Policing Project
NYU School of Law

NYU School of Law
40 Washington Square South
New York, NY 10012

info@policingproject.org
@policingproject
212.992.6950

December 5, 2023

Clare Martorana
U.S. Federal Chief Information Officer
Office of the Federal Chief Information Officer
Office of Management and Budget
725 17th Street, N.W.
Ste. 50001
Washington, D.C. 20503

Submitted electronically via www.regulations.gov

Re: Request for Comments on Draft Memorandum – *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (AI)*, OMB-2023-0020, Comments of Policing Project at New York University School of Law

The Policing Project is a nonpartisan center at New York University School of Law dedicated to promoting public safety through transparency, equity, and democratic engagement. We submit this comment in response to the Office of Management and Budget's (OMB) Request for Comments (RFC) on its draft memorandum titled "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence" or "the Guidance."

Our comment focuses on the application of the Guidance to law enforcement and makes six key points:

(1) If adopted as drafted, the Guidance would provide much needed transparency and accountability over federal law enforcement agencies' use of AI and would serve as a model framework for responsible and equitable use of this technology at all levels of government.

(2) The minimum practices identified for safety- and rights-impacting AI would establish essential safeguards, but the waivers provision should be revised to ensure waivers for law enforcement agency purposes are limited and subject to oversight.

(3) To ensure transparency and accountability over law enforcement agency compliance with the required minimum practices, the Guidance should require agencies to publicly document implementation in the AI use case inventory and should narrowly apply any law enforcement exemptions to this public reporting.

(4) The list of presumptive rights-impacting AI purposes should be revised to include law enforcement use of robots and drones and to ensure application to future rights-impacting law enforcement tools.

(5) For agencies with law enforcement components, the qualifications for the Chief AI Officer should include expertise in the protection of civil rights and civil liberties.

(6) OMB should encourage federal agencies to review the extent to which they can promote adoption of the Guidance by state and local law enforcement agencies that receive federal funding.

Where applicable, we note in the text of our responses the number of the RFC question to which we are responding.

I. Background on the Policing Project

The Policing Project is a nonpartisan, nonprofit center at NYU School of Law. We conduct research and also do work on the ground all over the country, with policing agencies and with the communities they serve, with the federal, state, and local governments, and with technology venders, to promote democratically-accountable and equitable policing. Our mission is to promote "front-end, accountability," which means the public has a voice in setting transparent, ethical, and effective policies and practices *before* policing agencies or government act.

One of our primary focus areas is the use of emerging technologies by policing agencies. Increasingly, this means AI-powered tools and systems. We have spent countless hours researching and discussing AI-powered policing technologies with racial justice and civil liberties advocates, technologists, and policing agencies themselves. We have developed numerous resources dedicated to promoting the sound governance of AI-powered public safety technologies, from a model statute regulating automated license plate readers (ALPRs) to regulatory frameworks for police use of facial recognition technology (FRT) and robots.[1] From 2019–2022, we staffed the Axon AI Ethics Board, an independent review board that guided and advised Axon, the developer of TASERS and the country's largest producer of body-worn cameras, around ethical issues related to the development and deployment of AI-driven policing technologies.

Our Comment draws on our deep study and past work on AI policing technology and our fundamental belief that adoption and use of these tools must be guided by democratic legitimacy, a commitment to racial justice, and an imperative to minimize harm.

II. The Guidance takes a major step toward establishing the federal government as a model of sound AI governance

As President Biden's AI Executive Order noted, AI "holds extraordinary potential for both promise and peril."[2] Perhaps nowhere is this more true than with regard to law enforcement use of AI. Although there may be real public safety benefits to law enforcement adoption of this emerging technology, sound governance is absolutely essential to mitigate the real risks to liberty, equity, and racial justice that are unique to policing agency use.

