FEDERAL AI UNLOCKED

PART 1 OF 3 | APRIL 2025

The Federal AI Imperative: Governance, Risk, and Compliance

NETWORK DESIGNS inc™

NetDes.com

THE ISSUE

Will the Federal government be able to strike a balance between adopting AI at the speed and scale our country demands?

The AI race is on, but speed alone won't win it. As the federal government continues to invest billions into AI, the real challenge emerges: Will the Federal government be able to strike a balance between adopting AI at the speed and scale our country demands and the need to do so responsibly and ethically? Join NDi in our Federal AI Unlocked series as we unpack how agencies are navigating this high-stakes frontier.

Introduction

Artificial Intelligence is rapidly transforming federal operations—from streamlining services to enhancing national security. The U.S. federal government is now the single largest purchaser of AI technologies and is predicted to spend more than $32 billion dollars in non-defense related AI R&D in 2026.[1] Federal agencies are deploying AI for everything from healthcare to disaster relief, evidenced by agency requests for $1.9 billion for AI R&D in FY2024 to further drive innovation.[2] With this expansion comes a heightened responsibility: to balance cutting-edge innovation with robust governance and compliance. Federal leaders recognize that as AI systems become more embedded in mission-critical tasks, they must be managed responsibly to protect security, ethics, and public trust.[2] AI can introduce serious risks if left unchecked. For example, models often embed biases that lead to discriminatory or unjust outcomes, and automated decisions can create a false aura of objectivity that obscures errors and biases.[3] In short, embracing AI's potential goes hand-in-hand with ensuring it is responsibly deployed with guardrails to mitigate risks to citizens' rights and safety.

Federal agencies are pursuing "Responsible AI"—championing innovation while creating governance structures to rein in AI's externalities. Join NDi as we explore that journey, outlining how government policies and frameworks provide guidance, how agencies are navigating AI risks (like bias and security vulnerabilities), and how compliance and oversight are evolving in the public sector. The goal is to tell the story of federal AI adoption through the lens of trustworthiness: how the government is working to harness AI's power for good while avoiding pitfalls and setbacks. **NDi's Take: Whether you're a policymaker or industry partner the narrative below offers insight into how the U.S. government is leading in the AI space—carefully, transparently, and with accountability.**

The AI Governance Landscape: Policies & Frameworks

The federal government has moved quickly to establish an AI governance landscape that marries **ambition with accountability**. A series of high-level policies and frameworks now guide how agencies approach AI. In late 2023, the White House issued Executive Order 14110 on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence laying out principles for ethical AI use across government. Building on that, a new Executive Order on **January 23, 2025** titled Removing Barriers to American Leadership in AI doubled down on U.S. AI lead-

REAL-WORLD APPLICATIONS

Operationalizing Responsible AI for Federal Impact

NDi brings real-world experience aligning AI solutions with federal governance standards like the NIST AI Risk Management Framework. Our work with the FAA applied Machine Learning and Natural Language Processing to surface regulatory relationships with explainability and transparency built in. We've enhanced federal chatbots with generative AI pipelines that prioritize ethical data handling and accuracy, and our internal AI incubator trains job-ready talent for high-stakes missions. From semantic search to bias-aware automation, NDi delivers AI that's engineered for trust, oversight, and compliance from day one.

93%

of federal agencies report having AI-use policies in place.

ership. This 2025 order calls for a comprehensive AI Action Plan to **maintain U.S. global AI dominance** in innovation and security, while ensuring AI systems are free from undue bias. It even directs officials to **revise or remove outdated regulations that inhibit AI innovation**.[4] Together, these directives signal a top-down commitment to foster AI advancement *and* implement guardrails so that progress doesn't come at the expense of American values.

Crucially, agencies aren't starting from scratch in interpreting these mandates—they are aligning with established frameworks for trustworthy AI. The **National Institute of Standards and Technology's AI Risk Management Framework (NIST AI RMF)**, released in January 2023, has become a foundational guide. NIST's framework is a "multi-tool" that helps organizations **design and manage trustworthy, responsible AI**.[5] It adds coherence to U.S. AI policy by defining what "trustworthy AI" entails and how to achieve it. Many agencies have embraced the NIST AI RMF to shape their strategies: for example, the Federal Aviation Administration (FAA) is developing an agency-wide AI governance framework aligned to **federal principles of trustworthy, accountable AI**.[6] Likewise, the Department of Justice's AI Strategy emphasizes **ethical and efficient AI governance in accordance with established laws and best practices**.[7] In short, federal organizations are weaving the NIST guidelines and White House principles into their own policies, ensuring a consistent government-wide stance on responsible AI.

