BEACON
GLOBAL STRATEGIES

THE DOWNLOAD

Get Ready to Red-Team: A New AI Executive Order Sets a High Bar for Industry

November 2023

What's New:
On October 30, the Biden administration issued its ambitious Executive Order (EO) on artificial intelligence (AI), the starting pistol in the U.S. government's race to set rules and regulations for the use of AI technologies at home and abroad. The directive is vast in scope, kickstarting a whole-of-government effort to craft policy that realizes the immense economic and societal promise of AI while managing the technology's very real risks to U.S. national security, cybersecurity, consumer protection and privacy, intellectual property, labor market stability, and civil rights.

Why It Matters:
In shying away from prescribing policy solutions, the EO's impact will largely depend on its implementation, Congressional funding, and international uptake in the coming months and years. But there is one group that will feel the effects of the mandate overnight: AI developers themselves. In a move that represents a tectonic shift in Washington's relationship with the private sector, the AI EO requires that a large swath of the American AI industry "red-team" its models for national security vulnerabilities and take mitigation measures accordingly. By placing the onus on companies to ensure Americans remain safe from the misuses of AI, the Biden administration has categorized the AI industry as part of the national security enterprise – and ushered in a new age of corporate responsibility that companies must meet head on.

> *"This isn't a mere bureaucratic exercise. It is a clarion call for a new era of responsibility … The government has essentially declared that the AI industry is part of the national-security apparatus, whether it likes it or not."*
>
> BGS Managing Director Klon Kitchen, *The Wall Street Journal*

Key Points:

- The 111-page EO tasks over one dozen departments and agencies with beefing up their AI workforce and issuing guidelines for incorporating the technology into the daily work of the federal government. To increase federal AI talent, the EO directs the Departments of State and Homeland Security to streamline visa processes for AI workers and establishes a campaign to attract global AI talent to the U.S.
- Certain agencies are charged with publishing best practices for the responsible use of AI in the industries they oversee, to include health care, financial services, and education.
- The EO invokes the Defense Production Act to require companies developing "potential dual-use foundation models" to report on the model's performance in AI red-team testing "on an ongoing basis" – and subsequent steps the company is taking to prevent these possibilities from becoming reality. **Reporting requirements kick in on January 28, 2024.**

EO Taskings by Executive Entity (Including Shared Taskings)

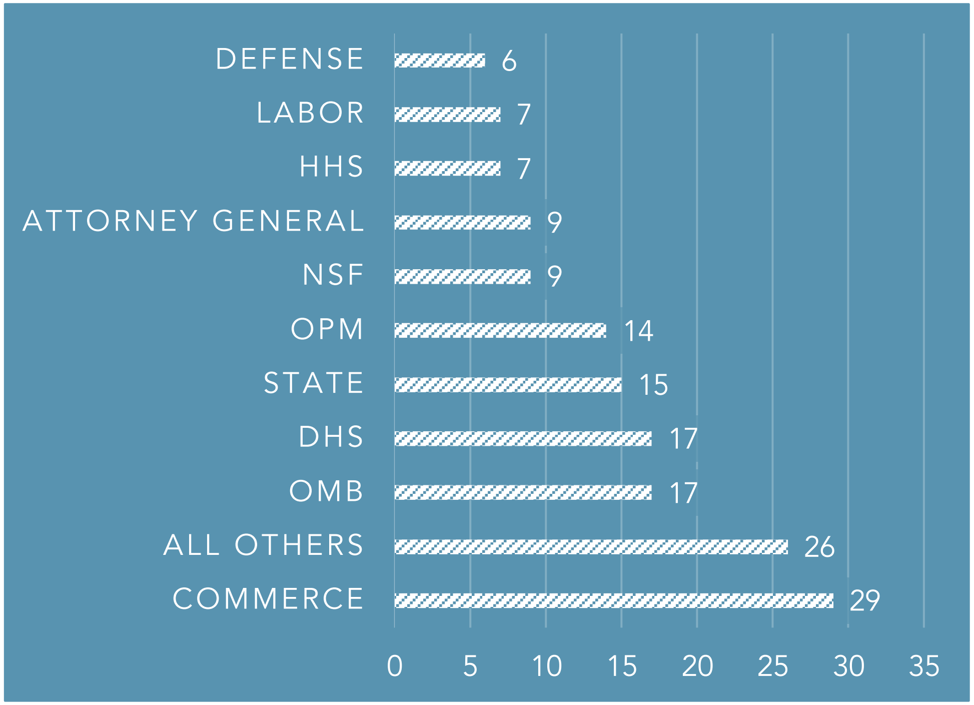

Key Points (continued):