Unfortunately, as is the case with many emerging technologies, law enforcement agencies have rushed ahead to use AI with little transparency and even less in the way of regulation or authoritative guidance on responsible use. Federal law enforcement is no exception – instead of serving as a model of accountable adoption and use of advanced technologies, it has deployed AI systems without adequate privacy protections or even any policy or training in place.[3] This unfettered use in a context as high-risk as law enforcement already has led to real harms to the public's civil rights and liberties – from false arrests to

[1] Regulating Use of Technology in Policing, Policing Project, https://www.policingproject.org/policing-tech-landing.
[2] Exec. Order No. 14110, 88 FR 75191 (2023), Sec. 1.
[3] *See, e.g.*, U.S. Gov't Accountability Off., GAO-23-105607, Facial Recognition Services: Federal Law Enforcement Agencies Should Take Actions to Implement Training, and Polices for Civil Liberties (Sept. 2023), https://www.gao.gov/assets/gao-23-105607.pdf (finding multiple federal law enforcement agencies use facial recognition technology without any policies or guidance in place to address civil rights and civil liberties, fail to fulfill privacy requirements, and do not require staff training on these systems).

excessive use of force.[4] What is needed instead of this rush-to-deploy model is rigorous study, stepwise adoption, public accounting of these technologies' benefits and costs, enforceable safeguards to mitigate risks to civil rights, racial justice, and civil liberties, and a commitment to abandoning systems and tools that do not advance public safety and equity. In other words, what is needed is a model of sound governance of AI.

This Guidance would represent a giant step in the right direction. It is precisely the sort of regulatory framework that long should have been in place around law enforcement use of emerging technologies, whether AI-driven or not. In the sections that follow, we highlight several ways in which its provisions would benefit from revision. But first we want to emphasize the critical protections it already establishes:

- It sets out a comprehensive list of AI applications presumed to be safety- and rights-impacting that explicitly includes key law enforcement use cases (pages 11-12).

- It establishes an appropriate baseline of minimum practices for safety- and rights-impacting AI that applies to federal law enforcement agencies, containing meaningful and enforceable safeguards that address the entire AI lifecycle (pages 15-18):

  ⇒ Pre-deployment, it requires agencies to conduct impact assessments, including cost-benefit analysis that requires demonstrating actual benefit as a condition to adoption.
  ⇒ It mandates real-world testing and independent evaluation of AI. As a result, the public would no longer need to rely on anecdotal evidence of AI's impact, but would have access to real data about whether or how well or poorly these systems work.
  ⇒ By requiring ongoing monitoring, it recognizes that compliance cannot be a static practice for these ever-changing systems.

- It also recognizes the need for additional requirements to ensure rights-impacting AI advances equity, dignity, and fairness and incorporates feedback from affected groups (pages 18-19).

  We were especially heartened by the call for engagement with affected groups. Our work in the policing context repeatedly has demonstrated that engaging with impacted communities is both necessary and practicable. For example, we run a grassroots program in Chicago that connects residents with patrol officers to provide feedback on neighborhood safety priorities.[5] We served as consultants to the federal monitor for the Cleveland consent decree.[6] We also have conducted community-based research on non-police alternate first response in cities across the country.[7] In all of these cases, and more, community feedback has proven invaluable and vital to improving public safety policy.

Crucially, these substantive requirements are enforceable. Agency compliance is mandatory and must be documented; use must stop if the AI system does not satisfy these conditions. This is what meaningful transparency and accountability look like.

[4] *E.g.*, Kashmir Hill, Eight Months Pregnant and Arrested After False Facial Recognition Match, N.Y. Times (Aug. 6, 2023), https://www.nytimes.com/2023/08/06/business/facial-recognition-false-arrest.html; Vanessa Romo, No Charges for Colorado Officers Who Held Black Children at Gunpoint, NPR (Jan. 8, 2021), https://www.npr.org/2021/01/08/955165485/no-charges-for-colorado-officers-who-held-black-children-at-gunpoint.
[5] Chicago Neighborhood Policing Initiative, https://chicagonpi.flywheelsites.com/mission.
[6] Cleveland Police Monitoring Team Overview, Policing Project, https://www.policingproject.org/cleveland-cpop-report.
[7] Reimagining Public Safety, Policing Project, https://www.safetyreimagined.org/about-rps.

Although we have suggestions for strengthening the Guidance, we fully appreciate the steps this Guidance represents to establish the federal government as a model of sound AI governance. We can only hope this model diffuses to state and local law enforcement agencies.

III. Strengthen the minimum practices for safety- and rights-impacting AI by adding oversight and narrowing the waivers provision (Question 6)

The Guidance's current minimum practices for safety- and rights-impacting AI would establish an impressive baseline of transparency and accountability for federal agency use of AI, but as applied to high-risk law enforcement use cases, the waivers provision must be revised to reduce discretion and ensure waivers only are granted in the presumably extremely rare circumstances in which they are warranted.

