America's AI Action Plan: Implications for Business Leaders

July 2025

Data & Trust Alliance

Executive Summary

The AI Action Plan, released on July 23, 2025 by the Trump administration, signals a turning point in the federal government's approach to artificial intelligence governance. The clear goal is the establishment of American AI dominance on the global stage. The Trump administration's framework recalibrates the government's role—shifting from top-down mandates to an emphasis on voluntary standards, public-private collaboration, and national competitiveness. It places greater responsibility for sound and safe AI deployment on industry actors, particularly those who bring AI into real-world applications.

This shift is reinforced by three complementary executive orders signed the same day and provides clearer direction for U.S. AI strategy. Together, these policies establish a broader deregulatory agenda that prioritizes rapid deployment, expands American AI globally, and constrains ideological content domestically.

1. Accelerating Federal Permitting of Data Center Infrastructure
2. Promoting the Export of the American AI Technology Stack
3. Preventing Woke AI in the Federal Government

While this largely deregulatory approach provides developers with multiple mechanisms to innovate, the spotlight—and operational responsibility—shifts towards deployers.

For business leaders navigating this evolving terrain, this brief highlights:

- **Key provisions of the Action Plan**—and where they diverge from the prior administration's posture
- **Implications for deployers**—who are responsible for integrating AI into real-world applications
- **Strategic openings for deployers**—to lead in this moment and shape the norms that will ultimately create a more viable and trusted ecosystem

Key Provisions and Implications of the AI Action Plan

The AI Action Plan builds on past executive orders while sharply pivoting in tone and scope. While it continues to acknowledge AI's promise and risk, the posture shifts toward a leaner federal government, industry-led governance, and geopolitical assertiveness. Here are the standout provisions—we map each key policy shift through the lens of what it means for deployer companies:

❶ Open Source AI & Competitive Pressure

The plan introduces:

- Encouragement of open-weight AI models
- Investment in R&D to foster further AI breakthroughs
- Expansion of the NAIRR (National AI Research Resource) Pilot to democratize access to open-weight models and data

How this is different:
This contrasts earlier frameworks emphasizing developer-side controls. The new model amplifies access to open-source and open-weight models but diffuses accountability.

Key implications for deployers:
The government's support for open-weight models and the NAIRR Pilot will accelerate access and innovation, but it also raises concerns about accountability, IP erosion, and safety. Deployers will bear the downstream risk if foundation models are misused or poorly governed. In addition, deployers may face growing pressure to adopt more open-source or open-weight models, particularly in federal procurement and public-facing contexts.

❷ From Regulation to Risk-Tiered Guidance

The plan formalizes a two-tiered framework:

- *Low-to-Moderate Risk AI*: Subject to voluntary standards and market-led transparency practices.
- *High-Risk / "Frontier" AI*: Subject to "targeted and proportionate" obligations such as safety evaluations, watermarking, and disclosures.

What's different: This represents a deliberate pullback from the Biden administration's approach, which emphasized mandatory safeguards across government. Instead, the Trump plan emphasizes discretion and market mechanisms over federal mandates.

Key implications for deployers:
Deployers become the de facto risk adjudicators. Gray-zone cases will force deployers to assess system criticality, determine applicable guidance, and ensure safe integration.

❸ Voluntary AI Standards & Evaluations

The plan introduces:

- Domain-specific efforts to develop voluntary national standards for AI systems (e.g., in healthcare, energy, and agriculture)—led by NIST—with a particular call-out to measure AI productivity gains in those domains.
- Support for the development of the science of measuring and evaluating AI models

How this is different:
Shifts from mandated federal safeguards to enterprise discretion and market-based governance.

Key implications for deployers:
Deployers will be the proving ground for federal guidance.

As CAISI and NIST build voluntary standards, implementation feedback from real-world enterprise use will shape what is practical, scalable, and globally relevant.

❹ AI Workforce & Talent Strategy

The plan introduces:

- Clarifications to IRS Section 132 to support tax-free skills training
- Expanded Pell Grant eligibility for AI-related fields
- New federal apprenticeship & job-matching programs
- Study of AI impact on labor market (via Census/BEA)

How this is different:
While deregulatory overall, the administration acknowledges public anxiety around job displacement and attempts to address it without imposing restrictions on automation or mandating corporate behavior. It shifts workforce planning from federal government-led to enterprise-anchored public-private collaboration.

Key implications for deployers:
Without clear federal guidance or support, deployers will need to drive AI workforce strategy and partner with state and local governments. They must define new competencies, co-create educational pathways, and support talent development from entry-level to executive. This includes long-term planning for skills development across diverse sectors and populations. The study of the impact of AI on the labor market can refine the work companies must do now to bridge the gap now.

