# MITRE's Response to the OMB RFI on Privacy Impact Assessments

## April 1, 2024

For additional information about this response, please contact:

Duane Blackburn
Center for Data-Driven Policy
The MITRE Corporation
7596 Colshire Drive
McLean, VA 22102-7539

policy@mitre.org
(434) 964-5023

# About MITRE

MITRE is a not-for-profit company that works in the public interest to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation. We operate multiple federally funded research and development centers (FFRDCs), participate in public-private partnerships across national security and civilian agency missions, and maintain an independent technology research program in areas such as artificial intelligence, intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience. MITRE's ~10,000 employees work in the public interest to solve problems for a safer world, with scientific integrity being fundamental to our existence. We are prohibited from lobbying, do not develop or sell products, have no owners or shareholders, and do not compete with industry—allowing MITRE's efforts to be truly objective and data-driven. Our multidisciplinary teams (including engineers, scientists, data analysts, economists, organizational change specialists, policy professionals, and more) thus dig into problems from all angles, with no political or commercial pressures to influence our decision making, technical findings, or policy recommendations.

MITRE has extensive experience supporting federal, state, local, and international government agencies as they analyze privacy risks and implement appropriate protections. MITRE's demonstrated privacy capabilities also include conducting research, development, and test and evaluation (RDT&E) activities that help government agencies better manage privacy risk, meet privacy compliance requirements, and strategically address privacy policy and technology challenges. RDT&E activities include investigating and reviewing privacy-enhancing technologies (PETs), developing a privacy threat model, and shaping privacy best practices to maximize the value of new and emerging technologies.

# Introduction and Overarching Comments

## Reimagine the Federal Government's Privacy Approach and Risk Model

The federal government's privacy approach and risk model need to evolve and mature beyond the current Privacy Impact Assessments (PIAs) framework. The traditional use of PIAs and the Fair Information Practice Principles (FIPPs) has proved to be restrictive and insufficient in addressing the full spectrum of privacy risks in the modern era. The Office of Management and Budget (OMB) needs to move beyond its currently stated plan of merely enhancing PIA processes and instead embark on a systematic and comprehensive reimagination of the approach to ensure robust and effective privacy protections. This approach should integrate Privacy by Design (PbD) and privacy engineering into all aspects of operations, fostering a culture of privacy awareness and accountability.

The current PIA-focused approach often results in a lack of accountability, as PIAs often concentrate primarily on semi-abstract impacts while neglecting the tangible harms to individuals. Additionally, the impacts in PIAs are often framed using the FIPPs as the sole risk model. The FIPPs are foundational and remain critically important but are not robust enough to contemplate how individuals may be affected. This can lead to a disconnect between the identified impacts and the actual harms to individuals.

The FIPPs highlight important aspects of privacy that organizations must address (e.g., accountability, transparency, purpose specification, and use limitation). However, they are not solely sufficient in the modern era due to their narrow and inelastic view of privacy. A proper risk model describes possible threats, identifies vulnerabilities that might be exploited by them, and lays out what would happen if each exploit were realized, including its likelihood and severity.[1] The FIPPs are not equipped to uncover privacy harms such as emotional distress, discrimination, or loss of dignity, nor do they identify the characteristics of data processing that may lead to those harms (e.g., appropriation, unanticipated revelation). They also do not surface potential knock-on effects or indirect cause and effect problems, or the social context of systems. The ideal risk model must supplement the FIPPs to ensure a more expansive analysis of the risks that technological systems could pose to privacy.

Even when effectively executed, PIAs serve as only a limited form of privacy risk assessment. They are not a substitute for a comprehensive approach to privacy that incorporates PbD and privacy engineering. PIAs should also not be standalone activities but rather ongoing processes that span the entire system development and operational life cycles.

MITRE strongly recommends that OMB move beyond its current plan to merely tweak PIA processes. Instead, a more systematic and comprehensive reimagination is necessary to ensure privacy protections are robust, effective, and truly serve the interests of impacted individuals. Furthermore, this new approach should not be siloed like PIAs, but rather integrated into the agency's broader risk management program and operational activities (e.g., the systems engineering life cycle), thus ensuring a comprehensive and effective management of all types of risks, including privacy.

## Thresholds for Privacy Analyses

OMB should mandate and provide guidance on performing initial privacy risk assessments that replace the current agency-specific Privacy Threshold Analysis (PTA) method. This advanced and nuanced approach would involve a genuine risk assessment to evaluate systems' potential privacy risks. The results of this initial evaluation would then determine the depth of the privacy risk assessment needed. This approach, which aligns with the proposed shift toward integrating privacy engineering into all aspects of operations, allows for a more robust and comprehensive assessment of privacy risks. Integrating this approach into the broader risk management program ensures a comprehensive and effective management of all types of risks, including privacy.

The traditional PIA process, as described in OMB Memorandum 03-22, often exhibits a circular logic. It requires the identification of high-risk systems for the purpose of risk assessment, which is essentially the process of identifying and evaluating risk. This circular reasoning has been perpetuated by agencies in the form of PTAs. Instead of continuing with this circular reasoning, OMB should adopt the reimagined privacy approach proposed above.