By granting sole authority to each agency's Chief AI Officer (CAIO) to waive requirements even for high-risk applications, the Guidance gives this official too much discretion. To ensure accountable decision-making, the Guidance should require additional sign off for waiver requests, especially for high-risk law enforcement applications of AI. At least in law enforcement contexts, the system-specific risk assessment required by the waivers provision should include a determination of the risk level presented by the use. If this determination reveals that the risk is anything other than low, then a CAIO waiver request should be subject to additional review, by, for example, the agency's civil rights office or OMB.[8] Because waivers outside of low-risk applications should be the exception not the rule, the administrative burden of this external check should be minimal.

OMB must narrow the qualifying conditions for issuing waivers, especially as applied to high-risk law enforcement applications. Currently, the Guidance permits waivers for any situation that represents an "unacceptable impediment" to "critical agency operations" without defining or limiting these key terms.[9] In our experience, law enforcement tends to think everything it does is "critical," and any safeguards are "unacceptable." We suggest revising this provision in two ways. First, the Guidance should constrict the qualifying criteria for high-risk law enforcement applications to permit waivers only when a requirement "*would create an insurmountable impediment to critical agencies operations and the system-specific risk assessment demonstrates that the risks of not fulfilling the particular minimum requirement do not outweigh the benefits*." In addition, OMB should require public documentation of the waiver determination to ensure transparent and accountable use of exemptions.

IV. Strengthen the AI use case inventory requirement (Question 8)

(1) Require agencies to publicly document implementation via the use case inventory

To ensure transparency and accountability over agencies' minimum practices implementation, OMB should require agencies to publicly document compliance in the AI use case inventory. Although the Guidance currently provides several options for how agencies might document their implementation, it does not require any particular method (page 13). This choose-your-own adventure approach to documentation will make it difficult for the public to assess compliance across agencies. Instead, OMB should mandate public documentation of implementation as part of agencies' annual AI use case inventory. Doing so will ensure that implementation of the key transparency and accountability measures set forth by the minimum practices are made public and will establish the use case inventory as a comprehensive public record of agencies' AI use.

[8] Civil Rights Offices of Federal Agencies, Dep't of Justice, https://www.justice.gov/crt/fcs/Agency-OCR-Offices.
[9] OMB AI Guidance, Sec. 5(c)(iii).

(2) Narrowly apply any law enforcement exemptions for safety- or rights-impacting AI disclosures in the use case inventory

The Guidance appears to take important steps to build on Executive Order 13960's annual AI use case inventory by requiring agencies to "report additional detail on how they are using safety-impacting and rights-impacting AI."[10] This is essential because federal agency compliance with EO 13960's AI use case inventory has been inconsistent at best and entirely deficient at worst.[11] And federal law enforcement agencies have been some of the biggest offenders on this score. For example, DOJ's most recent disclosures consisted of one page of information, listing a single use of AI by the FBI for a "threat intake processing system" to analyze crime tips.[12] This single page contains no disclosure related to the FBI's use of facial recognition technology despite the fact that the Bureau has been using this AI-powered technology for criminal investigations for almost a decade.[13] And there are zero disclosures for several other DOJ law enforcement agencies' use of facial recognition – from DEA to ATF to the U.S. Marshals – even though a recent Government Accountability Office audit reported significant use of this technology by each of these agencies. To effect the clear transparency and accountability goals of the Guidance, OMB should require that federal law enforcement agencies publicly disclose their use of safety- or right-impacting AI and their compliance with the minimum practices in the use case inventory.

And to ensure meaningful compliance by law enforcement agencies with AI use case disclosures—rather than the lackluster compliance we've seen to date—OMB should revise the equivocal language used to describe disclosure requirements and ensure law enforcement exemptions are interpreted extremely narrowly. For example, in a number of places, the draft Guidance establishes a "practicability" standard for whether or not to require use case disclosure.[14] Similarly, in several places, the Guidance specifically exempts "sensitive law enforcement" information.[15] This language creates too much leeway for law enforcement agencies to disclose the bare minimum or even entirely avoid disclosure. In our experience, law enforcement often employs a broad interpretation of what information is "sensitive" and a narrow interpretation of what is "practicable" to disclose publicly. As noted above, past practice specifically related to federal law enforcement compliance with the AI use case inventory supports this. But in truth, there should be almost no instance in which it is impractical to disclose a description of the use and capabilities of an AI system – including in the law enforcement context.[16] A use case inventory is a non-case-specific

[10] Shalanda D. Young, Proposed Memorandum for the Heads of Executive Departments and Agencies, Office of Management and Budget, Sec. 3(a)(iv) [hereinafter "OMB AI Guidance"].