❺ Federal AI Monetization & Procurement Strategy

The plan introduces:

- Reinforced appointment of Chief AI Officers at federal agencies
- AI use inventories with narrowed public disclosure obligations
- Replacement of M-24-10 with lighter-touch guidance based on "national criticality"

How this is different:
Reduces required transparency, increases sensitivity to content-based scrutiny in public-sector AI, including misinformation and mention of diversity, equity, and inclusion, and reduces climate safeguards.

Key implications for deployers:
Deployers selling into the federal government must align with evolving procurement filters, including new content alignment criteria. With reduced disclosure requirements, companies must take proactive steps to demonstrate credibility, especially where guidance remains narrow or undefined. These requirements largely differ from that of SLTT government and international partners.

❻ Supply Chain, Compute Access, & Infrastructure

The plan introduces:

- Investment in U.S.-based chip manufacturing
- Support for trusted partner supply chains
- A "Compute Marketplace" to enable equitable access to infrastructure

How this is different:
The Compute Marketplace is new and signals a recognition of structural inequities in compute access, especially for researchers and small firms.

Key implications for deployers:
While the Compute Marketplace opens the door to a dynamic, innovation-driven vendor ecosystem, the absence of enforceable requirements for vendors raises critical risks for deployers—chief among them, compromised data security and potential erosion of competitive advantage.

❼ Civil Liberties, Privacy, & Social Impact

The plan introduces:

- DOJ and DHS reviews of law enforcement use cases
- Encouragement (not requirement) of privacy safeguards
- No new federal civil rights or privacy protections
- Revision of the NIST AI Risk Management Framework to eliminate references to misinformation, diversity, equity, and inclusion, and climate change.

How this is different:
Moves from mandated impact assessments to decentralized, company-led accountability.

Key implications for deployers:
Deployers will take on any downstream risks and reputational damage of issues relating to bias, surveillance harms, or discrimination.

❽ Compute Infrastructure & Sustainability Expectations

The plan introduces:

- Streamlined permitting for data centers
- Removal of DEI and climate review guardrails

How this is different:
Expedites deployment while transferring responsibility for any environmental and social standards to industry.

Key implications for deployers:
Permitting reforms allow faster construction of data centers but remove climate and DEI review guardrails. This may create reputational and operational risks for deployers who do not proactively manage sustainability and equity expectations across communities and investors.

❾ Increasingly Fragmented Norms

The plan introduces:

- Fewer top-down mandates; emphasis on sector-specific, voluntary guidance
- Updates to the NIST AI Risk Management Framework that don't align to international perspectives

How this is different:
Moves governance downstream, creating space for leadership by the private sector and SLTTs but also complexity and fragmentation.

Key implications for deployers:
With federal mandates in retreat, state (SLTT) and international regimes are stepping in. Deployers will need to navigate increasingly divergent expectations, particularly around content alignment, civil rights, and trust benchmarks. Cohesive frameworks will be essential.

continued

A Few Strategic Openings for Deployers to Lead and Shape Norms

The Trump administration's AI Action Plan reframes the federal role in AI governance, effectively devolving many responsibilities to industry. This shift, while offering increased flexibility and innovation runway, also transfers the burden of operational discipline, reputational stewardship, and risk adjudication squarely onto deployers. However, this new landscape also presents strategic openings for members of the Data & Trust Alliance, as forward-thinking organizations, to shape norms, build trust, and set competitive benchmarks.

Shape standards through early adoption and influence. The plan elevates voluntary standards—especially those developed by NIST and CAISI—as the primary mechanism of AI oversight. Deployers should be the proving ground for this type of federal guidance—and hold the opportunity to define the contours of these standards by building and piloting draft frameworks, sharing real-world implementation feedback, and co-authoring best practices. *This framework-shaping is our flagship work at the Data & Trust Alliance and helps position our members as early movers to not only influence the rules of the road but also gain credibility with regulators, customers, and global partners— ultimately creating a more viable and trusted ecosystem.*

Leverage procurement as a governance lever. In the absence of binding federal rules, procurement preferences and content filters are becoming de facto governance tools. Deployers can use this to their advantage—by both meeting evolving government expectations and setting their own standards for vendors and partners. *The Data & Trust Alliance's current work on procurement diligence and guidance is an effective tool to reach this goal.*

Embed workforce planning as central, not supplemental, to AI strategy. The deregulatory stance on automation and labor displacement places the onus on companies to anticipate and address the human impact of AI adoption. *This can be addressed as part of the Data & Trust Alliance's new Peer Communities as they do the work of reimagining horizontal value chains across the enterprise.*

Use the Data & Trust Alliance's peer communities to determine where the real-world challenges and gaps are—and leverage them to build the necessary frameworks that deliver value and trust, safeguard against emerging risks to deployers, and align with government and market signals.