Atlantic Council

Commission on Artificial Intelligence

US leadership in the age of AI

October 2025 - June 2026

The GeoTech Center

The Atlantic Council GeoTech Center has a mission to shape the future of technology and data to advance society.

Editors

Graham Brookie

Tess deBlanc-Knowles

Raul Brens Jr

© 2026 The Atlantic Council of the United States. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Atlantic Council, except in the case of brief quotations in news articles, critical articles, or reviews.

Please direct inquiries to:

Atlantic Council
1400 L Street NW, 11th Floor
Washington, DC 20005

Honorary Congressional Co-Chairs

Sen. Todd Young (R-IN)

Sen. John Hickenlooper (D-CO)

Rep. Jay Obernolte (R-CA, 23)

Rep. Suzan DelBene (D-WA, 1)

Commission Co-Chairs

Ron Ash, *Accenture Federal Services*

Kemba Walden, *Paladin Institute and Atlantic Council Board Member*

Commissioners

Graham Brookie, *Atlantic Council (Executive Director)*

Amanda Craig Deckard, *Microsoft*

Brie Sachse, *Siemens USA*

Chris Massey, *The Bird's Nest Consulting Group*

Dave Levy, *Amazon Web Services*

Don Vieira, *Skadden, Arps, Slate, Meagher & Flom LLP and Affiliates*

Fred Kempe, *Atlantic Council*

John Goodman, *Atlantic Council*

Markham Erickson, *Google*

Matthew Gravis, *Atlassian*

Molly Montgomery, *Meta*

Nabiha Syed, *Mozilla Foundation*

Nathan Jokel, *Cisco*

Ned Finkle, *NVIDIA*

Rachel Gillum, *Salesforce*

Sarah Heck, *Anthropic*

Sridhar Ramaswamy, *Snowflake*

Thomas Zacharia, *AMD*

Tyson Lamoreaux, *Arista Networks*

Note from co-chairs

Across the full sweep of human history, technological change has been a constant and has redefined the rules of power. From the Stone Age to the Information Age, technological advancement changes how we live, what we can achieve, and how quickly we achieve the next breakthrough. Today, that acceleration has reached a massive inflection point. We are living through the greatest technological change in human existence.

Artificial Intelligence (AI) is the defining technology of this era. Its implications—for economic competitiveness, national security, democratic governance, and the human condition itself—are profound and reshaping the world around us.

The opportunity to harness AI for human advancement is enormous, and the consequences of failing to lead into this future are just as great. We convened the Atlantic Council Commission on AI with purpose, conviction, and urgency to accomplish something core to the organization's mission—to shape the global future together.

Throughout its 250-year history, the United States has stood at the cutting edge of successive technological revolutions—not by chance, but by design. We bring together builders, innovators, and mission sets. Discovery, creativity, openness, and a willingness to come together, adapt, and achieve something greater than the sum of our parts is part of the fabric of our country and critical to securing our national advantage. With the Commission, our goal was to build a comprehensive and connective approach for continued US leadership to match the speed of technological change.

To that point, over the course of the Commission's eight months of work, we witnessed significant change from when we launched our efforts just eight months ago.

During that period, leading AI labs released approximately eighteen frontier models—a new model roughly every two weeks. An American AI company became the first valued at five trillion dollars, having grown 500 percent within three years. The use of AI was confirmed in at least two US military operations in Venezuela and Iran. The adoption of AI agents performing complex, multi-step, multi-modal tasks expanded rapidly. The Pope has published an encyclical on AI. And, perhaps most consequentially, recent models like Anthropic's Mythos that are fundamentally changing the global security landscape across sectors were released.

The Atlantic Council Commission on AI brought together leaders from industry, civil society, and government to meet the moment. Two factors shaped the Commission's work and sense of urgency. First, positive policy solutions and partnerships once thought impossible are now withing reach. Second, the window to realize necessary change is narrowing. We must act now.

This effort of the Atlantic Council Commission on AI was the first of its scale and ambition since the release of Mythos, proving we need a new model of public and private partnership matched with a strong commitment to working together in the national interest.

While the findings of the Commission do not represent the viewpoint of any single individual or organization, we found cohesion, consensus, and shared purpose through our investigation of the elements needed to sustain US leadership in AI for years to come. We're especially thankful to each commissioner who contributed time, expertise, and leadership throughout our work together.

In this final report, the Atlantic Council Commission on AI sets forth the building blocks for a comprehensive strategy for AI leadership matched with a foundation for measuring our collective success. The Commission's recommendations are durable and require deliberate action. We all have more work to do, and we hope you'll join us in the effort. Let's get to work.

Ron Ash

Kemba Walden

Graham Brookie

Acknowledgments

The Atlantic Council gratefully acknowledges the leadership, time, and expertise of the commissioners, honorary congressional co-chairs, and commission chairs whose guidance shaped this report. The Council also thanks the experts, colleagues, and partners who contributed to the Commission's work through briefings, consultations, research support, and strategic input—including Matthew Rose, Ximena Cespedes, Kevin Klyman, Nicole Isaac, Michael Greenwald—whose perspectives helped sharpen the report's analysis and recommendations.

The report was authored by Atlantic Council experts and contributors across the organization. The Council also acknowledges the GeoTech Center team for their contributions, including Tess DeBlanc-Knowles, Raul Brens Jr, Coley Felt, Trisha Ray, Emily Sespico, Ryan Pan, Emma Hopp, and Olivia Torres.

The findings and recommendations reflect the Commission's collective deliberations and should not be attributed to any individual commissioner or institution. This report was written and published in accordance with the Atlantic Council's Intellectual Independence Policy, which requires all donors to agree to the Council maintaining independent control of the content and conclusions of its work.

Chapter Authors

Chapter 1: Introduction
Graham Brookie and Coley Felt

Chapter 2: Overall findings
Graham Brookie and Tess deBlanc-Knowles

Chapter 3: AI innovation
Tess deBlanc-Knowles and Ryan Pan

Chapter 4: Building talent and institutional readiness
Trisha Ray

Chapter 5: Governance of AI
Coley Felt and Ryan Pan

Chapter 6: Supply chain
Sara Ann Brackett

Chapter 7: Powering AI
Joseph Webster and Frank Willey

Chapter 8: Allies and partners
Mark Scott and Ryan Pan

Chapter 9: Onward – Measuring US success
Joseph Bonivel, Will LaRivee, Raul Brens, Tess deBlanc-Knowles, and Ryan Pan

Chapter 10: Onward – Navigating geopolitics of AI
Graham Brookie, Matt Petit, and Ryan Pan

Table of contents

Note from co-chairs 2

Acknowledgments 3

Chapter authors 4

Executive summary 7

Chapter 1: Introduction 9

- Overview 9
- A world defined by three converging trends 9
- The US policy response 9
- The commission's approach 10

Chapter 2: Overall findings 11

- Competitiveness requires public trust and optimism 11
- Innovation and integration are equal priorities 11
- The global security landscape has fundamentally changed and requires urgent action 11
- AI is a collective ecosystem that requires holistic strategy and management 11
- American strength scales with allies 12
- Durable leadership is grounded in long-term strategy 12
- Measurement is key 12

Chapter 3: AI innovation 13

- Flashpoints 13
- Findings and recommendations 14

Chapter 4: Building talent and institutional readiness 17

- Flashpoints 17
- Findings and recommendations 18

Chapter 5: Governance of AI 21

- Flashpoints 21
- Findings and recommendations 22

Chapter 6: Supply chain 24
Flashpoints 24
Findings and recommendations 25
Chapter 7: Powering AI 28
Flashpoints 29
Findings and recommendations 30
Chapter 8: Allies and partners 33
Flashpoints 33
Findings and recommendations 34
Chapter 9: Onward - Measuring US success 37
Why measurement is needed 37
Where current measurement falls short 37
Introducing a national AI scorecard 38
Baseline indicators 38
Conclusion 45
Chapter 10: Onward - Navigating the geopolitics of AI 46
Alliances as key ingredients to US competitiveness in AI 46
Building a practical consensus on AI alliances 46
Competitiveness at home and abroad: the allied tech stack 47
Meeting middle powers where they are 47
Managing systemic competition with China 48
Future questions to explore 48
Building the coalitions needed to sustain US AI leadership for a safer world 49
Conclusion 50

Executive summary

The global AI landscape is defined by three intersecting forces: intensified geopolitical competition, deepening interdependence, and accelerating technological disruption. Authoritarian states are leveraging AI to consolidate domestic control and extend influence abroad, while recent frontier model developments, particularly in cybersecurity, underscore the urgency of the United States maintaining the lead in capability development. The Atlantic Council Commission on AI was formed to meet this moment and establish a blueprint for sustained US leadership in AI across six critical domains: innovation, talent, governance, supply chain, energy, and allies and partners.

The Commission reached a unifying conclusion: US leadership in AI will not be determined by any single action, but by many actions pursued through a connected, coherent, and comprehensive approach. Critically, that leadership must be grounded in public trust as a foundational element across all six domains. Building citizens' confidence in AI, while demonstrating clear safeguards and tangible benefits, is a prerequisite for the United States to compete effectively.

The Commission identified seven cross-cutting imperatives that must underpin any coherent national AI strategy:

Competitiveness requires public trust and optimism. Policymakers and industry must demonstrate the net positive impact of AI on citizens' daily lives, providing clear guardrails and assurances around safety and security.

Innovation and integration are equal priorities. A strategy focused solely on frontier model development is insufficient. Adoption and deployment must be elevated as equal priorities to ensure broad social returns.

AI is a collective ecosystem that requires holistic strategy and management. Infrastructure, workforce readiness, and integration priorities must be addressed in parallel. Progress in any domain is constrained by the weakest link in another.

The global security landscape has fundamentally changed and requires urgent action. Cybersecurity implications of new AI models require structured public-private partnerships to enable coordinated awareness and preparation.

American strength scales with allies. Sustained leadership requires actively building AI development, deployment, and governance capacity among partners. The breadth of the United States' coalition is as important as the depth of its domestic capabilities.

Durable leadership is grounded in long-term strategy. The United States needs an enduring national AI strategy that establishes durable priorities while remaining explicitly designed for adaptation as the technology evolves.

Measurement is key. A national AI scorecard that continuously tracks key indicators across all six domains is as foundational to the success of a national strategy as the models themselves.

Within each domain critical to national AI competitiveness, the Commission identified the significant flashpoints that must be addressed and put forward a series of findings and recommendations intended to build on and accelerate government actions taken to date. The slate of recommendations reflects a nonpartisan approach, focused on advancing actions that, to be effective, should be sustained across political cycles and realized as part of a strategic whole.

Innovation. Sustaining US leadership in AI depends on the nation's ability to fund, build, and steward the foundations of innovation. Expanding federal R&D investment, including through novel funding mechanisms, and building out supporting infrastructure must be matched by actions to accelerate AI adoption in key sectors.

Talent. Without proactive policies, the US faces vulnerabilities in the form of an education system unprepared for the pace of technological change, uncertainty in immigration pathways for top talent, and undersupply of critical skills. A whole-of-society approach to workforce skilling and immigration reform are immediate priorities.

Governance. A durable federal governance framework must address risk tiering, data governance, transparency, and incident reporting. The federal government has the opportunity to leverage federal adoption to drive innovation across sectors, but realizing these ambitions depends on modernizing federal procurement systems.

Supply chain. A cohesive approach to hardware, cybersecurity, and data governance is essential. Strengthening the hardware foundation of the AI supply chain demands calibrated export controls and expanded domestic manufacturing. Securing AI systems requires both reinforcing federal cybersecurity institutions and developing AI-specific mitigations.

Energy. Building and maintaining the energy infrastructure necessary to power an era of rapidly escalating compute demands is a central ingredient in a strategic, national approach to AI. A pragmatic, all-of-the-above generation strategy must be paired with grid modernization and permitting reform.

Allies and partners. The United States should champion an allied AI stack built on interoperable standards, shared institutions, and a credible counternarrative to authoritarian

influence. Robust US leadership in AI safety and security will ensure that democratic values, rather than autocratic alternatives, define the emerging global AI order.

The Path Forward

Two areas warrant immediate follow-on action.

First, the Commission proposes establishment of a national AI scorecard: a taxonomy of baseline indicators to assess US competitiveness and guide a responsive national strategy. Rigorous evaluation is essential for technology evolving at this pace, as it enables governments to design smarter policies, measure return on investment, and course-correct in real time.

Second, the US approach to international partnerships must evolve. This requires resolving open questions around interdependence, interoperability, private-sector roles, mutual defense, and capability gaps. The rapid advancement of AI will continue to complicate and reshape geopolitics. Without a strategy grounded in the technical, economic, and geographic realities of the AI ecosystem, existing alliances risk fracturing, accelerating the fragmentation of the global order when cohesion is most needed.

Chapter 1: Introduction

Overview

The world is undergoing a fundamental shift driven by artificial intelligence (AI). In just a few years, AI has moved from a specialized research field to a foundational technology shaping economic productivity, national security, and the architecture of global power. The question facing the United States is no longer whether AI will reshape the international order but whether the United States and its allies and partners will be positioned to lead what emerges.

This report is the culmination of the work and expertise of the Atlantic Council Commission on AI, convened by the Atlantic Council's Technology Programs to bring together senior leaders from government, industry, and academia around a common framework for understanding AI competition and cooperation. The commission's work focused on six areas critical to US competitiveness that are distinct across an array of relevant sectors but act in concert to build, power, and apply AI that shapes everyday life. The goal of this report is not simply to describe the AI landscape but to translate its complexity into actions and continued leadership.

A world defined by three converging trends

The global AI landscape is shaped by three intersecting trends: intensified geopolitical competition, deepening interdependence, and accelerating technological disruption.

First, competition for AI leadership now spans the full technology stack—from semiconductors, high-speed networking, and data center infrastructure to models, applications, and governance norms. Countries are not simply competing to build the best models; they are competing to shape the ecosystems, standards, and supply chains that will underpin the global AI economy. Strategic competitors are closing capability gaps faster than anticipated and are gaining traction in global markets where governments and companies are making long-term decisions about which technology stack to build their digital infrastructure around.

Second, AI development is interdependent. Supply chains for semiconductors, critical minerals, data center components, and the talent that drives innovation in those supply chains cross borders in ways that make full technological separation unrealistic. At the same time, governance frameworks for managing these interdependencies—export controls, technical standards, and data governance—remain incomplete and uneven across regions.

Third, the rate of technological change is accelerating. Capabilities that once took years to develop now emerge in months, often outpacing the institutions designed to regulate, procure, and deploy them. Governments are increasingly forced to make long-term strategic decisions in an environment defined by short development cycles and high uncertainty.

Together, these three trends—competition, interdependence, and acceleration—define the strategic environment for AI policy and geopolitics.

Against this backdrop, authoritarian states including China and Russia aim to leverage AI to maintain absolute control at home and export their influence abroad. This approach is designed to undermine a free, open, secure, and interoperable global technology and political ecosystem, and it stands in stark contrast with actualizing AI's potential to benefit all of humanity. Countering authoritarian efforts to bend the arc of AI toward repression and surveillance, and to exploit the technology for national gain, requires more than a series of siloed tit-for-tat responses. It requires an integrated strategy that realizes the strength of collaboration and open systems, involves stakeholders and partners, and links how AI is designed, funded, powered, and governed. Failure to do so risks ceding US global leadership not only in technology but also in economic and national security.

The US policy response

The US AI policy has undergone significant reorientation in response to these dynamics. Recent policy efforts have increasingly prioritized competitive advantage, infrastructure buildout, and international adoption of the US AI technology stack. This includes major investments in domestic semiconductor manufacturing and data center infrastructure, efforts to rein in regulation affecting AI development and deployment, and a growing emphasis on exporting US AI technologies, standards, and governance approaches to allies and partners.

Access is a through line across these policy efforts. Sustained US leadership in AI will depend on enabling reliable access to the foundational inputs of the AI economy, including raw materials, compute capacity, education, skills, and the networks of innovators that drive the field forward. Expanding and securing this access is central to broader adoption of the US AI technology stack and to building the public trust needed to sustain it.

Competitive advantage, in this context, means more than producing the leading model. US leadership in AI will depend not only on frontier model development but on whether the United States and its allies can build, power, deploy, and govern the systems that will underpin the global AI ecosystem. This

approach inherently involves structuring and implementing policies and investments over a wide range of interconnected areas and doing so over time horizons that extend beyond one political cycle.

The commission's approach

The Atlantic Council Commission on AI was designed to engage, navigate, and lead through this complexity.

The commission's recommendations are intended as directional rather than prescriptive, aiming to build on and accelerate government actions taken to date. The slate of recommendations reflects a nonpartisan approach focused on advancing actions that, to be effective, should be sustained across political cycles and realized as part of a strategic whole. Recommendations represent the deliberation of the commission, as opposed to the views of each individual commissioner.

The six areas of focus considered by the commission were chosen because they are catalytic to US national success; however, each is managed in a silo or by a specific sector. Across all six areas, the commission's work reflects a central conclusion: US leadership in AI will not be determined by a single action but by many actions that must be made through a connected, coherent, and comprehensive approach. It will depend on innovation pipelines, resilient supply chains, energy access, talent development, effective governance institutions, and strong alignment with allies and partners. In short, AI leadership is not just about who invents the most advanced systems, but about who can build and sustain the ecosystem, infrastructure, and alliances that will shape the global economy and international order.

The commission's findings identify connections, strengths, and gaps that go beyond one policy or sector, as well as metrics for measuring national success. All form a foundation for a national strategy that can adapt at the speed of change in the age of AI.

Chapter 2: Overall findings

As the commission considered the six areas critical to US competitiveness in AI, certain findings were consistent and emerged in every domain. These make up the commission's foundational findings. Each represents a strategic imperative for a coherent, effective national approach to maintaining global leadership in AI and fully harnessing the potential of AI for the benefit of society, the economy, and national security.

Competitiveness requires public trust and optimism

In a democracy, the system is only as strong as its citizens' belief in it, participation in it, and ability to see tangible results from it. Realizing the full potential of AI in the United States is no different.

Building trust in AI will be essential for the US democratic system to compete effectively. Policymakers and technology companies alike will need to demonstrate the net positive impact of AI in citizens' day-to-day lives and provide assurances and clear protections around the safety and security of AI systems.