Key building blocks of federal AI governance include:

- **Executive Orders and Strategies–** High-level directives (e.g. the Jan. 2025 AI Leadership EO) set the vision by prioritizing innovation, security, and the removal of barriers. These policies urge agencies to **promote American AI leadership while upholding values of transparency and fairness**.[4] Agencies like DOJ translate these into action by providing "clear guardrails" for AI use in their missions.[7]
- **NIST AI Risk Management Framework–**A voluntary but influential framework that addresses **bias, explainability, robustness, and security** in AI systems.[8] By adopting NIST's guidance, agencies can systematically identify and mitigate AI risks. (Notably, even state laws such as Colorado's new AI Act give credit to organizations following the NIST AI RMF-underscoring its authority.[4])
- **Agency AI Policies & Boards–**Across government, agencies are crafting AI policies aligned with federal frameworks. As of 2024, **93% of federal agencies report having AI-use policies in place** (vs. 78% in state/local governments[9])— a clear performance indicator of policy adoption. Many have formed AI working groups or governance boards to enforce these policies internally.

This emerging governance latticework provides a roadmap for agencies. Executive Orders define *what* must be achieved (leadership, safety, ethics), and frameworks like NIST's AI RMF illustrate *how* to achieve it. The result is a federal AI ecosystem striving to encourage cutting-edge AI deployments that are **innovative yet principled**. By proactively establishing policies, the government aims to avoid a Wild West of AI—instead creating an environment where agencies can experiment with AI solutions **confident that they're staying within well-defined ethical and operational boundaries**.

Navigating Risk: NIST AI RMF & Beyond

Even with strong policies in place, implementing AI responsibly day-to-day is a complex challenge. Federal agencies must navigate a maze of risks from algorithmic bias and security vulnerabilities to reliability issues as they integrate AI into critical operations. Here, the NIST AI Risk Management Framework (RMF) has become an essential compass. The NIST AI RMF provides a structured approach for agencies to **identify, assess, and mitigate AI-related risks** across the AI lifecycle. It introduces key characteristics of trustworthy AI (for example, that systems should be safe, secure, explainable, privacy-enhanced, fair, and free from harmful bias[5]) and outlines processes to achieve those outcomes. Agencies are using this benchmark to double-check that their AI systems won't inadvertently violate civil rights, expose security loopholes, or undermine missions due to faulty data or design. In practice, this means doing things like rigorous bias testing of AI models, independent security evaluations, and continuous monitoring of AI decisions for errors.

"There was a sense that, 'Oh, let's just give it to the robots and there will be no more discrimination.' But that actually isn't true at all—bias is often built into the data."

— Rohit Chopra, Director of the Consumer Financial Protection Bureau

According to Rohit Chopra, Director of the Consumer Financial Protection Bureau, "There was a sense that, 'Oh, let's just give it to the robots and there will be no more discrimination.' But that actually isn't true at all—bias is often built into the data."[10] AI is not a magical solution immune to human flaws. Without careful oversight, it can **amplify biases or errors** present in training data or algorithms. Federal teams are therefore "baking in" risk mitigation steps alongside AI development. For example, the **NIST RMF's "Map, Measure, Manage, and Govern" approach** has agencies mapping out potential risks early (such as biases impacting decisions), measuring system behavior against benchmarks (like fairness metrics), and establishing controls and governance to manage these risks throughout deployment.[5]

Agencies are also sharing knowledge on AI risks through communities of practice and interagency councils. Using common frameworks means a **more unified defense** against AI pitfalls. A flaw discovered by one agency (say, a facial recognition bias) can lead to updates in best practices that benefit all. We are seeing concrete outcomes of this risk-focus: in 2024, agencies conducting AI risk assessments identified **227 use cases as "high-risk" (impacting rights or safety)** that required special mitigation plans.[11] This metric reflects a growing rigor in how agencies scrutinize AI projects. Rather than deploying blindly, agencies are pausing or redesigning AI tools that could harm the public. In fact, OMB reported that dozens of AI applications were halted or given extra oversight until they could meet new risk management requirements—a sign that the framework and oversight are catching issues before they escalate.[11]

