

JUNE 2023

# AI Governance Authority Options Memo

## Overview

Artificial intelligence (AI) should be governed by sector and outcomes. AI regulation also should focus on use cases that are of high consequence, whether beneficial or harmful to individuals and society. This approach requires two types of regulatory authorities, one each for:

1. Promulgating regulation or guidance that identifies high consequence AI use cases; and
2. Applying regulation to high consequence AI.

Existing regulatory bodies have the capacity to develop regulatory frameworks for identifying what constitutes high consequence in their particular domains and developing regulation for AI systems designated as high consequence within those domains. A new governance authority, however, could either directly regulate or support sector-specific regulators by providing guidance<sup>1</sup> so that specific regulatory frameworks are foundationally consistent.

Consequently, a new governance authority must possess the capacity to:

1. Either directly, or through a supporting role, establish the regulatory requirements for identifying high consequence AI use cases across departments/agencies and establish the regulation applicable for AI tools deemed high consequence; and
2. Either directly, or through a supporting role, conduct case-by-case reviews of AI tools submitted for high-consequence designation.

---

<sup>1</sup> The term “guidance” is used in a colloquial sense of conveying priorities, expectations, and standards. It does not indicate a directive with the force of regulation, as can be the case with some published department or agency guidance documents.

In doing so, the new governance authority requires a workforce with both expertise and bandwidth to perform this mission, while managing costs, preventing conflicts of interest, and avoiding actual or perceived regulatory capture.

This memo offers a menu of options for locating and organizing this new governance authority. It first outlines three potential locations for a new governance authority: (1) as a new standalone agency; (2) as a new office within an existing department or agency; and (3) as a new office or unit within the Executive Office of the President (EOP). The memo then provides organizational details for four models, evaluating relevant advantages and disadvantages of a new governance authority with: (1) a full-time permanent staff; (2) a mixed staff of full-time, permanent employees and rotating academic and private sector fellows; (3) a mixed staff of full-time, permanent employees and external contract-based analysts; and (4) a full-time permanent staff supported by a federally funded research and development center (FFRDC).

## Potential Locations for a New Governance Authority

### 1. New Standalone Agency

#### Description

Create a new standalone independent agency responsible for establishing foundational regulatory requirements for identifying high consequence AI use cases across departments/agencies and establishing the regulation applicable for AI tools deemed high consequence. The new agency would liaise with sector-specific departments/agencies in their building on that foundational regulation with their particular expertise. Additionally, the new agency would serve as a central collection point for submission of AI tools requiring review for high-consequence designation. The agency would determine which sector-specific department/agency, or combination of departments/agencies, would be the appropriate authority for full review.

#### Advantages

A standalone agency would offer a locus of expertise and institutional memory within the federal government. Dedicating a new agency to this task would avoid this mission falling through the cracks in an existing department/agency with a range of preexisting responsibilities and obligations. It also could provide a strategic-level perspective that allows the new agency to see and work across the overlapping jurisdictions and equities of sector-specific departments/agencies.

## Disadvantages

A new standalone agency could pull the center of regulatory activity away from sector-specific departments/agencies and toward the new agency. Sector-specific departments/agencies are best positioned to connect expertise in AI with their respective expertise to produce nuanced, practical regulation. This dynamic, therefore, could lead to redundancy and inefficient review processes as the new agency and sector-specific departments/agencies both work on these matters. Furthermore, if sector-specific departments/agencies will continue to operate as the primary regulators, a preceding step with a new agency could yield the sector-specific departments/agencies merely rubber-stamping the first round of work. Simultaneously, there are limits to which a new agency can set a common standard across the whole-of-government; the distinction between executive and independent agencies prevents either type of agency from compelling a position or action from the other type. Voluntary coordination across that divide would be the best achievable modus operandi.

Lastly, a new, independent entity would carry the highest setup cost, as it could not lean on existing administrative infrastructure, such as human resources, information technology or legal services, and would require more time for establishment.

## 2. New Office within an Existing Department or Agency

### Description

Establish a new office within an existing federal department or agency to serve as a central hub across the U.S. Government. For instance, an office within the Department of Commerce or the Federal Trade Commission (FTC) could play the role of facilitating government-wide work by establishing foundational regulatory requirements for identifying high consequence AI use cases across departments/agencies and other regulatory bodies and establishing the regulation applicable for AI tools deemed high consequence. This new office also would serve as the collection point for AI tools submitted for review, forwarding submissions on to the appropriate sector-specific oversight entity for the adjudication.