The United States lags other countries, including direct competitors like China, in terms of public trust and optimism in AI. According to Stanford's 2026 AI Index Report, large majorities (79–84 percent) in China, Indonesia, and Thailand are excited about AI-powered products and believe that AI will profoundly change their daily lives in the next 3–5 years (82–89 percent). In the United States, just 38 percent of respondents express excitement about AI, and 57 percent believe that AI will profoundly impact their lives. A 2025 poll conducted by Gallup and the Special Competitive Studies Project found that 60 percent of Americans somewhat or fully distrust AI. This is especially acute with young Americans. Another Gallup survey found that while more than half of Gen Z respondents (ages 14–29) used AI daily or weekly, marks for hopefulness and excitement about AI fell as marks for anger about AI increased.

Trust in AI is collective and cumulative. Put simply, the United States cannot compete effectively if Americans view AI primarily as something that threatens jobs, negatively impacts children, increases energy bills, strains the environment, and makes existential risks such as catastrophic or uncontrolled forms of warfare possible. These real concerns need to be understood and acknowledged by policymakers and managed collectively because they are felt collectively.

Government initiatives should prioritize return on social investment in AI and take demonstrable steps to be responsive to citizen concerns around social, economic, security, and environmental impacts from AI. Without that trust, it will be impossible to build national purpose and might significantly slow progress on national goals.

Innovation and integration are equal priorities

Innovation is only as effective as its adoption.

Much of the conversation around AI competitiveness in the United States focuses on the development of the next generation of AI models and the Silicon Valley-fueled race to artificial general intelligence (AGI) or superintelligence. However, to effectively compete globally and build overall social returns for the general population, the focus on innovation should be balanced with a concerted strategy for access to AI technologies and responsible integration of AI across society, the key components of which are outlined in this report.

Deployment and adoption of AI across sectors of the economy, if undertaken strategically and with societal benefit at front of mind, holds the potential to power economic growth essential for continued global leadership while also demonstrating to citizens positive, deliberative change that AI can drive.

Focusing solely on out-innovating competitors impedes the nation's ability to realize the full social, economic, and security benefits of AI and could, in fact, set the conditions for falling behind globally in critical areas of competition.

The global security landscape has fundamentally changed and requires urgent action

Recently released models from leading US AI labs demonstrate the potential of AI to transform cybersecurity, both as a defensive asset and as a vector for novel threats. Maintaining the US lead in AI capability development is therefore essential, buying time for US government, industry, and critical infrastructure operators and partners and allies to strengthen cyber defenses. A scenario in which adversaries achieve field game-changing AI-powered cyber capabilities first could prove catastrophic for US national security and economic stability.

Meeting this challenge requires robust, structured partnerships between the private sector and the federal government to enable proactive awareness of emerging capabilities and to facilitate coordinated preparation across sectors for the impact of such capabilities. Without deliberate coordination at the highest levels, the United States risks ceding both the initiative and the institutional knowledge needed to stay ahead.

AI is a collective ecosystem that requires holistic strategy and management

AI is a result of an intricate and interconnected ecosystem.

There is no single Sputnik moment or narrow Manhattan Project in the age of AI. Instead, there exists a chorus of many

happening at once—areas of cooperation and competition, opportunities and risks, and intertwined causes and effects occurring simultaneously.

AI leadership depends on the direct tech stack that produces the technology as well as human expertise, computational infrastructure, data, high-speed networking and connectivity, energy availability, and much more. Indeed, the commission chose its areas of focus because they are connected and catalytic but often managed in silos or by sector. To effectively harness AI's potential and maintain national competitiveness, the United States should govern and invest in AI as an integrated national ecosystem. This means addressing infrastructure dependencies, workforce readiness, and integration priorities in parallel rather than in sequence, recognizing that progress in any domain is constrained by the weakest link in another. Modernizing the electrical grid, enabling responsible government use of AI, and reforming educational pipelines are mutually reinforcing imperatives that succeed or fail together.

As an open market democracy, the United States' key strategic assets in competing effectively in an era of AI are diffuse and largely autonomous. This is a key strength but one that requires intention and connectivity. Taking an ecosystem approach requires forging alignment as a feature rather than a coincidence among government, industry, academia, and civil society.

American strength scales with allies

The architecture of global geopolitics is fundamentally changing, and AI will shape and underpin what comes next.

Sustained global leadership in AI will require developing the world's leading capabilities while also actively building capacity for AI development, deployment, and governance among partners across the globe. Doing so generates economic opportunity for US industry, fortifies the global AI supply chains the United States already leads and depends upon, and draws allied nations into a shared technological community anchored by common standards, values, and mutual accountability. Should the United States fail to equip its allies with the technological tools to participate in AI, its strategic competitors will not hesitate to fill the vacuum—and benefit as a result.

The age of AI presents an opportunity to renew and deepen the global network of alliances that has given rise to a period of unprecedented growth, security, and prosperity, and the ecosystem that underpinned breakthroughs in technology including AI. The fundamental shifts in the global landscape, caused or enabled in part by AI, also make it necessary to engage more countries as close partners. In an era of strategic competition, the breadth of the US coalition might matter as much as the depth of its domestic capabilities.

Durable leadership is grounded in long-term strategy

Breakthroughs in aspects of AI are happening every day.

The rate of change in the technology itself is dizzying, as AI has gone from being a novel scientific research field to a technology that impacts everyday life in a manner of a few short years. The first generative pre-trained transformer (GPT) model was released in June 2018. The computational power used to train AI models has been doubling approximately every six months since 2012. Performance benchmarks, adoption, investment, costs, and efficiency related to AI have all seen nearly unprecedented gains. Even if—as some researchers predict—the exponential nature of accelerating AI capabilities plateaus in the near term, the social impact of AI will not follow suit.

In this moment of rapid technological evolution, it is seductive to see everything as a short-term sprint. However, success will be measured in generations. A collection of technologies that will change the global architecture and geopolitics requires strategy and foresight beyond the current innovation cycle. AI policy to date has been largely reactive, responding to capability breakthroughs, competitive pressures, news cycles, and emerging risks as they arise. A long-term strategy, based on the elements outlined in this report, must go beyond one administration, one congressional budget cycle, one annual earnings report, or one group of stakeholders. US leadership requires establishing and pursuing a long-horizon AI strategy that establishes enduring national priorities while remaining explicitly designed for adaptation.

Measurement is key

The United States needs a national AI scorecard that continuously measures key indicators of success.

Effective measurement and evaluation are as foundational to US national strategy as they are to the latest AI models. The United States lacks a consistent, rigorous framework for measuring the full spectrum of factors that drive AI competitiveness. Without this, policymakers risk designing policy interventions based on incomplete data or anecdotal understandings of where the United States stands relative to competitors.

A democratic country cannot manage what it cannot measure (quantitatively) or explain (qualitatively). A comprehensive national scorecard that tracks the full AI ecosystem rather than isolated indicators is foundational for a coherent and effective national strategy. Consistent and transparent measurement accelerates public trust, allows for holistic management, and informs long-term strategy. It is a critical enabler for every element of national success in the emerging AI era.

Chapter 3: AI innovation

US leadership in AI innovation set the technology on its current trajectory. While the recent acceleration of capabilities has occurred over increasingly short time frames, the foundation on which it rests was constructed over the course of decades—fueled by long-horizon federal investments in basic research, a world-leading academic enterprise that attracted the best and brightest from across the globe to pursue their ideas in the United States, and a vibrant private-sector environment that provided funding and pathways for new ideas to reach the market.

The rest of the world took note. Competitors and allies alike have modeled national AI strategies on US strengths. China, in particular, has pursued a series of structured national plans aimed at achieving global leadership in AI innovation and adoption. The results have been substantial. Chinese universities now rank among the top institutions in the world, academic publishing on AI out of China has surged, and a vibrant AI start-up ecosystem has emerged. In 2025, China's approach evolved to make bets on open-source AI and prioritize AI integration. Based on usage data from the OpenRouter platform, over the course of 2025, Chinese open-source models grew to account for 30 percent of global AI use, up from 1.2 percent at the end of 2024. According to analysis published by Hugging Face, the most popular platform for open-source AI development, over the course of 2025 open-source models from China surpassed those from the United States in both monthly and overall downloads, swiftly accounting for the plurality of downloads on the platform.

While US frontier labs continue to release the most advanced, capable models that exist, recent reports illustrate how once-authoritative gaps continue to narrow. This state of affairs demands the United States assess its current approach and take decisive steps to address the flashpoints and cracks that are emerging across the nation's innovation landscape.

Flashpoints

Weakening academic base

US innovation has long been sustained by the wellspring of ideas and talent developed through the nation's academic enterprise, the strength of which has been enabled through sustained federal funding for research and education initiatives. The Information Age would not have existed without the creation of the internet in a narrow science community funded by the US Advanced Research Projects Agency (now the Defense Advanced Research Projects Agency) and the US National Science Foundation (NSF) before the technology became commercialized for the broader public. The techniques of deep learning and reinforcement learning that underpin the impressive capabilities of today's AI models came out of federally funded academic labs. However, academia's role in the future of AI has been overshadowed by that of the private sector in recent years. And while the importance of AI has accelerated, federal funding for AI research and development (AI R&D) over the course of the past four years has leveled out at around $3.3 billion annually. This plateau comes in spite of efforts, such as those authorized in the CHIPS and Science Act of 2022, to boost federal support for research, development, and education in technical areas of strategic importance such as AI. Since the passage of the act, appropriated spending for agencies such as NSF, the US Department of Energy (DOE), and the US National Institute for Standards and Technology (NIST) have fallen well short of the levels authorized in the legislation.

Limited funding for AI research in academia means a shortage of technical resources that are critical for AI innovation and academic research applications—most importantly, access to advanced computing resources. This has set up a widening gap between private-sector R&D initiatives that have access to the computational resources needed to innovate on the cutting edge of AI and academic endeavors that do not. Per a Stanford University report: "Among academics with access to compute, it is typical for them to have access to between 1 and 8 AI chips, whereas industry researchers may have access to thousands." A federal initiative, the National AI Research Resource (NAIRR), was launched as a public-private partnership in 2024 with the goal to connect researchers with the necessary computational, software, and data resources (including both federal and commercially available options) to sustain AI innovation in academia and small businesses. However, the research community continues to face challenges in accessing computational resources. This creates barriers to sustaining essential R&D that might not have direct commercial incentive but sets the groundwork for future breakthroughs, including novel approaches to more generalizable capabilities and improvements related to AI assurance, interpretability, and security. In addition, this lack of access slows the pipeline of talent development on which the private sector relies for growth.

Availability and accessibility of data

High-quality data enable AI systems to generate analysis and predictions accurately and efficiently, which becomes particularly salient when developing AI models for specialized tasks in high-impact domains such as healthcare, manufacturing, and science. The Trump administration's AI Action Plan sets a goal for the United States to "lead the creation of the world's largest and highest quality AI-ready scientific datasets." In November 2025, Trump launched the Genesis Mission, which aims to build an integrated AI platform to leverage scientific

datasets to drive scientific discovery through DOE's national laboratories. However, many factors still hinder access to data for public benefit, including the restriction of access to corporate-owned data, regulatory uncertainty around data governance, and the downsizing of statistical agencies. Within organizations, data readiness remains a prerequisite for AI innovation and adoption. However, challenges related to data siloing, governance, quality, and interoperability often limit the ability of organizations across sectors to harness AI for their benefit. Increasing the availability of data is instrumental in optimizing task- and sector-specific AI applications and fueling innovation in both academia and the private sector. Yet questions remain about how to responsibly expand use of federal datasets, incentivize open access, and better leverage technical options such as synthetic data and other privacy-enhancing technologies.

Uneven AI adoption across sectors

China recently issued its "AI Plus" initiative, aimed at promoting integration of AI across various sectors including science and technology, industrial development, consumption, social welfare, governance, and global cooperation, with the overarching goal of reaching AI penetration of 70 percent across its economy by 2027. This stands in contrast to the US model, which to this point has centered more around pursuit of AGI and a race for the next generation of model rather than the integration of current AI capabilities into economic sectors. Recent Census data show a picture of highly uneven AI usage rates by sector in the United States: while the information sector has a 38 percent adoption rate, the manufacturing sector has 13 percent usage, and the transportation sector's rate is even lower at 8 percent. Uneven application of AI across the economy prevents innovative experiments and use cases from being developed in diverse sectors, undermining US competitiveness in AI and weakening the ability to build public trust through the demonstration of promising AI capabilities and innovations.

Infrastructure readiness

Compounding infrastructure deficits across the US economy present challenges for national efforts to push forward on AI adoption. As detailed in Chapter 7 of this report, the nation's aging energy grid requires modernization to meet the demands of data centers. Simultaneously, the country's connectivity infrastructure must evolve to support the high-bandwidth, low-latency networks demanded by AI workloads, a challenge that will become more pressing as AI increasingly moves into edge devices and into the physical world through robotics. In addition, existing digital divides between rural and urban areas across the country will deepen if not addressed. Across US industries, accumulated technology debt, legacy hardware, and flawed or immature data infrastructure leave many organizations across critical sectors without the foundational capacity required to deploy and scale AI effectively.

Risks of frontier models and AGI

As AI capabilities rapidly progress, so does the need to address risks related to safety, security, and transparency. The increasing rate of change from the relatively narrow use of large language models to more complex, autonomous deployment of agentic AI and multimodal AI adds urgency to the need for technically grounded approaches to responsible deployment and governance. Organizations such as the Frontier Model Forum represent industry-led efforts to collaborate on understanding and managing risks to public safety, such as from the intersection of AI with chemical, biological, radiological, nuclear weapons and advanced cyber threats. Governments around the world have established national institutes focused on various elements of the AI safety and security paradigm. In the United States, the Center for AI Standards and Innovation, housed within NIST, collaborates with the private sector on evaluations and assessments of potential AI security vulnerabilities and has carried out more than forty evaluations of frontier AI models. It recently launched an AI Agent Standards Initiative focused on "ensuring a trusted, interoperable, and secure agentic frontier." Globally, the International Network for Advanced AI Measurement, Evaluation, and Science, previously led by the United States and now led by the United Kingdom' AI Security Institute, promotes the exchange of knowledge and best practices among national institutes. Such efforts are paramount as AI capabilities progress rapidly.

Findings and recommendations

Finding: The US government's role as a promoter of AI innovation remains essential. US industry relies on the ideas and talent that come out of the US university ecosystem—the strength of which is predicated on sustained federal support. Over-reliance on the private sector to drive AI innovation will constrain the opportunity space for new approaches and applications, weakening the nation's ability to innovate at the pace of global competition. Action is needed to ensure that the foundations of US innovation can continue to power progress.

Recommendations:

- **Strengthen funding for AI research and development.** Federal support is the linchpin of the US national research enterprise. Federal funding should focus on strengthening associated components of the broader AI ecosystem and supply chain, in addition to advancing AI capabilities. This includes fostering innovation in AI hardware design and manufacturing, cloud infrastructure, next-generation networking, and advancing novel AI applications in critical areas such as healthcare, energy, manufacturing, and transportation. The commonly cited recommendation from the National Security Commission on AI to double

annual federal funding for non-defense AI R&D year over year until it reaches $32 billion remains a worthy goal.

- **Explore novel funding structures.** The pace of technological advancement and the nature of AI ecosystems mean that the standard approaches to federal support might no longer suffice. The government should explore and expand funding mechanisms such as large-scale funding for top investigators in the model of the Howard Hughes Medical Investigator program (which has led to more than thirty Nobel Prizes), support for focused-research organizations, rapid grants to allow for quick reaction to technological developments, and joint research programs that can link students directly with private-sector AI projects to enable graduates to contribute immediately to frontier innovation. It is also worth sustaining NSF efforts to pioneer promising approaches to spur regional innovation and build nimble research teams that can accelerate innovations to market.

Finding: Investing in the infrastructure for innovation and adoption will power progress. Funding for AI and AI-related R&D is necessary but insufficient on its own to fully harness the potential of AI at the national level. Recent federal initiatives, such as the Genesis Mission and the NAIRR, represent promising steps toward establishment of the requiredinfrastructure to sustain innovation, but additional action is needed to achieve the envisioned impact and to support broader integration into the economy.

Recommendations:

- **Resource promising federal initiatives like the Genesis Mission and the NAIRR.** If adequately resourced with federal funding that transcends political cycles, the Genesis Mission and the NAIRR hold the potential to supercharge AI innovation: the Genesis Mission from the perspective of harnessing AI for advances in scientific discovery, and the NAIRR for ensuring that the US research community can pursue big, bold ideas in AI that set the groundwork for the next generation of applications and capabilities. The task force set up by Congress to develop a roadmap for the NAIRR recommended that it be resourced at a level of $2.6 billion over six years. As of fiscal year 2026, it has received annual funding of about $30 million per year.

- **Launch public-private partnerships to support research in critical, hardware-intensive domains such as robotics and manufacturing.** In addition to the computational and data infrastructure enabled through initiatives such as the Genesis Mission and NAIRR, there is potential to accelerate progress in areas such as AI-enabled robotics and manufacturing, in which the United States is falling behind China. China has five times as many robots working in its factories as the United States, and more than the rest of the world combined. The United States should pursue public-private partnerships in which the private sector can provide researchers access to high-cost hardware infrastructure and robotics platforms, including networking and other testbeds, and the federal government can fund the research initiatives—working toward the shared end goal of innovation and advancement within critical sectors. Models such as the Platform for Advanced Wireless Research program, which is supported by NSF and an industry consortium of thirty companies and associations, provide a blueprint for impactful collaboration to accelerate breakthroughs in critical technology areas such as advanced wireless networking.

- **Enhance efforts to expand broadband access.** The federal government should accelerate efforts to connect all Americans to high-speed internet—such as through the Broadband Equity, Access, and Deployment (BEAD) Program—ensuring that connectivity requirements reflect the demands of AI workloads, particularly in terms of uplink capacity.

Finding: Signals and incentives from the government can accelerate AI adoption in key sectors. Integration of AI in key sectors of the economy holds the potential to strengthen the overall US economy and build trust among the population as people begin to see the benefits from the technology in their lives. The history of general-purpose technologies shows that the rate of diffusion is a key determinant of their capabilities, and the United States should fight to win the race to diffuse AI throughout society. The government can play an outsized role, particularly in certain sectors that rely on federal funding or involve significant federal regulatory oversight, but concerted actions will be needed.