- If your company is building or deploying AI – or even helping others to do so – you are likely on the hook for national security red-teaming. The EO defines "dual-use" with a broad brush, to include any model "that is trained on broad data; generally uses self-supervision; contains at least tens of billions of parameters; is applicable across a wide range of contexts; and that exhibits, or could be easily modified to exhibit, high levels of performance at tasks that pose a serious risk to security, national economic security, national public health or safety, or any combination of those matters."
- Red-teaming exercises must eventually comply with guidelines being developed by NIST over the next nine months. Until then, red teams should – at a minimum – test a model's ability to democratize biological weapons manufacturing, exploit cyber vulnerabilities, influence "real or virtual events," and self-replicate.
- The Administration is not encouraging differential regulation of open versus closed source AI models – for now. The EO directs the Commerce Department to solicit public input and submit policy recommendations to the President on the benefits and risks of dual-use models with widely available weights.
- AI developers aren't the only ones with reporting requirements. U.S. cloud service providers will be required to report "foreign persons" training large AI models that could be used in "malicious cyber-enabled activity."
- The EO also encourages the Federal Trade Commission to ensure fair competition in the AI marketplace, raising questions about how Commissioner Lina Khan's known distrust of "Big Tech" will translate in the AI regulatory space.

What We're Thinking:

- *Something's gotta give.* The EO's scope reflects the fact that Large Language Models (LLMs) and other AI technologies carry implications for virtually every aspect of federal governance and American life. In practice, finite resources and busy political calendars mean the Biden administration will be forced to prioritize certain parts of the EO over others. Thornier provisions that may be pushed to the back burner include those aimed at limiting workforce disruptions and reforming immigration processes.
- *Congress still holds the purse.* The Biden administration cannot implement its landmark EO without the financial backing of Congress, which appears to be headed toward passing stopgap measures to avoid a government shutdown. Neither should we expect significant AI legislation in the near term, with Senate Majority Leader Chuck Schumer (D-NY) saying recently that the chamber is months away from introducing a comprehensive AI bill. Whether a divided Congress can pass AI legislation is an open question.
- *Don't hold your breath.* The bulk of the EO tasks departments and agencies with conducting studies and stakeholder outreach on the promise and perils of AI, only after which they will begin drafting rules and regulations and submitting them for interagency review. Now is the time for companies to engage, offering the government expertise, information, and feedback to shape what's coming. A proactive posture now can prevent frustration and pain later.
- *The AI industry is now a national security partner …* The AI EO is unprecedented in designating the private sector as the first line of defense against threats to U.S. national security. While U.S. tech companies have long advanced national security objectives – including by combatting state-sponsored disinformation and maintaining internet connectivity in Ukraine – the EO marks a formal recognition by the U.S. government that it cannot protect America's interests without the help of its tech industry. Corporate responsibility just took on an expansive new meaning.

What We're Thinking (continued):

- *...It should act like one.* Having been given a stewardship role over Americans' security, it is imperative that companies rise to the occasion. Industry must engage the administration and the broader national security enterprise to align with government priorities and shape expectations about what AI red-teaming can and cannot do. Moreover, AI developers that fall under the EO's definition of "dual-use" must swiftly stand up a robust red-teaming operation. Moving forward, if nefarious actors exploit a model that a company has failed to red-team in good faith, that company faces an existential reputational risk.
- *No one can or should do this alone.* Companies must pair their internal red-teaming capabilities with outside expertise versed in the vast array of national security threats magnified by AI. Red teams staffed by AI engineers alone won't cut it – they will require experts ranging from nuclear chemists and radiologists to ethicists and international arbitration lawyers. And with national security on the line, it is critical that companies share best practices with one another. Prioritizing one's competitive advantage over the security of the nation imperils us all. The time for siloed thinking is over.
- *Beacon can help.* We're helping many of the world's leading companies navigate all of this and we can help you too. We can assist business leaders who are ready to red-team, who need to red-team but don't know how, or those who simply need help tracking and understanding their evolving responsibilities. Our clients rely on our technical expertise and deep national security influence as they seek to understand, shape, and comply with the host of AI rules and requirements that will soon rollout. For more on our AI policy offerings and support, reach out to the **Global Technology Policy Practice** at bgs@bgsdc.com.

More about *The Download*

The Download is for leaders in government, industry, and civil society who operate at the intersection of national security and global business. Produced on an as needed basis, this product concisely explains important issues and places them in context with data and analysis – with the specific intent of enabling action.

More about Beacon Global Strategies

Beacon supports clients across defense and national security policy, geopolitical risk, global technology policy, and federal business development. Founded in 2013, Beacon develops and supports the execution of bespoke strategies to mitigate business risk, drive growth, and navigate an increasingly complex geopolitical environment. Through its bipartisan team and decades of experience, Beacon provides a global perspective to help clients tackle their toughest challenges.

Learn more at bgsdc.com or reach out to bgs@bgsdc.com

insight.
strategy.
action.

BEACON
GLOBAL STRATEGIES