This revised approach has several benefits. An appropriate risk assessment, which includes an analysis of potential harms to individuals, provides a more holistic view of privacy risks. This enables agencies to more effectively address risk before actual problems occur and determine when not to move forward with a system when privacy risks cannot be adequately addressed. Furthermore, a fuller risk assessment of an appropriate scope and depth would align resources to

[1] S. Shapiro. "Modernizing Privacy Risk Assessment." *Issues in Science and Technology* 38, no. 2 (Winter 2022).

the efforts where they are most needed: the assessment and oversight of systems that pose the highest risk to the privacy of the data subjects. This would be an improvement over the current paradigm, which requires significant analysis and resource expenditures on all systems, even those with minimal PII risk.

## Enhanced Training Is Required

There is a pressing need to enhance training—both under the existing PIA paradigm but especially so in the more systematic approach we propose. This training is particularly crucial for senior federal officials who play a pivotal role in implementing privacy policies. These officials need to have a deep and comprehensive understanding of privacy risk analysis. This knowledge is indispensable for the proper execution of PIAs or our proposed replacement approach. Without this foundational understanding, the effectiveness of privacy risk management could be significantly compromised.

One of the existing challenges with PIAs is their abstract invocations of privacy risk. For these assessments to be effective, the individual completing the PIA must have a conceptual understanding of risk in general and a focused understanding of privacy risk specifically. A review of existing PIAs across the federal government will reveal wild variations in this understanding, which has fundamentally undermined the effectiveness of many PIAs.

To address this issue, OMB should develop and provide basic training materials on privacy risk and privacy risk assessment. These materials should include training on basic risk concepts as well as on how these manifest regarding privacy. This initiative will equip federal employees with the necessary knowledge to conduct robust and comprehensive privacy risk assessments.

Moreover, the training should emphasize the importance of leveraging information artifacts generated even in agile systems engineering life cycles. These artifacts, such as architecture and entity relationship diagrams, and more textual elements such as data dictionaries, can provide useful inputs to the PIA process and contribute to the system description aspects of the PIA. This will reduce PIA process friction while enabling more analytically robust PIAs. While less detailed versions of some of these artifacts may be more appropriate for publication, the more granular originals should be leveraged for the purposes of analysis.

Just as privacy risk models should extend beyond the FIPPs to include a broader view of privacy risks, so too should training. Training should equip practitioners to recognize privacy risk in the context of the individuals affected by the system, their organization, the broader ecosystem within which their process or system operates, and their social intersections.

Finally, while effective training is critical for enhancing the quality of PIAs and ensuring they meet their stated goals of identifying and managing risk and articulating it to the public, it is not sufficient generically. Because privacy is contextual, government agencies should refine the OMB training to address their particular use cases.

## OMB Analysis of PIAs

More deliberate oversight is needed to ensure PIAs meet their goal. Currently, there are hundreds of existing PIAs that arguably do not identify privacy risk, let alone enable agencies to manage the risk or provide the public with useful insight. Some agencies have articulated that they view the PIA as a checkbox exercise. Having meaningful review or random auditing of PIAs by OMB across agencies will help ensure that the PIAs serve their purpose and build public trust.

# Answers to Questions Posed in the RFI

## 1. A wide range of privacy risks are associated with the creation, collection, use, processing, storage, maintenance, dissemination, disclosure, and disposal of personally identifiable information (PII). What improvements to OMB guidance on PIAs as analytical tools and notices to the public would assist agencies in identifying, addressing, and mitigating these risks, including when an agency:

### a) Develops, procures, or uses information technology to handle PII;

A wide range of privacy risks are associated with the handling of PII. To assist agencies in identifying, addressing, and mitigating these risks when developing, procuring, or using information technology to handle PII, several improvements to OMB guidance on PIAs as analytical tools and notices to the public are recommended.

First, the FIPPs approach to PIAs, while foundational, is overdue for a major overhaul. It worked well for the first decade of the E-Government Act, but with the advent of advanced technologies such as artificial intelligence (AI), cloud computing, and the complexities of cross-border data issues, it is necessary to transition from a compliance-based approach to a risk mitigation and reward maximization approach. The PbD principles of including privacy throughout the life cycle of a system and underlying data sets should be mandated.

Second, while the FIPPs provide a sound basis for informational privacy requirements, they do not completely address all issues. For example, the FIPPs assume that the authorities for collecting and using personal data are appropriate. However, in some cases, laws and regulations allow organizations to use personal data in ways that are not necessarily appropriate in some environments. Organizations should look beyond compliance requirements to include social and ethical considerations when addressing privacy. The FIPPs do not necessarily advocate for this, nor do they address associated potential privacy harms. Ideally, OMB guidance on PIAs as analytical tools and notices should request that agencies look to the FIPPs for guidance but that they also address social and ethical considerations regarding the handling of personal data when they are analyzing privacy impacts and identifying privacy risks and risk mitigation approaches.

Third, it would be beneficial for agencies for OMB to share use cases demonstrating how completing a PIA could lead to the identification of specific vulnerabilities and related controls to mitigate them. This would provide a clearer understanding of how PIAs can directly contribute to risk mitigation.

Finally, OMB could increase the usefulness of the PIA approach with a greater emphasis on triage and prioritization. The National Institute of Standards and Technology's (NIST's) privacy risk assessment methodology (PRAM) provides a useful model in this regard. PRAM walks system owners through characterizing the overall environment of a system, evaluating data flows, identifying areas of privacy concern, prioritizing which privacy concerns to address, and identifying measures to address the prioritized privacy concerns. Incorporating such a methodology into OMB guidance would provide much-needed connectivity between the process of evaluating risks and the process of deciding how to address them.