Recommendations:

- **Implement federal innovation incentives in key sectors to encourage AI adoption.** In sectors such as education, healthcare, or transportation, federal initiatives that encourage AI integration and modernization of legacy technical infrastructure through existing funding mechanisms hold the potential to drive the development of high-impact AI applications and accelerate tangible, positive impacts of AI on Americans' everyday lives. Such efforts could be realized by leveraging tax credits, low-interest financing, or long-standing grant programs by building in opportunities and requirements related to AI innovation and infrastructure modernization. Federal funding can also support pilot programs to test AI applications in high-impact settings, developing best practices and guidance for integration that can facilitate scaling across industries.

- Advance consistent regulatory policy to facilitate responsible AI adoption. As discussed in more detail in Chapter 5, clear and consistent policies at the federal level around the use of AI in specific sectors can both provide clarity to companies in terms of compliance and accelerate responsible adoption of AI.

Finding: Safety, security, and assurance are all necessary ingredients for continued innovation and leadership. Without trust in AI systems—built through rigorous safety testing, robust security measures, and independent assurance mechanisms—public and institutional adoption could stall, effectively slowing the pace of innovation regardless of technical capability. Significant safety or security failures coming from AI systems risk setting the industry back by years. Steps should be taken to reduce this risk.

Recommendations:

- **Double down on international cooperation.** International cooperation on AI safety and security can help develop and enforce shared safety and evaluation best practices and standards while building common awareness of emerging risks. The United States should lead in developing these AI safety and evaluation standards through NIST, engaging through international standards development organizations. Such cooperation becomes more urgent as AI capabilities rapidly evolve and transform the landscape in areas such as cybersecurity.

- Fund continued innovation in development of AI evaluation infrastructure, new benchmarks, and assurance methods. The US government should allocate AI R&D support to advancing the science of AI evaluation and developing new approaches to assuring AI model performance, which is particularly essential for high impact uses such as national security.

Chapter 4: Building talent and institutional readiness

Winning in AI requires developing technical skills, strong public institutions, and programs to support workforce transitions. The United States' AI leadership depends on three imperatives: developing homegrown talent from grades K–12 through graduate education; attracting the world's best researchers and builders; and adapting the existing workforce, including that inside government, to function effectively in an AI-driven economy. Together, these imperatives will define whether the United States can maintain its technological edge.

Put simply: sustaining US leadership in AI depends on the American people. This includes a full-spectrum talent strategy that develops world-class domestic talent, attracts the best global innovators, and continually integrates and upskills the existing workforce to keep pace with rapid technological change. In April 2025, the White House Executive Order Advancing Artificial Intelligence Education for American Youth made AI literacy a core US policy goal and serves as a baseline.

The United States' innovation potential stems from its world-class workforce. This workforce, in turn, has enabled the country's technology sector to produce the world's five most valuable global companies, including Microsoft, Google, Apple, and Amazon. However, while the United States has historically been the destination of choice for top-tier talent, it is gradually losing ground to countries offering attractive incentives and investing in their education systems. Based on robustness of education systems, labor force participation and innovation, the United States currently ranks fifth in the world behind Japan, China, Sweden, and Singapore. It ranks twenty-fourth out of thirty-seven Organisation for Economic Co-operation and Development (OECD) countries in math, and twelfth in science. The 2025 Computer Science (CS) Rankings featured just one US university in the top ten.

AI development requires both software skills and hardware skills. The former category includes skillsets traditionally associated with CS and AI: programming languages, machine learning, and natural language processing, as well as a strong foundation in mathematics and statistics. Data engineering talent is also critical to maintaining data pipelines for AI. Hardware skills include the talent required to build and operate data center infrastructure—including networks, optics, semiconductors, mechanical and electrical, and clean room engineers—core to AI. Also critical are mid-career practitioners who translate models into production systems, manage AI deployments at scale, and operationalize AI across sectors. These practitioners are as strategically important as the researchers who produce breakthroughs.

Beyond building a technical workforce, human soft skills such as ingenuity and creativity have traditionally made up a significant component of US leadership driving innovation—albeit one that is more difficult to measure. Sustaining AI's diffusion requires the capacity to identify where AI is genuinely useful, but also the imaginative capacity to see where it is not being applied at all. Here, the arts and the humanities—scholars, designers, ethicists, and domain experts who can bridge AI possibilities with social complexity—are key for durable AI adoption. An emerging AI talent base must also include capable policymakers. Federal and state governments alike must be able to anticipate, and not just react to, the workforce shifts AI will drive. The US Department of Labor's (DOL) AI Literacy Framework identifies five foundations of AI literacy: understanding AI principles (core capabilities and concepts); exploring AI uses; directing AI effectively (prompt engineering); evaluating outputs; and using AI responsibly. Literacy means understanding how AI works; fluency means knowing how to use it to solve real problems. For the United States to be truly AI ready, AI fluency—not just literacy—will be required across the nation's workforce and government.

Flashpoints

Evidence-supported policymaking

The narrative around AI-driven workforce changes is diffuse and speculative, and the scope of impact appears to demand urgent action on all fronts. However, if everything is urgent, nothing is: "policymakers," as one report put it, "cannot manage what they cannot measure." In the near term, the most salient measurement challenge is the impact on entry-level opportunities, paired with factors that are affecting new college graduates such as post-pandemic labor market corrections and older cohorts delaying retirement. A 2025 survey shows that 59 percent of recent college graduates have found it "very challenging" to secure a full-time, entry-level job. Many employers report plans to eliminate entry-level roles as AI adoption automates routine tasks, shrinking the pool of traditional "starter" jobs for early-career workers. At the same time, according to data from LinkedIn, AI can also create new jobs, with 1.3 million jobs added in 2024-25.

An educational system unprepared to meet rapidly growing AI skills demand

The United States cannot meet its AI-ready workforce goals through top-ranked institutions and higher education alone. The 2021 Atlantic Council Commission report spotlighted alternative pathways for an AI-ready workforce, including through career and technical education (CTE) programs and community colleges. However, hiring faculty and staff has become increasingly challenging as academic institutions face funding cuts and compete both with one another and with higher-paying opportunities in industry.

The talent pipeline begins before the higher education level. China's Ministry of Education, for example, has instituted a comprehensive, tiered AI education system across all K–12 levels. Unless the United States modernizes its own K–12 teaching approach, it will be outpaced in developing the talent base essential for AI competitiveness.

Attracting and retaining world-class talent

Highly skilled, foreign-born innovators and entrepreneurs have historically brought great economic value and innovation to the United States. A US National Bureau of Economic Research study found that immigrant inventors produce nearly a quarter of total innovation output (patents and economic value generated), despite representing just 16 percent of all US-based inventors. Sixty percent of US start-ups in the 2025 Forbes Artificial Intelligence 50 list have foreign-born founders. Four of the seven most valuable US public companies are led by immigrants, including Nvidia's Jensen Huang (Taiwan), Google's Sundar Pichai (India), and Microsoft's Satya Nadella (India).

Many countries have instituted visa policies to attract AI talent. China introduced a visa for young science and tech talent and has incentivized Chinese nationals to return home after their foreign education. Similarly, the United Arab Emirates (UAE) has begun offering a Green Visa that allows skilled foreign workers to self-sponsor. Canada's Global Skills Strategy enables Canadian employers to hire highly skilled foreign talent, with a processing time of only two weeks. In the United States, discussions about the importance of and ongoing need for such talent have become politically divisive. The challenge is nuanced, with a growing perception that US-born individuals are increasingly falling behind.

The federal government's talent challenge

Public institutions need AI talent to effectively govern, acquire, and adopt this technology. In 2024, the US federal government recruited two hundred AI experts from the private sector through an "AI talent surge." In 2025, however, many of these experts departed federal service as a result of US Department of Government Efficiency (DOGE)-led downsizing initiatives. While DOGE's initiatives created immediate pressures on government capacity, long-term, structural issues such as lengthy hiring timelines, salary caps, and rigid qualifications systems have long affected the federal government's ability to attract and retain talent. The Trump administration launched the Tech Force initiative in December 2025, with the aim of recruiting one thousand elite engineers who will "receive technical training, engage with industry leaders, and work closely with senior managers from companies partnering with the Tech Force."

The government needs in-house expertise for AI-fluent policymaking that supports desired workforce outcomes. The recruitment strategy for most technically demanding roles in government is finding people who are driven by mission. Models for attracting technical talent to public service include programs like the CyberCorps Scholarship for Service, which has recruited and trained more than five thousand individuals in cybersecurity since 2002. Such programs have the added benefit of creating a cadre that brings an understanding of how government works to their future private-sector roles, serving as unofficial bridge builders between sectors.

Findings and recommendations

Finding: AI talent is a critical enabler of national competitiveness and prosperity. The current US policy environment sends mixed signals. The Trump administration has articulated AI dominance as a paramount national goal through its AI Action Plan, built on the pillars of accelerating innovation, building AI infrastructure, and leading in international diplomacy and security. Yet the workforce policies that undergird this ambition are relatively weak. AI talent should be seen as a national security interest, acknowledging its dual role in maintaining a modern, agile government and as a cornerstone of sustained US economic leadership.

Recommendations:

- **Create a national interest designation for AI to recognize the technology, and the fostering of related talent, as a strategic priority.** This designation, enacted through legislation, should carry concrete operational weight, including expedited processing for AI-related visa petitions, preferential fellowship eligibility, and a clear legislative mandate for cross-agency coordination on talent strategy.

- **Measure and disclose AI talent trends.** To address the lack of robust, longitudinal data, the US Office of Science and Technology Policy (OSTP), in collaboration with DOL, should publish an annual report on the state of AI talent in order to track domestic supply, international competition benchmarks, and federal talent gaps.

Finding: Education reform and stronger nontraditional pathways for AI skilling are urgently required. Broad-based AI fluency—building on a baseline of basic literacy—across the US population is essential for enabling trust. AI is often perceived negatively by Americans who feel they have little control over how the technology is used in their lives. The Trump administration has moved in the right direction with its April 2025 Executive Order on Advancing AI Education for American Youth, establishing a policy of promoting AI literacy and proficiency by integrating AI into education, providing comprehensive AI training for educators, and fostering early exposure to AI concepts to develop an AI-ready workforce. However, additional actions should be taken.

Recommendations:

- **Fully fund and scale the White House AI Education Task Force's K–12 mandate.** The task force is executing a presidential AI challenge for US youth, establishing public-private partnerships for K–12 AI education, and strengthening existing agency programs that can support K–12 education initiatives. These are important initial steps but must be backed by dedicated and sustained federal appropriations.

- **Establish funding for the integration of AI core topics in traditional college education.** Congress should establish dedicated federal funding, administered through NSF and DOE, to support the integration of foundational AI topics (AI literacy, computational thinking, statistics, foundational machine learning concepts) into undergraduate curricula across all disciplines, not just computer science programs. Similarly, relevant humanities credits (philosophy, ethics, sociology, law) should be integrated into CS/AI degrees. Funding should support curriculum development and institutional capacity building, with priority for state-funded universities that serve student populations at scale.

- **Prioritize teacher training as an enabling step.** The bottleneck in K–12 AI education is educators' capacity to deliver it. NSF and DOE should expand funding for educator credentialing in AI literacy.

- **Expand and formalize nontraditional AI skilling pathways.** The establishment of public-private partnerships to provide lifelong learning and reskilling initiatives is core to cultivating a broader AI-ready workforce and reorienting the existing workforce toward AI-enabled workflows. Examples include the National Applied AI Consortium, a federally funded program that partners with tech companies through community colleges to build capacity for AI education and training nationwide. Such partnerships can also help align curricula with skills the market demands. This includes the expansion of pathways such as vocational training, micro-credentials, community colleges credentialing, and apprenticeships.

Finding: Attracting and retaining talent is foundational to US technological leadership. The United States must be ready to compete with other countries seeking to attract talent. However, it is currently undermining AI dominance goals through an immigration posture that deters the talent it needs. For example, the Trump administration's $100,000 fee for new H-1B visa petitions went into effect in September 2025, disrupting high-tech talent pipelines, particularly for small businesses and venture-backed start-ups. Instead, the US government should institute sensible immigration reform that balances concerns about immigration inflows with the crucial role foreign talent plays in the innovation ecosystem.

Recommendations:

- **Reform the H-1B program to better serve AI talent goals.** Reform of the H-1B program will incentivize foreign nationals to join the US workforce at the cutting edge of innovation through sensible sponsor fees, higher visa caps, and predictable transfer, renewal, and adjudication timelines. The H-1B program as it is currently designed also constrains workers, making it challenging to change employers without jeopardizing status. On the sponsor side, the system creates significant burdens for small firms and start-ups, allowing larger incumbents to derive the most benefit. For fiscal year 2027, US Citizenship and Immigration Services (USCIS) is replacing the random H-1B lottery with a weighted selection process. USCIS should work with OSTP to develop a weighting rubric that favors AI skills within the H-1B selection process.

- **Extend Optional Practical Training (OPT) and science, technology, engineering, and mathematics (STEM) OPT for AI-relevant PhD graduates.** Currently, post-graduate work authorization for international students in STEM fields caps out at three years before an H-1B is required. Given the low H-1B lottery odds (less than a 30 percent approval rate in recent years), many top AI researchers complete their training at US universities and then leave for countries with more predictable pathways to residency and citizenship. Studies by the Center for Security and Emerging Technology (CSET) and the National Academies have found bottlenecks in the US immigration system are an extremely relevant factor in these researchers' decisions to leave. Extension of the post-graduation work authorization for PhD candidates in AI-related fields is critical for retaining top talent during these individuals' peak years of innovation potential.

- **Create a dedicated AI innovator visa.** Congress should authorize a new visa category specifically for individuals with demonstrated expertise in AI research, development, or governance. This category can be modeled on the O-1 visa for extraordinary ability but with streamlined adjudication timelines, lower evidentiary burdens for early-career researchers, and a direct pathway to permanent residency.

Finding: Building a skilled federal AI workforce is a prerequisite for effective governance and deployment of AI. The federal government cannot effectively govern, acquire, or deploy AI systems without meaningful in-house expertise. This is both a capability gap and a credibility gap. Agencies without AI-literate staff are poorly positioned to hold vendors accountable, set meaningful procurement standards, or anticipate second-order effects of AI adoption on their missions.

Recommendations:

- **Strengthen Tech Force with sustained employment guarantees and mid-career tracks.** Following an exodus of AI talent from government agencies, the Tech Force program should create sustained pathways to ensure mission-driven individuals are retained and trust in federal government career pathways is rebuilt.

- **Revive and expand mission-driven pipeline programs.** Programs like the CyberCorps Scholarship for Service demonstrate that structured, mission-first talent pipelines can work at scale. A parallel AI Scholarship for Service program, administered through NSF in partnership with universities, should be established to provide full tuition coverage and a stipend to students pursuing AI-related graduate degrees in exchange for a minimum two-year federal service commitment. This would create a valuable pipeline of AI talent into government.

- **Mandate AI fluency training for all senior federal acquisition and procurement officials.** Acquisition roles are an underappreciated aspect of the federal AI talent gap. Federal contracting officers who lack AI literacy cannot write effective requirements, evaluate vendor proposals meaningfully, or hold contractors accountable for performance. The Office of Management and Budget (OMB), in coordination with the Federal Acquisition Institute, should develop and mandate a standardized AI certification for all contracting officials working on AI-adjacent acquisitions.

Finding: Understanding the impact of AI on the US workforce writ large is necessary to build sound policy responses. The United States currently lacks a consistent, standardized government-produced assessment of how AI is reshaping the labor market in real time, including disruptions for mid-career workers. The White House AI Action Plan directed DOL to establish an AI Workforce Research Hub "to lead a sustained federal effort to evaluate the impact of AI on the labor market and the experience of the American worker."

Recommendations:

- **Launch a quarterly AI workforce impact assessment.** DOL should publish a quarterly assessment, modeled on the rigor and accessibility of the monthly jobs report, tracking job displacement and creation attributed to AI adoption by sector. The report should track shifts in wages for AI-adjacent skills, as well as effects on entry-level jobs. Proposed legislation would direct the Departments of Labor, Education, and Commerce to assess which demographic groups are most likely to benefit or be harmed by AI-driven changes. Standardized, recurring measurement will enable policymakers to detect and more proactively manage emerging impacts.

- **Operationalize the AI Workforce Research Hub.** The hub directed in the AI Action Plan could serve as the institutional home of the quarterly assessment. It should establish data-sharing partnerships with AI labs (which have begun publishing their own economic impact surveys). Additionally, the US Bureau of Labor Statistics (BLS) incorporates AI into its ten-year employment projections, but notes, "The timing and scale of many potential impacts of GenAI [generative artificial intelligence] are too uncertain to be reflected in BLS projections." The hub should collaborate closely with BLS to structure projection methods related to fit-for-purpose AI for rapid technological change.

Chapter 5: Governance of AI

Governments have become primary partners in the AI economy, not only as regulators but as major users, integrators, investors, and market shapers.

In the United States, the scale of federal AI adoption has expanded rapidly. As of April 2026, federal agencies had reported **3,611 individually reported AI use cases**, including 445 high-impact use cases, across forty-one agency submissions spanning internal operations, healthcare applications, and service delivery. This growth has been reinforced by procurement and integration efforts led by the US General Services Administration, which have **facilitated access** to frontier AI systems, including large language models, from leading private-sector providers such as Anthropic, Google, OpenAI, and xAI. As agencies increasingly rely on private-sector AI systems, new governance and procurement challenges have emerged.

Policy guidance has evolved in parallel. In April 2025, the OMB issued Memoranda **OMB M-25-21** and **OMB M-25-22**, which together established a more structured approach to federal AI adoption. These directives emphasized responsible innovation, streamlined acquisition, and risk management, while mandating the designation of chief AI officers within agencies to oversee governance, deployment, and compliance. In December 2025, OMB issued Memorandum **OMB M-26-04**, "Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles," which introduced new procurement and transparency requirements for generative AI systems, particularly large language models used by federal agencies. At the same time, national security applications of AI have accelerated. The US Department of Defense (DOD) **launched** the AI Acceleration Strategy in January 2026, underscoring a shift toward integrating advanced and agentic AI capabilities into military operations.

