
# *NATIONAL PRIVACY RESEARCH STRATEGY*

*A Report by the*

PRIVACY RESEARCH AND DEVELOPMENT
INTERAGENCY WORKING GROUP

SUBCOMMITTEE ON NETWORKING AND INFORMATION
TECHNOLOGY RESEARCH AND DEVELOPMENT

*of the*

NATIONAL SCIENCE AND TECHNOLOGY COUNCIL

JANUARY 2025

## About the Office of Science and Technology Policy

The Office of Science and Technology Policy (OSTP) was established by the National Science and Technology Policy, Organization, and Priorities Act of 1976 to provide the President and others within the Executive Office of the President with advice on the scientific, engineering, and technological aspects of the economy, national security, health, foreign relations, the environment, and the technological recovery and use of resources, among other topics. OSTP leads interagency science and technology policy coordination efforts, assists the Office of Management and Budget (OMB) with an annual review and analysis of federal research and development in budgets, and serves as a source of scientific and technological analysis and judgment for the President with respect to major policies, plans, and programs of the federal government. More information is available at https://www.whitehouse.gov/ostp.

## About the National Science and Technology Council

The National Science and Technology Council (NSTC) is the principal means by which the executive branch coordinates science and technology policy across the diverse entities that make up the federal research and development enterprise. A primary objective of the NSTC is to ensure science and technology policy decisions and programs are consistent with the President's stated goals. The NSTC prepares research and development strategies that are coordinated across federal agencies aimed at accomplishing multiple national goals. The work of the NSTC is organized under committees that oversee subcommittees and working groups focused on different aspects of science and technology. More information is available at https://www.whitehouse.gov/ostp/nstc.

## About the Subcommittee on Networking & Information Technology Research & Development

The Networking and Information Technology Research and Development (NITRD) Program has been the nation's primary source of federally funded work on pioneering information technologies (IT) in computing, networking, and software since it was first established as the High-Performance Computing and Communications program following passage of the High-Performance Computing Act of 1991. The NITRD Subcommittee of the NSTC guides the multiagency NITRD Program in its work to provide the research and development (R&D) foundations for ensuring continued U.S. technological leadership and for meeting the nation's needs for advanced IT. The National Coordination Office (NCO) supports the NITRD Subcommittee and its Interagency Working Groups (IWGs) (https://www.nitrd.gov/about/).

## About the Privacy Research & Development Interagency Working Group

The Privacy R&D IWG coordinates federal R&D aimed at protecting privacy during information processing, including R&D of privacy-protecting information systems and standards. This R&D supports advances in large-scale data analytics that can improve healthcare, eliminate barriers to education and employment, and increase efficiencies in the transportation and financial sectors while minimizing risks to individual privacy and possible harms such as discrimination, loss of autonomy, and economic losses. The Privacy R&D IWG reports investments to the Cyber Security and Privacy (CSP) Program Component Area.

## About This Document

This strategy establishes objectives and priorities for federally funded privacy research, provides a framework for coordinating privacy research and development, and encourages multidisciplinary research that recognizes privacy needs of individuals and society and the responsibilities of the government. The science and technology advances established by this strategy will enable individuals, commercial entities, and the government to benefit from technological advancements and provide meaningful protections for personal information and individual privacy.

## Copyright



## Disclaimer

Any mention in the text of commercial, non-profit, academic partners, or their products, or references is for information only; it does not imply endorsement or recommendation by any U.S. government agency.

## National Science and Technology Council

### Chair

**Arati Prabhakar,** Assistant to the President for Science and Technology; Director, Office of Science and Technology Policy (OSTP)

### Acting Executive Director

**Lisa E. Friedersdorf,** Office of Science and Technology Policy

## Subcommittee on Networking and Information Technology Research and Development (NITRD) Co-Chairs

**Joydip Kundu,** Deputy Assistant Director, Computer and Information Science and Engineering, U.S. National Science Foundation (NSF)

**Craig Schlenoff,** Director, NITRD National Coordination Office (NCO)

## Privacy Research & Development Interagency Working Group (IWG) Co-Chairs

**David Kuehn,** Program Manager, Federal Highway Administration U.S. Department of Transportation

**Angela Robinson**, Mathematician, Cryptographic Technology Group
National Institute of Standards and Technology

**Anna Squicciarini,** Program Director, Secure and Trustworthy Cyberspace (SaTC) Division of Computer and Network Systems (CNS)
National Science Foundation (NSF)

### Technical Coordinator

Olachi Onyewu, NITRD NCO

### Writing Team Members

James Joshi (NSF)
TJ Kasperbauer (NIH)
Angela Robinson (NIST)
David Kuehn (FHWA)
Alexandria Phounsavath (DHS)
Anna Squicciarini (NSF)

