



Cohere Response to the Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

[FR Doc. 2025-02305] (90 FR 9088)

March 15, 2025

Introduction

Cohere appreciates the opportunity to provide input on the National Science Foundation and White House Office of Science and Technology Policy's (OSTP) Request for Information (RFI) on the National AI Action Plan. As a developer of state-of-the-art AI foundation models purpose-built for enterprise use, Cohere brings a unique perspective grounded in practical experience deploying secure AI systems across multiple industries.

We strongly support the Administration's efforts to enhance U.S. leadership in AI. An effective AI Action Plan must:

- **Make AI adoption a centerpiece:** Promote policies that accelerate government and private sector adoption of AI.
- **Reform and modernize government procurement:** Place technologies from innovative startups in the hands of the government to make a difference in the lives of everyday Americans.
- **Prioritize regulatory flexibility:** Pursue risk-based governance and a sectoral approach to ensure security and enable innovation.
- **Invest in AI R&D Infrastructure:** Make the government an enabler of technology and innovation through increased access to public data and compute, and talent programs.

This submission outlines Cohere's core policy priorities and recommendations to foster a competitive, innovation-friendly AI ecosystem in the United States and allied nations. Embracing policies that unleash innovation while managing risks is necessary. We stand ready to collaborate across government departments to realize this vision, underscoring the necessity of a comprehensive national approach.

About Cohere



Founded in 2019, Cohere builds enterprise-grade AI models and end-to-end products addressing real-world business challenges, prioritizing privacy, security, multilingual support, and verifiability. Our CEO and co-founder, Aidan Gomez, contributed significantly to the invention of Transformer architecture, foundational to today's widely used large language models (LLMs) that power generative AI.

As one of the fastest-growing **independent** startups globally, Cohere has a diversified set of investors, and is hardware and cloud agnostic. Our independence allows flexible deployment of our models across technologies, cloud platforms, and even private, air-gapped environments, making our solutions highly accessible. We partner with allied government agencies and leading global companies such as Oracle, RBC, and Fujitsu, focusing on seamless integration, deep customization, and accessible solutions for businesses to deliver immediate practical value. This adoption-focused mindset – honing in on solutions that organizations can realistically integrate, test, and scale – underscores our practical and applied approach to AI. The recent launch of our secure AI agents platform [North](#) is an example of this work.

A Practical Overview of the State of AI

Today's AI models, particularly LLMs like Cohere's [Command A](#), demonstrate transformative capabilities for society as a whole. These capabilities will continue to grow, become more accurate, and easier to use. However, it is important to recognize and prioritize policy through which many of these improvements have stemmed.

Cohere's CEO and co-founder Aidan Gomez [penned an open letter](#) in December, highlighting that in the past 18 months, AI breakthroughs have originated not from raw compute, but rather from advances in data quality, synthetic data, and novel training techniques like reinforcement learning. Recent model releases by companies worldwide, notably by DeepSeek in China, have affirmed that these methods offer clear opportunities to enhance model performance and cost efficiency. In fact, Cohere already employs many of the techniques recently showcased to improve efficiency and performance. These developments confirm that **solely focusing on scaling compute will not ensure dominance in future AI advancements.**

This highlights an important reality: model performance on general use cases and benchmarks are converging across model developers. In other words, the performance of the most advanced versions of LLMs are trending about the same. This type of homogeneity, while powerful, is inconsistent with the needs of most everyday



businesses. General academic benchmarks don't showcase how models will tackle real-world tasks on an assembly line or analyzing financial research at a local bank.

More and more, businesses and governments require that models – and how we assess their performance – be customized to their **specific needs over general-purpose solutions**. Focusing on practical value with tailored solutions – not theoretical milestones like Artificial General Intelligence (AGI) measured by notional benchmarks – will win the day. The real frontier lies in crafting intuitive, secure, and sector-specific solutions that people actually want to use, and enabling enterprises to adopt this technology and realize its benefits throughout the economy. This “AI 2.0” will transform business and government processes, making every company an AI company, provided we create the right conditions for seamless adoption and integration.