Throughout this expansion, federal AI programs face new—but expected—operational and governance challenges. Existing frameworks for cybersecurity, data privacy, and internet technology acquisition might not always be agile enough to manage the unique risks of generative AI. For example, a privacy policy can **impede** data sharing among agencies and prevent effective collaboration on shared AI applications. Agencies must also contend with risks unique to AI systems. **NIST's Generative AI Profile (AI 600-1)** identifies a range of these risks including hallucinated outputs, information integrity, and data privacy and security risks, which in federal contexts raise particular concerns about preserving distinct classification controls across data inputs and outputs in AI systems. Agencies also face emerging governance gaps around the integrity and interoperability of data and the absence of common standards for tracking model provenance, including where models originated, what data they were trained on, and what modifications have been made. At the same time, the federal government faces intense competition from the private sector in attracting and retaining AI talent, as well as in securing access to the computational infrastructure required to develop and deploy advanced systems. These challenges are compounded by risks unique to generative AI, highlighting the need for updated governance approaches that better align security, usability, and innovation.

More broadly, AI governance in the United States has emerged through a patchwork of policies and **voluntary guidelines and best practices** rather than a unified legal framework. Federal agencies rely on existing authorities, executive guidance, and procurement rules to govern AI systems, while Congress has yet to pass comprehensive AI legislation. Meanwhile, states are moving forward with their own AI laws and regulations, producing a fragmented governance landscape that complicates compliance, innovation, and national policy coordination. Recent federal actions have sought to address this fragmentation. The December 2025 **White House executive order** on national AI policy reflects an effort to establish a more unified federal framework, including potential preemption of state-level regulations viewed as barriers to innovation. However, this approach introduces new tensions between national consistency and state-level experimentation, particularly in areas such as safety and civil rights.

Beyond domestic governance, the United States is also shaping the global AI ecosystem through industrial policy and export controls. Since 2022, **restrictions** on advanced semiconductors and related technologies have aimed to limit adversaries' access to critical AI infrastructure. At the same time, the 2025 **AI Action Plan** outlines a complementary effort to expand the global adoption of US AI technologies and infrastructure, particularly through partnerships with allies and partners. Together, these efforts reflect a strategy focused both on restricting access to critical inputs and scaling the US AI ecosystem internationally.

Flashpoints

Balancing rapid adoption with risk management

Federal AI policy is attempting to accelerate adoption while also ensuring responsible use. **OMB M-25-21** and **OMB M-25-22** encourage agencies to reduce barriers to innovation and streamline procurement, while also requiring risk management practices for high-impact AI systems. More recent guidance, including **OMB M-26-04**, adds new requirements for generative AI systems related to transparency, documentation, and unbiased outputs, particularly in procurement. These policies show that the government is trying to move quickly on adoption while simultaneously building governance and

oversight mechanisms, which can slow implementation and create uncertainty for agencies and government vendors. The recent Anthropic dispute with the DOD over military use of its Claude AI models underscores the real-world tensions between federal procurement, AI safety policies, and national security requirements.

Fragmentation between federal and state AI policies

The United States does not yet have a unifying body of law governing artificial intelligence. Federal oversight is spread across sector-specific regulations, agency guidance, procurement policies, and voluntary standards. At the same time, states have started to pass their own AI laws, producing a fragmented and uneven regulatory landscape. Recent polls show low public optimism toward AI in the United States. Comprehensive and forward-looking AI governance is essential to restoring public trust. The AI Action Plan marks a positive step by outlining national objectives for AI innovation, infrastructure, and diplomacy, but it lacks the force of law. Meanwhile, poorly designed federal preemption could override state regulations that aim to protect civil rights and public safety, which are areas of legitimate concern for many Americans.

Insufficient AI workforce inside the government

Federal agencies continue to face challenges in recruiting and retaining AI talent, competing directly with the private sector for technical expertise. Agencies also face difficulties in terms of providing education and training to the current workforce, competition with the private sector for AI-skilled professionals, compensation limitations, and slow hiring processes. These challenges have been compounded by the restructuring or elimination of some government technology units in recent years. The US Tech Force is intended to bring approximately one thousand technical experts into government for limited terms to help implement AI projects. However, long-term workforce and institutional capacity challenges remain.

Findings and recommendations

Finding: Federal leadership in AI governance is necessary, but it is currently fragmented and inconsistently implemented. The federal government has issued substantial AI governance guidance—including OMB Memoranda M-24-10 and M-25-21—but implementation remains uneven across agencies, and no statutory framework exists to enforce consistent standards or define permissible use cases at a national level. Action is needed to establish durable, enforceable federal leadership in AI governance.

Recommendations:

- **Establish a coherent federal regulatory framework.** The federal government should develop mandatory and voluntary standards addressing risk tiering, procurement standards, data governance, transparency requirements, and incident reporting to prevent fragmented governance, reduce reliance on diverging state-level or industry-led frameworks, and avoid repeating challenges seen in privacy regulation. It should collaborate with private-sector developers to leverage technical expertise while safeguarding public trust, particularly for high-risk applications. This framework should ensure consistent implementation of existing guidance, close gaps where OMB authority is limited, and create a foundation for consistent national policy.

- **Act as a catalytic adopter of AI.** As a complement to efforts outlined in preceding chapters to expand access to compute, data, and educational resources, the United States should use federal adoption strategically to drive broader innovation across states, sectors, and industry. Effective federal AI deployment can set standards, demonstrate best practices, and encourage wider adoption downstream. By showcasing how AI can be integrated safely and efficiently in public programs, the federal government can accelerate adoption and improve outcomes across government, private-sector, and state-level initiatives.

Finding: Federal agencies lack the procurement systems and workforce capacity needed to adopt AI effectively. Agencies face significant barriers to integrating AI into operations—from outdated procurement processes to a workforce that has not kept pace with the demands of overseeing and deploying AI systems. Addressing these structural gaps is essential to realizing the benefits of AI in government.

Recommendations:

- **Reform procurement systems to enable rapid acquisition of emerging AI capabilities.** Agencies should implement adoption cycles that allow for continuous evaluation, adaptation, and scaling of AI tools while maintaining operational security. Procurement reform should include structuring acquisitions around outcomes rather than inputs and should be supported by readiness assessments such as data audits and governance reviews. Agencies should also collaborate with AI evaluation experts (e.g., at NIST) to streamline and standardize assessments of AI tools and platforms, including for security and auditability.

- **Invest in training and professional development for administrative and executive staff.** Agencies need people capable of responsibly adopting and overseeing AI systems and should revitalize technology-focused positions reduced by prior budget cuts.

Finding: AI's geopolitical significance demands that national security and international competitiveness are central to federal AI strategy. The United States must develop governance approaches that balance innovation, security, and strategic influence—ensuring that AI adoption strengthens, rather than undermines, the United States' global position. Policy and investment decisions should align AI deployment with strategic objectives, support defense readiness, and enhance the United States' position as a leader in emerging AI technologies.

Recommendations:

- **Ensure AI adoption supports national defense and strengthens partnerships with allies**. Federal AI strategy should prioritize applications that reinforce defense readiness and maintain US competitiveness in AI development and deployment. This includes joint AI infrastructure investments with allies, coordinated export control regimes, and shared testing and evaluation arrangements for defense-relevant AI systems.

- **Develop governance approaches that balance innovation, security, and strategic influence.** Policy frameworks should account for AI's dual-use nature, ensuring that the drive for innovation does not come at the cost of security or strategic positioning.

Chapter 6: Supply chain

The coordination of chips, data, infrastructure, code, and human ingenuity produces AI systems. The AI supply chain consists of the physical and virtual resources required to develop and deploy AI models, including raw materials, specialized technology, and manufacturing capacity. Defining the scope of the AI supply chain is a pressing challenge for policymakers, as overly broad definitions risk diluting prioritization efforts and overly limited definitions fail to anticipate relevant risks. Certain critical components—such as advanced semiconductors, networking infrastructure, software and computing resources, data, and models—form the core of the AI supply chain, but each component retains external dependencies. This complex chain faces growing threats of disruption, coercion, and strategic dependency, which threaten the security, reliability, and availability of AI systems.

As AI and machine learning have transitioned from niche academic topics to strategic national interests, the AI supply chain has become both a focus of policy debates and a multifaceted geopolitical competition. Export control debates have fixated on the semiconductors used to develop and deploy AI models, while other deliberations focus on the copyright implications of training data and the security and sustainability of data center buildouts. Meanwhile, the Trump administration's AI Action Plan calls for the expansion of US AI infrastructure and semiconductor manufacturing capability, as well as an acceleration of AI innovation through measures such as removal of regulations and development of datasets. Governance frameworks such as the European Union (EU) AI Act impose controls on AI applications, while voluntary commitments from AI labs have led to pre-release model evaluations by governments and nonprofits. The recent US National Cyber Strategy also identified securing the AI stack as a priority, recognizing the importance to US leadership of data, infrastructure, and models. Initiatives from the US Department of State such as Pax Silica, which frame allied coordination on semiconductor and critical mineral supply chains as crucial components of AI competitiveness, reflect the growing reality that securing the AI supply chain is inseparable from broader industrial and geopolitical strategy.

As the global consensus converges on the status of AI as a technology of critical national importance, strengthening and securing the nation's AI supply chain in the face of nation-state and non-state threats will continue to be a pressing challenge. Disruptions affecting key US cybersecurity institutions compound these challenges. The dissolution of the Cyber Safety Review Board (CSRB), as well as the backlog and funding challenges affecting the US Cyber and Infrastructure Security Agency (CISA) and NIST's National Vulnerability Database (NVD), have all reduced the US government capacity to address emerging and existing risks to the AI supply chain, even as the cyber capabilities of frontier AI models grow. Policymakers and companies will need to establish and verify the trustworthiness of the entire supply chain, rather than individual components. Concerns regarding espionage, sabotage, and the exfiltration or proliferation of advanced capabilities will continue to captivate policy and industry audiences. But without meaningful consensus on evaluation criteria for trusted AI infrastructure and supply chains, policy initiatives will languish in the implementation phase.

Flashpoints

Supply chain and manufacturing sovereignty

In response to the strategic vulnerability of US reliance on overseas manufacturing, successive US administrations and Congress have encouraged the establishment and expansion of domestic manufacturing facilities for advanced semiconductors. Concentration in the chip industry has resulted in a handful of companies, particularly Taiwan Semiconductor Manufacturing Company, holding outsize importance to the manufacturing of chips essential to the AI supply chain. Governments must ensure visibility into the potential strategic risks of concentration, dependence, and delegated control across the AI supply chain. International debates about digital sovereignty will overlap with AI-specific concerns about compute access, leading to proposed geographic restrictions on access to AI infrastructure and manufacturing capabilities. The escalating costs of AI development will drive competition and create strategic dependencies, forcing countries and enterprises to make difficult choices between leveraging the advantages and redundancies of global capabilities or maintaining domestic control and visibility.

Software and compute insecurity

Managing the software used to develop, test, and deploy AI models will be a critical challenge. Third-party packages and infrastructure, including open-source software, are an essential feature of most enterprise AI use cases, and establishing best practices for security and reliability will be a prerequisite for robust use. AI models increasingly possess advanced cybersecurity capabilities, including for vulnerability detection and exploitation, which will test the ability of organizations to respond at speed and scale to secure their software. Computing providers will need to manage ever-evolving hardware, including specialized chips, memory, and networking infrastructure. Providers must balance availability, resilience, and security considerations, particularly as they configure their global infrastructure footprint. Governments will need to weigh the benefits of incentivizing robust security practices with the risks of imposing prohibitive costs on enterprises.

Model openness

After the release of the Chinese DeepSeek-R1model, debates burst into the public sphere about the practice of "open-sourcing" AI models making their weights public to allow convenient, self-hosted, and largely unrestricted use. US companies, including frontier labs, continue to strategically release open-source models in addition to their flagship proprietary models, and the Trump administration's AI Action Plan included a provision emphasizing the strategic importance of US leadership in open models to become global standards. Organizations will continue to seek information on the sources, training data, and safeguards used in both open- and closed-source model development. The tension between demonstrating capabilities and limiting adversary misuse of models will continue, and companies and policymakers should continue to assess whether open-source models create meaningfully distinct risks and should be assessed differently than closed-source models.

Data access

The sheer amount of content used to train flagship AI models has resulted in both a litany of court cases over the limitations of copyright and alarms about a plateau in model quality without continued increases in access to quality training data. Governments, consumers, and enterprises have expressed concerns about advantages in data access that specific countries, including China, could derive from their centralized prioritizations of data gathering and permissive regulatory environments for data use. As AI adoption continues, these same stakeholders will continue to seek clarity about the data captured, accessed, and transformed by AI tools as well as verify the fidelity of data from both adversarial influence and injection. Ensuring data access while managing data protection needs will require both amplifying existing protections, such as encryption, and developing new techniques to prevent novel threats, such as model distillation and training data leakage. One size will not fit all, and policymakers should continue to evaluate which sensitive data components deserve enhanced scrutiny and protection, and which data elements should be fully utilized in model training and tuning.

Export controls have slowed, but not handicapped, China's AI ambitions

As a dominant player in global AI markets, the United States has used export controls on advanced AI chips, semiconductor manufacturing tools, and related personnel support to shape international AI access. The effectiveness of these controls in limiting China's AI capabilities depends on enforcement, the pace of Chinese domestic alternatives, and the willingness of allied jurisdictions to align on parallel restrictions. These measures have been reinforced by congressional proposals, such as the SAFE Chips Act and the Comprehensive Outbound Investment Security Act, to tighten controls on semiconductor exports and outbound investment to China. In early 2026, the US government approved licensed exports of Nvidia H200 AI chips to China under conditions intended to balance national security with commercial interests—a move that sparked substantial debate in Washington. At the same time, smuggling, enforcement challenges, and China's push for domestic alternatives underscore persistent limitations in how comprehensively these controls can curb Chinese AI capabilities, as evidenced by the growth of Chinese companies such as DeepSeek, Alibaba, and Moonshot.

Findings and recommendations

Finding: Developing a standard definition and conception of the AI supply chain will lend coherence and connectivity to policy initiatives. Policy efforts to surge capacity and resources for emerging technology require defined scopes and clear prioritization. Before the US government can establish lines of effort and workstreams to effectively collaborate with allied and partner governments and industry partners, action is needed.

Recommendation:

- **Drive alignment around a definition of the AI supply chain.** The government should delineate which components of the AI supply chain are in focus for each of the US government's AI initiatives. Without this clarity, the interconnected nature of the AI supply chain will lead to redundant efforts and insufficient resources to meet overarching challenges.

Finding: Semiconductor manufacturing and critical mineral supply chains require sustained, long-term commitments. The concentration of advanced chip manufacturing in a small number of companies and geographies is a risk that initial policy actions are far from resolving. The supply chain challenge extends beyond the production of cutting-edge chips and must address advanced packaging, other semiconductors, advanced manufacturing equipment such as extreme ultraviolet (EUV) and deep ultraviolet lithography equipment, and critical mineral supplies. Policy attention has focused disproportionately on leading-edge graphics processing units (GPUs), while other categories of chips, such as field-programmable gate arrays (FPGAs), which face distinct competitive dynamics and vulnerabilities, have been largely overlooked. Without sustained efforts across the full manufacturing value chain, early investments risk being stranded.

Recommendations:

- **Align export control and industrial policy objectives.** Export controls on advanced semiconductors and related technologies are critical to maintaining US advantages in computing, but they must be calibrated to effectively constrain adversary capabilities without undermining the

commercial viability of US companies or alienating allied nations. Enforcement challenges and China's investments in domestic alternatives demonstrate why export and import controls should be calibrated as one element of a broader industrial policy strategy. Export controls must also take into consideration the tertiary impact on capital available, especially among hyperscalers, to invest in domestic energy and other infrastructure.

- **Sustain and expand incentives for domestic semiconductor manufacturing, assembly, testing, and packaging**. Congress should extend incentives for advanced packaging, testing, and assembly operations. Advanced packaging is increasingly recognized as a bottleneck to domestic capabilities, and federal incentives should explicitly target the buildout of domestic capacity in this area.

- **Address the distinct vulnerabilities of specialized semiconductors, including FPGAs.** Current semiconductor policy has assumed that the vulnerabilities and strengths in leading-edge chip markets apply uniformly to specialized silicon. FPGAs, which are critical for defense systems, telecommunications, and AI inference, depend on lagging-edge manufacturing nodes in which China has built massive state-backed capacity. Incentives for allied assembly, testing, and packaging capacity should be part of supply chain diversification efforts for these segments, in which China already holds a significant downstream advantage.

Finding: The security of AI software and compute infrastructure requires both applying existing cybersecurity best practices and developing AI-specific mitigations. AI systems inherit the security challenges of the software and cloud ecosystems in which they are built and operated, but they also introduce novel risks that existing frameworks do not adequately address. Meanwhile, the institutional infrastructure for cybersecurity in the United States has been weakened at precisely the moment AI is raising the stakes. These institutions are the foundation on which any AI-specific security effort must be built. Creating new AI security mechanisms without repairing the underlying infrastructure will retread old obstacles and prevent rapid advancements.

Recommendations:

- **Restore and strengthen federal cybersecurity institutions for the AI era.** Congress should permanently reauthorize the Cybersecurity Information Sharing Act of 2015, adequately fund the NVD, and ensure CISA has the workforce and authorities it needs to support AI supply chain security. The dissolution of the CSRB left the United States without a critical investigative mechanism for analyzing major cyber incidents; a successor capability should be established. Without these foundational institutions operating at full capacity, AI-specific security initiatives will lack the operational infrastructure they need to succeed.

- **Map AI supply chain components to existing cybersecurity best practices and identify gaps.** Developers, deployers, and policymakers should avoid reinventing the wheel. They should systematically map existing cybersecurity controls and best practices to components of the AI supply chain. Where existing best practices fall short for AI-specific risks, standards bodies such as NIST and AI-focused entities such as the Center for AI Standards and Innovation should develop guidelines for targeted mitigations. This approach of assigning general cybersecurity responsibilities to established bodies, while directing AI specific analysis to focused organizations, avoids burdening AI safety-focused organizations with responsibility for the security of the entire software and data ecosystem while ensuring AI-specific gaps receive dedicated attention.