### Advantages

Creating a new office within an existing federal entity would avoid the costs and time delays of standing up a full new agency. It also could allow the new office to draw on existing authorities and competencies of the host entity both in promulgating relevant regulations and in enforcement.

## Disadvantages

Locating the new governance authority in an existing federal entity could result in legacy structures and imperatives limiting or shaping the new office's mission. There could be host pressures to conform the new mission to pre-existing mandates and capacities. Such dynamics could lead other departments, agencies, and commissions to view the new office not as a neutral facilitator among a whole-of-government effort. Additionally, as with the preceding option, a multi-layer approach of a new office providing an initial review before transmitting a case to the sector-specific oversight authority could result in rubber-stamping concerns. Finally, as with the preceding option, the distinction between executive and independent agencies places a cap on how much a new office in an existing federal entity could enact government-wide guidance.

### 3. New Office or Unit within the EOP

#### Description

Establish a new office or unit within an existing office within the EOP, such as within the Office of Management and Budget (OMB),<sup>2</sup> to serve as a coordinating function among sector-specific oversight entities. The new office or unit would possess a core staff to both establish and maintain the baseline high consequence use case guidance and collect review submissions that would be distributed to sector-specific departments, agencies, and commissions for the rulemaking process.

#### Advantages

An EOP-based governance authority's inability to regulate would self-limit its scope, pushing the regulatory activity to the sector-specific departments, agencies, and commissions with the best relevant expertise. The office or unit also would sit on more central and neutral ground within the interagency process, positioning it to facilitate among the sector-specific departments and agencies with equities in rulemakings and specific case-by-case reviews.

#### Disadvantages

An EOP-based governance authority would face constitutional constraints in that it could not directly regulate; it would have to play a supporting role to sector-specific departments and agencies. Additionally, an EOP-based entity cannot mandate independent agencies and commissions comply with issued guidance.

---

<sup>2</sup> This option could be similar to the placement and operations of the Office of Information and Regulatory Affairs (OIRA) within OMB.

## Possible Organizational Models for a New Governance Authority

In addition to where a new governance authority is located, its organizational structure could take several forms that either focus on building in-house governmental talent or using public-private mechanisms to draw on external expertise.

### 1. Full-time Permanent Staff

#### Description

In any of the three locations outlined above, the new governance authority could rely on a permanent staff of full-time employees (FTEs). Depending on a permanent FTE status, staff could include the use of excepted service billets to offer more competitive packages to attract talent to federal service. Depending on which of the previously outlined locations selected, the FTE staff would either be concentrated in a new agency or office, or distributed across a set of departments, agencies, and commissions.

#### Advantages

This option would build internal government capacity. A long-term permanent staff removed from the private sector could lessen conflicts of interest and mitigate against perceptions of regulatory capture compared to alternatives reliant on private sector expertise. However, it would not wholly eliminate those concerns, as some may see the incentive for staff to depart for the private sector as possibly influencing their work.

#### Disadvantages

This option entails high budgetary costs. It requires budgeting for a larger permanent staff that would need market competitive compensation to attract the required talent. Depending on the location selected, those costs could be centralized or distributed across a range of departments, agencies, and commissions, with tradeoffs in either case. Larger, centralized budgetary costs could daunt authorizers and appropriators, while more distributed costs could involve a broader set of authorizing and appropriating committees. A fee structure can partially defray these costs; however, it must remain balanced so that the new governance authority is not seen as dependent on private actors in a way that would incentive it to cater to the private sector. An in-house FTE staff with less public-private circulation also could cut the new authority off from knowledge of some of the rapid developments occurring in the private sector. Finally, multiple layers of coordination and review can contribute to creating a cumbersome and time-consuming process, as well as result in a rubber-stamping effect by the final sector-specific oversight department, agency, or commission.

## 2. Core Full-time Permanent Staff Supplemented by Rotating Fellows from Private Sector and Academia

### Description

In any of the three previously outlined locations, a new governance authority could employ a small, core leadership and administrative staff, but rely on a rotating cadre of private sector/academic fellows to perform the analytical work. Depending on the selected location, fellows could sit centrally in the new entity or be distributed. For instance, creating a new agency or office within an existing department or agency would align with centralizing fellows in that new entity. Alternatively, an EOP-based governance entity could host rotating fellows or detail them to sector-specific departments, agencies, and commissions to perform analytical work, similar to the White House Fellows model.

