explains the US lead. In the Letta report on the future of the European single market, closing the gap is "essential for maximizing public interest sharing and limiting the concentration of private value from data collection and profiling". Such recommendations are echoed by other reports calling for more pan-European funds and measures to improve institutional investors' access to venture funds.

The gap is one of scale and structure. US-based venture capital funds total roughly €930 billion, about six times the €150 billion held by EU funds. The European Central Bank finds the gap widens at later funding stages, partly because US monies draw far more on pension funds and foundations. In 2025, 75% of global AI venture-capital deal value went to US firms, compared with 6% to the EU. European venture capital investment is still rising, approximately doubling in the first half of 2026 compared to a year earlier. But a public-interest approach cannot reduce the challenge of closing the venture capital gap, especially since the US venture capital ecosystem is not necessarily a model of efficiency in its allocation of resources. And the current focus on AI may be crowding out other needed investments. The venture capital model also encourages firms to sell to large, wealthy incumbents, which pushes EU firms toward the United States.

The prominence of private-sector investment in AI should not overshadow the significant role of public-led industrial policy in the broader financing environment. The EU's trajectory on AI from a largely regulatory focus to the AI Continent Action Plan and various industrial policy efforts to build out cloud and data-center capacity, such as in the Tech Sovereignty Package, speaks to this trend. Such tools do not necessarily mean AI development is being steered in a public interest direction, according to the six key criteria listed in the introduction. The tools instead point to the way existing mechanisms could be harnessed and the funding possibilities that exist beyond private-led investment in AI. Public institutions such as state investment banks, can also provide financing for investments with longer-term timelines and greater uncertainty.

The Broader Needs

A public-interest innovation policy should set clear guardrails, address identifiable market and coordination failures, and create the conditions for diverse ecosystems to develop and compete. Expanding innovation where societal value exceeds immediate commercial returns requires a careful balance and a mix of policy tools to avoid counterproductive predetermination of technological outcomes.

The Transatlantic Dimension

The transatlantic opportunity lies in complementary experience and convergence of innovation policy without resulting in one model. The United States combines deeper private capital markets with established technology ecosystems, while Europe has increasingly deployed public investment and industrial-policy instruments alongside regulation. Mutual learning should focus less on replicating each other's models than on identifying which enabling conditions allow innovation ecosystems to generate and diffuse broader public value.

Recommendations

Use public funding, procurement, and approvals to fill gaps left by private capital. Governments should identify applications and capabilities with clear societal value but weak commercial incentives, including transparency, evaluation, privacy-preserving systems, public-service delivery, and workforce adaptation. Public support should be conditional on measurable public-interest outcomes, and governments should offer credible demand, appropriate data access, and predictable rules.

Create the conditions for existing ecosystems to become stronger through broader participation rather than recreating those conditions from scratch. Studies show that over 80% of green and digital technology in Europe emerges from regions that are already leading economically, raising concerns that AI could further concentrate resources and opportunity. If successful ecosystem strategies rely on pooling existing strength and sets of actors, policies should reinforce existing ecosystems while expanding participation to under-resourced regions. Policies should also support intermediary organizations that connect researchers, entrepreneurs, investors, governments, and communities, and ensure that the benefits of innovation go beyond established technology hubs. European and US states, regions, cities, universities, regulators, incubators, and financial institutions should exchange operational models and develop joint initiatives for public-interest procurement, ecosystem development, and sector-specific AI.

Align public and private capital with public interest objectives. Simply scaling up venture capital is insufficient to ensure that AI investments benefit the public interest. As Europe seeks to deepen and integrate its capital markets, policymakers should strengthen transparency for AI investment flows and identify areas that remain underfunded despite their societal value. Public financing institutions, including development banks and other public investment vehicles, should help de-risk and co-finance mission-driven projects that private investors may overlook, while guarding against excessive concentration of capital in a narrow set of firms, technologies, or regions. Better tracking of AI-related financial flows would also help policymakers assess if investments are advancing public-interest goals or merely reinforcing existing market dynamics.

Conclusion

The United States and Europe approach AI from different starting points. They differ in their regulatory traditions, their attitudes to privacy and risk capital, and in their technological capacity, with the United States holding a clear lead in frontier AI. Still, across the value chain they increasingly face the same challenges: How to build AI infrastructure that benefits local communities, how to access data without eroding rights, and how to direct innovation toward societal needs.

As some of their approaches also start to converge, learning from and collaborating with each other becomes an even greater opportunity. Many of the solutions identified in this report are still emerging, often at the state, regional or city level, and each side can draw on the other's experience, rules, and initiatives to test, adapt and improve them. In a field that evolves faster than Washington or Brussels can legislate alone, this exchange gives both the ability to remain competitive, strengthen their shared security, and show that democracies can lead in AI while serving their citizens.

Aligning AI with the public interest is not a constraint turning into a necessary condition. Citizens must have a say in a technology that is reshaping their communities, workplaces, and public services. At the same time, AI's gains must be broadly shared to justify the public resources it draws on. For companies, public trust secures permits, customers, and a durable license to operate. Rising opposition shows that a buildout without public support will not last. AI will deliver on its promise only if citizens trust how it is built and deployed, a challenge allies can meet better together.