## Table of Contents

## List of Abbreviations and Acronyms

| Acronym | Definition |
|---|---|
| AI | Artificial Intelligence |
| CCPA | California Consumer Privacy Act |
| CPRA | California Privacy Rights Act |
| DP | Differential Privacy |
| ECPA | Electronic Communications Privacy Act |
| EO | Executive Order |
| FIPPs | Fair Information Practice Principles |
| FISA | Foreign Intelligence Surveillance Act |
| FTAC | Fast-Track Action Committee |
| FTC | Federal Trade Commission |
| GDPR | General Data Protection Regulation |
| HIPAA | Health Insurance Portability and Accountability Act |
| HPC | High Performance Computing |
| IT | Information Technology |
| IWG | Interagency Working Group |
| LLMs | Large Language Models |
| ML | Machine Learning |
| NIST | National Institute of Standards and Technology |
| NITRD | Networking and Information Technology Research and Development |
| NPRS | National Privacy Research Strategy |
| NSF | U.S. National Science Foundation |
| NS-PPDSA | National Strategy to Advance Privacy-Preserving Data Sharing and Analytics |
| NSTC | National Science and Technology Council |
| OMB | Office of Management and Budget |
| OSTP | Office of Science and Technology Policy |
| PDaSP | Privacy-preserving Data Sharing in Practice program |

| | |
|---|---|
| PETs | Privacy Enhancing Technologies |
| PPDSA | Privacy Preserving Data Sharing and Analytics |
| R&D | Research and Development |
| SaTC | Secure and Trustworthy Cyberspace program |
| SBIR/STTR | Small Business Innovation Research/Small Business Technology Transfer |
| U.S. | United States |

## 1. Executive Summary

People's lives are inextricably interconnected with cyberspace and information systems. The computing revolution has enabled advances in many sectors of the economy, while social interactions have been profoundly affected by the rise of the Internet, mobile communications, and rapid advances in artificial intelligence (AI), including recent fast-paced growth of large language models (LLMs) or foundational models trained on large amounts of data. Ever-increasing computational power, hyperconnectivity online, and exponential growth of powerful data collection devices and techniques in a wide range of application domains, such as transportation, education, health care, and finance, are accelerating these trends. Massive data collection, storage, processing, and retention in the digital era challenge long-established privacy norms and introduce significant risks of privacy harms to individuals with negative consequences to communities and society at large. The increased ability to conduct data analytics at scale, including training large and powerful AI models, is indispensable to progress in science, engineering, medicine, and social good. However, when information about individuals and their activities can be tracked, combined, inferred, and repurposed without their knowledge or understanding, risks emerge that these data actions could result in such individuals experiencing physical harm, unfair discrimination, loss of autonomy, financial loss, and loss of dignity. The presence of these privacy risks can have a devastating and chilling effect on people's behaviors, diminish public trust in cyberspace, and exacerbate potential harm to both individuals and society.

The federal government is mindful of these privacy risks and the critical need for foundational, use-inspired and translational privacy research and development (R&D). The Executive Order on Safe, Secure and Trustworthy Development of AI (EO 14110) emphasizes that "Americans' privacy and civil liberties must be protected as AI continues advancing," and that the agencies "shall use available policy and technical tools, including privacy-enhancing technologies (PETs) where appropriate to protect privacy." [1] The Executive Order on Ensuring Responsible Development of Digital Assets (EO 14067) highlights the need for protecting consumer and other stakeholders by ensuring privacy protection and safeguards against "unlawful surveillance."[2] Similarly, the Blueprint for an AI Bill of Rights[3] emphasizes the need for data privacy reinforcing the message to individuals: "you should be protected from abusive data practices via built in protection and you should have agency over how data about you is used."

This National Privacy Research Strategy (NPRS) updates the 2016 NPRS and outlines the strategic priorities for privacy R&D to be pursued by researchers and practitioners from public and private sectors. It establishes objectives for federally funded (both extramural and government-internal research) as well as industry-funded privacy R&D, provides a common direction for coordinating R&D in privacy-preserving technologies, and encourages multidisciplinary research that recognizes the responsibilities of public-private stakeholders and the needs of society at large.

---

[1] The White House. (2023, October). *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*. https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

[2] The White House. (2022, March). *Executive Order on Ensuring Responsible Development of Digital Assets.* https://www.whitehouse.gov/briefing-room/presidential-actions/2022/03/09/executive-order-on-ensuring-responsible-development-of-digital-assets/

[3] The White House Office of Science and Technology Policy (OSTP). (2023). *Blueprint for an AI Bill of Rights*. https://www.whitehouse.gov/ostp/ai-bill-of-rights/

The overarching goal of this strategy is to promote innovative privacy research and privacy-preserving technology while advancing the well-being and prosperity of individuals and society.