- **Secure the open-source software ecosystem.** Nearly every AI system, including commercial frontier models, depends on third-party open-source libraries and packages. For AI infrastructure components for which reliability and security are paramount, the US government should incentivize migration to memory-safe programming languages like Rust, which prevent developers from introducing a class of vulnerabilities into programs, and fund tooling that makes this transition practical. Targeted federal investment in independent security audits of critical open-source packages would improve ecosystem-wide resilience and could serve as a blueprint for the prioritization of neglected resources that are both increasingly critical to AI development and increasingly at risk due to the rapidly advancing cybersecurity capabilities of frontier models.

- **Establish structured communication mechanisms between industry and government for AI supply chain threat intelligence.** The US government should work with industry to develop standardized reporting frameworks, including common taxonomies for threat classification, vulnerability disclosure, and standard protocols. Government bodies should share actionable intelligence with industry where possible, deepening trust and collaboration between critical entities in the AI supply chain. Establishing and resourcing an AI Information Sharing and Analysis Center (AI-ISAC), modeled on existing ISACs as recommended in the Trump administration's AI Action Plan, would be one pathway to the partnership and collaboration the government must establish with the private sector.

Finding: Model openness is emerging as an arena of strategic competition that demands a coherent US approach. Chinese open-source models provide users performance close to that of the leading proprietary models at a fraction of the cost and without restrictions on use that accompany many US open-source models. In 2025, the volume of downloads of models from Chinese providers on the platform Hugging Face surpassed that of US counterparts. According to leadership at Andreessen Horowitz, about 80 percent of the start-ups it sees leveraging open-source stacks are using Chinese models. US companies such as Airbnb, Pinterest, and Notion have all built open Chinese models into their AI offerings. The rapid global adoption of Chinese open-source AI models has implications that extend beyond the innovation landscape into the supply chain domain, creating dependencies on Chinese-origin technical ecosystems that could prove difficult to reverse. The United States must compete more effectively in the open-source space while developing frameworks to manage associated risks.

Recommendations:

- **Close the gap in open-source model disparities.** US government and industry should come together to spur more robust open-source offerings from the US AI ecosystem. This includes providing adequate compute to academic and nonprofit developers of open-source models, supporting the development of open-source consortia to pool resources, reducing regulatory and legal hurdles to the open publication of model parameters and training data, and facilitating diffusion of open models. The US federal government should also carry out evaluations of Chinese open-source models, as articulated in the AI Action Plan.

- **Develop risk assessment frameworks for open-source AI models.** Open-source models can be modified in ways that closed models cannot, enabling new vectors for experimentation and for risk. Companies and policymakers should assess whether open-source models create meaningfully distinct supply chain risks and should be treated differently than closed-source models. Risk frameworks for both open- and closed-source models should account for the full lifecycle of model use, including fine-tuning, deployment, and ongoing maintenance, rather than focusing narrowly on the moment of model release.

- **Treat model ecosystem influence as a supply chain question.** Policymakers should integrate model openness into broader supply chain and industrial policy thinking. The question of which nation's models are embedded in global AI applications carries strategic weight analogous to which nation's telecommunications equipment is embedded in global and domestic networks. Broader conversations about supply chain security should reflect this reality, and federal procurement guidance should address model-origin considerations for AI systems used in sensitive government applications.

Finding: Fragmented approaches to data governance have left the United States without a coherent framework for securing and ensuring the integrity of the entire set of data in the AI supply chain. Policymakers have repeatedly shifted focus from one data component to the next, from training data to model weights to inference outputs, without building a comprehensive framework for securing the full spectrum of data in the AI supply chain. This pattern produces policy that is reactive, incomplete, and vulnerable to the emergence of new threat vectors that fall outside the scope of which component happens to be in focus. Overconfidence about which data element will most drive AI development can lead policymakers to skip past important questions, wrongly treating them as resolved. A more systematic approach is needed.

Recommendations:

- **Develop tailored data governance approaches that balance access and protection.** The United States should pursue data governance approaches that enable access to high-quality data for AI R&D while providing meaningful protections against novel threats to the integrity, availability, and confidentiality of essential systems, such as through model distillation, training data extraction, and inference attacks. This includes investing in privacy-enhancing technologies, establishing clear frameworks for access to federal datasets, and engaging with industry on standards for model robustness to adversarial attacks. The United States should also explore the feasibility of establishing a platform for exchanging data between government and commercial sources for the purpose of AI development.

- **Address the competitive implications of asymmetric data access.** The advantages in data access that countries like China derive from centralized data gathering and permissive regulatory environments represent a structural challenge that the United States cannot and should not replicate. Instead, the government should pursue policy strategies that leverage the strengths of the US model, including the depth of the private-sector data ecosystem, the quality of federally curated datasets, and the trust that robust privacy protections engender in international partners, while ensuring that regulatory uncertainty does not unnecessarily constrain US competitiveness.

Chapter 7: Powering AI

Electricity supply has shifted from being of peripheral interest to central for AI and technology companies. As AI companies release new, larger models and AI adoption accelerates, potential electricity shortfalls are looming larger. Accordingly, US AI leadership will depend, in part, on how the country meets surging AI-related demand. Three challenges are critical for policymakers and industry: building out the power system with new infrastructure and innovative technologies, making more efficient use of power systems, managing the gap between projected and actual AI load, and meaningfully addressing social license concerns.

A primary goal for US AI is to build an electricity system that can deliver sufficient electricity to power AI workloads. Electricity is already a bottleneck for US data centers, and insufficient or unreliable supply could slow or derail US AI ambitions. Although AI's trajectory and ultimate impact remain uncertain, the technology will likely reshape US electricity infrastructure and markets. With the existing grid already due for replacement, the system is primed for a period of fundamental change.

The electricity sector has little historical experience managing rapidly growing demand and dynamic operating conditions, but it is now contending with immense uncertainty and transformation. From 2010 to 2021, US electricity demand was relatively flat, as demand rose only 1.4 percent to 3,805 terawatt hours, for a compound annual growth rate of only 0.13 percent. In contrast, US power demand has risen consistently in recent years, driven by electrification, industrial revitalization, and data center consumption for cloud computing needs and AI. While residential and industrial electricity demand are up 3 percent and 4 percent, respectively, from 2021 levels, commercial electricity demand rose 12.4 percent from 2021 to 2025. Data center hubs unsurprisingly outpaced average national growth rates: Virginia, which hosts the largest data center market in the world by operational capacity, accounted for 15 percent of the increase in total commercial demand growth, while other emerging data center hubs in states such as Texas and Georgia also outpaced the national average.

Under the right conditions, data centers can have a net positive impact on electricity affordability. The large and consistent power demand profile of data centers, along with hyperscalers' ability to support investments in grid services and upgrades, means that new data center load can lower average retail electricity prices when costs are allocated according to the "beneficiary pays" principle. Furthermore, in US states with high demand growth over the past five years, average inflation-adjusted retail prices have generally declined or risen less than in states with flat or falling demand, where prices have generally increased. The net impact of data center load growth depends largely on the characteristics of the existing system and new load, as well as how utilities design rates, conduct system planning, and allocate costs.

Even as electricity demand is rising, supply and transmission are straining to keep pace. Much of the grid is at end of life: more than half of US coal plants were built in the Jimmy Carter presidency or earlier; coal plant maintenance has been deferred ahead of retirements; 55 percent of in-service distribution transformer units are more than thirty-three years old and are approaching the end of their service life; and much of the US electricity grid, including transmission, was built around the time of the moon landing.

The United States is struggling to bring new generation and transmission capacity onto the grid, especially in comparison to China. In 2025, the United States deployed 55 gigawatts (GW) of new generation to the grid. Conversely, China deployed 543 GW of new generation capacity that year, or more than the United States has cumulatively installed since 2008. According to Chinese sources, in 2023, China built about 40,445 kilometers (or 25,133 miles) of new transmission lines greater than 220 kilovolt (kV); the US built 888 line miles of high-voltage (more than 345 kV) transmission lines in 2024, the last full year for which data are available. The United States is not falling behind China in the electricity race; it is not even competing.

Taking all this together—growing demand, constrained supply, the need to replace aging infrastructure, and insurance risks from climate change—US electricity prices are rising and might begin soaring. While electricity prices changed little from 2010 to 2021, in real terms, they are now outpacing inflation and have risen nearly 21 percent from 2021 to 2024. In addition, if inflation and interest rates rise amid global supply chain outages or other crises, especially those surrounding oil, then capital costs will climb, which ultimately increases the total cost of the infrastructure buildout and can cause projects still under development to become unviable. With power demand projected to rise due to AI and other factors, incremental generation capacity deployment facing growing headwinds, and the grid nearing its end of life, most analysts believe prices will continue their march upward.

Rising electricity prices could spark backlash, constraining AI development. While electricity prices have traditionally not been as impactful for consumers as gasoline or groceries, nor as psychologically salient, public opinion polls suggest consumers are increasingly sensitive to electricity prices. Importantly, the effects of rising electricity will be uneven, falling most acutely on individuals earning lower incomes and living in rural areas. Electricity costs will figure increasingly prominently in political debates over AI.

US AI efforts will hinge, to a significant degree, on the nation's ability to provide reliable and affordable electricity and to adopt pragmatic approaches. If the United States fails to pay sufficient attention to the energy landscape on which future AI growth rests, it runs the risk of higher prices, greater outages, more public backlash to AI, and insufficient energy to effectively compete with China.

Flashpoints

Electrons as an AI bottleneck

Energy is a potential bottleneck for AI but not a major cost driver. Amortization of hardware and labor account for the bulk of costs for both AI training and inference workloads; energy comprises only 2–6 percent of costs for training. Nevertheless, impressive technical advances in AI, including longer model "thinking," suggest that inference workloads will become more electricity intensive. Whether for training or inference workloads, data centers cannot run without electricity. If US electricity supply-demand balances are disrupted, then AI will suffer.

Can data centers curtail electricity consumption at key moments? Some level of data center curtailment might not be catastrophic, especially for inference workloads. A Duke University study indicated that data centers could curtail inference operations during peak-of-the-peak demand periods, while some companies at the AI-energy nexus are aiming to reposition data centers as "flexible, intelligent grid assets." Some AI models already appear to be headed toward surge pricing, while some consumer-facing inference workloads can tolerate greater latency without severely degrading the product. When training AI models, however, power interruptions could hold severe consequences, especially for frontier models.

AI inference and training workload capacity factors might converge. Due to AI advances, models can solve more complex tasks. In turn, inference workloads can extend across multiple days or weeks. If inference workloads indeed prove more electricity intensive or compute constrained, AI model providers might employ tiered or surge pricing: some are already offering two times capacity outside of certain peak demand windows. Computational resources, memory chips, and EUV machines are key AI bottlenecks. But with total planned US compute capacity in 2030 thirty times larger than mid-2025 levels, electricity constraints could also throttle both model training and consumer inference usage, constraining US AI efforts.

Social license to operate: Electricity affordability and environmental impacts

Increasing demand, constrained supply, and grid replacement needs could lead to rising prices and reliability concerns. Political concerns around AI, data centers, and electricity prices might sharpen if affordability concerns intensify.

Because most social license concerns revolve around affordability, it's worth contextualizing electricity prices alongside other consumer expenditures. Gasoline for vehicles cost the average consumer unit—as defined by the US Bureau of Labor Statistics—about $200 per month in 2024. Electricity costs totaled about the same amount. Food at home ran the average household about $500 per month during the same year. If electricity prices rise, the US AI buildout could face a political backlash. Electricity affordability is moving to the center stage in US politics: the Trump administration unveiled its Ratepayer Protection Pledge, while Pennsylvania Governor Josh Shapiro sued the Pennsylvania–New Jersey–Maryland Interconnection (PJM), the regional transmission organization that oversees transmission system reliability and coordinates electricity markets across thirteen states in the eastern United States—including Virginia, home to Data Center Alley, and Washington, DC—to speed up its interconnection queue. Regulatory and market reform, strong contracts, and effective rate design are essential to ensure that data centers pay their fair share of electricity costs while contributing positively to grid reliability and lowering per-customer system costs.

Electricity costs will also be impacted by natural gas dynamics, especially the interplay between exports and domestic use. US exports of liquefied natural gas are expected to more than double from 2025 levels by 2029, and they will constitute more than 20 percent of all US dry gas production by 2030. Pipeline exports to Mexico are also expected to rise, albeit more gradually. Rising export volumes will lift natural gas prices, all else equal, raising fuel prices for electricity. Ultimately, shifts in natural gas demand have limited effects on the data center buildout, as even a massive 4–6 gigawatt-sized data center will only consume about 1 billion cubic feet per day of natural gas, depending on turbine efficiency, or about 1 percent of total US natural gas consumption. Nonetheless, relying solely or even primarily on natural gas to power AI data centers amplifies exposure to fuel price volatility, especially if energy-inefficient turbines like single-cycle peakers are employed.

Attempting to meet AI power needs via coal might also prove contentious. Coal is politically unpopular: only 27 percent of Americans say the United States should support coal mining,

versus 65 percent and 45 percent, respectively, favoring production of renewables and nuclear energy. Tech companies might have a difficult time navigating the political demands of their clients and user base, which largely oppose coal, and the Trump administration, which favors the technology.

Finally, using coal to power AI poses technical risks. The median age of a US coal plant is forty-five years, while many have delayed maintenance. Ramping up throughput at these facilities could strain hardware and potentially degrade reliability, especially because coal plants are unable to quickly adjust output, a feature that will likely become more important in matching AI inference workloads' unpredictable demands.

Energy infrastructure and a (potential) market correction

AI is already a powerful tool. Still, if its short-term benefits are overhyped, if AI workloads make dramatic gains in energy efficiency, or if interest rates rise due to energy shortages, a market correction or even a bubble could emerge, likely lowering projections of energy demand. Technological changes are uncertain but could upend forecasts in either direction: the reality is energy infrastructure could be overbuilt or underbuilt. In addition, some sort of AI model consolidation, such as a "winner take most" scenario, could lead to a single AI model dominating the global marketplace, with uncertain impacts for energy demand. The United States might continue to provide a large share of global data center capacity, in which case global AI adoption could increase overall US energy supply requirements. Given the inherent uncertainty surrounding AI, policymakers and industry actors must consistently adjust assumptions as developments unfold.

Data centers' water impacts are overstated but socially salient

Larger data centers can each "drink" up to 5 million gallons of water per day, or about 1.8 billion gallons annually—the equivalent of a town with a population of about ten thousand to fifty thousand residents. While this might seem large, the United States uses about 118 billion gallons of water each day for irrigation. Put another way, the average hamburger requires more than five hundred gallons of water to produce. The average ChatGPT search requires somewhere between 0.003 and 0.007 gallons, or less than two tablespoons. In other words, the water needed to produce a single hamburger would power approximately 71,000 ChatGPT searches, conservatively.

Nevertheless, data center water consumption has emerged as a focal point of public interest, and water scarcity can be a real regional concern, especially in the developing world. In certain water-scarce areas of the country such as California or Arizona, AI might compete for limited water resources with other use cases. In Arizona, for example, agriculture and golf courses account for about 72 percent and 2 percent of the state's daily water usage, respectively.

Findings and recommendations

Finding: Too little attention is paid to less glamorous areas of the electricity system, such as transmission. The US electricity grid is aging, and the United States dramatically lags China in building out new transmission. If the United States cannot sustain its current grid, much less build the incremental generation and transmission needed for AI, then its ability to capture the global AI market will be hampered—potentially severely. Deliberate action is needed.

Recommendation:

- **Overhaul aging grid infrastructure and expand the power grid.** Energy companies and public utilities must invest not only in replacing or retrofitting aging transmission lines, distribution lines, and supporting grid infrastructure, but also in building new generation, storage, transmission, and distribution projects. Upgrading and expanding the grid must consider both physical and digital infrastructure upgrades, including advanced transmission technologies, Internet of Things (IoT)-enabled system connectivity and monitoring, and AI-driven grid optimization and control. System efficiency improvements often generate positive returns in the long term and utilities can leverage hyperscalers' access to capital to make these investments, thus improving reliability, mitigating cost increases, and accelerating US AI development.

Finding: An "all of the above" approach to powering AI's electricity needs is critical, but policymakers must also recognize the reality of what is feasible over different time horizons. Speed is critical. Technologies that quickly increase the generation and transmission capacity of the grid will enable rapid development of more powerful AI models. Natural gas will remain integral to meeting growing demand while solar, onshore wind, and battery storage can be deployed cheaply and quickly.

Other fuel sources—nuclear energy, geothermal, hydro, wind, and even coal—will play an important role in maintaining reliability of the existing grid. Nonetheless, in terms of incremental generation, these fuel sources will not be constructed at the scale and speed sufficient to supply AI's incremental power needs due to infrastructural and political constraints. Nuclear energy construction timelines often span more than a decade. Geothermal holds immense promise but has yet to reach commercialization. Hydro is a valuable resource, but its potential has already largely been captured. Onshore wind has yet to be fully tapped; offshore wind, however, faces transmission

constraints, market challenges, and political opposition. The United States hasn't built a coal power plant since 2013 due to poor economic fundamentals, and the most recent facility now faces a nearly two-year outage through March 2027.

Recommendation:

- **Adopt a diversified approach to building electricity generation capacity, implementing both short- and long-term solutions.** For the next five years, US near-term incremental electricity generation capacity growth will overwhelmingly consist of solar, natural gas, onshore wind, and batteries working in tandem with one another.

Finding: Non-energy costs matter: ensuring unconstrained US and allied access to AI accelerator chips holds major energy implications.