### b) Initiates, consistent with the Paperwork Reduction Act, a new electronic collection of information that contains PII;

When an agency initiates a new electronic collection of information that contains PII, consistent with the Paperwork Reduction Act, there are a wide range of privacy risks associated with the creation, collection, use, processing, storage, maintenance, dissemination, disclosure, and disposal of that PII. To assist agencies in identifying, addressing, and mitigating these risks, improvements to OMB guidance on PIAs as analytical tools and notices to the public could be made.

One such improvement would be to align all related regulatory and compliance requirements into a cohesive Plan of Action and Milestones. This would ensure that all relevant aspects of privacy protection are considered and addressed in a systematic and comprehensive manner. This Plan of Action and Milestones could include the following required outcomes:

- System Authorization: Agencies would strengthen the partnership between privacy and security by verifying and documenting that the system has the necessary security and privacy measures in place to address privacy risk. Agencies could, for example, document the comprehensiveness of privacy and security alignment in a System Security and Privacy Plan, as required by both OMB A-130 and NIST 800-53, rev. 5.[2]
- FedRAMP Authorization: Agencies would confirm and document that they follow a standardized approach to cloud product and service security and privacy assessments, authorization, and continuous monitoring.
- Paperwork Reduction Act Compliance: Agencies would confirm they followed all requirements for verifying their authority to collect PII, including determining that no process exists for the agency to collect the information in a manner less burdensome to the public.
- Section 508 of the Americans with Disability Act Compliance: Agencies would confirm that they make information and data accessible to persons with disabilities and that privacy protections are equally effective for all users.
- Civil Rights/Civil Liberties Risk Assessments: Agencies would evaluate the potential impact of the system on civil rights and civil liberties and implement measures to mitigate any identified risks.
- System of Records Notice Publication and Updates: Agencies would reflect a practice of regularly updating the system's System of Records Notice to reflect changes in the system or its use of PII, and ensuring that these notices are clear and accessible to the public.
- AI Risk Assessments: As AI technologies become increasingly prevalent, agencies would follow a proscribed method to assess the specific privacy risks associated with their use and implement appropriate mitigations.

[2] OMB Circular A-130. *Managing Federal Information as a Strategic Resource*, July 28, 2016. See Section 4.c.9.: "Agencies shall ... Develop and maintain a privacy plan that details the privacy controls selected for an information system that are in place or planned for meeting applicable privacy requirements and managing privacy risks, details how the controls have been implemented, and describes the methodologies and metrics that will be used to assess the controls." See also NIST SP 800-53 Rev. 5. *Security and Privacy Controls for Federal Information Systems and Organizations*. September 2020: Control PL-2, System Security and Privacy Plans.

### c) Uses a third-party website or application that makes PII available to the agency; or

Agencies must continue to conduct relevant risk assessments when they engage vendors and third parties, and must account for the privacy risks of new and developing technologies such as cloud computing and AI. We recommend that risk assessments for third-party websites and applications that make PII available to agencies include the same improvements for improving privacy risk analyses recommended throughout this response.

### d) Engages in a relevant cross-agency initiative that involves PII?

Privacy risk assessments and data sharing agreements continue to be a core component of shaping the needs and expectations of each organization that participates in cross-agency initiatives that involve PII. Organizations need to:

- Define the purposes of the initiative.
- Document their understanding of the context of the initiative and how it fits into the broader data processing ecosystem between and beyond the agencies involved.
- Document their understanding of the details of the technologies and data.
- Ensure consent follows the data.
- Maintain the necessary data provenance.
- Define data quality standards, including understanding when changes to the data must be propagated.

Additional measures may also be necessary depending on the context of the initiative.

## 2. What other models or best practices for conducting and documenting PIAs or similar analyses could improve agencies' PIAs?

### a) Are there approaches to analyzing and documenting how an entity addresses and mitigates privacy risks used by non-federal government entities, specific sectors or industries, academia, or civil society that OMB should consider?

There are several models and best practices that could improve agencies' PIAs or similar analyses. OMB should consider overhauling its overall approach by leveraging appropriate risk models and a model-agnostic analytical methodology. This would enable OMB to effectively address the list of specific contexts and issues provided in this RFI. A variety of models and methods are available across academia, industry, and other governments. A survey or compendium of these could serve as a highly useful resource for departments and agencies.

One key area to consider is privacy threat modeling. Risk is a combination of threats that exploit vulnerabilities, leading to negative outcomes. In privacy risk management, privacy consequences are frequently identified and analyzed. However, privacy threats are not well understood, and many organizations, including those involved in the PIA process, do not actively incorporate privacy threat modeling into their risk management procedures. Nevertheless, comprehending and effectively evaluating privacy threats is vital in privacy operations. Privacy risk models need to broaden their scope beyond consequences, on which they typically concentrate, to also model threats and vulnerabilities.

Furthermore, much current threat modeling is cybersecurity focused and does not consider privacy at all. If privacy concerns are considered, they are often treated solely as a confidentiality

concern. While security is one of the FIPPs, the security aspects of privacy are broader than confidentiality and include integrity and availability. Furthermore, privacy issues beyond security that involve other rights and interests of the data subjects are not considered. As we previously commented, the FIPPs may no longer be considered a robust and comprehensive model of privacy interests. Nonetheless, considering only confidentiality threats fails to address basic FIPPs such as the authority to collect, purpose specification, and minimum necessary data collection, and well as other privacy issues (e.g., those that lead to harms like financial loss and embarrassment). Privacy risk models must include privacy-specific threats and vulnerabilities.