To achieve these goals, this strategy identifies the following priorities for privacy research:

- Foster multidisciplinary approaches to privacy research and solutions;
- Understand and measure privacy preferences and impacts;
- Develop system design methods that incorporate privacy preferences, requirements, and controls;
- Increase transparency of data collection, sharing, use, and retention;
- Ensure that information flows and use are consistent with privacy rules; and
- Reduce privacy risks of data analytics and AI, including the potential of re-identifying anonymized data.

# 2. Introduction

## 2.1 Privacy Research Purpose

Networking and information technology is transforming life in the 21st century, changing the way people, businesses, and government interact at scale. Vast improvements in computing, storage, and communications technologies, including rapid progress in AI and advanced analytics, are creating unprecedented opportunities for enhancing individuals' social wellbeing; improving health and health care; eliminating barriers to education and employment; and increasing efficiencies in many sectors, such as manufacturing, transportation, finance, and agriculture.

Advances in information technology have mixed results. For example, the promise of these new systems and applications often stems from their ability to create, collect, store, transmit, process, and archive information on a massive scale. However, the exponential growth in the quantity of personal information that is being collected and retained, combined with the increased ability to analyze it and combine it with other information, is creating valid concerns about privacy risks and about the ability to manage these unprecedented volumes of data responsibly. The presence of such risks can create a chilling effect on people's behaviors and rapidly reduce trust in cyberspace.

However, the progress of privacy understanding, and experience under legal and regulatory protections, has not kept pace with the exponential increase in data collection, processing, and storage, and the resulting risks to privacy. Today, information exists in a complex and dynamic ecosystem that includes:

- Individuals whose information and data elements are collected, processed, or stored;
- Data collectors and data brokers, who buy, repackage, and sell collected information;
- Analytics providers, including AI model developers, who create systems for processing such information to extract valuable insights from data;
- Data managers, who maintain data that may include sensitive records or when combined with other data could risk privacy harms; and
- Data users, who make decisions based on the data analytics.

The decreasing cost of storage has enabled organizations to collect large amounts of data and save the data in long-term repositories, making such data available for unanticipated or even unforeseeable future use. Meanwhile, there is a growing array of ubiquitous consumer devices, environmental sensors, and tracking technologies designed to collect, process, and archive information continuously, often

without the data subjects knowing exactly what is being collected about them and how it will be used.[4] This exponential growth in the type and amount of data collected, the increasing analytic capabilities and AI training at massive scale, and the lack of appropriate mechanisms to control use are among the many factors driving increased privacy concerns. The availability of disparate datasets is setting the stage for a "mosaic effect,"[5] where carrying out data analytics or AI model training across datasets can result in linkages that reveal privacy-sensitive information or generate inaccurate and potentially problematic inferences, even though in isolation the datasets may not raise privacy concerns.

In response to ongoing rapid technological progress to leverage data while preserving privacy, several federal initiatives have been taken to provide leadership across a broad range of privacy issues. The National Strategy to Advance Privacy-Preserving Data Sharing and Analytics (PPDSA)[6] was released in 2023 to establish a roadmap for advancing foundational, use-inspired, and translational research for developing technologies that focus on enabling data sharing and analytics in a privacy-preserving manner. This National Privacy Research Strategy complements the National Strategy to Advance PPDSA. The EO 14110[7] has pushed forward a PETs development agenda, especially within AI technologies, that has resulted in several new initiatives. Previous and current federal R&D activities promoting privacy are summarized in the table below.

| **Agency R&D Activities Promoting Privacy** |
|---|
| Census Bureau and National Institute of Standards and Technology (NIST) U.S. PETs Lab |
| United States Department of Homeland Security (DHS) Data Autonomy Framework |
| DHS Office of University Programs Center for Accelerating Operational Efficiency Data Analytics – Privacy Enhancing Technology projects |
| DHS Privacy Enhancing Technology Applications and Evaluation |
| DHS Science & Technology Directorate Test and Evaluation of Facial Recognition and Facial Capture Technology in accordance with International Organization for Standardization/International Electrotechnical Commission and NIST technical guidance |
| DHS Silicon Valley Innovation Program Synthetic Data Generation topic call |

[4] Federal Trade Commission. (2024, March) *FTC Cracks Down on Mass Data Collectors: A Closer Look at Avast, X-Mode, and InMarket*. https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/03/ftc-cracks-down-mass-data-collectors-closer-look-avast-x-mode-inmarket

[5] White House Office of Management and Budget. (2013, May). *Open Data Policy-Managing Information as an Asset.* https://www.whitehouse.gov/wp-content/uploads/legacy_drupal_files/omb/memoranda/2013/m-13-13.pdf

