Palantir

The Honorable Lori Trahan
2233 Rayburn House Office Building
Washington, DC 20515

Greetings,

We welcome the opportunity to provide feedback on ongoing efforts to reform the Privacy Act of 1974. Respect for individual liberties is central to the American way of life, and privacy rights have always been a crucial part of these freedoms. That is why a commitment to safeguarding privacy shaped Palantir's founding more than two decades ago, and it's why privacy has always been a defining and fundamental characteristic of our work.

We approached our response to your request for input with respect for the delicate balance that must be struck between ensuring the foundational systems of our government are equipped to best serve the American people efficiently and effectively, while at the same time honoring the privacy rights that are foundational to our nation's system of government and way of life. Our recommendations are based on insights gathered over 20 years of experience building technology to improve institutional mission outcomes while upholding American values in the use of our software products, including AI enablement tools and platforms.

We stand ready to assist in future efforts to advance this initiative, and we look forward to working together.

1. General questions.

    a. **What are your biggest concerns with the federal government's collection, maintenance, use, or dissemination of personal information?**

        - We believe the government plays a critical role in safeguarding the personal data and upholding the privacy rights of all Americans. However, we share the concerns of many that large-scale data collection by the federal government poses additional risks to the fundamental values that define Western liberal democracy.

        - While new, ever-advancing methods and technologies that enable data management and data analysis play an essential role in supporting the efficient functioning of our government institutions, they also can become incompatible (even in the absence of malign intent) with the outdated organizational paradigms that underlie legacy privacy legislation like the Privacy Act of 1974.

        - Modern processing and analytics technology has rendered former (and in many places, current) models of data storage and management obsolete. In the past, fragmented, poorly secured, and unknown data systems may have served as a partial mechanism — intentional or otherwise — for upholding data privacy by virtue of their systemic inefficiency, achieving privacy through obscurity. The lessons of the terrorist attacks of September 11, 2001, however, demonstrated how such fragmentation can also pose grave national security risks. But in breaking down information sharing barriers that may have served as an unnecessary impediment to the lawful and desirable sharing of select, mission-critical elements of data, our government institutions also needed to simultaneously address the privacy risks that were created or exacerbated by the shift away from a reliance on privacy by obscurity. Investing in data privacy reform and tools that can safeguard data privacy in modern contexts while also enabling the delivery of institutional missions and mandates therefore becomes a central challenge of informational privacy in the modern age.

b. **How should the federal government balance securing privacy with other priorities, especially promoting security, reducing waste, fraud, and abuse, and improving service delivery (for example, through the use of a public identity verification platform)?**

- Too often discussions on this important theme begin with a zero-sum assumption, i.e., that improvements in security, waste reduction, fraud, abuse, service delivery, etc. *must* come at the expense of undermining privacy interests. This starting point should be treated as a false dichotomy and operational demand should focus - to the greatest extent possible - on rejecting a tradeoff between offering greater data privacy safeguards and reducing waste, fraud and abuse, or improving the essential services of the federal government. In our experience, these seemingly competing considerations can and should be treated as complementary, not exclusive - with each improving as a function of better design of systems of data management that enhance security controls while making data more accessible in agency-specific demands and contexts. The Computer Matching Amendments to the Privacy Act, while significantly limited and obsolete today, illustrate a way to balance the search for fraud, waste, and abuse with the need for due process rights.

- As one industry example, Palantir Technologies' success working in some of the most highly regulated applications, industries, sectors, and markets around the world serves as a concrete demonstration that these considerations can be thoughtfully navigated in a manner that shifts the presumption away from zero-sum to mutually optimized outcomes. This is one of the reasons why Palantir has historically supported - and continues to support - robust and evolving legal and regulatory safeguards that are positioned to adapt to ever-changing threats to data privacy. We know from experience that the dual objectives constitute a legitimate and largely achievable engineering challenge that technology providers *should* be made to treat as a foremost optimization goal. For example, our tools make it practical to provide for data use and sharing for specified purposes while preventing or discouraging use and sharing for other purposes. This mindset is central to the principles and technologies that comprise Palantir's focus on Privacy and Civil Liberties Engineering.

c. **What are the unique privacy risks created by the government's use of artificial intelligence? How can Congress mitigate those risks?**

- The use of AI – and especially Generative AI (GenAI), including Large Language Models (LLMs) – introduces novel privacy risks. For example, the vast training data for modern GenAI systems will surely contain personal data, so a GenAI system may respond to an unrelated user prompt with personal data "memorized" from training.