To ensure the adoption and customization that small businesses and the US technology ecosystem seeks, we must preserve the ability to build technologies that are fit for purpose. For example, **focusing solely on [model size](#)** ignores the reality that adoption needs to extend beyond growing superclusters. Individual businesses and government departments will require solutions that are both cloud, chip, and [scale-agnostic](#). This has been a major focus for Cohere. In fact, our latest model release, [Command A](#), is able to run on just two GPUs while beating leading state-of-the-art models running on 32 or more GPUs in performance on enterprise tasks.

Public policy should therefore be focused on enabling these opportunities, rather than the highly speculative hazards of superintelligent machines. To date, much of the policy debate has focused on addressing speculative risks, possibly posed by broadly available models intended for individual consumers. Enterprise risks are different, and the opportunities for productivity and growth are infinite. Solving for real-world hurdles that businesses and government agencies face today is essential to winning in AI. Just as critical will be not hamstringing these technologies, avoiding fragmented policies and one-size-fits all-approaches. A balanced and thoughtful approach to AI regulation is essential for America and allied nations to win the global AI race. Overregulating the industry with unnecessary burdens risks slowing technological advancement and potentially conceding further ground to China, undermining America's position in the global tech arena.

Cohere's Core Priorities for U.S. AI Leadership

1. Make AI Adoption a Centerpiece

Cohere urges OSTP to make federal and private sector adoption of AI a centerpiece of the Action Plan, backed by policies that incentivize agencies to experiment and rapidly scale successful AI deployments.

The United States must lead by example in harnessing AI's potential – both to advance national security and capture its economic benefits. Government adoption of AI is fundamental to upholding the country's national security imperative, and implementing secure AI solutions on-prem is crucial. There is a critical window to invest in AI for defense, intelligence, and homeland security applications, staying ahead of adversaries racing to exploit AI capabilities. By integrating secure AI across federal agencies—from the Pentagon to civilian agencies—the U.S. will be better equipped to detect threats and maintain its technological edge.

Beyond advancing our national security interests, AI adoption in government can vastly improve government operations and help meet citizens' needs. AI can automate compliance and reporting, improve infrastructure and public safety, speed up government services, and enhance informed decision-making. Federal agencies spend over [\\$100 billion on IT](#) annually, with much of it maintaining outdated legacy systems. Cohere urges OSTP to make federal and private sector modernization these systems a centerpiece of the Action Plan, backed by policies that incentivize agencies to experiment, spend their time and resources more impactfully, and rapidly scale successful AI deployments.

Recent U.S. AI regulatory proposals originate from fear about the largest models, with rigorous requirements that would entrench large, existing incumbents. Instead, policy should be focused on enabling AI adoption: addressing uncertainty around how existing law applies, a lack of clarity around definitions, and enabling small and medium enterprises to innovate and participate in government initiatives.

Agencies across the government should treat AI as a strategic priority and a mission-enabler. The AI Action Plan should include initiatives that:

- Set clear goals and funding for government AI adoption and modernization programs, such as pilot projects in each agency to solve specific operational challenges with AI.

- Promote the embedding of AI systems in federal operations – from automating rote processes and augmenting analytics, to enhancing cybersecurity and decision support – with appropriate safeguards.
- Include dedicated funding for inter-agency collaboration platforms to share data and AI tools, and streamlined approval processes for new AI uses.
- Clarify existing laws that might be blockers to private sector adoption – including for heavily regulated critical infrastructure sectors, which are the most in need of these technologies.
- Create common definitions – ideally legislatively – that allow the public and private sector to work together effectively. Currently, US states and agencies across the federal government use differing definitions of AI models, systems, and risks, creating a patchwork for AI suppliers and public sector entities.
- Give equal weight to smaller AI companies in advisory boards and legislative discussions as you consider public policy. Innovative startups are closer to the technology and often have more at stake in finding creative solutions for these issues.
- Support copyright and other policies that allow AI training: Copyright is designed to promote creativity and human progress – not to impede the development of new technologies. We caution against imposing new obligations that would effectively push AI progress offshore or into the hands of only the largest corporations. Instead, we urge preserving open, technology-neutral frameworks that allow lawful uses of data for innovation.