[6] The White House. (2023, March). *National Strategy to Advance Privacy-Preserving Data Sharing and Analytics.* https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Strategy-to-Advance-Privacy-Preserving-Data-Sharing-and-Analytics.pdf

[7] The White House. (2023, October). *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*. https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

| |
|---|
| Federal Trade Commission (FTC) Privacy Research Workshops |
| National Institutes of Health (NIH) Data Linkage Governance Metadata Project |
| NIH iDASH Secure Genome Analysis Competition |
| NIST Collaborative Research Cycle |
| NIST Privacy Engineering Collaboration Space |
| NIST Privacy Framework |
| National Science Foundation (NSF) Privacy-preserving Data Sharing in Practice (PDaSP) program |
| NSF Secure and Trustworthy Cyberspace (SaTC)program |
| NSF U.S.-UK Privacy Enhancing Technologies Challenge |

A significant amount of NITRD's earlier privacy research investment includes explicit efforts in health care, regulation compliance, foundational and multidisciplinary research explorations, as well as research on privacy as an extension of or complementary to other research on cybersecurity. While the 2016 NPRS heralded more coordinated efforts among federal agencies to pursue R&D activities, the computing and information landscape has evolved significantly since then, creating a need to recognize emerging privacy risks and harms and further streamline our national strategy for privacy research.

This 2025 NPRS establishes strategic objectives for federally funded as well as public sector-driven research in privacy and provides guidance to federal agencies for developing and sponsoring R&D activities in this area, while identifying a broader strategic research roadmap for the R&D community. This research strategy emphasizes pursuit of new knowledge and scientific foundations and socio-technical solutions that identify and mitigate emerging privacy risks and harms. The strategy recognizes that research for technologies must also be accompanied by an understanding of privacy risk and appropriate privacy protections for the research itself. There are no one-size-fits-all approaches for privacy protections. The end goal of this strategy is to help society realize the benefits of computing, communication, and information technologies with appropriate contextual privacy protection policies and practices[8] to minimize or mitigate privacy risks and harms to individuals and society both on- and offline. Strategies for minimizing potential risks to privacy must consider a range of opportunities, from minimizing data collections to proper safeguarding of data throughout its lifecycle.

This NPRS calls for research along a continuum of challenges, from how people understand privacy in different situations and how their privacy needs can be formally specified, to how these needs can be respected and how mitigation and remediation can be accomplished should privacy expectations or

[8] Department of Health and Human Services. (2024, March). *Federal Policy for the Protection of Human Subjects ('Common Rule')*. https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/index.html

requirements and interests be violated. Finally, the NPRS highlights the need to transition research results into practice so that all stakeholders can benefit from value that can be extracted from data in a privacy-preserving manner. Appendix A summarizes the main steps in the development of the strategy.

## 2.2 Privacy Characterization

Privacy is difficult to characterize, despite legal and policy definitions.[9] A full treatment of privacy requires a multidisciplinary approach, including perspectives from ethics and the humanities, the social sciences, laws and governance, and STEM fields. Embodying such broad considerations, the federal government's approach has been guided by the Fair Information Practice Principles (FIPPs), a framework for understanding stakeholder considerations utilizing concepts of fairness, due process, and information security. The 2012 Consumer Privacy Bill of Rights is based on the FIPPs, supplemented importantly by the concept of "respect for context."[10] Research is needed to help bridge the gap between statements of principles and effective implementation in information systems. The 2014 report on big data[11] provides that privacy "addresses a range of concerns reflecting different types of intrusion into a person's sense of self, each requiring different protections." Privacy can be defined in multiple ways, depending on whether one highlights aspects such as confidentiality, the control of dissemination of personal information and their use, the control of one's identity, or the negotiation of boundaries of personal spaces. Indeed, privacy definitions and characterizations continue to evolve and themselves constitute an open research question. Privacy R&D should not be limited by any view or definition of privacy and should be explored from many perspectives. Research examining the usefulness of different approaches and their applicability to general or specific privacy challenges should accompany such explorations.

The research priorities outlined in this document are based on a privacy characterization that is a combination of four key concepts: *subjects*, *data*, *actions*, and *context*. As a coarse characterization, *subjects* encompass an individual or a group of individuals, the identity (as well as pseudonymity and anonymity) of individuals and groups and their rights, autonomy, and privacy preferences. *Data* encompasses the data and derives information about these individuals and groups, also referred to as data subjects. *Actions* encompasses the various data collection, storage, processing, analysis, and retention practices, controls that constrain such practices, as well as impacts (negative and positive) of the collection and use of data on individuals, groups, and society. The interactions among *subjects, data, and actions* that enable flow of information, the interpretation of those interactions, and the risk of harm are influenced and conditioned by the *context*.