- The core privacy risks of AI, however, are not unique to the recent advancements in AI systems. GenAI may exacerbate some legacy data privacy challenges and introduce other novel issues, but many of the attendant risks relate to enduring challenges of data management and digital infrastructure. In that vein, we have observed that foundational investments in tools and capabilities that support core data protection and privacy considerations — what we've described in a separate article as 'basic' Privacy-Enhancing Technologies (PETs) — are necessary for governing any AI/ML system in consequential settings.

- Nor are the risks of AI systems purely related to privacy. As AI systems are used to assist – if not to automate – some forms of decision-making and agentic activity, it is imperative that the federal government is mindful of the risks of infringements on individual liberties, including procedural requirements such as due process in a law enforcement or administrative context. Here too, the federal government can use technology to implement guardrails on AI systems, including robust testing & evaluation, monitoring, effective human oversight and control interfaces, and other AI Governance best practices. For example, privacy and civil liberties impacting risks in the use of AI systems can be addressed through the following approaches:

  - Integrating trusted 'ground truth' data sources and institutionally calibrated Ontologies (i.e., the software transposed model of data and logic that organizations treat as canonical for their disciplines and operations) can serve as critical capabilities for mitigating model 'hallucination' risks.

  - Structured guidelines for directing when GenAI-dependent systems should 'hand off' specific classes of operations to better suited logic-based tools (e.g., calculators for mathematical operations).

- Chain-of-Though (CoT) prompting to provide better interpretability and explainability as to the underlying methodology by which GenAI-based systems produce specific outcomes such as recommendations, decisions, or other outputs.

- Unit testing, perturbation testing, LLM-as-a judge paradigm, benchmark evaluations, etc. represent only a narrow subset of broad field of AI testing and evaluation methodologies that supplement or go well beyond red-teaming approaches often regarded as the limited standard for assessing GenAI reliability. (See our series on responsible AI in practice for additional details on these and other approaches to GenAI risk mitigation in practical, operational contexts.)

  - We have long advocated for the federal government to promote the effective and conscientious use of AI, especially when AI challenges the privacy and civil liberties of Americans. Moreover, we know from first-hand experience that AI innovation can not only be pursued in a way that is compatible with security, privacy, data protection, and related fundamental rights, but AI is also often most effective when it supports rights-protective outcomes. In this vein, we reiterate our earlier recommendation on the importance of – as a starting point for addressing the risks of AI use by the government – renouncing the presumption of an unavoidable tradeoff between offering greater data privacy safeguards and employing AI to carry out government agency missions. By explicitly rejecting this false dichotomy, Congress can encourage AI developers and deployers to pursue innovations and innovative uses that minimize AI risks.

d. **How can the federal government most effectively leverage privacy-enhancing technologies (PETs)?**

  - The federal government should invest in the promotion, adoption, and application of privacy-enhancing technologies (PETs), which protect privacy, civil liberties, and related fundamental rights.

  - Utilization of PETs alone, though, is not enough. PETs can be narrowly constrained to more 'exotic' techniques such as Differential Privacy and Homomorphic Encryption, which are rooted in statistical or mathematical techniques with nuanced privacy guarantees. While we certainly recommend the application of such PETs where appropriate, they should be used in conjunction with more fundamental architecture and 'basic' PETs that advances privacy, data protection, and civil liberties, which we detail further in response to questions below.

2. **Modernizing the Privacy Act of 1974.**

a**. Definitions.**

i. **How can the Privacy Act's core definitions, including "individual, record," and "system of records" be modernized to reflect the federal government's current information management practices? How should these definitions take into account the Office of Management and Budget's incorporation of the term personally identifiable information into recent guidance, including OMB Circular A-130?**

- When the Privacy Act was signed into law more than five decades ago, its framework focused on the "record" - a discrete collection of information on an individual - as the fundamental unit of interest for a privacy framework. The advent of complex data structures, metadata constructs, granular data management, and myriad other data innovations though, has rendered the very concept of a "record" obsolete. A revised Privacy Act should abandon the record concept and focus instead on PII as a general and perhaps more fluid concept that applies as much to the context of data usage as to the atomic unit of analysis. This shift would help to address the complexities of data flows and enable more robust protections for individuals in an era of pervasive data collection and processing.

ii. **Should the Privacy Act address privacy concerns faced by organizations, including businesses and nonprofits? If so, how?**

- The challenge of passing comprehensive consumer privacy legislative reform has proved daunting enough as its own matter. In the interest of focus and tractability, this effort may be better served with a still-limited (but sizable and important) focus on reforming the Privacy Act to address government agencies and their activities. Non-governmental privacy issues should certainly be addressed, but the best vehicle for this is separate legislation. We have offered thoughts on such legislation in the past.