For instance, MITRE developed the PANOPTIC Privacy Threat Model™, which provides a standard structure for mapping privacy attacks that can be used to model privacy threats and facilitate privacy risk management. While security risk modeling typically focuses on confidentiality-based threats to information about individuals (e.g., data breaches), PANOPTIC enables identification of threats beyond those threats (e.g., threats related to consent; notice; and inappropriate use, sharing, or retention of information about individuals). OMB should provide guidance to agencies that integrates privacy threat modeling into the PIA process as part of broader privacy risk modeling, by requiring the identification of privacy threats via the use of PANOPTIC or other privacy threat models and appropriate responses to specific threat patterns.[3]

In addition, the current wording of OMB Memorandum 03-23 permits agencies to describe the PII held and used by a system in general ways. However, privacy risks sometimes emerge when an analyst considers (1) the full data set, including each data type, not merely data categories; (2) the exact data flows from point to point; (3) the exact use of each data type; and (4) how data changes state as it is processed. Regarding the last point, for example, data elements that are not PII on their own or are publicly available PII may combine in a way that introduces privacy risk. The mere fact that PII is publicly available does not automatically imply the absence of privacy interests. (This highlights the importance of employing a robust preliminary privacy risk assessment to determine whether a more in-depth risk assessment is warranted.) For high-risk systems, a greater degree of granularity could lead to better privacy threat identification and mitigation. A greater degree of specificity would better meet the requirements set out in Circular A-130 for evidence-driven policy.

Finally, the use of diagrams is crucial in PIAs. Informational privacy risk is grounded in data flows, which are themselves grounded in system architecture. Diagrams are fundamental for describing both, yet PIAs as currently executed invariably eschew such diagrams. This not only reduces PIA readability but also hampers analysis of the target system(s). Data flow diagrams and system architecture diagrams, which have been used for decades, should be included, if only in abbreviated forms, in any PIA, even if they must be developed exclusively for that purpose. This would enhance the descriptive and analytical capabilities of PIAs.

[3] PANOPTICTM Privacy Threat Model. 2023. MITRE, https://ptmworkshop.gitlab.io/#/panoptic. Last accessed March 25, 2024.

### b) Are there similar approaches to analyzing and documenting how an entity addresses and mitigates other risks in information governance (e.g., security risks) that OMB should consider from other federal guidance or frameworks?

There are several models and best practices that could improve agencies' PIAs or similar analyses. In particular, approaches used in other areas of information governance, such as security risk management, could provide valuable insights.

In recent years, NIST has integrated privacy into the Risk Management Framework (RMF) to support evolving guidance, such as OMB Circular A-130. The RMF, developed as a life-cycle approach to managing cybersecurity risks, is most commonly applied at the system level. The latest versions of several RMF-related Special Publications (SPs) now discuss privacy and the relationship between cybersecurity and privacy risk management, including NIST SP 800-37 rev. 2 and NIST 800-53 rev. 5.

However, while privacy is included in the RMF, in many cases "and privacy" has simply been added after mentions of "security" without elaborating on privacy nuances. This means the RMF remains primarily focused on cybersecurity. To address this issue, more privacy-specific guidance that aligns with each of the RMF's steps (Prepare, Categorize, Select, Implement, Assess, Authorize, and Monitor) should be developed. This will help organizations better understand privacy considerations for each phase. Additionally, the use of the RMF as a privacy risk management tool will need to ensure privacy is adequately covered and not usurped by a cybersecurity focus. It is also worth noting that some organizations may still be transitioning away from treating the RMF as a compliance process and toward encouraging risk management.

The NIST Privacy Framework also provides privacy risk management tools. Its purpose is to provide a common approach to help organizations manage privacy risk by:

- Considering privacy and effects of products, systems, and services on individuals
- Communicating about their privacy practices
- Encouraging cross-organizational workforce collaboration

Figure 2 in the Privacy Framework (pictured below[4]) depicts the relationship between cybersecurity and privacy risk in an effort to highlight that privacy risk can arise from any type of data processing activity and points to the need to identify and manage those events beyond cybersecurity-related privacy events (e.g., breaches, unauthorized activities). Privacy risks from data processing can occur in a multitude of ways throughout the information life cycle. For example, some privacy risks during data processing may arise from governance failures, such as collection without consent, excessive collection, data inaccuracy, misuse, excessive retention, and lack of notice regarding information about individuals. Other types of privacy risks may arise from failures during system design and development, such as using poor coding practices or implementing privacy-invasive technologies over more privacy-supportive options. Some technologies and capabilities inherently introduce privacy risks while addressing other challenges (e.g., surveillance capabilities can be useful for anti-theft and physical safety but inherently introduce privacy risks by recording people). Regardless of the source or failure that

[4] NIST Privacy Framework: A Tool for Improving Privacy Through Enterprise Risk Management. 2020. NIST, https://www.nist.gov/privacy-framework. Last accessed March 20, 2024.

introduces the risk, privacy risk assessments must holistically consider the sources of privacy risk in their specific context, not just those that stem from cybersecurity-related events.