2. Reform and Modernize Government Procurement for AI

Current federal procurement rules and lengthy contracting cycles favor established vendors, creating high barriers for startups to provide innovative services to the government. The AI Action Plan should encourage updates and reform to the Federal Acquisition Regulation (FAR) and agency procedures to streamline the acquisition of AI systems, and reduce entry barriers for non-traditional contractors.

This means adopting flexible, agile procurement models – such as challenge-based solicitations and pilot programs – that allow agencies to rapidly test and integrate cutting-edge AI solutions. Similarly, streamlining compliance requirements and vendor qualification criteria that unintentionally exclude startups is also essential.

By making federal procurement more accessible, the government can leverage its "power of the purse" to catalyze AI innovation and avoid vendor lock-in. We urge OSTP to:

- Direct agencies to implement interoperability requirements in AI contracts to ensure solutions from different providers can work together seamlessly.
- Simplify solicitations and compliance paperwork to get to the use case and need more directly.
- Ensure procurement evaluations place appropriate weight on innovation, security, and performance rather than just business size or past contract volumes.
- Avoid single-source, vendor lock-in that only supports large legacy players and instead structure government contracts and procurement processes to encourage diverse participation.
- Consider challenge-based solicitations and pilot programs:
 - Include programs that would give project managers more latitude and ability to begin work on certain pilots prior to full certification or procurement completion, provided certain industry standards are met.
 - Promote "regulatory sandbox" provisions for testing innovative approaches.
 - Expand the use of Other Transaction Authorities (OTAs) for AI projects generally.
- Preference and provide support for smaller companies with innovative technologies trying to navigate the morass of different certifications.
- Ensure government contracting or public policy doesn't unfairly support open or closed source over the other.
- Focus on finding commercial solutions to address government needs first before building models or applications from scratch. This can speed adoption while leveraging existing external expertise. Cohere specializes in customization and collaborating to build specialized models.

Not only will these recommendations deliver direct benefits to taxpayers, but they will also signal to the industry and the world that the U.S. government supports and drives AI innovation. Fostering a fair, competitive AI marketplace for government contracts will ensure the best and most cost-effective AI technologies are deployed in federal missions, not just those from the largest incumbents. As President Trump's Executive



Order 14141 emphasizes, the U.S. must "support a fair, competitive AI ecosystem for businesses of all sizes" to maintain its AI leadership.

3. Prioritize Regulatory Flexibility through Risk-Based & Sectoral Governance

Regulatory Flexibility

Policymakers should avoid overly prescriptive mandates or one-size-fits-all oversight regimes that could inadvertently stifle innovation. Instead, the AI Action Plan should champion agile, principles-based governance focused on the actual risks and contexts of AI use.

Overly rigid or broad-brush regulation poses a real danger: regulatory ambiguity and over-prescription will risk inhibiting investment, innovation, and adoption—especially for small- and medium-sized enterprises (SMEs) and start-ups that are powerful engines of AI innovation.

Unfortunately, the current AI regulatory landscape suffers from a fundamental misalignment: the attempt to govern AI through broad, technology-focused frameworks ignores the nuanced realities of different industries, evolving technology and shifting responsibilities in the supply chain. This approach has led to overlapping legislation, inconsistent standards, and unnecessary compliance burdens that impede innovation and misdirect the work of governments and regulators, without effectively addressing applied sector-specific risks. State-level legislation, in particular, poses an existential compliance risk for startups who cannot keep up with dozens of differing definitions and standards.