b. **Disclosure requirements.**

i. **Should the law's provision that requires agencies to only maintain "only such information about an individual as is relevant and necessary to accomplish a purpose of the agency," or data minimization provision, be strengthened? If so, how?**

- We recommend against including provisions that are dependent on broad standards that can be interpreted liberally, such as “relevance.” Instead, we feel that while data minimization should be required, it should also be accompanied by agency-specific requirements for granular data access controls and selective revelation approaches to reinforce proportionality, minimization, and use limitations most applicable to the context of data usage. Modern technology has made it feasible to help ensure that only certain individuals who serve particular functions have access to sensitive data within well-defined bounding conditions (e.g., temporal, purpose-based, classification-based, etc.). One broad middle ground here is to combine general statutory standards with directions to OMB to issue more detailed guidance on applying the standards in context.

ii. **How can the requirements regarding individuals' access to and ability to amend their information be improved? Furthermore, how can agencies' implementation of this requirement be modernized?**

- While we agree an individual’s access to their private data is important, given our expertise we feel this would not be an appropriate place to offer advice. We would be happy to address it as part of future discussions of privacy rights, though.

iii. **Should Congress consider requiring that agencies provide individuals a "right to be deleted," in which individuals may request that an agency delete their records? If so, how should providing such a right be balanced against other governmental interests, including promoting national security, improving service delivery, and reducing waste, fraud, and abuse?**

- While we support a broad right to deletion, we note that there are significant challenges associated with intended forms of appropriate deletion. For example, deletion may run a spectrum from “soft deletion” (rendering data difficult to recover or link) to “hard deletion” (rendering data — or even the hardware storing the data - permanently destroyed, irretrievable, or unlinkable). Different deletion techniques may provide varying degrees of assurance that data could not be reconstituted or resurrected for unwarranted purposes. But deletion approaches with firmer destruction of data guarantees may also carry unintended consequences. We are concerned that data deletion may weaken due process protections and accountability measures by undermining audit trails, as well as compromise the integrity of existing datasets, creating conflicts with legal or regulatory requirements with specific data retention mandates and preventing judicial or congressional review. We stress the need to reconcile the individual’s right to deletion with the demands of transparent and accountable governance. We recommend pursuing measures such as anonymization and comprehensive, rigorous access controls to preserve appropriate privacy protections for individuals while keeping agencies accountable for their actions.

c. **Written consent requirement.**

- While we recognize that issues relevant to written consent are important to address, we do not feel that advising on these particular issues fall within the scope of matters on which it would be appropriate for Palantir to offer its advice, given our work as a software company is not particularly relevant.

d. **Exceptions to the written consent requirement.**

i. **5 U.S.C. §552a(b)(1) provides an exception, known as the need to know exception, "to those officers and employees of the agency which maintains the record, who have a need for the record in the performance of their duties."**

1. **Should the need to know exception be narrowed, clarified, or otherwise modified? If yes, how?**

- The Need to Know exception is a good foundation for access controls, but it should be bolstered by measures that add contextual limitations while ensuring the varied and complex missions of individual agencies are not compromised. For example, need to know could be further parameterized around temporal or micro-purpose-based considerations that could then be transposed into rigorous technical access control measures.

2. **How can Congress improve the transparency around agencies' granting of need to know exceptions?**

- One technique for improving transparency with respect to need to know exceptions is utilizing a more refined paradigm for data access, such as Purpose Based Access Controls (PBACs). PBACs grant access to individuals based on specific purposes that are defined narrowly and apply only to certain portions of datasets. The utility of PBACs in a government agency use case is clear: Because purposes are set by data governance teams to contain data specifically scoped to help the user meet their goal, granting access to entire datasets for the completion of specific tasks is no longer necessary. Utilizing PBACs is one solution to address the problems inherent in navigating access controls for sensitive datasets. Other examples draw upon other methods of parameterizing data access, such as temporal constraints.