Within this characterization, "privacy" concerns the proper and responsible collection, creation, use, processing, sharing, transfer, disclosure, storage, security, retention, and disposal of information about

[9] Department of Justice. (2022, January). *The Freedom of Information Act, 5 U.S.C. § 552*. https://www.justice.gov/oip/freedom-information-act-5-usc-552

[10] The White House. (2012, February). *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy*. https://obamawhitehouse.archives.gov/sites/default/files/privacy-final.pdf

[11] The White House. (2014, May). *Big Data: Seizing Opportunities, Preserving Values*. https://obamawhitehouse.archives.gov/sites/default/files/docs/20150204_Big_Data_Seizing_Opportunities_Preserving_Values_Memo.pdf

people. This includes decisions by entities about when *not* to collect or store, *not* to create, *not* to transfer, and *not* to permit certain uses of information to protect legitimate privacy interests.

## 2.3 Key Challenges for Privacy

The following challenges motivate the research priorities for this strategy.

### 2.3.1 Influence of Context on Privacy

Individuals tend to share personal data with others or organizations within a particular community for specific purposes. For example, individuals may share their medical status with health care professionals, product preferences with retailers, legal interests with law firms, spiritual concerns with religious organizations, and trip plans with travel agents. The community and the application domain provide a context for sharing data. When information shared with one community shows up in another outside of the intended context, this may violate individuals' expectations of privacy. Context matters greatly in privacy. The content of the personal data, the relevant ethical or social norms, relationships between the parties, and other factors all matter.

The contextual nature of privacy creates a challenge for designing privacy-preserving systems because people will consider privacy from varied viewpoints, may use diverse terminologies to express their privacy concerns and preferences, perceive privacy-related harms differently, and vary their privacy requirements with circumstances. Moreover, system designers may lack adequate mechanisms to specify the properties that comprise privacy and to establish that such properties are satisfied by some deployed system. While there has been rapid development in privacy research over the last decade, techniques for specifying information and computing systems as well as the computational models are still lacking adequate capabilities to address privacy challenges. In addition, there is a need for sociobehavioral and ethics research on privacy expectations and behaviors associated with perceived and actual differential harms from disclosure of personal information across groups and contexts, as well as research on assessing, ensuring, and communicating the balance between individual or group privacy and societal good for various data-driven applications.

### 2.3.2 Transparency in Data Collection, Use, and Retention

Current government and business approaches to provide transparency about data collection, storage, and use practices have fallen short. The traditional notice-and-choice framework, in which data collectors and users set forth practices in lengthy privacy policies and deem individuals to have read, understood, and consented to them, has its limits. Privacy notices that are sufficiently detailed become too long and difficult to understand for individuals to read and give meaningful consent, while notices that are phrased broadly in order to cover all anticipated future uses lack sufficient details for consent to be meaningfully informed. Today, there are so many organizations seeking to collect and use information for analytics and training AI models that individuals realistically do not have the ability to evaluate each collection notice and associated data use and cannot track unconsented uses of data. Looking forward, as people are increasingly surrounded by sensors that continuously collect data in domains such as transportation, building energy management, or public safety, and as organizations increasingly adopt advanced analytics or powerful AI systems, it is becoming more challenging to ensure privacy through existing data control such as consent and disclosure mechanisms. Beyond existing laws and regulations, more effective, efficient, and scalable solutions are needed to support transparency, accountability, consent, and choice, for data control by individuals, oversight by regulators, and legal and regulatory compliance.

Users also have little public understanding of data storage and retention practices and their potential implications, due to a lack of transparency and common practices from data providers. For instance, the availability and persistence of collected data can contradict people's expectations that minute details of their past will not be forever available. It can also contradict individuals' expectations about how collected data will be used, because providing notice of prospective changes in data-handling or data sharing practices can be challenging. Data longevity and complex data provenance and lineage also makes it difficult for individuals to withdraw or change their consent regarding particular data uses. Accordingly, research and development on methods to communicate data use and retention in plain language given an increasingly complex environment are needed.

### 2.3.3 Data Aggregation, Analysis, and Release

Increased capabilities in data collection, aggregation, analysis, and machine learning are fueling the discovery of new patterns, correlations, and knowledge about the world, and an increasing use of classification, generative, and predictive algorithms. Some large commercial AI models may include privacy-sensitive information. The sensitive information can be revealed if the algorithm is poorly designed or exploited by an adversary. Individuals are often unaware of the outcomes of these algorithms and technologies. Government and industry have difficulty evaluating the privacy-preserving elements of systems that use these complex algorithms or scoring them under a certification schema. Organizations increasingly rely on non-public algorithms to make a variety of decisions or take a direct action. However, there is a risk that predictive algorithms could, for example, enable decisions that result in (perhaps unintended) consequences such as bias and discrimination that can lead to potential harms. Unintended consequences can and do occur, and the actual scale and impact on privacy are not known.