To align regulatory approaches and target these risks, we instead propose a federal sector-led regulatory approach that leverages the deep expertise of existing agencies and regulatory frameworks. Agencies should be encouraged to focus on near-term barriers to adoption and deployment – taking into account use case, deployment context (API, public cloud, or private deployment), and access control (consumer vs. enterprise, freely available vs. controlled distribution) – while prioritizing high risk applications that influence consequential decisions or access to essential services. The success of this approach requires agencies to first clarify how existing regulations apply to AI before creating new ones. Governments must also develop these clarifications with industry at the table and ensure as much consistency of terminology and frameworks as possible across sectors.

This approach recognizes a crucial reality: AI applications in healthcare, for example, face fundamentally different challenges than those in financial services, energy infrastructure, or transportation. Each sector has its own risk profiles, operational requirements, and established regulatory mechanisms that can be adapted for AI governance as needed. We encourage the AI Action Plan to:

- Promote a sectoral approach to public policy where model companies work with deployers who are specialized in the industries they serve.
- Designate primary regulatory authority to agencies with sector-specific jurisdiction.
- Focus on high risk use cases, deployment contexts, and access controls.
- Direct agencies to clarify existing regulations or consider how existing regulations apply to the use of AI in regulated activities under their purview before creating new law.
- Use existing cross-sectoral risk assessment frameworks and build sector-specific risk assessments only for specific deployment contexts.

The U.S. should likewise prioritize harmonization of definitions and standards with other agencies and jurisdictions, grounded in U.S. initiatives. Existing voluntary, risk-based workstreams at the National Institute of Standards and Technology (NIST) are well-positioned to drive this harmonization. American leadership in global standards will advance American leadership in AI.

Testing and Assessment

Developing effective testing and assessment frameworks for AI systems requires a delicate balance between ensuring security and enabling innovation or experimentation. Industry has led the way and pioneered testing standards that have shown to be both cutting-edge and effective. Public policy in this space should focus on supporting these efforts, providing additional data, standardization, and highlighting best practices rather than building new preemptive and burdensome regimes on top of this work. This approach scales with the potential impact of AI applications while supporting domestic innovation, economic growth, and ensuring global competitiveness.

The AI Action Plan should maintain an innovation-friendly environment by:

- Avoiding blanket licensing or certification requirements for AI models.
- Setting performance-based standards instead of pre-market approvals.

- Focusing on risks that are known, measurable or observable for frontier, high-risk AI use cases.
- Avoiding compulsory sharing or public disclosure of trade secrets and other proprietary information.

We support the Administration’s focus on AI security. Cohere’s recent [Security Framework](#) outlines our perspectives on the practical questions raised by AI that are not covered by terms like “safety.” Our AI security framework addresses key dimensions, including infrastructure and model security, robustness, and privacy, ensuring these considerations are integrated into our model design and engineering workflows. Cohere’s experience underscores that risk management must be context-specific: the intended use of an AI system informs what risks are most salient and what controls are appropriate. We recommend the AI Action Plan encourage agencies to work with industry on sector-specific AI assurance techniques (for example, financial sector standards for algorithmic non-discrimination in lending).

While countries around the world are creating (or considering creating) AI Safety/Security Institutes (AISIs) to help build government expertise in understanding and evaluating AI risks, we believe the most pertinent role for these types of bodies is to serve as centers of scientific excellence and technical expertise rather than disseminate regulation. Their primary function should be supporting sectoral agencies and industry players in advancing measurement science with technical guidance, harmonized definitions, research, and standard-setting. These institutes can develop testing methodologies and conduct research on existing and emerging risks, and provide technical assistance to both regulators and companies for applied risks. In this way, institutes can leverage the unique expertise, resources and credibility of governments to advance the field and provide guidance and tools for best practices. Just as space agencies realize the government’s commitment to advance space science, so too can institutes advance the science of AI security.

