The Next Move: Special Edition

Regulatory and policy developments in tech — November 2023

White House mobilizes bold push for responsible adoption of AI

By Sean Joyce, Rohan Sen, Jennifer Kosar, Tim Persons, Ilana Golbin, Jocelyn Aqua and Ege Gurdeniz

The issue

The Biden administration issued its long-awaited executive order (EO) on artificial intelligence (AI). The order is the government's biggest step toward regulating the fast-moving technology. It calls for new standards, funding, training and enforcement to mitigate AI risks, while also paving the way for the technology's widespread adoption.

The EO builds on the administration's 2022 Blueprint for an AI Bill of Rights, which noted the technology's many risks — including those to the workforce, privacy, critical infrastructure, national security and democracy — but also its potential for good. The order also builds on the voluntary commitments the White House secured more recently from executives of major AI companies, who pledged to conduct internal and external "red-teaming" (simulated attacks) on their AI models, share data about safety to third parties and develop technology to identify AI-generated content, among other measures. The EO seeks to address these and many other risks.

Who needs to pay attention? While companies that create "foundation" AI models and those that serve the federal government are most immediately affected, the EO is a bellwether of how AI may be regulated in the future. This will likely affect all companies. In addition, the order may touch them in other ways, for example, by influencing how AI models and their use can drive healthcare advances and climate solutions, shape consumer expectations for transparency and accountability, and affect the workplace.

The administration's focus

Noting AI's "extraordinary potential for both promise and peril," the Biden administration cited the urgent need for a partnership between and among government, industry, academia and civil society to mitigate the technology's substantial risks and harness its capacity for good.

The EO represents a coordinated approach spanning the entire federal government to lead in this frontier space. Drawing heavily from the NIST AI risk management framework, the order reflects a policy of advancing AI's development and use according to eight guiding principles and priorities. By focusing in equal parts on innovation and safeguards, the EO shows the government's stance on the matter, one that reinforces the criticality of responsible AI.

- **Technical aspects.** The order addresses a wide spectrum of issues connected to the more technical aspects of AI, from safety and security to testing, privacy, transparency into technical capabilities of the model and safeguards against misuse.
- **Competition aspects.** There is further acknowledgement that AI is on track to become a core technology, almost akin to a utility service. The EO envisions measures to support a free and fair market for these technologies to benefit all.
- **Societal aspects.** In addition to technical impacts, the directive references impacts to society such as climate, education, workforce (re)training and retooling, labor disruptions, criminal justice, etc. These provisions have the potential to impact more industry sectors in the long run as the administration tries to accommodate the use of AI in these domains.
- **Sector reach.** The breadth of issues the EO addresses means that most, if not all, industry sectors will feel its impact, for example, through the privacy, consumer and fraud protection, and civil rights provisions. The order does emphasize certain roles in the AI ecosystem, however, which will affect some organizations more directly.
    - Entities that create general purpose "foundation models," especially those with wide and powerful capabilities, must consider their role in how these models are used for purposes the EO is seeking to protect against.
    - Healthcare sector companies must promote the use of responsible AI practices in drug discovery and veteran care, and develop specific practices for assurance of AI systems in healthcare technology.
    - Financial services firms can expect additional guidance from regulators with respect to preventing their AI systems from violating their obligations regarding customer information, fair lending and housing, and cybersecurity risks, to name a few.
    - Companies that provide services to the government will need to pay particular attention to how they demonstrate responsible practices related to their AI systems, given the unique provisions for their vendors.

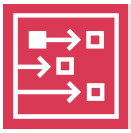

Your next move

While targeted in many ways and lacking the force of law, the EO extends far beyond the government's use of the technology. It sets policy objectives for federal regulatory agencies and will likely have far-reaching consequences.

To prepare for its implementation, companies should understand the potential direct and secondary impacts, identify the gaps and opportunities, and plan accordingly. Consider taking these steps.

- **Determine the EO's impact on your organization.** Identify all potential intersections between the many required federal actions and your business or sector — including those of your suppliers. From there, determine your potential exposure and develop a plan to manage it. Recognize that some of the EO's effects may be indirect, for example, by setting the tone for consumer sentiment regarding AI more broadly or by directing global leadership and coordination on future standards and frameworks. Use these opportunities to collaborate with your regulators on evolving practices, where relevant.
- **Monitor for and comment on rulemaking.** Over the next several months, the charged departments will be developing and issuing further, more specific guidance to measure compliance with the EO provisions. Your regulatory compliance teams will need to be ready when rules are issued. In addition, consider participating in comment and feedback periods to shape rule development.
- **Assess and manage risks of your AI technologies.** Follow the principles of responsible AI that the EO embodies.
    - Create a risk taxonomy that reflects the guiding risk principles that enable AI risks to be measured, managed and, if necessary, transparently reported to others over time.
    - Develop an enterprise governance model that considers the described obligations and responsible AI principles more broadly. A critical and foundational step to developing a governance model is aligning the roles and responsibilities of existing teams, as well as defining new ones, to support oversight.
    - Incorporate practices to assess and test your AI systems according to their risk. This may include practices like AI red-teaming described in the EO.
    - Upskill your staff to adopt and use AI responsibly and be prepared for changes to their workflow in the face of increased AI adoption, including process retooling and reskilling to maintain critical institutional knowledge.
    - Consider adopting a comprehensive privacy program, including frameworks for privacy-enhancing technologies (PETs).