It is essential for all stakeholders to understand privacy concerns to effectively address the challenges posed by today's digital environment. Whether they are policymakers, legal advisers, technology developers, business leaders, or consumers, stakeholders play a critical role in shaping the landscape of privacy protections. Without a thorough grasp of the privacy issues at stake, stakeholders may implement solutions that are inadequate, overly restrictive, or fail to align with societal values. Understanding privacy concerns enables stakeholders to make informed decisions that balance the need for innovation with the imperative to protect individual rights.

Furthermore, the government faces unique responsibilities to avoid harms from inappropriate or unexpected privacy disclosures and at the same time making data collected using public funds available as broadly as possible in support of societal benefits. The growing attention in publishing statistics, analyses, and raw data held by the federal government in order to provide public benefits from public investments in research raises privacy concerns as well. There are limitations to existing approaches for protecting privacy, such as the removal of personally identifiable information (PII) or use of de-identification techniques, to address the privacy risks of large-scale data collection, analytics, and release. For example, differential privacy (DP) as a statistical disclosure limitation technique has shown promise in providing privacy guarantees, but there is significant challenge in deploying DP in broader practical applications. Similarly, homomorphic encryption allows for the performance of operations on encrypted data without ever needing to decrypt it and thus has shown promise in providing privacy guarantees, but the performance of fully homomorphic encryption can be a barrier in practice. As more information about individuals is collected and shared, often as a source of profit as compared to government's focus on data access for public good, data analytics can often be used to

link sensitive information back to individuals, despite efforts to de-identify data. This situation creates opportunities for personal information to be misused.

## 2.4 Desired Outcome

The goal of this NPRS is to produce scientific foundations, knowledge, and technology that will enable individuals and groups, commercial entities, and the government to benefit from technological advancements and data use while proactively identifying and mitigating privacy risks. Solutions are needed that will establish trustworthy boundaries in the process and in the outcomes.

Privacy creates opportunities for political expression and choice. Privacy protections provide a space for negotiation between consumers and businesses about data practices. When privacy is not protected, individuals and society suffer from harms, including erosion of freedom, discrimination, loss of trust in institutions, or reduced innovation from self-censoring by the population.

Sustaining privacy requires technologies targeted for particular use, as well as foundational science and engineering to analyze contexts that might lead to privacy harms and produce technologies to prevent or mitigate them. Much work to date on privacy has focused on specific narrow technologies and applications. This strategy seeks:

- To prioritize and promote fundamental multidisciplinary research to develop scientific foundations for privacy that would enable rigorous analysis of the drawbacks or limitations, risks, and potential benefits to privacy and society from data collection, storage, processing, and analysis systems; and
- To promote use-inspired and translational research to foster the development of socio-technical, inclusive solutions that better protect privacy and allow more robust and precise negotiation of privacy expectations and preferences, and their representation or specification, and enforcement in specific application contexts.

These R&D outcomes are important to ensure a safe and secure data-centric digital future where data access, sharing, and use in a privacy-preserving, transparent and equitable manner is easily achieved, while affirming democratic foundations, uplifting human dignity and well-being, and ensuring national security.

This strategy does not attempt to set privacy standards or norms; however, the research outcomes of this strategy should support individuals, communities, and the federal government in achieving privacy protections. Likewise, this strategy does not address privacy policy issues associated with law enforcement or national security (although the research under this strategy should help clarify many related issues). Appendix B further discusses the legal and policy context for privacy in the United States. Moreover, this strategy does not address how to rectify poor computer security or information protection practices, which can cause privacy harms. Federal research strategy for improving computer and cyber security is presented in the 2023 Federal Cybersecurity Research and Development Strategic Plan[12].

[12] NITRD. (2023, December). *Federal Cybersecurity Research and Development Strategic Plan*. https://www.nitrd.gov/pubs/Federal-Cybersecurity-RD-Strategic-Plan-2023.pdf

# 3. National Privacy Research Priorities

The following research priorities, jointly established by the information technology research funding agencies, focus on critical capability gaps in the privacy domain. The priorities provide a strategy; federal agencies will be responsible for making tactical decisions about how to structure, fund, and execute specific research programs based on their missions and capabilities, so that the overall research portfolio is consistent with this strategy. This national strategy is intended to inspire a range of parallel R&D efforts in both the public and private sectors.