3. **Should there be limits on those "officers and employees" who can receive need to know exceptions? If so, should this access differ by type of federal employee (for example, political appointees vs. civil servants)? Furthermore, should access differ by the relative risk or scope of a particular system of records?**

    - Determinations of legitimate system and data access, including for exceptional circumstances, are better left to government agencies and their oversight functions (including relevant Congressional authorities). In making these determinations, however, policymakers should be aware of the existence and proven capabilities of state-of-the-art data management systems that provide highly granular access controls to enable dynamic data sharing decisions at extremely fine-grained levels of detail. Data access decisions can be managed on an individual by individual, data point by data point basis, to ensure extremely precise data minimization. Managing data by role, risk, purpose, classification, timing, or other salient factors is manageable from a technical perspective.

ii. **5 U.S.C. §552a(b)(3) provides an exception for an established routine use identified in the system of records notice (SORN) that has been published in the Federal Register."**

1. **Should the definition of "routine use" as "a purpose which is compatible with the purpose for which [the information] was collected" be narrowed, clarified, or otherwise modified? If yes, how?**

    - We acknowledge the importance of this issue. While it is not a core priority for us to provide specific advice, we are happy to include it as part of future discussions surrounding these legislative efforts.

2. **Is the information included in the SORN, and the medium of publication via the Federal Register, sufficiently effective to notify individuals about the use of their information? Could the SORN and the process by which it is made public be improved?**

    - SORNs are more than likely too broad to provide meaningful notice. At the same time, the process to update and amend SORNs (via the APA) is lengthy, cumbersome, and may be subject to ambiguity and decision on the appropriate triggers for significant systems changes warranting SORN revisions. An updated Privacy Act should provide greater transparency of agency systems and uses, but it should also provide for an expedited process for making this information available, as well as clearer and more actionable guidance on the form, severity, and significance of systems changes that demand corresponding refinements to SORNs.

    - The SORN process may also benefit from refined guidance on the level of notification and detail mandated to provide sufficient notification on not just in-system developments, but also inter-system or -agency data developments, such as data sharing between agencies and with third-parties (including researchers).

e. **Data sharing between agencies and with third-parties (including researchers).**

    i. **It is widely known that anonymized data can sometimes be combined to potentially identify individuals. How can the Privacy Act be updated to mitigate against the risks of de-anonymization in large datasets?**

        - The Privacy Act should include language directing agencies to maintain robust, context-specific access controls, audit logs, and data rights management systems (DRMs) that prevent undesirable or unauthorized reidentification of deidentified data, even when certain elements of disparate datasets are combined together to advance specific government objectives, such as epidemiological research or pathogen surveillance. But perhaps more critically, the Privacy Act should reflect a more sophisticated view of the range of approaches to deidentification that may be applied in different scenarios to address varying degrees or reidentifcation risk (as we've outlined in detail here).

- An updated Privacy Act could also contemplate more sharing under formal data usage agreements that expressly prohibit data recipients from trying to reidentify data shared with them. The provisions of such agreements could be backed up with technical capabilities enabling selective revelation, obfuscation, and audit controls that help to ensure implementation of their requirements.

ii. **How can the government share personal information—with other agencies, researchers, states and localities, and other entities—in ways that are effective and privacy-preserving?**

- We acknowledge the importance of this issue. While it is not a core priority for us to provide specific advice, we are happy to include it as part of future discussions surrounding these legislative efforts.

iii. **Should Congress consider imposing restrictions on intra-agency data sharing? If so, how?**

- Yes. Indeed, requirements for additional restrictions on intra-agency data sharing should be implemented to reinforce strong notions or proportional data access and to place the onus on agencies to utilize well-established data protection PETs and other selective revelation and accountability tools to ensure only the data required for a given authorized initiative is shared on any given occasion. Such capabilities enforce these concepts include:

- Access Controls - Any system holding data collected or used by the federal government should maintain robust access controls. Integrating data into centralized systems can greatly improve the federal government's efficiency in providing services to the American people, but aggregating information originally held in disparate systems raises the risk of misuse and challenges to privacy. This is why it is imperative that any such system can implement granular access controls to ensure that data can only be viewed by personnel with the appropriate authorization.

    - This may require access controls at more granular levels. For example, when data is integrated, column- and row-level access controls might be necessary to reflect a user's authorization to view certain sensitive entries or fields.

- Purpose Justification - Another related challenge to privacy is ensuring that data is only used for the purposes for which it is collected. Capabilities and processes to oversee compliance with purpose specification can better prevent misuse or repurposing of data.

- Data Minimization - Sensitive data should be minimized by default. Data minimization provides privacy by design when users might legitimately need certain sensitive data, but only for specific purposes and specific times.

    - Beyond access controls – Access controls are ultimately binary - they either grant or deny the ability to interact with specific data. There are other, more nuanced data minimization technologies that help preserve privacy by (among other approaches) obfuscating/minimizing data that a user may technically be authorized to use. Further, alongside purpose justification tools, each act of de-obfuscation can be formally acknowledged with corresponding metadata that is tracked in an auditable ledger to help reinforce accountability.