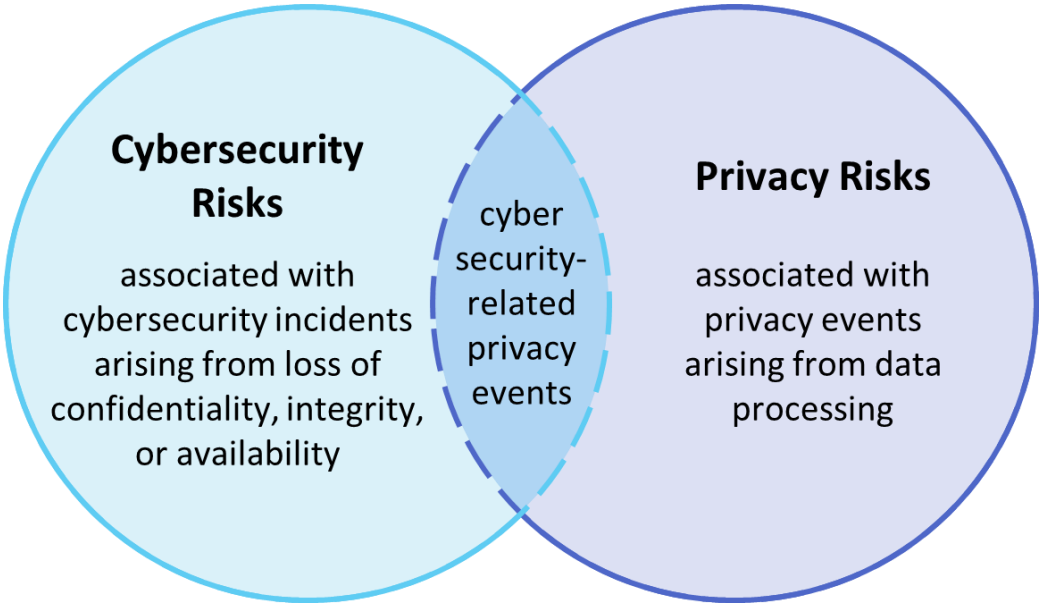


The Privacy Framework Core provides 100 privacy risk management outcomes (called "Subcategories") that can help organizations assess their privacy posture and determine a path forward. For example, the Privacy Framework includes Subcategories that encourage organizations to understand their roles, and by extension the roles of their systems, in the data processing ecosystem, by inventorying and mapping data and data flows as well as the myriad components of an environment that play a role in assessing and managing risk. The Privacy Framework also encourages implementing risk management strategies to foster an environment of ongoing risk management.

### 3. What guidance should OMB consider providing to agencies to help reduce any duplication that may arise in preparing PIAs along with other assessments focused on managing risks (e.g., security authorization packages or the AI impact assessments proposed in OMB's Draft Memorandum on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence) and to support these assessments' different functions?

OMB should consider improving guidelines for identifying systems that require a full PIA. Presently, although the Privacy Threshold Analysis approach is employed, a full PIA is still required for any system containing PII, irrespective of the scope of the PII and its exploitability. A more nuanced approach could be developed to identify higher-risk collections of PII and tailor the extent and level of detail of the analysis to the risk posed by the PII in the system.

While privacy has its own specific elements and issues, a PIA is fundamentally a risk assessment. As such, it should reflect and leverage what is generally understood about risk and risk assessments, regardless of the domain. Risks involve threats exploiting vulnerabilities, leading to adverse consequences with varying degrees of likelihood and severity. The range of

these potential threats, vulnerabilities, and consequences is defined by a risk model. Ideally, this risk model should consist of separate models for threats, vulnerabilities, and consequences.

Current PIAs often refer to risks, but in practice this is a largely *pro forma* invocation lacking a systematic analytical structure. Moreover, the risk model (which is actuality is a consequence model) currently employed (and so intertwined with PIA instruments as to mask its role), the FIPPs, is inadequate for addressing the kinds of privacy risks inherent to complex socio-technical systems.[5]

The FIPPs must be augmented or combined with other privacy risk model elements, including privacy threat models such as MITRE PANOPTIC™ and broader consequence models such as Solove's taxonomy of privacy.[6] This requires explicitly distinguishing between the privacy risk assessment methodology employed (and its supporting instrument) and the privacy risk model leveraged by the model-agnostic methodology. Although this approach may initially seem more burdensome due to the need for new and substantial guidance, treating PIAs more explicitly as risk assessments will set the stage for tailored, context-specific privacy risk models and innovative assessment methodologies (e.g., STPA-Priv[7]). Doing so will be far better aligned with agency environments and missions and ultimately much more effective than the current one-size-fits-all approach.

This kind of more substantive and sophisticated analysis can support system privacy plans (SPPs) and vice versa. PIAs and SPPs should exhibit a high degree of synchronization and be mutually supportive. Other types of assessments (e.g., civil liberties, data ethics, AI) may also intersect with PIAs, depending on context. Therefore, explicit coordination between PIAs and other potentially relevant artifacts should be part of any updated PIA guidance provided by OMB. Such coordination can enable different assessments to leverage one another, potentially reducing the overall effort required.