[11] *See, e.g.*, Christie Lawrence et al., Implementation Challenges to Three Pillars of America's AI Strategy, Stanford HAI (Dec. 2022), https://hai.stanford.edu/sites/default/files/2022-12/HAIRegLab%20White%20Paper%20-%20Implementation%20Challenges%20to%20Three%20Pillars%20of%20America%E2%80%99s%20AI%20Strategy.pdf (surveying federal agency compliance with the AI use case inventory and finding it "problematic" with 78% of federal agencies failing to publish an inventory); Bowman Cooper, Like Looking for a Need in an AI-Stack, Center for Democracy & Technology (July 21, 2023), https://cdt.org/insights/like-looking-for-a-needle-in-an-ai-stack/ (reviewing federal agency AI use case inventory disclosures and finding the "information provided by each agency is inconsistent and unclear, making it difficult for the public to understand exactly how the use of AI impacts them").

[12] AI Use Case Inventory Submission on Open Data, U.S. Dep't of Justice, https://www.justice.gov/open/page/file/1517316/download.

[13] U.S. Gov't Accountability Off., GAO-19-579T, Face Recognition Technology: DOJ and FBI Have Taken Some Actions in Response to GAO Recommendations to Ensure Privacy and Accuracy, But Additional Work Remains 2 (June 2019), https://www.gao.gov/assets/gao-19-579t.pdf.

[14] OMB AI Guidance, Sec. 3(a)(iv) n.8; Sec. 5(c)(iv)(H).

[15] *Id.*

[16] *See* Barry Friedman & Maria Ponomarenko, Democratic Policing, 90 N.Y.U. L. Rev. 1827, 1884-85 (2015) (discussing law enforcement objections to transparency and observing that "the need for secrecy is not nearly as acute as it may seem. . . ." and distinguishing between details related to specific investigations, which have a rightful claim to secrecy, and details related to use of tools or techniques which "can be made public and publicly debated without undermining law enforcement interests").

accounting of system functionality. It is not a revelation of tactical plans or sensitive information about a particular investigation. Therefore, substantive public disclosure of the information required by the minimum practices should be the default.

We make two final recommendations to ensure the use case inventory is useful and used. First OMB should ensure these inventories contain sufficient detail to enable the public to evaluate the impact of an agency's AI use. This means agencies should be required to disclose the vendor(s) for any safety- or rights-impacting AI. We also join the recent recommendations by a coalition of civil rights and tech policy groups that the AI use can inventory include additional detail about demographic information on AI system use and outcomes.[17] Second, we recommend that OMB direct agencies institute a public complaint process that individuals or organizations can use if they believe an agency has not sufficiently documented, reported, or implemented a required minimum practice.

V. Supplement the law enforcement purposes presumed to be rights-impacting (Question 5)

The Guidance's list of presumed safety- and rights-impacting purposes is a strong start. The explicit inclusions of law enforcement use of facial recognition technology and other biometric tools, person-based and place-based predictive policing, and license plate readers are especially important. As noted above, these AI tools already have been the source of false arrests, unnecessary uses of force, and due process violations.[18]

To ensure the Guidance adequately captures high-risk law enforcement applications, we recommend two additions to the list of rights-impact AI. First, OMB should amend the rights-impacting list to include high-risk law enforcement applications of aerial vehicles (or drones) and robots. Although robotics (though not drones) are listed under "safety-impacting" purposes, the risks presented by law enforcement use of these tools stretch beyond the physical risks that this category seeks to protect. Drones and robots both have significant AI-powered surveillance capabilities that can present serious risks to the public's civil rights and liberties from privacy to free expression to racial discrimination. For example, policing agencies—including federal agencies—have used drones to surveil peaceful protests and to disproportionately monitor communities of color.[19] And police use of these tools is only increasing.[20] Both tools merit explicit inclusion in the "rights-impacting" list along with the AI-powered surveillance tools already present.