## 3.1 Foster Multidisciplinary Approaches to Privacy Research and Solutions

This priority is overarching, speaking to how privacy research is conducted. It aims to advance research that will improve organizations' ability to protect privacy while executing their missions and responsibilities, and more generally, improve the ability of individuals, groups, and all members of society to create systems that collect and process information in a manner that respects privacy, ensures fairness and equity, and prevents unfair discrimination. To achieve these objectives, this priority calls for multidisciplinary research involving disciplines such as computer and information sciences, engineering, social and behavioral sciences, biomedical science, psychology, economics, law and policy research, and ethics. Multidisciplinary approaches are necessary to:

- Characterize privacy goals, and potential risks and harms,
- Understand privacy events (acts and actions with the potential to compromise privacy; privacy events may or may not be considered privacy violations and may or may not result in privacy harms)[13] and
- Build robust socio-technical foundations to engineer privacy-protecting systems that can mitigate privacy risks and minimize privacy harms.

Multidisciplinary and holistic approaches are needed to understand how the adoption of privacy protections is advanced or impeded by policy and regulatory factors, organizational and business aspects, market competition, and economic and social incentives or disincentives, as well as technologies. Multidisciplinary research is needed to gain insight into whether and when privacy protections are addressed best technologically, through law, ethics and policy, or some combination of all methods.

Transitioning research into real-world technology implementations requires stakeholders at many levels to understand the privacy risks and trade-offs to make decisions that protect privacy while meeting the needs of the organization. A controlled environment allows discovery of privacy issues, new privacy concepts, and effective risk assessments to inform design and implementation of appropriate privacy solutions.

Decisions addressing privacy concerns require multiple perspectives, beyond technical, policy, and legal experts to organizational leadership and staff who understand the critical business functions of an organization. Everyone involved in using or overseeing the use of the data has a responsibility to understand privacy concerns.

[13] NIST. (2020, January). *NIST Privacy Framework: A Tool For Improving Privacy Through Enterprise Risk Management, Version 1.0*. https://csrc.nist.rip/publications/detail/white-paper/2020/01/16/nist-privacy-framework-version-10/final

Efforts to maximize adoption of privacy protections must also consider potential market inefficiencies, such as asymmetric information between consumers and producers, where there is an aspect of product quality which consumers have much less ability to evaluate compared to producers. Asymmetric information may lead to underinvestment in producing that aspect of quality—in this case, in producing privacy-preserving features. Learning from fields with similar potential market inefficiencies (e.g., product safety, environmental pollution) may provide some research models for the types and combination of solutions to address and support adoption of privacy protections.

Key research questions include:

- How are privacy requirements, policies and practices for research understood by people with different roles and at different levels of the organizations who are responsible for the business functions?
- How can organizations effectively and consistently design, implement, and adopt context-based privacy policies and practices?
- What novel models, methodologies, and approaches are needed to enable modeling of socio-technical dimensions of privacy from multiple disciplinary perspectives?
- How can multidisciplinary perspectives enhance the understanding of privacy challenges in the digital age?
- How can a multidisciplinary approach to privacy research be used to anticipate future privacy challenges?
- What are the best approaches and practices for integrating socio-technical, legal, and ethical considerations into privacy research?
- How can multidisciplinary research help in creating more inclusive and equitable privacy policies and practices?
- How can interdisciplinary dialogues be fostered to address privacy issues comprehensively?
- How do socio-economic, cultural, and other group differences influence privacy expectations, and how can a multidisciplinary approach address these variations?
- How can multidisciplinary perspectives accelerate the development and deployment of privacy-preserving technologies that address technical, ethical, and societal challenges?

## 3.2 Understand and Measure Privacy Preferences and Impacts

Privacy preferences are often diverse, context-specific, dynamic, difficult to predict, and difficult to measure. Research is needed to develop methods and technologies that provide the capabilities to characterize the various and evolving preferences, expectations, norms and rules for activities, socio-political considerations, information disclosure, and data flows in the digital realm that involve personal information. Effective techniques to understand privacy preferences, in context and over time, may not guarantee meeting privacy preferences. Further, there are cases where responding to preferences does not fit with system design or objectives. However, they will support the development of techniques that empower individuals to make informed privacy-related choices, identification of minimal privacy requirements for the purposes of building privacy-preserving systems, understanding how various privacy preferences impact participation with new technologies, and suitable remediation of privacy harms. Research is also needed to understand the overall impacts of privacy for society and the potential tensions between individual or group privacy preferences and societal goals or public good.

Explicit in the improved understanding of privacy preferences and impacts must also be the ability to define various privacy objectives (e.g., individual control, accountability, respect for context, and

transparency) and the ability to measure how information systems meet or do not meet those objectives. The measurements should aid individuals in helping them make informed privacy-related decisions as well as support machine-based analysis and reasoning. System designers and developers need to better understand what individuals, groups, and communities value regarding privacy; what people's privacy preferences and expectations are; in what ways privacy might be infringed upon; and expectations for remediation and recovery in case of privacy infringement, in order to develop systems that are more respectful of peoples' privacy choices. A better understanding of individuals' privacy awareness is also important for understanding what types of information must be given to individuals to enable them to make informed choices about their activities. Furthermore, greater awareness of privacy preferences and perceived deviations from them can inform social, legal, and system design policy.