Finally, one limiting factor for robust privacy risk assessments under the current PIA model may be the requirement to publicly publish PIAs. Agencies tend to be concerned about introducing other types of risk by being candid or revealing detailed information about the innerworkings of its systems (e.g., increasing concern over its activities, increasing likelihood of a cybersecurity event) in PIAs. OMB may need to consider in its guidance encouraging maximum transparency so that individuals are fully informed and able to understand risks to them, and to increase agency accountability. Additionally, OMB should consider providing guidance regarding artifacts that must be evaluated when conducting a privacy risk analysis and the information that must be summarized from them in a PIA. For example, agencies should review and understand a system's architecture to adequately assess privacy risk, but a detailed architecture diagram may not be necessary or effective to include in the published PIA because it may be too complex for individuals without system architecture and engineering expertise to comprehend, rendering it useless or confusing rather than clarifying. At the same time, a simple, high-level representation

---

[5] S. Shapiro. Time to Modernize Privacy Risk Assessment. 2021. National Academies of Science, Engineering, and Medicine, https://issues.org/modernize-privacy-risk-assessment-fipps/. Last accessed March 18, 2024.

[6] D. Solove. A Taxonomy of Privacy. *University of Pennsylvania Law Review* 154, no. 3 (2006), p. 477, http://ssrn.com/abstract=667622. Last accessed March 25, 2024.

[7] S. Shapiro. Privacy Risk Analysis Based on System Control Structures: Adapting System-Theoretic Process Analysis for Privacy Engineering. 2016. IEEE Security and Privacy Workshops (SPW), pp. 17-24, https://ieeexplore.ieee.org/document/7527748. Last accessed March 25, 2024.

of system components and their interactions can be helpful for understanding what a system actually does, supplementing the typical narrative format of many PIAs.

### 4. What role do PIAs play in your search for information about how agencies handle PII and address privacy risks? For what purpose(s) do you read agencies' PIAs?

OMB should recognize that some privacy professionals read PIAs for the purposes of understanding privacy issues at the agencies that sponsor their work. These PIAs may also be instructional in understanding how various agencies interpret their privacy obligations and how thorough they are in documenting their privacy controls. While we expect and believe that many individuals read PIAs to determine their rights and to ensure government activity is being conducted legally and ethically, privacy experts benefit from PIAs in additional, substantive ways that benefit agencies and the public.

### 5. What improvements to PIAs would help you better understand agencies' assessment of privacy impacts and risk mitigation strategies?

#### a) What improvement(s) would you recommend to make it easier to find and access agencies' PIAs?

The E-Government Act of 2002 specifically required agencies to make PIAs available to the public in a clear, unambiguous, and understandable format. While many agencies adhere to the spirit of the E-Government Act in posting the PIAs on their website, it can often be challenging both to locate the PIAs and to understand them due to their complex and non-transparent language.

To enhance accessibility, MITRE recommends that all federal government PIAs be consolidated and posted on a single site, similar to the approach used for AI systems on ai.gov. This centralized approach would simplify the process for individuals seeking PIAs, particularly the PIAs of programs involving multiple agencies.

Additionally, PIAs are often named based on the system they address, which can make it difficult for individuals to locate a specific PIA unless they know the exact name of the system. To overcome this issue, OMB should provide guidance to agencies to organize and name their PIAs based on the business process handling the PII, similar to the naming convention used for AI use cases on ai.gov.

Finally, PIAs are often written using technical language, including terms that are specific to the agency and its practices, and acronyms that are not spelled out and not immediately obvious to the public. Agencies that use passive voice construction may be obscuring accountability and responsibility, which may reflect that the agency itself has not made clear assignments concerning privacy activities. While MITRE argues elsewhere in this response for a degree of technical specificity when discussing issues of engineering, PIAs are also a tool for accountability and public consumption. Even technical language may be expressed in terms that are grammatical, well organized, and clear. PIAs or successor assessments must be accessible to provide transparency; without it, the PIA is not valuable or useful for this purpose. By consolidating PIAs on a single site and naming them based on the relevant business process, ideally using a consistent taxonomy, such as an updated version of the Federal Enterprise

Architecture, individuals will find it easier to locate the PIAs they are seeking. This approach would significantly enhance the accessibility and transparency of PIAs.

b) What improvement(s) would you recommend to make it easier to read and understand agencies' PIAs?

As noted in our response to question 2(b), the absence of diagrammatic information in typical PIAs poses a significant challenge. Without visual aids, readers are left to mentally construct a system model based on potentially dense textual descriptions of data and processing. This task can be difficult even for privacy practitioners with a technical background, and it can be insurmountable for non-experts attempting to make sense of the information. Where appropriate, PIAs should contain supporting visuals to ensure an easier read and understanding of an agency's new system or program.

## 6. How can agencies increase awareness of PIAs among stakeholders?

To enhance awareness of a PIA for a new system or program, agencies can consider implementing several strategies:

- Stakeholder Meetings: Organize stakeholder meetings specifically to discuss the PIA. This approach could be particularly beneficial for the most complex and contentious systems or programs, because it would provide an additional layer of transparency and foster a deeper understanding of the respective privacy risks and mitigations.
- Training Sessions: Conduct training sessions or workshops to educate stakeholders about PIAs. These sessions can provide detailed information about what PIAs are, why they are important, and how they are conducted.
- Informative Content: Create and distribute informative content such as brochures, flyers, or online resources that explain PIAs in an easy-to-understand manner.
- Case Studies: Share case studies of how PIAs have helped organizations increase awareness. These real-life examples can provide a practical understanding of the benefits of conducting PIAs.
- Feedback and Surveys: Conduct surveys or ask for feedback to understand the level of awareness among stakeholders. This can help in identifying areas where more awareness is needed and inform future communication and education efforts.

## 7. AI and AI-enabled systems used by agencies can rely on data that include PII, and agencies may develop those systems or procure them from the private sector.

a) What privacy risks specific to the training, evaluation, or use of AI and AI-enabled systems (e.g., related to AI system inputs and outputs, including inferences and assumptions; obtaining consent to use the data involved in these activities; or AI-facilitated reidentification) should agencies consider when conducting PIAs?