AISIs should also focus as much on practical, immediate challenges that present barriers to safe adoption of AI today as they do on speculative risks. This includes developing standardized testing protocols for applied issues like discrimination in the context of AI applications subject to anti-discrimination laws, cybersecurity, and privacy. It is critical that the government focuses on domains where it has access to relevant data. This is essential in avoiding duplication of efforts and providing mutually beneficial outcomes between the public and private sector. For example, the government is best positioned to receive and maintain threat intelligence to model security. It could provide



research and sharing of this data that companies do not possess such capabilities. Confidentiality guarantees – from joint research on assessment to threat emerging databases – can foster trust and voluntary sharing.

We urge OSTP to ground the AI Action Plan in a risk-tiered framework that directs the most scrutiny and safeguards toward high-risk AI applications (e.g., security-critical or rights-impacting uses) while streamlining requirements for lower-risk uses. This proportionate approach aligns with global best practices and ensures compliance efforts focus where they matter most.

By leveraging existing privacy, consumer protection, and security regulations – and filling gaps with targeted measures only as needed – the government can avoid "overlapping or conflicting rules" that create needless burden. We advocate regulatory pragmatism: use tools like the NIST AI Risk Management Framework and sector-specific guidance to address real-world risks, rather than imposing sweeping new licensing or compliance regimes.

4. Invest in AI R&D Infrastructure – Data, Compute, and Skills

America's leadership in AI is not assured – our competitors are investing heavily in AI. The Chinese government has made it a national goal to surpass the U.S. in AI by 2030 and is "positioning Chinese firms to become the next AI leaders" through massive state-led efforts. The U.S. and like-minded nations need to move decisively to retain our technological edge by making strategic investments to expand access to fundamental resources that fuel AI development. To do this, we must:

- Accelerate federal open data initiatives and incentivize privacy-preserving data sharing.
- Implement the OPEN Government Data Act fully and launch new public-private data trusts for AI research.
- Fund the proposed National AI Research Resource (NAIRR) to provide shared computing resources.
- Expand NSF and DOE programs that grant AI researchers time on supercomputers.
- Consider tax incentives or credits for companies (especially SMEs) investing in AI-critical infrastructure.
- Support AI workforce development through education grants, visas for AI experts, and reskilling programs.

These recommendations mirror efforts by bipartisan members in the Congress, such as Sens. Rounds and Heinrich. By leading in AI research, the U.S. can set standards rather than ceding that role to competitors. We encourage OSTP to explicitly recognize AI leadership in research as a national security priority in the Action Plan, aligning efforts across the Department of Defense, the Department of Homeland Security, the Department of Energy, the Intelligence Community, and beyond to accelerate adoption of reliable AI.

Conclusion

Cohere strongly supports the development of a National AI Action Plan that positions the United States and its allies to win the AI innovation race, while managing risks responsibly. Our recommendations focus on practical, actionable steps. This balanced approach – promoting competition and flexibility alongside targeted safeguards – is the surest path to maintain American leadership in AI. It will further unlock AI's immense potential for economic growth.

Cohere is optimistic about an AI-powered future and committed to its realization. We stand as a willing partner to federal agencies and policymakers in implementing the AI Action Plan's vision – whether through providing technical expertise, participating in pilot programs, or sharing our methodologies. Should you wish to engage further, please do not hesitate to reach out to our government affairs lead, **A.J. Bhadelia at [aj \[at\] cohere.com](mailto:aj@cohere.com)**.

Together, government and industry can cultivate an AI ecosystem that is innovative, secure, and aligned with our values. The time to act is now – with the right policy framework, the United States can usher in a new era of AI-driven prosperity.

Respectfully submitted,

A handwritten signature in black ink, appearing to read 'A.J. Bhadelia', with a long horizontal flourish extending to the right.

A.J. Bhadelia
North America Government Affairs and Policy
Cohere