    - Sensitive Data Discovery & Management - The federal government needs systems that can be used to help understand where sensitive data might reside – and how it flows across different use cases and applications – in order to apply the correct privacy controls. While the collection and use of sensitive data should be planned at the outset of each project, technologies to identify, catalog, and manage sensitive data are crucial component of privacy protection to ensure that privacy controls are robust and comprehensive.

    - Deletion - When sensitive data is not needed, the federal government should not look to PETs, access controls, or data minimization. Rather, the federal government should ensure that agencies and organizations have tools for comprehensive data deletion. However, deletion must be used with great care and forethought so as not to jeopardize other critical functions.

- Audit logging - Robust audit logs for both use and disclosure are essential for ensuring that the federal government is equipped to oversee the use of technology. To help ensure long-term accountability, the federal government could consider extending audit log retention beyond the current 18-month default.

- As with the adoption or promotion of any new technologies, it's equally critical that the federal government prioritize investments in education, upskilling, and training users of technology systems on data governance and security best practices. In addition, it is critical that organizations have the capacity to carry out oversight and audit responsibilities to ensure that the Americans' privacy is upheld.

- Some critics might say that using such privacy protective approaches will slow down the pace of innovation. Our position, however, has long been that technologies that protect privacy and civil liberties can accelerate innovation in part by directing engineering efforts towards better defined optimal outcomes. Engineering efforts that reject zero-sum assumptions (noted in our General Questions remarks above) out of the gate help yield technology breakthroughs and capabilities that allow organizations working with sensitive data to carry out their missions efficiently and responsibly.

f. **Civil remedies.**

i. **Should Congress consider strengthening the Privacy Act's private right of action to seek injunctive or compensatory relief? If so, how?**

- We acknowledge the importance of this issue. While it is not a core priority for us to provide specific advice, we are happy to include it as part of future discussions surrounding these legislative efforts.

g. **Privacy leadership, innovation, and oversight.**

i. **What role should the Office of Management and Budget (OMB), especially its Office of E-Government & Information Technology and Office of Information and Regulatory Affairs (OIRA), play in promoting privacy across the federal government, including through standards-setting?**

- NIST is the gold standard for developing technical guidance, frameworks, benchmarks, and tools to help federal agencies improve their privacy practices. Given its expertise in creating standards for cybersecurity, data management, and technology, NIST is uniquely positioned to drive consistent, robust, and adaptable privacy protections across agencies, especially in the context of modern data challenges like granular data management and dealing with Personally Identifiable Information (PII).

ii. **What role should the National Institute of Standards and Technology (NIST) play in developing technical guidance, frameworks, benchmarks, and tools for agencies to improve their privacy practices?**

- We acknowledge the importance of this issue. While it is not a core priority for us to provide specific advice, we are happy to include it as part of future discussions surrounding these legislative efforts.

iii. **What role should agency Chief Information Officers (CIOs) play in promoting privacy at agencies? What role should Senior Agency Officials for Privacy (SOAPs) play? How should these two officials work together?**

- CIOs and SOAPs are critical stakeholders for understanding and mitigating privacy concerns at agencies - these two officials must work together to understand and synthesize technical, legal, and operational realities that impact the privacy of the American's their agencies serve. They cannot be stove-piped from one another. More generally, breaking down barriers between legal assessments of systems and the technical provision of those systems requires cooperations with mission owners as well.

iv. **What role should independent officials, councils, and boards—including Inspectors General, the Federal Privacy Council, and the Privacy and Civil Liberties Oversight Board—play in overseeing the federal government's privacy practices?**

- We believe it is crucial for there to be external third parties that are appropriately empowered and capable of monitoring the federal government's data privacy practices, however we also caution that such bodies, being external to ongoing agency operations, would be best placed to conduct timely reviews and provide guidance, rather than monitor the day-to-day activities of an agency on a granular level.

3. **How can related laws, including but not limited to the Computer Matching and Privacy Protection Act (CMPPA), the E-Government Act of 2002, and the Federal Information Technology and Modernization Act (FISMA), and the Foundations for Evidence-Based Policymaking Act of 2018 be similarly modernized to better secure Americans' privacy?**

- We recognize that there are many existing laws that touch on government use of data and welcome the opportunity to provide input on them. However, for the purposes of this response, we will only reiterate that building state-of-the-art data privacy architecture and governance into the systems our agencies use every day to serve the American people is essential to modernizing or rewriting each of these laws.