Beyond NIST, agencies draw on sector-specific expertise to manage AI risks. The Department of Defense, for instance, has developed

REAL-WORLD APPLICATIONS

Semantic Search Built for High-Stakes Environments

NDi's work at the FAA advanced Machine Learning (ML) and Natural Language Processing (NLP) to map regulatory relationships across over 1.5 million documents. We engineered explainable AI tailored to aviation safety, delivering intelligent search, bias-aware data structuring, and transparent results visualization. The system enables real-time access to critical compliance information, supporting federal missions with clarity and accountability. It's AI that informs, not obscures.

The federal government more than doubled its cataloged AI use cases from 710 to 1,757 in one year, yet implemented new risk controls that paused or adjusted dozens of potentially harmful systems. This indicates that AI adoption is accelerating alongside a commitment to risk mitigation.[11]

detailed principles for **ethical AI in warfare** and toolkits for testing AI robustness under battlefield conditions.[12] Civil agencies like the Department of Health and Human Services must weigh privacy and safety heavily when using AI in healthcare decisions. Yet, despite differences in mission, the connective thread remains: **robust risk management** is non-negotiable. By institutionalizing processes to mitigate bias, secure AI systems against threats, and validate algorithms, agencies align AI deployment with their core objectives and values.

The payoff is twofold. First, it protects citizens, ensuring an AI that helps decide benefit eligibility, for example, does so fairly and accurately. Second, it bolsters innovation because teams and the public have greater confidence in AI that has been through proper vetting. In the federal space, responsible AI isn't seen as a check-box compliance exercise; it's viewed as **critical to mission success.** | **NDi's Take: As agencies continue to refine their use of the NIST RMF and related tools, we can expect the federal AI ecosystem to mature in its ability to spot risks early and address them, allowing the government to push the envelope with AI in a safe, controlled manner.**

Regulatory Compliance in AI Adoption

Hand-in-hand with risk management is the need for strict **regulatory compliance** when federal agencies adopt AI. Simply put, agencies must ensure that introducing AI into government services **does not violate laws or public trust.** A robust compliance approach covers ethics, security, and privacy requirements that AI systems must meet, or agencies face legal and reputational consequences. In recent years, the federal oversight community has sharpened its focus on AI compliance. The White House Office of Management and Budget (OMB) has issued clear directives to agencies on this front, notably **OMB Memo M-24-10 (March 2024)** entitled "Advancing Governance, Innovation, and Risk Management for AI." This memo requires agencies to integrate AI into existing compliance structures rather than treat it as a Wild West. Concretely, it told agencies to **publish AI compliance plans**, conduct impact assessments for AI that affect rights or safety, and ensure independent testing of AI systems.[3] Each agency's plan must outline how their AI use aligns with ethical principles and legal obligations (covering areas like privacy, civil rights, and cybersecurity). This move by OMB essentially forces agencies to **show their homework**. Before scaling up AI, they must prove they've addressed the question, "Is this AI use lawful, ethical, and secure?"

Compliance also means adherence to baseline statutes and regulations that predate AI. Laws around privacy (Privacy Act, HIPAA), civil rights and non-discrimination, and administrative law all still apply—AI or not. As former FTC Chair Lina Khan bluntly put it, "There is no AI exemption to the laws on the books."[10]

REAL-WORLD APPLICATIONS

Generative AI, Deployed with Precision and Oversight

For the FAA's CE Chatbot, NDi integrated generative AI to improve public engagement and responsiveness—without compromising ethical boundaries. Our Azure based internal tool, NDiGPT, curates and validates training content before deployment, ensuring accuracy and eliminating model drift. This curated approach earned national recognition and demonstrated how mission-specific AI can be scaled safely. When it comes to AI transparency, we lead with control, not convenience.