When conducting PIAs for AI and AI-enabled systems, agencies should consider the privacy risks of data they use to train AI systems, data they enter into operational systems, and data that the operational systems produce. Some risks are common across each, while others are specific to each process.

**Common Risks Across All Processes**

- Informed Consent: Obtaining informed consent from individuals whose data is used in AI systems can be challenging. Users may not fully understand the implications of their data being used in this way, making it difficult to ensure that consent is truly informed.
- Data Sharing: AI systems often involve sharing data with third parties, such as cloud service providers. This can pose a privacy risk if these third parties do not have adequate privacy and security controls in place.
- Transparency and Accountability: The complexity of AI systems can make it difficult for individuals to understand how their data is being used and whether it has resulted in unfair treatment, posing a transparency and accountability risk.
- Repurposing of Data, Models, and Model Output: Repurposing data, models, and model output is common practice in AI, which can lead to privacy risk. Downstream use cases may not fully understand the limitations of the data, model, or model outputs due to lack of documentation. AI developers make assumptions that lead to inappropriate data use, model, or model outputs that are built on false assumptions and can harm users.

**Data Used to Train the System**

- Data Collection and Storage: AI systems require vast amounts of data for training. This data often includes personal information, which, if not properly anonymized, encrypted, or otherwise secured, can pose a significant privacy risk.

**Data Entered into Operational Systems**

- Data Loss: AI systems could inadvertently leak sensitive information. For example, an AI model trained on medical records could reveal sensitive health information about individuals.
- Reidentification: Even when data is anonymized, AI systems can sometimes reidentify individuals. This is particularly a concern when multiple data points can be used to identify an individual.

**Data Developed by the System**

- Inferences and Assumptions: AI systems can make inferences and assumptions about individuals based on their personal information. These inferences may lead to discrimination based on protected classes such as age, religion, gender, or ethnicity. Even when these attributes are excluded from training data, AI systems may find and utilize surrogates that produce similar results.
- Hallucinations: Generative AI systems occasionally respond to prompts with facts about individuals that are not true ("hallucinations"). These hallucinations may include attributions of speech acts the person did not make, crimes and unethical actions they did not commit, and statuses they do not have. Privacy officials must anticipate and address the risk of agencies or individuals taking action that affects individuals based on these hallucinations.
- Long-Term Storage and Use: AI systems often require long-term storage of data for training and evaluation. This can pose a privacy risk if the data is not properly secured, or if it is used for purposes other than those for which it was originally collected.

By considering these risks within the context of these categories, agencies can conduct more targeted and actionable analyses, leading to more effective privacy risk management in AI and

AI-enabled systems. Likewise, by recognizing that each application requires its own thorough analysis across the AI life cycle, privacy risk can be more accurately identified and managed.

Finally, it is important that privacy training include AI-specific training for the risk managers who will be working with AI. AI is not 100% accurate (nor is it designed to be). Programs need to factor that into their risk calculus and response (e.g., redress).

### b) What guidance updates should OMB consider to improve how agencies address and mitigate the privacy risks that may be associated with their use of AI?

When considering guidance updates to improve how agencies address and mitigate the privacy risks associated with their use of AI, OMB should again consider the process categories from our response to part (a):

**Common Guidance Across All Processes**

- Training and Education: Establish training and education for agency staff on the privacy risks associated with AI and how to mitigate them, including data ethics.
- Roles and Responsibilities: Establish clear guidelines to identify the responsibilities of senior agency leaders, program managers, AI developers, and parties with vested interests in the AI systems, tools, and outcomes.
- Regular Updates: Given the rapid pace of AI development, OMB should regularly update its guidance to keep pace with new technologies and emerging privacy risks.

**Data Used to Train the System**

- Context-Specific Guidelines on Using PETs: Provide guidance on using PETs to secure personal information used in AI systems, recognizing that effective anonymization is inherently contextual. Emphasize the need to tailor the approach to the data and context, acknowledging that there are no one-size-fits-all solutions for de-identification. This guidance should encourage agencies to consider the specific context and potential risks when deciding on the most appropriate de-identification techniques.
- Informed Consent: Provide guidelines on how to obtain informed consent from individuals whose data is used in AI systems. These guidelines should include clear explanations of how the data will be used and the potential risks involved.

**Data Inserted into Operational Systems**

- Regular Audits with Measurable Standards: Develop self-audit and independent audit standards for AI systems to ensure they are not inadvertently revealing sensitive information or perpetuating discriminatory practices. AI-based systems typically require specialized auditing techniques (i.e., algorithmic auditing) in addition to standard methods. NIST 800-53 provides families of security and privacy controls that can be used to build audit and assessment criteria. The MITRE Privacy Maturity Model[8] recommends the development of measurable and demonstrable standards that can form the basis of audits.

[8] The MITRE Privacy Maturity Model provides a framework for developing, implementing, maintaining, and evaluating privacy programs within organizations. This document may be used to assess both the completeness and maturity level of a privacy program. The framework was developed based not only on comprehensive research of relevant laws and guidance, but also on practices that have been assessed as effective in many organizations. It is publicly available free of charge. See MITRE Privacy Maturity Model. October 20, 2019. The MITRE Corporation, https://www.mitre.org/sites/default/files/2021-11/pr-19-3384-privacy-maturity-model.pdf.