- **Apply a trust-by-design approach in your AI systems.** Although the EO calls for standards to be developed, it does not define any requirements for organizations to demonstrate externally validated trust in AI systems or their outcomes. This will likely fall to individual agencies — including regulatory ones — as they act on the EO's directives. These may include audit, regulatory or management attestation requirements, all of which require management to first demonstrate their own compliance and effective practices. Companies keen to understand what they'll need to attest to will have to wait longer as these initiatives take shape. What is clear, however, is that a growing number of interested parties will likely need some documentation on AI systems and their associated data and testing practices, especially those used for critical decisions.
- **Consider segmenting your AI models.** With increasing scrutiny of models used for national security, national economic security and national public health, it'll be advantageous to segment models supporting these types of activities from models for broader, more generalized use. For example, a single enterprise-wide LLM may no longer work if different functions and use cases become subject to different provisions. A multiple-model approach may be more practical for diverse organizations looking to manage evolving compliance expectations.
- **Explore ways to deploy AI defensively.** The EO contains examples of harnessing the benefits of AI to mitigate risk, such as developing AI tools to find and fix vulnerabilities in critical software or automatically detecting cyber incidents. Consider ways to follow that example in your own operations.
- **Prepare for transparency.** If your business will face direct obligations — performing risk assessments, sharing red-teaming results — document your processes and controls and assess their readiness for external reporting. Make sure your public statements and internal practices are aligned, given the increasing scrutiny by regulators, customers and the media.

The EO on AI is a clear signal on how active the government intends to be in this space. Over the coming weeks, we can expect more specificity as individual agencies provide further resolution to policy directives. In the interim, adopting sound risk management based on responsible AI principles will position you to meet the requirements envisioned in the order.

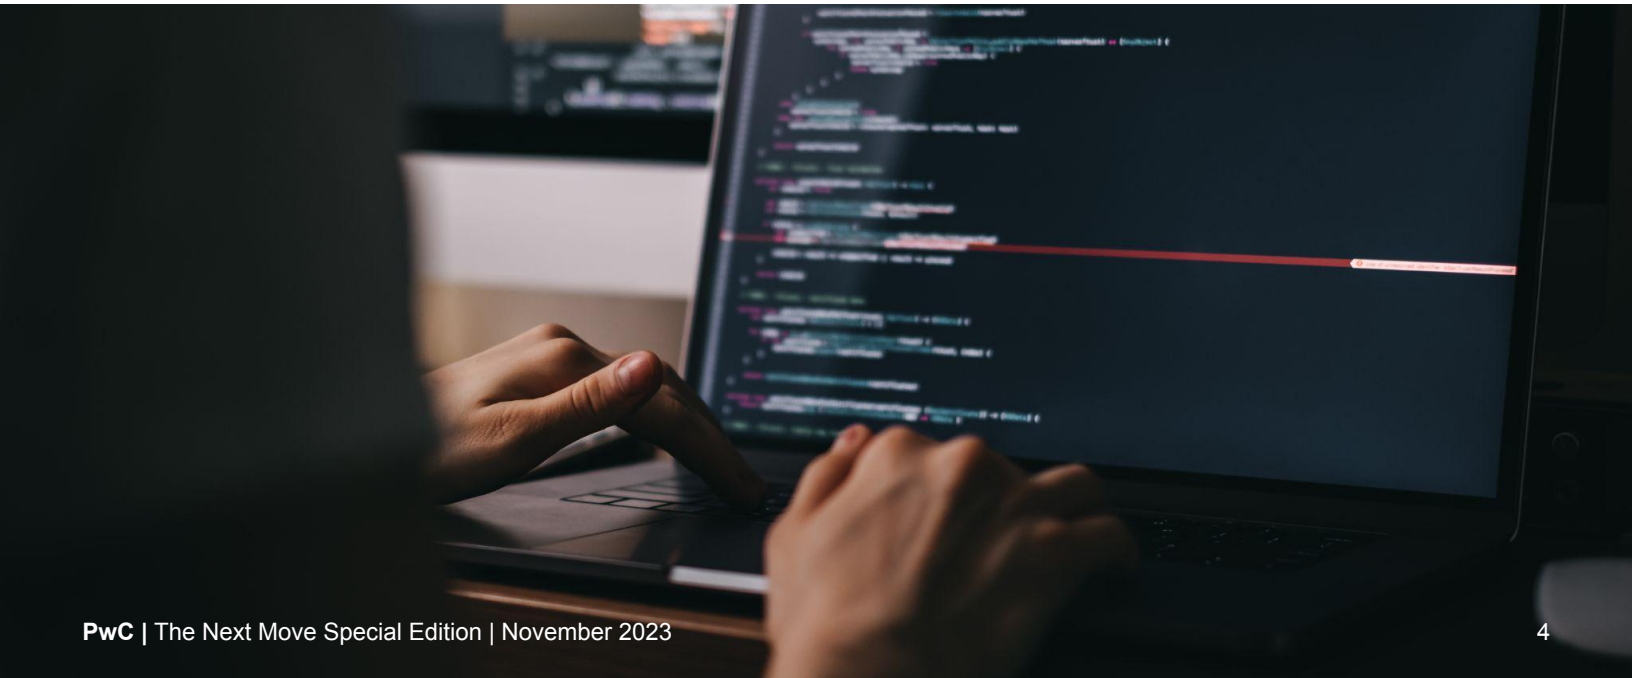

Explore more AI insights

- Generative AI: Transform the future of business and lead with trust
- Managing the risks of generative AI: A playbook for risk executives — beginning with governance
- Tech Effect: Your digital guide to growth in a people-led, tech-powered world