Federal AI systems must therefore be designed and deployed in ways that **uphold existing laws**. For example, an AI used in hiring by a federal contractor must not discriminate based on protected attributes, or an AI that processes personal data must follow privacy and security protocols. Agencies are embedding these requirements into their procurement and development processes. The **Department of Justice's Civil Rights Division** and other oversight bodies have signaled they will scrutinize government AI uses for any disparate impacts or biases that could violate rights.[10] Likewise, Inspectors General and the GAO (Government Accountability Office) are auditing agency AI projects to ensure compliance with cybersecurity standards and proper data handling.

Despite these mandates, compliance remains a work in progress. Early checks show mixed results. By late 2024, **only about half of federal agencies had published the required AI compliance plans** per OMB's directive.[3] Those plans that were published also varied in quality and transparency. This highlights a familiar tension: regulations can be issued from above, but agency follow-through takes time and effort. The encouraging news is that the mere act of measuring compliance is spurring action. Agencies that fell short are now under pressure to catch up, and the spotlight on compliance rates is motivating the laggards. Meanwhile, federal procurement rules are evolving to enforce compliance down the supply chain. For instance, OMB's guidance (and a forthcoming update M-24-18) is expected to require that **vendors certify their AI products meet certain ethics and security criteria**.[4] By holding contractors accountable, the government can prevent "black box" AI solutions from slipping in without proper vetting.

In practical terms, robust AI compliance means agencies doing things like conducting **privacy impact assessments** whenever a new AI system will handle personal data, ensuring AI decisions can be **explained and documented** to satisfy due process in public services, and maintaining **human oversight** for AI-driven processes to intervene if something goes wrong. It's the unglamorous side of AI adoption, but absolutely essential for sustaining public trust. Federal leaders know a high-profile AI failure—say an AI system denying benefits unjustly or a data breach via an AI tool—could set back broader adoption. Thus, they are working proactively to build compliance into the AI lifecycle. When done right, compliance measures don't stifle innovation; rather, they act as **safety nets**, giving agencies the confidence to scale effective AI solutions knowing they are on firm legal and ethical ground.

The Future of AI Policy in Government

Looking ahead, federal AI governance is poised to evolve rapidly as both the technology and public expectations advance. We can anticipate a wave of **new regulatory developments** aimed squarely at the toughest challenges of AI: bias, security, and ethics. One likely focus area is **algorithmic bias and fairness**. Thus far, agencies have managed bias largely through guidelines and internal reviews, but future policy could formalize this—for instance, requiring routine **algorithmic impact assessments or bias audits** for any high-stakes AI system (similar to environmental impact statements for major projects). There are growing calls for legislation that ensures

If done right, the payoff is not just cutting-edge government AI, but AI that Americans can trust knowing it's governed with their best interests, rights, and security in mind.

AI decisions in areas like hiring, lending, or criminal justice are free from unlawful discrimination. We're already seeing movement at the state level (e.g., **Colorado's AI law governing high-risk algorithms to prevent discrimination** and Illinois' law banning biased AI in hiring),[4] and these may serve as templates or pressure for federal standards. A federal law or regulation in the near future could mandate agencies to **proactively test AI for fairness** and publicly report results injecting greater transparency into government AI use.

The federal journey with AI is still in its early chapters, but the narrative is clear. By embedding governance, risk management, and compliance into the heart of AI adoption, the U.S. government is striving to **set the gold standard for public sector AI**. The transformation driven by AI in federal operations can indeed be revolutionary—improving efficiency, enabling data-driven policymaking, and enhancing citizen services—but only if accompanied by the guardrails that prevent unintended harm. With strong frameworks, vigilant risk management, and a culture of accountability (through roles like Chief AI Officers and oversight councils), federal agencies are writing a story of responsible innovation. The next few years will be critical in filling out this story, as policies tighten, and technologies mature. **NDi's Take: If AI is done right, the payoff is not just cutting-edge efficiency, but AI that Americans can trust—knowing it's governed with their best interests, rights, and security in mind. By proactively adapting to future challenges (bias, security, ethics), the U.S. public sector will continue leading in AI, proving that even in a rapidly changing tech landscape, good governance and innovation can go hand in hand.**

REAL-WORLD APPLICATIONS

Human-Centered AI Talent, Ready for Deployment

Through our Power Up Program, NDi cultivates AI/ML professionals equipped for federal mission needs—trained on real projects in predictive analytics, model development, and secure cloud environments. Interns transition into billable roles quickly, reducing onboarding time and elevating workforce readiness. This program directly supports responsible innovation by embedding ethical practices and compliance from the ground up. It's how we scale trustworthy AI—by building it into the talent pipeline.