Second, because policing technology always is evolving, we recommend adding language to subsection 5(b)(ii)(B) that focuses on the capability of concern rather than just the specific device. Specifically, we would add a clause to cover law enforcement use of "any AI to identify criminal suspects and track or monitor individuals." Today, we rightly are concerned with law enforcement's ability to use specific AI-powered surveillance tools like facial recognition technology or license plate readers. But there is no predicting which new tools are coming down the pike tomorrow that will serve similar purposes. It is essential that the Guidance is flexible enough to cover future developments.

[17] Algorithmic Justice League et al., Letter to the President re: Advancing Anti-Discrimination Testing in Artificial Intelligence Executive Order (Aug. 7, 2023), https://www.upturn.org/work/letter-to-the-biden-harris-administration-on-their-forthcoming-ai-executive.

[18] *See, e.g.*, *supra* note 4.

[19] Policing Project at NYU Law, Civil Rights and Civil Liberties Audit of Baltimore's Aerial Investigation Research (AIR) Program (Nov. 2020), https://static1.squarespace.com/static/58a33e881b631bc60d4f8b31/t/5fc290577acac6192a142d61/1606586458141/AIR+Program+Audit+Report+vFINAL+%28reduced%29.pdf; Jason Koebler et al., Customs and Border Protection is Flying a Predator Drone Over Minneapolis, Vice (May 29, 2020), https://www.vice.com/en/article/5dzbe3/customs-and-border-protection-predator-drone-minneapolis-george-floyd.

[20] Ese Olumhense, The Tech at 'Cop Con,' The Markup (Nov. 3, 2023), https://themarkup.org/news/2023/11/03/the-tech-at-cop-con.

VI. Ensure CAIOs have civil rights and civil liberties expertise (Question 1)

OMB should amend the Guidance to require that the CAIO in agencies with law enforcement components have expertise in civil rights and civil liberties. When it comes to a high-risk context like law enforcement, this official must be qualified to assess the unique ways in which AI in policing can amplify harms to racial justice and individuals' civil rights and liberties. In policing, AI's oft-touted productivity gains have the potential to translate to more state-sponsored surveillance and over-enforcement in certain communities, often communities of color. To guard against these unique risks, we recommend the Guidance explicitly include expertise in civil rights and civil liberties as a necessary qualification for agencies with law enforcement components.

VII. Support sound AI governance through grantmaking (Question 3)

If implemented as currently drafted, the Guidance eventually should make the federal government a model of AI governance. This is no small accomplishment, but limiting its scope to federal agencies would represent a missed opportunity – especially when it comes to law enforcement. Law enforcement in particular remains largely a creature of the states. There are over 18,000 law enforcement agencies at the state and local levels compared to 83 at the federal.[21] And, relevant to this discussion, two things are true about many of these state and local law enforcement agencies: (1) many of them have been using AI systems ranging from facial recognition to ALPRs to predictive policing for years without any regulation or guidance in place; and (2) many receive federal funding to support their operations.

The Guidance should encourage federal agencies to use their grantmaking authority to apply these principles to state and local law enforcement agencies that receive federal funding and support. The federal government regularly uses funding as a carrot to achieve state and local compliance with rules like these. Just last year, President Biden's Executive Order 14074 on policing required the Attorney General and Secretary of Homeland Security to review and exercise their authority to award federal grants and provide support to state and local law enforcement based on adoption of the policies established in that Order.[22] OMB should direct agencies to engage in a similar review process to determine how the Guidance's provisions might be extendable to state and local law enforcement. By supplementing the Guidance with this requirement, OMB has an opportunity to ensure that its sound AI governance practices will apply to some of the most frequent users of safety and rights-impacting AI: state and local law enforcement agencies.

Thank you for the opportunity to comment.

Respectfully submitted,

Barry Friedman
Founder and Faculty Director

Max Isaacs
Senior Staff Attorney

Katie Kinsey
Chief of Staff

[21] Duren Banks et al., National Sources of Law Enforcement Employment Data, Office of Justice Programs, U.S. Dep't of Justice (Apr. 2016), https://bjs.ojp.gov/content/pub/pdf/nsleed.pdf; Connor Brooks, Federal Law Enforcement Officers, 2016 – Statistical Tables, Office of Justice Programs, U.S. Dep't of Justice (Oct. 2019), https://bjs.ojp.gov/content/pub/pdf/fleo16st.pdf.

[22] *See* Exec. Order No. 14074, 87 FR 32945 (2020), Sec. 20.