Recommendation:

- **Factor energy considerations into decision-making around export controls.** In addition to the national security concerns of sharing advanced chips that strengthen AI capabilities of adversaries, US and allied data centers need rapid access to compute capacity. Epoch AI researchers in 2024 estimated that AI accelerator chips comprised roughly 44 percent of amortized hardware capital expenditures and energy costs. AI accelerator chips likely comprise a larger share of costs now. If domestic AI companies face higher semiconductor prices due to exports of advanced chips to adversary states, hyperscalers will be more capital constrained, thus limiting the buildout of data center capacity and the energy infrastructure that will serve it. Hyperscaler power purchasing agreements and procurement at scale hold promise in scaling new technologies, such as small modular reactors, geothermal, innovative battery chemistry, and transformers. If fewer data centers can be built domestically or in allied nations because of chip sales to adversary states, then it will be more difficult to achieve economies of scale for new energy solutions.

Finding: Maximize the utility of existing generators and grid assets by installing advanced transmission technologies, improving energy management, and investing in energy efficiency measures. Maximizing the capacity of existing infrastructure offers a cost-effective and timely complement to building new transmission lines and energy generation assets. The following actions should be prioritized to optimize grid performance.

Recommendations:

- **Invest in advanced transmission technologies.** Dynamic line ratings, advanced conductors, power flow controllers, topology optimization, and AI- and IoT-driven optimization software can increase the throughput of existing transmission lines, reducing grid congestion and the need for new infrastructure. In addition, advanced transmission technologies are strategically advantageous to the United States as they can be deployed relatively quickly along existing corridors and are not subject to the arduous permitting and study processes that often delay new infrastructure projects.

- **Install high-voltage direct current (HVDC) transmission lines where suitable.** HVDC lines have many advantages over alternating current lines and are particularly useful for long-distance power transmission.

- **Pursue innovation and strengthen efficiency in energy management of AI workloads.** Improving AI workload efficiency, more efficient query processing, storage management, cooling systems, and other innovative tools can support data center operations, lower energy inputs, and help hyperscalers manage their systems with greater precision. Grid managers should also explore options to harness AI software for grid infrastructure. Regulators and policymakers must ensure that the cost of these software solutions can be recovered by utilities through the rate base or through targeted funding or incentive opportunities.

- **Advance efficient use of electricity for AI and other sectoral use cases to help taper energy demand growth.** No common measurement or benchmark exists for AI electricity usage. Additional inquiries by hyperscalers and independent researchers yielding a shared framework for measuring efficiency, with metrics such as queries per watt or energy consumption per inference operation, could serve as a foundation for improving AI workload efficiency that could inform government and industry infrastructure planning and development or the implementation of performance-based incentives. Energy-efficient technologies in residential, commercial, and industrial sectors reduce energy consumption and can unlock power for AI applications without compromising the quality or usefulness of energy services for those users. Demand-response programs can help customers lower their power consumption during periods of peak electricity demand. In turn, reducing consumption at peak periods both lowers users' bills and makes additional power available in the system.

- **Incentivize the installation of energy-efficient technologies in utility rate design, implement supportive policies that enable the building of new infrastructure, and reward the modernization of inefficient legacy systems.** Investments in energy efficiency can yield long-term cost savings, improve grid reliability, and lower power system costs. Utilities can create incentive programs for energy-efficient appliances for consumers and leverage efficiencies in grid infrastructure and operations. Utilities should also conduct system planning to provide

clarity on long-term system needs to which infrastructure developers can respond. Regional transmission organizations and government agencies can complement utility planning frameworks with regional and national plans, respectively. Policymakers can also jump-start efficiency opportunities with targeted policies, such as agency-directed technical assistance, grant programs for critical infrastructure projects, and incentives for both consumer products and utility assets.

Finding: Expediting permitting and engaging communities effectively would accelerate the pace of the energy infrastructure buildout. Approval processes for energy infrastructure projects can last for more than a decade in some cases, and the current patchwork of permitting authorities is creating a system that is insufficient to meet the demands of AI and consumers.

Recommendations:

- **Pass comprehensive permitting reform legislation.** This would set enforceable deadlines for permit decisions, increase transparency on decision timelines, centralize permit coordination in a single agency, enable concurrent review of multiple project permits, and set reasonable limits on judicial review.

- **Engage stakeholders, including the grid operator and communities, early and often in development efforts.** Effective engagement reduces the likelihood of community opposition that can halt or significantly delay project development. To avoid costly litigation and project delays later in development, community benefit agreements and transparent, consistent communication with local communities should become the standard.

Finding: Engaging in interregional coordination and proactive planning of the electricity grid holds major benefits. Interregional transmission lines can have tremendous value in both lowering electricity costs and bolstering reliability. Making proactive decisions about the infrastructure upgrades that will be needed in the future and identifying suitable sites for new projects can unlock opportunities to expedite permitting and interconnection to the power grid.

Recommendations:

- **Identify and prepare suitable corridors for power transmission projects and hubs for data centers and generators by local, state, and federal government entities and grid operators.** These measures reduce the need for bespoke consideration of impacts on a project-by-project basis.

- **Ensure that data centers bear the full cost of infrastructure built solely to serve them.** They should also pay a just and reasonable share of the costs of projects that also benefit other users, consistent with the "beneficiary pays" principle.

Finding: Engaging on electricity costs and contextualizing water consumption relative to other use cases are critical for AI development. AI's electricity challenges are real. Fears about rising energy costs exacerbated by AI is a legitimate consumer concern, and the US AI complex should help address this problem head on by strengthening the grid, expanding transmission, easing the natural gas turbine shortage, and investing in generation sources with zero fuel cost. At the same time, water consumption issues are largely overhyped.

Recommendations:

- **Pursue closed-loop water recycling systems that starkly lower consumption.** Treating and reusing wastewater on-site drastically reduces the need for freshwater intake. This shift also significantly reduces the volume of discharge, cutting both environmental impact and operational utility costs.

- **Engage in more vigorous public education campaigns.** Contextualize water demands relative to other use cases, including agriculture, cryptocurrency, and golf courses.

Chapter 8: Allies and partners

The United States, its allies, and partner countries are at an inflection point on AI. There is a need to build cross-border coalitions, interoperable standards, and shared public institutions that can promote innovation, align governance structures, and elevate democratic values at a moment when decades-old norms and alliances are fraying and there is increased economic competitiveness between countries. Trends of increasing geopolitical competition, interdependence, and technological change are intensifying while traditional mechanisms—and trust—in a rules-based order are eroding as the United States has retrenched in many areas of global engagement and areas of leadership established in the post-World War II era.

The United States remains central to the commercialization and governance of AI worldwide. But a series of domestic and international policymaking shifts now challenges that position, which is based on the country's world-leading research and development, industrial expertise, and technical capacity.

The Trump administration's AI Action Plan includes leading international AI diplomacy and security as a core pillar. The plan sets out priority actions to export the US AI tech stack to global partners, counter Chinese influence in international bodies, strengthen export controls, and coordinate on national security risks. Since the issuance of the plan, the administration's attention within the international context has primarily focused on AI exports. At the 2026 India AI Impact Summit, Michael Kratsios, director of the White House Office of Science and Technology Policy, unveiled the details of the American AI Export Program, including: a national champions initiative that will integrate the leading AI companies of partner nations with the US AI stack to help build domestic tech industries; new financing programs managed through the World Bank, Export-Import Bank, and US International Development Corporation; and a new tech corps within the Peace Corps to provide import partners support in deploying AI for public services.

US officials have also promoted bilateral agreements to strengthen US industries' position in the burgeoning global AI market, including via federal funds. Recent agreements include the AI Opportunity Partnership with India—focused on "pro-innovation regulation" and digital infrastructure—and the US–Japan Framework Agreement—focused, in part, on AI critical materials.

Longtime international partners, particularly those within the EU and other Western countries, have shifted their attention from creating new regulatory and legislative oversight. Non-US politicians and officials have embraced the need to jump-start their countries' sluggish economies via the potential economic benefits and efficiencies associated with AI. In that regard, there is growing alignment between the US and non-US approaches to AI, including through proposals such as the EU's so-called Digital Omnibus, which proposes a delay in the implementation of the bloc's AI Act; the European Commission's dual Apply AI Strategy and AI in Science Strategy; Canada's AI Strategy taskforce; and Singapore's AI National Strategy.

This alignment poses opportunities and challenges for the United States.

A growing consensus that AI advances must be harnessed for economic development provides scope for multinational and bilateral agreements to meet such policymaking objectives. Global AI governance efforts, particularly those related to the United Nations and the Group of Seven (G7), will continue, though a rush to regulate the emerging technology has so far abated. China, however, continues to pursue its own geopolitical interests via international governance structures.

On the other hand, many countries, including close US partners, are seeking to embed principles of "digital sovereignty" into the development of domestic AI industries. That includes efforts to use public procurement to support domestic companies rather than US competitors; restrict data access via localization principles; and impose additional controls that might restrict the ability of US firms to operate in countries from Brazil to Australia to Germany. Here, again, the issue of trust—not of citizens in their government but rather among partner countries—remains paramount to US competitiveness in AI.

Flashpoints

Infrastructure buildout

Governments worldwide are expected to spend, collectively, tens of billions of dollars over the next five years on AI infrastructure, including high-performance computing centers and domestic cloud computing networks to underpin economic growth. This represents both an opportunity and a risk for US industrial players operating domestically and internationally. The current US administration has backed the export of the US AI stack as a mechanism to embed US industrial expertise, values, and norms at the center of how allies develop their indigenous AI industries. Some of these countries, however, have shown greater interest in decoupling themselves from potential reliance on world-leading US AI infrastructure. That includes potentially using government procurement contracts

to preference local competitors to US companies as a means of supporting national enterprise. As of early 2026, this remains more a risk than a reality.

Governance and regulation

While countries are pulling back from specific AI regulation, many are still seeking greater ownership of how these systems are rolled out within their borders. That includes ongoing engagement with international bodies like the OECD, despite the recent pullback from the current US administration. This represents an opportunity for US industry, given many companies' global footprints and long-term relationships with both non-US governments and commercial partners worldwide. Demonstrating strong governance structures is a competitive advantage at a time when non-US policymakers are balancing the need to show citizens, they are in control of the emerging technology and the economic imperative of harnessing AI for growth.

Digital sovereignty trends

This topic is still poorly defined, but its underlying principle, in which non-US policymakers are pursuing more agency over AI-enabling infrastructure and applications, has become more pronounced over the last twelve months. Data residency and sovereignty requirements are among the most important frictions in allied AI collaboration. Allied and partnered nations have imposed strict data localization laws. Data sovereignty trends might cause problems for US companies looking to bid for non-US contracts, particularly those associated with public institutions, as well as making joint AI development and collaboration more complex. US government and industry must balance two somewhat contradictory policy objectives: those from the current administration seeking to support US companies, at home and abroad, as part of its AI Action Plan; and those from non-US governments that are equally looking to jump-start national AI industries. However, significant differences remain among non-US countries, including those in more developed parts of the world with existing AI-focused expertise and those from the so-called Global Majority where AI infrastructure is sparse. In addition, as more non-US companies begin to accelerate AI adoption, industrial use cases that embed proprietary datasets within US infrastructure are raising questions around data governance and data location. Addressing concerns about control and security over often sensitive data will be critical as sectors across the globe navigate the integration of AI into novel use cases.

Chinese influence

Outside the Western world, Beijing is making a play for both economic and political partners via the prism of AI. Its companies and governance efforts represent a rival structure to those offered by the United States and its partners. Beijing has centered its pitch, primarily to Global Majority countries, via the provision of open-source technologies that stand in contrast to proprietary AI from the United States. China has also developed extensive connections within emerging economies through the promotion of likeminded, broader digital governance initiatives. That includes co-authoring submissions to the United Nations' (UN) Global Digital Compact, the stated objective of which includes creating "a comprehensive global framework for digital cooperation and governance of artificial intelligence."

Findings and recommendations

Finding: The advantage of the US AI stack multiplies when considering "allied AI stacks." Equipping allies and partners with solely "American-made" AI stacks is not an effective approach for transatlantic cooperation on AI. Looking at recent geopolitical tensions, non-US countries are reasonably concerned that an AI system that is fully designed and deployed with US-led technologies could increase their dependence on the United States and lead to a loss of agency. The fear of the United States shutting down any AI systems with a "kill switch to force any diplomatic agreement results in Europe's and other partners' growing inclination toward digital sovereignty. This drive has resulted in proposals such as the establishment of "AI Factories", the Digital Omnibus to simplify the EU AI Act and create a single market, and the European Commission's release of the Tech Sovereignty package that includes multiple AI and cloud components. An AI ecosystem consisting of both US and allied stacks is advantageous: it can multiply the United States' comparative advantage, especially in areas that are the current weak spots, such as semiconductor fabrication, open-source AI models, high-bandwidth memory, and critical minerals.

Recommendations:

- **Strategically align with allies and partners on AI export frameworks.** The AI Action Plan directs federal agencies to devise a plan for an "AI global alliance" to ensure export control alignment and policy levers across governments in adopting US-led AI protection systems. The administration should move forward to establish this construct

and set up a global AI export framework centered on encouraging investments in a diverse set of frontier models and applications in the global marketplace among US allies and partners, with appropriate security controls to protect US interests. Importantly, these products should meet safety and security standards developed through an agreed-upon framework co-created by the United States, allies, and partners, and promoted in international standards development organizations. Such standards should be reinforced by provable technical controls to establish a verifiable trust architecture. The export framework should include supporting the allied AI stack as a throughline from the industrial layer up, including energy equipment, semiconductor fabrication, and advanced manufacturing.

- **Leverage the Tech Prosperity Corps for global partnerships.** The White House's announcement of the Tech Prosperity Corps at the India AI Impact Summit was a welcome development. In parallel to the traditional US Peace Corps, the Tech Prosperity Corps is meant to export technical talent to deploy AI applications in public services in the developing world. Although the initiative states potential collaboration will include import partners, the emphasis remains on the US AI stack. The next step is to ensure the allied AI stacks developed through the frameworks recommended above are also fully embedded in the initiative, with substantive input from cross-border industry partners. Sending AI talent from both the United States and allied partners to the developing world can bolster buy-in for a global AI ecosystem that can effectively counter autocratic influence.

- **Create diplomatic positions to champion US and allied AI stacks.** The United States should establish an AI Attaché Initiative, designating attachés at priority US diplomatic missions to advocate for policies that will enable exports of the US and allied AI stacks to the world, increasing the coherence of the AI infrastructural building in the global AI ecosystem through US and allied support.

- **Anchor allied AI cooperation in shared public interest institutions.** The credibility of an allied AI stack depends on more than interoperable hardware and aligned export controls. It depends on shared institutions that allies trust to steward the technology in the public interest, such as research consortia, standards bodies, civil society networks, and cross-border fellowships. Taking such steps to strengthen these components of cooperation moves beyond building coalitions among allied governments to instil public trust across borders.

Finding: Strategic communication is key to successful alliances and partnerships. Communication plays a vital role in successful foreign policy. While the United States has signaled a strong interest in incorporating partner countries' leading AI companies into the American AI Export Program, in practice, allies and partners remain skeptical of US intent, particularly at a time when the US government has increasingly involved itself in territorial disputes and other geopolitical stand-offs.

Recommendations:

- **Reinforce US reliability.** The United States should focus messaging not on US superiority but on US reliability as a partner of choice to provide safer and greater market access than adversaries like China, enabling more capable and robust AI models that are not subject to authoritarian government censorship. It should actively coordinate with allies and partners on contingency plans and provide better assurances to maintain the continuity of their infrastructure and services in response to geopolitical events.

- **Build common purpose to counter Chinese influence.** The United States must also reinforce the message of unity with allies and partners in countering Chinese influence and preserving cooperative and peaceful use of AI. The AI Action Plan highlights this objective but focuses on Chinese influence in international bodies, without elaborating on the United States' vision for these bodies. Just as in 1968, when fifty-nine nuclear and non-nuclear states joined together to sign the Treaty on the Non-Proliferation of Nuclear Weapons under US leadership to prevent a nuclear war, the US government today has a responsibility to articulate a similar unifying purpose. It should reclaim the power of international organizations and, if necessary, prevent the global AI ecosystem from being influenced by autocratic regimes and harmful practices, including preventing non-state actors such as terrorists from abusing AI systems.

Finding: Interoperable data-sharing standards add a technical muscle to powering allied AI coordination. There is currently a lack of standards for secure data governance and coordination among allies and partners. This was highlighted in the recent meeting convened by the NATO Communication and Information Agency. As data are a critical piece in the architecture of AI, it is necessary to create an interoperable ecosystem for data governance, particularly for data used to train AI models. This includes the provision of agreed-upon frameworks and standards for data, which can effectively reduce administrative burden and increase trust and support US allies and partners in key domains that involve AI capabilities, particularly in national security, global health, humanitarian crises, and other areas that require international collaboration.

Recommendation:

- **Create a data governance ecosystem.** The United States should co-create interoperable standards for data classification, evaluation, and sharing protocols with allies and partners. This can reduce friction in allied collabora-

tion while increasing mutual trust. To address allies' and partners' legitimate concerns about stronger control over their data in certain use cases, such an ecosystem should facilitate privacy-preserving information sharing and verifiable technical controls over data access.

Finding: The US leadership in global AI safety and security is a crucial ingredient for AI leadership on the world stage. Without it, adversarial nations will fill the vacuum, seeking to export their own AI frameworks, standards, and technologies, which reflect values that run counter to democratic governance and individual rights. Concerns about AI safety and security risks are preoccupying allies and partners. For example, agentic AI has created new threats that produce harmful content and weaken system security with highly automated decision-making processes that risk undermining human control to a greater degree than previous AI models. Forging international approaches to manage shared AI risks and create updated security principles to keep pace with technological change is essential to ensuring the democratic values remain at the core of the global AI order.