About the Author: *Matthew Hopkins is a seasoned AI leader focused on mission-aligned artificial intelligence in the federal government. A recent Georgia Tech graduate, he conducted research on human-machine interaction focused on AI and holds credentials in AI product design and robotic process implementation from MIT, Duke, and Columbia. As a leader at Network Designs, Inc. (NDi), he has developed AI/ML solutions for secure access, regulatory intelligence, and chatbot optimization. He also built NDi's Artificial Intelligence Hub and Power Up Program, an incubator for next-gen AI/ML talent. His work bridges technical depth with a strong grasp of federal compliance and responsible innovation.*

Citations

1. Stewart, N. (2024, January). Funding for the Future: The Case for Federal R&D Spending. https://www.scsp.ai/wp-content/uploads/2024/01/RD-White-Paper-2.pdf
2. Government Accountability Office. (2024, September 9). Artificial Intelligence: Agencies are implementing management and Personnel Requirements. https://www.gao.gov/products/gao-24-107332#:~:text=Why%20GAO%20Did%20This%20Study
3. Tobin-Miyaji, M. (2024, November 21). *Federal agencies largely miss the mark on documenting AI compliance plans as required by AI Executive Order*. EPIC. https://epic.org/federal-agencies-largely-miss-the-mark-on-documenting-ai-compliance-plans-as-required-by-ai-executive-order/#:~:text=The%20U,that%20must%20be%20carefully%20scrutinized
4. Gilbert, N. (2025, January 23). *Executive order: Removing barriers to American leadership in Artificial Intelligence: Insights*. Holland & Knight. https://www.hklaw.com/en/insights/publications/2025/01/executive-order-removing-barriers-to-american-leadership-in-ai#:~:text=-directs%20White%20House%20officials%20to,the%20Biden%20AI%20EO
5. Kerry, C. F. (2023, February 15). *NIST's AI Risk Management Framework plants a flag in the AI debate*. Brookings. https://www.brookings.edu/articles/nists-ai-risk-management-framework-plants-a-flag-in-the-ai-debate/#:~:text=The%20National%20Institute%20of%20Standards,about%20AI%20policy%20and%20development
6. *Use of Generative AI Tools and Services*. U.S. Department of Transportation. (2024, October 3). https://www.faa.gov/documentLibrary/media/Notice/N_1370.51_Generative_AI_Tools_and_Services.pdf
7. *Artificial Intelligence Strategy for the U.S. Department of Justice* . U.S. Department of Justice. (2020, December). https://www.justice.gov/d9/pages/attachments/2021/02/04/doj_artificial_intelligence_strategy_december_2020.pdf
8. *Comparing NIST AI RMF with other AI frameworks*. RSI Security. (2025, February 12). https://blog.rsisecurity.com/comparing-nist-ai-rmf-with-other-ai-risk-management-frameworks/#:~:text=Comparing%20NIST%20AI%20RMF%20with,Request%20a%20Free
9. Fox-Sowell, S. (2024, September 19). *Federal government outpacing state, local agencies on AI adoption, survey finds*. StateScoop. https://statescoop.com/federal-government-state-local-ai-adoption-2024/#:~:text=The%20survey%20found%20that%20the,of%20federal%20agencies
10. SolasAI. (2023, June 28). *"Sounding the alarm": 3 quotes from government officials about responsible AI*. Medium. https://pub.solas.ai/sounding-the-alarm-3-quotes-from-government-officials-about-responsible-ai-8bc976762913
11. Adler, M. (2024, December 18). *Federal government discloses more than 1,700 AI use cases*. FedScoop. https://fedscoop.com/federal-government-discloses-more-than-1700-ai-use-cases/#:~:text=Per%20the%202024%20consolidated%20inventory%2C,year%20documented%20710%20use%20cases
12. Olay, M. (2024, October 29). *Hicks highlights DOD's commitment to responsible AI use*. U.S. Department of Defense. https://www.defense.gov/News/News-Stories/Article/Article/3949441/hicks-highlights-dods-commitment-to-responsible-ai-use/#:~:text=,Hicks%20said