### Advantages

This option would lower, but not eliminate, costs. While the new governance authority would have to offer reasonable compensation packages to attract talent, additional motivations (public duty, reputation, etc.) could inspire experts to serve short “tours of duty.” The work’s digital nature also could defray costs as shorter-term fellows would not have to relocate to Washington, D.C. for their service. The model also would support a robust exchange of quality and up-to-date talent from the private sector or academia to perform the work.

### Disadvantages

This option raises concerns of conflicts of interest that may inhibit reviewers and chill private actors from submitting their AI tools for review. Fellows would recuse themselves on submissions from their home institution. But this option would have to mitigate concerns that a fellow returning to a competitor would not abuse insights gained from the intellectual property reviewed. Non-disclosure agreements (NDAs) could help with this challenge, but may deter individuals from becoming fellows if it might curtail their future professional pathways post fellowship. Furthermore, the rotating fellow model comes at a cost to institutional memory as fellows could cycle through on relatively short timelines. Lastly, if fellows are located in a centralized office, rather than detailed to the sector-specific oversight entities, the dynamic contributes to a similar cumbersome process and potential for rubber-stamping that was present in the preceding option.

### 3. Core Full-time Permanent Staff Supplemented by Nongovernmental Contract Experts

#### Description

In any of the three previously outlined locations, a new governance authority could employ a small, core leadership and administrative staff, but rely on external part-time contractors to review submissions of AI tools and provide analysis to inform regulation. Experts would remain at their nongovernmental home institutions and could be paid per review or at an annual rate for a certain number of hours or reviews. As with the preceding option, experts would recuse themselves from home institution submissions and sign NDAs for their work. Additionally, as with the preceding option, depending on the selected location, the external experts could be affiliated with a new agency/office or managed by the new governance entity, but detailed to the specific departments and agencies, or commissions as required.

#### Advantages

This option further reduces costs by relying on external expertise on contract. It defrays costs as experts would not relocate to Washington, DC. It also most directly plugs into private sector and academic talent, requiring the least disruption to these experts' lives in order to incentivize participation.

#### Disadvantages

This option may not deliver a sufficient workforce for the mission. Even with its flexibility, the total number of work hours may not be sufficient to perform quality reviews, particularly in the cases of more complex submissions. There also is a question whether a side contract (or even volunteer) capacity will deliver quality work. Additionally, as with the previous two options, contractors connected to a new agency or central office, rather than sector-specific departments agencies, or commissions, could result in drawn-out process and rubber-stamping concerns. Furthermore, an NDA's constraints may limit desire to participate, or result in too many experts conflicting out. Finally, the contracting model has less of an institutional memory problem compared to rotating fellows, as experts could return over cycles; but it still would involve more turnover than a full-time professional staff.

### 4. Core Full-time Permanent Staff Supplemented by an FFRDC

#### Description

In any of the three previously outlined locations, a new governance authority could employ a small, core leadership and administrative staff, but rely on an FFRDC to supply the necessary analytical capacity.<sup>3</sup> The new governance authority would receive and direct AI tool submissions to the appropriate regulator — whether that is to the governance authority itself as a new agency or office, or to the appropriate sector-specific oversight entity. In each case, the regulator would task the FFRDC to conduct the review and make a recommendation. The FFRDC would deliver that recommendation to the departments, agencies, and commissions for appropriate rulemaking. Funding for the new entity would be limited to the core staff. More funding would be dedicated to the collection of departments, agencies, and commissions that are clients to the FFRDC to provide a baseline budget, which could then be supplemented by fees levied on those submitting AI tools for review.

### Advantages

This option better distributes costs across several actors, focusing less attention on a single point for appropriations. The expert capacity is government adjacent in the FFRDC, but still depends, to a degree, on service-based funding from the departments. Additionally, creating an FFRDC provides a more lasting government-connected set of expertise, supporting institutional memory. It also reduces potential conflicts compared to preceding options that rely heavily on private sector expertise; nor does it preclude using private sector expertise via contract in certain cases.

### Disadvantages

This option does not necessarily reduce costs, though the fee structure could defray it significantly. Legislation would have to harmonize that fee system with a traditional FFRDC ability to manage costs by taking on additional contract work for nongovernmental clients. Permitting contract work for nongovernmental clients allows for more funding sources, but the space for conflicts grows. This option also raises questions of the cost and ability to attract quality talent to the FFRDC. Finally, as is the case with the preceding options, multiple layers of coordination and review can contribute to a cumbersome and time-consuming process, as well as a rubber-stamping effect by the final sector-specific oversight entity.

---

<sup>3</sup> This model draws on the relationship between the Office of Science and Technology Policy and the Institute for Defense Analysis' Science and Technology Policy Institute.