Recommendations:

- **Build on existing AI safety and security frameworks.** The United States should continue constructive conversations with allies and partners and accelerate agreements on existing frameworks that are beyond privacy concerns and extend into high-risk AI use cases, especially around national security use cases that might undermine domestic civil rights protections. The EU AI Act provisions related to high-risk AI systems are set to enter into force in August 2026, and the time is ripe for the United States to engage in these discussions and build consensus with allies early in the implementation phase. Additionally, the Trump administration has created the AI Agents Standards Initiative, under NIST, which seeks to foster industrial-led standards and open protocols to ensure security, interoperability and adoption of agentic AI. The next course of action is to collaborate with allies and partners to expand this initiative at a global level by working with the International Organization for Standardization (ISO) and other international organizations to enhance multisectoral input and the consistency of global AI safety and security standards and a strong US presence in the standard-setting process.
- **Create interoperable information-sharing and security-assurance ecosystems.** Creating a multilateral system between the United States, allies, and partners for real-time information-sharing and shared security-assurance baselines around AI risks, vulnerabilities, and capabilities can help elevate mutual trust, address digital sovereignty concerns without defaulting to hard localization, and meaningfully tackle emerging global AI safety and security risks. These security baselines can include privacy-enhancing mechanisms such as encryption and access controls, auditability, and incident response. Current efforts, including the International Network for Advanced AI Measurement, Evaluation, and Science, provide an excellent starting point. Countries such as the United States and United Kingdom, which regularly carry out advanced AI evaluation exercises, should share the results of these exercises with other countries. The next step is to create shared physical and digital mechanisms through which relevant findings can be accessed and analyzed, while generating proactive responses, all in real time. In addition, these networks should be expanded to include additional countries whose digital ministries are carrying out AI measurement work, and the United States should leverage other alliances such as the Five Eyes for similar information sharing. These information-sharing ecosystems can sustain long-term collaboration to mitigate AI safety and security risks to the global economy, security, and individual rights.

Chapter 9: Onward - measuring US success

Realizing AI's potential to advance the United States' economy, society, and position in the world requires a deliberate, structured measurement architecture. It requires a system that enables qualitative and quantitative assessments to facilitate evidence-based decision-making around how AI is developed, deployed, and governed. Without a consistent, rigorous methodology that tracks the full suite of elements contributing to AI competitiveness, policymakers risk designing ineffective, short-sighted policy interventions that could handicap the US national ability to compete effectively.

The United States needs a national scorecard to measure AI consistently and cumulatively.

Today's AI measurement landscape is not lacking in data; it is fragmented. Metrics are developed and tracked across separate domains: technical benchmarks measure model capability; safety and security metrics track incidents and risks; adoption data capture diffusion across firms and workers; public surveys reflect trust and sentiment; government inventories track usage and compliance; and infrastructure metrics assess energy and compute capacity. Each of these provides valuable insight, but they operate in isolation, often using different definitions, time horizons, and units of analysis. As a result, policymakers can see rapid progress in capability or adoption while missing underlying risks in trust, infrastructure constraints, or governance effectiveness.

This chapter proposes a taxonomy that can effectively contextualize and measure US success through a comprehensive framework. Such a framework would enable a more systemic approach to track adoption, assess societal and economic impacts, and ensure ethical and responsible use. More importantly, it would create a feedback loop for continuous improvement, allowing strategies to evolve in step with technological advancements. By institutionalizing this approach, the US government can optimize outcomes, strengthen public trust, and position the nation competitively on the global stage.

Why measurement is needed

Robust measurement is foundational to an effective national strategy. Without clear, credible metrics, evidence-based policymaking becomes speculation. Establishing rigorous evaluation frameworks enables governments to design and iterate smart policies to deliver real public value while measuring the results of actions taken and return on investment. Using metrics as performance indicators and early warning systems, these benchmarks signal whether the US national approach is succeeding or requires adjustment. For a technology as dynamic as AI, this level of visibility understanding is essential. Transparent, data-driven evaluation strengthens accountability, builds public trust, and enables continuous improvement.

In the age of AI, measurement that keeps pace with technological change can also target pain points and scale solutions more quickly. Just as frontier models improve through a combination of increased scaling of data and compute and regular cycles of expert feedback, the US national strategy for AI should also mirror that approach through a data- and feedback-driven mechanism—such as a national scorecard—that measures key indicators with continuous policy data as a catalytic enabler for tailored policy measures and sustained competitive advantage in the global AI landscape.

Where current measurement falls short

Three weaknesses stand out in the current measurement landscape. The first is that too much time is spent analyzing how AI models score against specific benchmarks rather than how they perform in the real world. AI model leaderboards, benchmarks, and investment figures are important for understanding and comparing model performance against each other, but they do not tell us how a model functions when being used by humans, whether AI is being adopted in sectors that matter the most to national competitiveness, or if public institutions are prepared to oversee those systems. This gap is increasing day by day, with Stanford's 2026 AI Index warning that AI capabilities are advancing faster than the frameworks needed to govern and evaluate them.

The second weakness is less apparent but just as important. The indicators and data needed to assess national competitiveness are scattered across agencies, surveys, and reporting systems that do not speak with each other. The mix of quantitative and qualitative datasets created to meet specific needs within respective agencies—such as compute access for researchers, business adoption, and public attitudes—does not amount to a clear picture of national competitiveness. As it stands, policymakers need to piece together a picture of national competitiveness from a wide range of sources such as Census and federal adoption data, federal agency AI use case inventories, and academic sources such as Berkeley Lab's interconnections queue analysis and Stanford's annual AI Index. That fragmentation makes it harder to analyze how compute, energy, talent, trust, and institutional capacity are reinforcing one another or drifting apart.

The third weakness is that policymakers do not have a clear way to distinguish between early warning signs (leading indicators) and downstream outcomes (lagging indicators). Some indicators can signal stress in a system well before it becomes a strategic problem and the consequences become vi-

sible. For example, a lack of educator training would show a strain on a system that is trying to increase AI education. Access to shared compute and growing delays in the power interconnection queue are other examples. Other indicators emerge later, after the effects have started to set in, as with productivity gains, wider adoption across the economy, or public backlash. A national measurement framework needs both kinds of indicators but also needs to treat them differently.

Introducing a national AI scorecard

The critical shortcomings in the current state of measurement make it clear that the United States needs a national scorecard to measure direction, systemic signals, and overall success in the age of AI. Policymakers are often searching for a picture of where the United States stands vis-à-vis China or for the impact of a specific set of policies, but the answers they receive end up being flawed because they are based on only one metric or a small handful of metrics. A national scorecard would integrate essential indicators to paint a fuller picture of US competitiveness in AI.

The indicators recommended in this chapter are extracted from the commission's insights based on its six areas of focus. Theyare not meant to be exhaustive. Rather, they provide a view of the direction the United States is moving in areas critical to maintaining global leadership on AI. The scorecard concentrates a core set of indicators in one place to systematically assess the value and progressive development of the United States' current AI path. The national scorecard could be extrapolated with similar measures among allies and partners to build an effective diffusion strategy, measuring the allied AI stack to scale US competitiveness globally.

The goal of creating a set of baseline indicators is to enable assessment and prioritization of strategic measures rather than isolated benchmarks. These measures reflect the commission's recommendations, focusing on enabling evaluation of whether the United States is making positive progress in terms of strengthening public trust, managing AI as a collective ecosystem, and effectively integrating across that ecosystem.

Public trust

In democratic societies, the success of AI policy hinges on trust at two interconnected levels: confidence in the technical performance of AI systems and confidence in the institutions—especially governments—responsible for deploying them. Reliable, accurate, and safe systems form the foundation of that trust—if systems fail to perform as intended, no governance framework can compensate. Simultaneously, public trust is shaped by how transparently, ethically, and responsibly the systems are implemented and how members of society can participate and strengthen the system. Trust and optimism in AI are not abstract ideals but measurable indicators of both system integrity and trust in government. Policymakers must lead with metrics that underpin the trustworthiness of an AI system, capturing overlapping areas of AI infrastructure, education, governance, and global leadership to boost technological reliability and citizen proficiency. Without this cross-cutting lens, even the most advanced AI capabilities risk limited adoption, erosion of public trust, and, ultimately, a failure to translate innovation into sustained national advantage.

Collective ecosystem

AI competitiveness is often framed through model performance and algorithmic breakthroughs, but the reality is far more multifaceted and grounded in strengthening the entire AI ecosystem shaped by multiple variables. Resilient, scalable, and efficient infrastructure, human capital, and effective government use are all crucial factors to AI deployment at a meaningful scale. From data centers and semiconductor supply chains to talent cultivation and government expertise, improvements across sectors and industries are the keys that enable AI to move from concept to scaled capability. There are metrics that directly correlate to model competitiveness, governance resilience, and deployment scalability; there are also crucial issues that policymakers must navigate to maximize the robustness of the AI ecosystem and manage negative social impact. Having robust measurements on hand across these areas, as outlined below, will enable management of AI as an interconnected ecosystem.

Effective integration

Measuring AI competitiveness requires thinking beyond sharpening the edge of innovation in a narrow technical community. Non-commercial or technical sectors—including the government, academia, civil society groups, and individual citizens—must be able to directly experience the benefits that AI brings. This consideration requires measuring integration of AI across industries and sectors of the national economy. Tracking effective integration of AI ensures national competitiveness is not limited to the technical level but is instead expressed throughout different levels of society, with individual citizens' balanced capabilities to apply AI to everyday life, contribute to responsible AI development, and elevate the nation's innovation ecosystem via the strengths of a diverse citizenry.

Baseline indicators

Organized along each of the critical areas for AI competitiveness studied by the commission, the baseline indicators outlined below will be essential for assessing the current state and relative strength of US competitiveness and enabling a strategic, national approach. This framework is by no means exhaustive but is intended to sketch out an integrated baseline that could be used to measure AI competitiveness moving forward.

Public trust

In democratic systems, public trust directly affects adoption, political sustainability, workforce acceptance, and ultimately national competitiveness. A technically advanced AI ecosystem that lacks public confidence will struggle to scale effectively, particularly in high-impact domains such as healthcare, education, finance, law enforcement, and public services. As a result, policymakers should incorporate trust metrics that measure both confidence in the AI systems themselves and confidence in the institutions deploying them.

Table 1: Public trust metrics

Competitiveness dimension	Metrics
Public perception of AI's impact	• Public opinion on whether AI will have a net positive impact on essential sectors, including education, employment, healthcare, and banking • Public opinion on whether AI will have a net positive impact on personal life, such as entertainment, news consumption, wellness, social life, and task completion • Percentage of the population expressing confidence in AI-assisted decisions across sectors (e.g., healthcare, education, and finance)
Public optimism about AI	• Public opinion on whether AI is providing, or will provide, positive impact to individuals and to the country in the next five years • Public opinion on whether AI will have a positive impact in the world
Public anger or skepticism about AI	• Public opinion on whether AI is being governed effectively • Public opinion on whether AI is negatively impacting individual well-being • Public opinion on risks related to AI (e.g., safety concerns, job loss, and energy prices)
Trust in the government's use of AI	• Percentage of respondents who believe government AI use is transparent and accountable • Percentage of respondents who support different government AI use cases (e.g., to allocate public benefits, assist in writing reports, and in defense) • Percentage gaps between AI use cases deployed by the government, the private sector, and through public-private partnerships
AI inclusion and equity perception	• Trust disparities across demographic groups • Perceived fairness of AI outcomes • AI adoption rates across socioeconomic groups • Public perception of whether AI benefits are broadly distributed
AI harm and incident reporting	• Number of reported AI-related harms and incidents (including by media) • Level of bias and discrimination in evaluation of AI models • Percentage of systems independently evaluated
Rates of adoption	• Percentage of AI adoption by sector • Percentage of AI adoption for individual purposes (e.g., entertainment, news consumption, wellness, and task completion)

Innovation

Maintaining a US advantage in AI development and employment requires sustained action across funding and policy mechanisms to capture economic value at scale. Particularly because these efforts compete with other national priorities, policymakers need clear measures built around several economic and technical dimensions to capture results in actionable ways.

Table 2: Innovation metrics

Competitiveness dimension	Metrics
Federal funding effectiveness	• Volume of federal funding supporting AI-related research and development • Volume of patents from federally supported efforts • Volume of papers accepted in top journals from federally supported efforts • Follow-on private investment rates linked to federally funded R&D
Academic sector vitality	• Position of US universities among global rankings in AI and computer science • Share of top-cited AI scholars from US institutions • Share of top-cited AI publications from US institutions
Model performance	• Performance of US models on key general capabilities benchmarks • Performance on sector- and function-specific model benchmarks
Economic competitiveness	• Percentage of leading AI labs based in the United States • Level of private investment in the AI sector, as compared internationally • Start-up formation and survival rates • Adoption rates in key economic sectors • Global adoption rates of US versus foreign closed and open models • Revenue growth of US AI firms in key sectors

Talent and institutional readiness

Cultivating an AI-ready workforce presents a fundamental policy challenge, as the deliberate pace of sound labor force management often lags the speed of technological change. To meet this challenge, a measurement framework should be designed to establish a baseline view of whether the United States is developing domestic talent, retaining high-value international talent, and preparing the federal workforce for AI adoption.

Table 3: Talent and institutional readiness metrics

Competitiveness dimension	Metrics
Domestic training and education	• K–12 math and computer science proficiency benchmarks • AI, CS, and statistics course availability in US high schools • Instructor AI credentialing rates in US high schools • Percentage of students completing an AI or data literacy module before high school graduation • Number of US adults completing federally supported AI upskilling programs (e.g., apprenticeships, community colleges, and Career and Technical Education (CTE) certificates)
Attracting global talent	• International student enrollment in US AI-related graduate programs, particularly at the PhD level • Percentage of top foreign-born AI students who study in the United States • Citizenship conversion rates from visas of foreign-born AI professionals • Number of H-1B, O1, and other visas given to foreign-born AI professionals • Number of new AI start-ups founded by immigrants, including unicorns
Federal workforce	• Number of new federal AI positions created • Federal AI position vacancy and retention rates • Number of federal workers completing federally supported AI upskilling programs across all agencies • Number of federal workers with access to state of the art AI systems
Workforce displacement	• Number of workers per sector displaced by AI • Number of workers per sector who shift roles due to the integration of AI

Governance

Effective AI governance requires more than policy articulation. It demands measurable capacity to adopt, oversee, and align AI systems at speed and scale. A measurement framework should assess whether governments are not only setting rules but also building institutional capability to implement them. One of the clearest indicators of this capability is the agility of government AI adoption, particularly through procurement and deployment timelines and workforce capacity.

Table 4: Governance metrics

Competitiveness dimension	Metrics
Agility of government AI adoption	• Average time from need identification to deployment authorization for AI projects • Average time from technology availability to federal agency deployment • Comparison of federal and private-sector deployment times • Amount of time required to procure a new AI system
Institutional readiness and policy maturity	• Number and percentage of agencies with published AI policies, "appropriate use" frameworks, and data governance policies • Frequency of policy updates (e.g. annual or biannual)
Federal AI workforce capacity	• Number of technology-focused positions filled • Time to hire for technical positions
Investment in AI assurance and evaluation	• Federal budget for AI safety testing and evaluation software and hardware • Federal funding for research on AI testing and evaluation

Supply chain

Key metrics on AI supply chain span all inputs to the AI ecosystem, including compute capacity, hardware and software supply chains, and network connectivity, providing a comprehensive view of national readiness for AI expansion.

Table 5: Supply chain metrics

Competitiveness dimension	Metrics
Compute and accessibility	• Total FLOPs (floating-point operations) available within the United States • Access to computing resources across public, private, nonprofit, and academic sectors • Amount of compute made accessible for researchers at subsidized rates by government or industry
Hardware supply chain	• Share of advanced chip production in the United States and allied nations • Share of AI-relevant semiconductor production in the United States • US and allied share of assembly, packaging, lithography, etch, deposition, and critical minerals processing • Number of active supply chain agreements • Investments in new domestic chip production capability
Software supply chain	• Number of key open-source repositories underlying AI and other software scanned for novel vulnerabilities • Vulnerability response time (e.g., common vulnerabilities and exposures (CVE) enrichment timelines)
Network and connectivity	• Broadband penetration • Fifth- and sixth-generation (5G/6G) coverage and amount of spectrum available for each • Average upload and download speed on mobile devices per metro area • Rural-urban connectivity gap

The table below illustrates that robust national AI efforts must extend beyond the AI models themselves. It requires a holistic approach that includes securing the supply of advanced hardware, modernizing network infrastructure, and ensuring access to computing resources.

Energy

Continued progress in AI development, innovation, and integration comes with significant energy demands and implications. Metrics such as total power capacity , geographic distribution, and the rate of infrastructure buildout provide a direct indication of whether a nation can scale AI systems in line with growing demand. High capacity and rapid, geographically distributed expansion signal the ability to deploy AI efficiently while minimizing regional strain and latency. Conversely, limited capacity, slow buildout, or overconcentration in specific regions can create deployment bottlenecks, constrain innovation, and increase vulnerability to localized infrastructure disruptions.

Table 6: Energy metrics

Competitiveness dimension	Metrics
Data centers	• Total power capacity (megawatts) • Geographic distribution • Buildout rate • Data center efficiency (measured in performance per dollar or trillion operations per second (TOPS) per watt)
Energy availability	• Electricity cost, normalized for incomes and gross domestic product (GDP) • Grid reliability • Energy generation capacity
Grid and transition	• Transmission line buildout • Interconnection backlogs • Grid congestion

Allies and partners

US AI competitiveness also depends on whether the United States can translate domestic strengths into durable influence across allied and partner ecosystems. In this domain, measurements should track cooperation in academia and industry across allies, international standards, and international collaboration on AI safety and security.

Table 7: Allies and partners metrics

Competitiveness dimension	Metrics
Diplomatic partnerships	• Number of US-led (or involved) international agreements and frameworks focused on AI • Number of US-led (or involved) alliances or international bodies working on AI • Number of countries and breadth of activities in US-led AI alliances (e.g., the International Network for Advanced AI Measurement, Evaluation, and Science) • US participation in global AI initiatives (e.g., via the UN and OECD)
Economic and academic partnerships	• Number of research papers in top AI and CS venues co-produced by US and allied experts • Number of allied or partner companies in US AI ecosystems • Number of US companies in allied or partner AI ecosystems
Safety and security collaboration	• Level of support for international agreements and frameworks on AI safety and security with US involvement • Participation in global standards bodies (e.g., the number of US-led AI proposals advancing in ISO) • Participation in global AI safety initiatives (e.g., the International AI Safety Report and joint AI testing exercises between countries) • Number and cadence of meetings the United States and allies conduct

Conclusion

A strong US AI strategy will be judged not only by the sophistication of its models but by whether it can strategically implement effective policies with a cross-cutting lens. A durable national advantage requires moving beyond aspirational goals and isolated benchmarks toward a disciplined system of metrics to expose gaps across systems, support focused strategies, and enable faster policy adjustments. By institutionalizing this measurement architecture, the United States can turn ambition into sustained technological leadership.