Audits should focus on concrete and documentable criteria, rather than generalities or abstract concepts, to ensure effective measurement and documentation.
- Data Sharing Policies: Develop best practices for sharing data with third parties, including requirements for data protection measures.
- Cross-Border Data Transfers: Provide guidance on how to transfer AI data sets into and out of foreign countries.

**Data Developed by the System**

- Long-Term Storage and Use: OMB should coordinate efforts with the National Archives and Records Administration to develop guidelines on how to securely store data for long-term use in AI systems, and on limitations for the use of this data.
- Transparency and Accountability: Update guidance to include strategies for improving transparency and accountability in AI systems, such as providing clear explanations of how the AI system works and how it uses data. Such guidance should take into account ongoing research on effective strategies for engaging with the public.
- Incident Response Plan: Agencies should update their incident response plans to address potential breaches, privacy concerns, or AI system failures promptly and effectively.

### 8. What role should PIAs play in how agencies identify and report on their use of commercially available information (CAI) that contains PII?

#### a) What privacy risks specific to CAI should agencies consider when conducting PIAs?

Privacy risk models applicable to other sources of PII are equally relevant to CAI. However, the nature of CAI introduces additional risks due to the opacity and sometimes questionable ethics of the underlying data ecosystem. The general public typically lacks a comprehensive understanding of this ecosystem and, even when aware, may have limited options to exercise control over their data and mitigate associated privacy risks. Unlike direct collection of PII, CAI inherently poses risks related to transparency and choice, as well as potential violations of contextual integrity (i.e., the social norms under which information is provided and presumed to be used).[9] To identify these risks, an oversight body such as OMB must generate a privacy risk model that extends beyond the FIPPs.

#### b) OMB M–03–22 requires PIAs "when agencies systematically incorporate into existing information systems databases of information in identifiable form purchased or obtained from commercial or public sources," while noting that "[m]erely querying such a source on an ad hoc basis using existing technology does not trigger the PIA requirement." What guidance updates should OMB consider to improve how agencies address and mitigate the privacy risks that may be associated with their use of CAI that contains PII?

The decision to conduct a PIA for the use of CAI that contains PII should be context based, and the necessary analysis should be explicitly integrated into PIAs. The traditional binary distinction between public and private information has always been somewhat blurred, and this is even more the case in today's socio-technical environment. The fact that personal information is publicly

[9] H. Nissenbaum. Privacy in Context: Technology, Policy, and the Integrity of Social Life. Palo Alto: Stanford Law Books, 2009.

accessible via web scraping and commercial data brokers does not negate the existence of privacy or civil liberties interests when agencies propose to collect or use this information. At a minimum, such use will likely constitute violations of contextual integrity, creating the potential for a variety of privacy harms, including those NIST identifies as exclusion, stigmatization, and power imbalance. In the current era, informational privacy is not solely about what agencies directly collect, but also about how they use data, regardless of its origins. If the use of PII is integral to a business process, this alone warrants a privacy risk assessment.

### 9. What guidance updates should OMB consider to improve how agencies address and mitigate the privacy risks that may be associated with their use of other emerging technology and data capabilities?

For most agencies, PIAs are performed separately from other agency risk management activities. PIAs document information about the types of PII collected, their uses, sharing, retention, and so on. However, privacy risks are not always explicitly discussed, so privacy risks may not be adequately identified, tracked, and mitigated. To address this challenge, PIAs should be integrated into agency risk management efforts across domains, including via engagement with data management and cybersecurity processes, explicitly feeding into enterprise risk management functions. PIAs should provide structured analytical mechanisms for identifying privacy risks and mitigations, rather than leaving these as free-form exercises, as is typical.

### 10. What else could help promote greater effectiveness and consistency across agencies in how they approach PIAs?

Enhancing the effectiveness and consistency of PIAs across agencies can be achieved through several strategies:

- Implementing Privacy Continuous Monitoring: This strategy can provide regular assessments of privacy capabilities and risks, ensuring that many potential threats are promptly identified and addressed.
- Use of Data Inventory and Data Flow Diagrams: These diagrams can offer a clear visual representation of what happens to data throughout the information life cycle, including how and where data is collected, processed, transferred or disseminated, stored, and disposed within and beyond an organization, making it easier to identify potential privacy risks.
- Implementing AI Chatbots: AI chatbots can help correct basic errors in records and assist in producing data in response to simple Freedom of Information Act requests. This not only enhances the accuracy of data but also promotes transparency and accountability, which are crucial elements in Privacy Impact Assessments.

### 11. What else should OMB consider when evaluating potential updates to its guidance on PIAs?

Many agencies are unsure about whether the use of user credentials alone constitutes PII within a system, thereby requiring a full PIA. OMB should consider clarifying that, while user credentials technically meet the definition of PII, they are most commonly used for system administration. User credentials do not always trigger the same type of privacy interests that the processing of information about citizens does, or about agency employees when the information is related to their roles as employees. In many cases, current security requirements are adequate protection

for user credentials. By exempting such systems from the requirement to complete PIAs, agencies could redirect resources to focus on developing PIAs for systems that pose much higher risk. It is important to note, however, that there may be exceptions to this approach, such as systems that correlate and analyze system user activities for the purpose of identifying insider threats. Any guidance that permits a PIA exception for user credentials must stress the importance of evaluating data processing activities for privacy risk before making the determination that a PIA is not required.