Chapter 10: Onward - navigating the geopolitics of AI

US strength scales with allies and partners.

This fundamental notion extends into an immediate future in which the existing global architecture is fraying, geopolitical competition is increasing, interdependence and hyperconnectivity persist, and the rate of technological change is compounding. AI is squarely at the intersection of each trend. The future is contested and no outcome is guaranteed. Nations that lead in AI will not only grow faster economically but will govern differently, fight differently, and project influence in ways that are difficult to predict.

The United States is already navigating these developments, seeking to maintain global leadership in AI while countering the growing influence of China. However, the complex nature of an AI ecosystem influenced by the six overlapping pillars (and others) investigated by the commission shapes a nation's technological advantage, defying the traditional definition of power from a purely economic or national security lens.

Strengthening AI competitiveness in this emerging landscape demands a new approach to international cooperation that is designed for a new global paradigm.

This effort will be measured in generations; however, immediate action is necessary given the acceleration of global trends and AI development. It lends itself to another effort at the scale and ambition of this commission's work on US competitiveness, thus this chapter is future facing—introducing key themes and hard questions that need to be investigated to evolve international partnerships to harness the opportunities and manage the risks of AI.

Alliances as key ingredients to US competitiveness in AI

Traditional alliances enable the United States to maintain military bases or personnel in allied countries and deploy forces in response to conflicts as well as in support of humanitarian aid operations. As AI is an increasingly central component underpinning both economic cooperation and security architecture, it brings into question whether new modes of power projection should be prioritized.

AI will also require new standards and approaches for interoperability among partner countries, including in areas such as AI safety and security, data and information sharing, and other areas noted in Chapter 8. As AI is integrated into legacy systems and more AI-based tools and systems emerge, new capabilities, concepts of operation, data requirements, and infrastructure investments will be required. This requires consistent US engagement and leadership in standards development organizations matched with active participation by allies and partners.

Building a practical consensus on AI alliances

The United States has an opportunity abroad—and a responsibility at home—to build a practical consensus and criteria for international partnership organized around AI competitiveness.

Since ChatGPT was released in November 2022, the Biden and Trump administrations pursued various policies designed to align more international partners with the US AI strategy. While the approach has varied, both administrations focused policy around AI as a national security priority, lessening critical dependencies, shaping global governance, and countering China's advancements in AI technologies and its ability to project power beyond its borders through technology. From the Biden administration's Diffusion Framework to the Trump administration's AI Action Plan, AI export program, and Pax Silica policies, there is consensus and some degree of continuity emerging to shape US global AI leadership. This can be built upon.

Rethinking the US approach to allies in the age of AI depends on consensus that partnership scales strength, defines critical needs, identifies mutual interests with international partners, creates criteria for alignment, and articulates the stakes of systemic competition.

On defining the United States' critical needs, AI leads to heightened priorities around building more resilient technology supply chains, particularly in areas in which the United States has less access but can lean on the strength of allies. The White House has recognized this, acknowledging that partnering with other countries on the AI supply chain will be essential. The US State Department's Pax Silica initiative seeks to forge stronger partnerships across the AI supply chain, from critical mineral refining and processing, to network infrastructure and semiconductor fabrication, to software and frontier models. The supply chain demands reveal the strategic importance of diverse geographies, raising the question of whether supply chain contributions hold the potential to substitute for traditional military commitments in bilateral and multilateral alliance structures. This also points to the need to consider new criteria when identifying and prioritizing partners.

While interdependencies can create critical vulnerabilities for the United States, they can also bind alliances and drive global leadership. Balancing both remains an imperative.

The concept of AI diffusion—strategically managing and monitoring AI adoption among countries and creating incentives for advancement with the United States—is useful in creating criteria for alignment. A necessary first step is distinguishing between core allies and partners, emerging partners, unaligned or neutral countries, and adversaries. Two fundamental areas for growth in AI diffusion are clear standards for advancement among international partners, including growing from an emerging partner to a core ally, and avoiding missteps in partner management, such as excluding existing allies or significant economic partners.

In every element of the commission's work, managing AI as a collective ecosystem emerged as a continuous priority. The approach to international partners is no different; it requires a holistic approach.

Competitiveness at home and abroad: the allied tech stack

Beyond building domestic capabilities, US AI competitiveness should also include establishing allied strength to maintain US global leadership and a democratic world order.

The United States has traditionally led, and is still leading, global AI innovation. However, the geopolitical landscape of AI competition is rapidly shifting. Recent international perception of US reliability has declined due to increased concerns over digital sovereignty, rising national ambition of AI capability around the world, the challenge from China to influence the global AI ecosystem, and growing dialogues about AI safety and security risks. An effective approach to national AI competitiveness should take a long-horizon view and prioritize aligning allies' and the United States' technological advantage, particularly expanding allies' strengths on which the United States can lean in areas where it lacks critical and immediate access.

The national scorecard for AI could be further extrapolated with similar measures among allies and partners to effectively scope the allied tech stack, measure interdependencies and interoperability, and identify capability gaps among partners. Meaningful indicators could measure interoperable infrastructure, international policy agreements and frameworks, robust national AI ecosystems that collaborate with the US commercial and innovation ecosystems, and resilient supply chains excluding authoritarian, adversarial, or non-rights aligned technologies.

The United States will increasingly need to account for variance in AI capabilities among partners and allies. However, the United States also has an opportunity to serve as a partner or "AI accelerator" among aligned countries in connectivity, capability, and governance. With the plurality of the world's computing resources, the United States is in a strong position to supply computational power to allies and partners. As with the other national approaches recommended in this report, an effective diffusion strategy requires accurate metrics to scale US competitiveness.

Meeting middle powers where they are

Working with middle powers—in part because of their economic, systemic geographic, and cultural scale—is a critical component of the United States' continued AI leadership and global leadership.

The United States and China remain in an escalatory pattern of systemic competition between open market democracy and an authoritarian model of state-led capitalism, and China seeks to ultimately remake the existing international system in its own image. Many other countries are attempting their own path on AI through tech sovereignty, determined to avoid being caught in bipolar competition among hegemons or dependency on a handful of multinational companies based beyond their borders.

The fast-evolving geopolitical landscape of AI demands deeper engagement with advanced middle powers, many of which are formal allies like Canada and the United Kingdom, matched with significant attention to emerging middle powers across Latin America, the Middle East, Central and Southeast Asia, and Africa. Many have actively developed national AI agendas or formed regional frameworks outside direct US or Chinese influence. They are ambitiously carving out their own paths toward digital and AI sovereignty, sometimes with regional collaboration, to avoid becoming overly dependent on powerful actors. For example, the United Arab Emirates has carefully balanced trade relationships with the United States and China to secure advantages in AI development and maintain its strategic position for global AI investors, while the Association of Southeast Asian Countries (ASEAN) has also recently announced the establishment of an ASEAN AI Safety Network to strengthen the region's role in global conversations on AI safety and security.

Taken collectively, middle powers represent a large—and growing—share of the global economy. In an increasingly multipolar world, these countries and regional groupings have greater leverage to shape the direction of global AI development, governance, and investment.

This trend from middle powers carries a clear strategic implication: the United States must build a broader and more flexible AI coalition to multiply its advantages, while demonstrating thoughtful leadership grounded in a sharp understanding of its allies' and partners' national interests.

Managing systemic competition with China

Evolving the US approach to alliances will be crucial to ensuring competitiveness in AI and the broader digital ecosystem, particularly China has advanced a model that centers its strategic interests abroad.

China has positioned engagement with Global South countries through supporting their digital and AI infrastructure development priorities with a Chinese tech stack bundled as part of the Digital Silk Road initiative. It has framed this initiative as capacity building and sovereignty consolidation in multilateral forums and bilateral dealings with a group of seventy-seven countries. However, these systems incorporate centralized architectures that reflect domestic governance priorities, embedding Chinese governance models into how partner countries manage data and digital systems over time. Shared China-hosted infrastructure and standards can create forms of interdependence that carry significant economic and strategic implications for the United States. Indeed, this strategy is central to China's goal of creating an alternative global order in the Global South to challenge US technological leadership in the global system. China has termed its preferred global order for AI the Global AI Governance Initiative and has made open-weight AI models the core of its Global AI Governance Action Plan.

The US response in the global AI ecosystem starts with reclaiming the narrative and proactively engaging with countries by leveraging its own technological and ideological advantage. Incorporating the US AI stack into allied AI stacks as recommended earlier in this report is only the first step toward a broader rethink of alliance restructuring. The bigger aim is to enable the United States and allies to be ready to develop and deploy AI in a coordinated way when responding to economic and national security needs. Cooperation must be further legitimized through global forums with the goal of advancing free, open, transparent, and interoperable technology ecosystems globally. In advancing this approach, the United States also advances the principles of freedom by not systemically censoring or biasing AI models. This requires engaging consistently in international standards development organizations and working from a shared agenda and set of priorities.

Rather than retreating from the United Nations and other multilateral platforms, the next step is thinking through how to evolve and renew the alliance structure to address the realities of today's geopolitical environment driven by technological change. By focusing on delivery, credibility, and alignment with partner needs, the United States can reinforce a model of digital governance that is responsive to a diverse and evolving global landscape while ensuring the Chinese model of digital dependence and state-centric systems does not come to dominate the digital ecosystems of the majority of the global population.

Future questions to explore

1. Interdependence: How can nations balance national AI ambitions with collective security needs?

High geopolitical instability at the current moment poses challenges to forming effective and coherent alliances. Nations are focusing inward and structuring their AI development agendas around meeting their citizens' needs, building sovereign tech capabilities, and competing for global AI leadership. The result is an increased polarity of AI powers and a fragmented world order. Disagreements are more likely to escalate, vulnerable countries are more likely to be harmed by authoritarian technologies, and AI safety and security risks might continue to rise without oversight. How should countries balance national ambitions and collective solutions to these global crises,especially given the widespread technical diffusion of AI and associated safety concerns? And how should the United States lead in maintaining this balance without compromising its own competitiveness?

2. Private-sector partners: What is the optimal role of industry in strategic alliances?

Private industry is a significant and dynamic asset for the United States given its open market and democratic system. Compared to other eras of technological change in which government and academia led innovation and scientific breakthroughs, private industry is playing a leading role in investment, innovation, and deployment in the age of AI. Companies are shaping the governance of AI in major ways, from safety testing to international standards to open technical protocols. The geopolitical significance of AI also means that the tech industry bears a larger role beyond developing new models or providing capital. The rapid deployment of AI to gain operational advantage gives industry unprecedented leverage in determining the stakes of national and global security. Hyperscalers and other large AI companies built and headquartered in the United States are increasingly engaged in geopolitics and have multinational equities beyond the narrow parameters of national competitiveness. What responsibilities—and opportunities—do US companies with equities beyond borders have in aligning with national priorities and serving as an extension and product of US global leadership? And how can multinational companies enable an emerging allied stack ecosystem with the technology that underpins alignment?

3. Interoperability: How should United States and its allies collaborate on collective data security and data architecture?

Access to data has long been an important advantage in advancing AI, but a safe and secure architecture in which data are collected, managed, and analyzed is becoming an even bigger key to national and economic security. As critical data flow across borders and systems through increased AI deployment, there is a need to collaborate with allies on collective security of data to prevent cyber espionage, data breach, or data poisoning by adversaries. The United States needs to build ways to work with allies on designing safe and secure data architectures. What would these architectures look like, and what are the mechanisms needed to allow for allied and partnered data collaboration amid the world's increased focus on data sovereignty even among allies?

4. Scaling security: What should mutual defense mechanisms look like in an AI-driven world?

While AI deployment rapidly enhances the precision and efficiency of military operations, its limitations need to be addressed with allies. This includes high-quality data gathering, careful human-machine interactions to avoid miscalculation, and the prevention of accidental harm to civilians or critical infrastructure. These risks can pose new questions to global collaboration. When automated systems, chatbots, and agents make life-and-death decisions and even communicate with each other at a much faster rate than political consultation allows, does that change the calculus of mutual defense commitments? How could the United States and allies update agreements to account for an increasingly automated environment?

5. Capability gaps: How do technological asymmetries affect alliance cohesion?

A successful alliance requires a certain baseline of technological competency among allies. However, there is an AI capability gap between the Global North and Global South, as well as among countries in both groups. While high-income countries are in a generally good position to benefit from AI, low-income countries in parts of Asia, Latin America, and Africa still struggle in terms of access to electricity and digital infrastructure. Meanwhile, Europe faces uneven development of AI across its member states, which weakens the region's AI capability. The concentration of talent, capital, hardware, software, and R&D output remains in only a handful of countries and companies. How might such technological asymmetries between allies and partners create new forms of dependence or tension? And how might that challenge cohesion for competitiveness?

Building the coalitions needed to sustain US AI leadership for a safer world

The rapid advancement of AI will continue to complicate and reshape geopolitics. Without a clear strategy that reflects the technical, economic, and geographic realities of the AI ecosystem, existing alliances and partnerships risk becoming easily dissolved, which would lead to further fragmentation of the world order. As authoritarian powers are advancing their global ambitions with advanced technology, the United States must think more creatively about how to structure existing and new alliances around technology. Sustaining US leadership will require a coalition strategy that strengthens competitiveness, advances trusted technology ecosystems, and supports a freer, more secure, and open future.

Conclusion

Artificial intelligence is reshaping the foundations of economic growth, national security, public trust, and global influence. The United States has significant advantages, but those advantages will endure only if they are matched with strategy, investment, institutional readiness, and close alignment with allies and partners. This report offers a framework for action across the systems that will determine AI leadership: innovation, talent, governance, supply chains, energy, measurement, and international cooperation.

The work of the Commission is a starting point, not an end point. Sustaining US leadership in AI will require continued engagement among government, industry, academia, civil society, and allies. The choices made now will shape whether AI strengthens open societies, expands opportunity, and supports a more secure and prosperous future.

Atlantic Council Board of Directors

CHAIRMAN
*John F.W. Rogers

EXECUTIVE CHAIRMAN EMERITUS
*James L. Jones

PRESIDENT AND CEO
*Frederick Kempe

EXECUTIVE VICE CHAIRS
*Adrienne Arsht
*Stephen J. Hadley

VICE CHAIRS
*Robert J. Abernethy
*Alexander V. Mirtchev

TREASURER
*George Lund

DIRECTORS
Stephen Achilles
Elliot Ackerman
*Gina F. Adams
Timothy D. Adams
*Michael Andersson
Alain Bejjani
Colleen Bell
Sarah E. Beshar
*Karan Bhatia
Stephen Biegun
Linden P. Blue
Brad Bondi
John Bonsell
Philip M. Breedlove
David L. Caplan
Samantha A. Carl-Yoder
*Teresa Carlson
*James E. Cartwright
John E. Chapoton
Ahmed Charai
Melanie Chen
Michael Chertoff
George Chopivsky
Wesley K. Clark
*Helima Croft
Ankit N. Desai
*Lawrence Di Rita
*Paula J. Dobriansky
Joseph F. Dunford, Jr.
Richard Edelman
Stuart E. Eizenstat
Tara Engel
Mark T. Esper
Christopher W.K. Fetzer
*Michael Fisch
Alan H. Fleischmann
Jendayi E. Frazer
*Meg Gentle
Thomas H. Glocer
John B. Goodman
Sherri W. Goodman
Marcel Grisnigt
Jarosław Grzesiak
Murathan Günal
Michael V. Hayden
Robin Hayes
Tim Holt
*Karl V. Hopkins
Kay Bailey Hutchison
Ian Ihnatowycz
Deborah Lee James
*Joia M. Johnson
*Safi Kalo
Karen Karniol-Tambour
*Andre Kelleners
John E. Klein
Ratko Knežević
C. Jeffrey Knittel
Joseph Konzelmann
Keith J. Krach
Franklin D. Kramer
Laura Lane
Almar Latour
Yann Le Pallec
Diane Leopold
Jan M. Lodal
Douglas Lute
Jane Holl Lute
William J. Lynn
Mark Machin
Marco Margheri
Michael Margolis
Chris Marlin
William Marron
Roger R. Martella Jr.
Judith A. Miller
Dariusz Mioduski
*Richard Morningstar
Georgette Mosbacher
Majida Mourad
Mary Claire Murphy
Julia Nesheiwat
Edward J. Newberry
Franco Nuschese
Joseph S. Nye
*Ahmet M. Ören
Ana I. Palacio
*Kostas Pantazopoulos
David H. Petraeus
Elizabeth Frost Pierson
*Lisa Pollina
Daniel B. Poneman
Robert Portman
*Dina H. Powell McCormick
Michael Punke
Ashraf Qazi
Laura J. Richardson
Thomas J. Ridge
Gary Rieschel
Charles O. Rossotti
Harry Sachinis
C. Michael Scaparrotti
Ivan A. Schlager
Rajiv Shah
Wendy R. Sherman
Gregg Sherrill
Jeff Shockey
Kris Singh
Varun Sivaram
Walter Slocombe
Christopher Smith
Clifford M. Sobel
Michael S. Steele
Richard J.A. Steele
Mary Streett
Nader Tavakoli
*Gil Tenzer
*Frances F. Townsend
Melanne Verveer
Tyson Voelkel
Kemba Walden
Michael F. Walsh
*Peter Weinberg
Ronald Weiser
*Al Williams
Ben Wilson
Maciej Witucki
Neal S. Wolin
Tod D. Wolters
*Jenny Wood
Alan Yang
Guang Yang
Mary C. Yates
Dov S. Zakheim

HONORARY DIRECTORS
James A. Baker, III
Robert M. Gates
James N. Mattis
Michael G. Mullen
Leon E. Panetta
William J. Perry
Condoleezza Rice
Horst Teltschik
William H. Webster

**Executive Committee Members*

List as of March 24, 2025

The Atlantic Council is a nonpartisan organization that promotes constructive US leadership and engagement in international affairs based on the central role of the Atlantic community in meeting today's global challenges.

© 2026 The Atlantic Council of the United States. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Atlantic Council, except in the case of brief quotations in news articles, critical articles, or reviews. Please direct inquiries to:

Atlantic Council
1400 L Street NW, 11th Floor
Washington, DC 20005