Privacy preferences are subjective and can vary by generation, cultural subgroup, socioeconomic status, and other factors. These variations can make it difficult to draw general conclusions about current privacy norms or predict how these norms may develop over time. Similarly, privacy impacts are shaped by and evolve with the use of newer technologies and in new contexts. Some privacy impacts may occur as the result of a specific event (for example, a data breach) or the introduction of a new process or technology that helps to protect privacy or increases privacy risk. Other impacts may occur as the result of an accumulation of disparate data over time that, when combined, reveal privacy-sensitive information or result in inferences not possible in isolation. It is necessary to systematically identify and assess privacy impacts and desires and consider how they interact with other goals of individuals, organizations, and society as a whole such as convenience, cost savings, and utility for public health and safety.

Various privacy events occur when there are deviations from privacy rules, norms, desires, or expectations of a group or individuals. Research is needed to develop methods and technologies for better understanding, detecting, and assessing such deviations and privacy harms that may occur as a consequence. In particular, socio-technical solutions are needed to detect privacy violations when they are not directly identifiable by individuals. . Research addressing privacy preferences should also aim to clarify the wide range of effects of technology on individuals and society, including the chilling effects of data collections and unauthorized uses.

Recovering from privacy violations requires remedies. Frequently, there is no legal recourse or even legal recognition that a privacy violation has taken place. Existing recovery mechanisms are limited and are inconsistent in their efficacy. In addition to the technical difficulty of recovery, it is important to consider the varying privacy preferences and needs related to remediation and recovery. Privacy events may impact groups or individuals in varying levels of severity. Understanding the needs and expectations for remediation and recovery is necessary as the development and adoption of data collection and analysis continues to expand.

Innovative research and technology may promise to advance the well-being of society but disproportionately harm some groups or individuals in the case of a privacy violation. It is essential to the success of this Strategy to understand how privacy preferences impact participation with new technologies and how new technologies may introduce new privacy issues. For example, machine-learning models trained on vast sets of medical data can improve medical research and arguably advance the well-being of society. However, some individuals may not wish to participate, may prefer a particular approach to obtaining their permission to have their data included in the training of the model, or may wish to revoke their consent.

For these reasons, this strategy calls for research to develop fundamental techniques for understanding and measuring privacy preferences and impacts. Such research should include techniques for assessing the emergence, codification, and evolution of societal practices, attitudes, and beliefs regarding privacy and harms from privacy events, and suitable remediation of privacy harms. Addressing these issues must involve technological, behavioral, economic, cultural, social, educational, psychological, ethical, and historical perspectives and related analyses.

Key research questions include:

- What methods can organizations use to track and measure the effectiveness of their privacy policies and practices? What actions can they take to address challenges and resolve issues?
- What research methods are most reliably and validly sample, measure, and represent people's privacy preferences, expectations, attitudes, beliefs, and interests in one or more communities?
- To what extent do privacy preferences, expectations, attitudes, beliefs, and interests vary by generation, cultural subgroup, socioeconomic status, or other socio-demographic demarcations?
- How and why do privacy preferences, expectations, attitudes, beliefs, and interests change or evolve? Among groups or subgroups, do certain factors or experiences influence the emergence of privacy expectations and beliefs regarding privacy more than others, and if so, why?
- What incentives can effectively promote privacy and the adoption of privacy-preserving technologies and relevant policies and practices?
- What measurable impacts have privacy incentives had on a range of social values, such as social justice, economic growth and security, and innovation?
- To what extent do incentives, such as sharing personal data for access to "free" services, modulate privacy expectations, attitudes, beliefs, and interests?
- What methods and technologies could identify privacy incidents and other privacy impacts effectively and efficiently? What methods would be effective for disclosing this information to affected parties and systems, and where appropriate or required, to the government?
- How do privacy events become regarded as privacy harms by individuals or groups? How can privacy harms be recognized, measured, and assessed?
- How do privacy events affect peoples' behavior? How can the effects of privacy events be measured?
- What information and methods can effectively inform and enable decisions regarding people's privacy preferences in the policy, regulatory, and legislative domains?
- To what extent does the public understand how technological and economic factors affect their privacy, and to what extent do people understand power and information asymmetries between individuals and data collectors/users?
- How do different privacy preferences, expectations, attitudes, beliefs, socio-political context, and interests in other countries (if they exist) drive any differences in privacy laws and regulations in the U.S.?
- What kinds of formalisms could define privacy objectives and impacts, and what techniques and metrics could be used to measure how information processing systems meet those objectives?
- How can the relationship of privacy objectives and other objectives of individuals, organizations, and society (for example, the objective to ensure health and safety through sharing epidemiological data) be understood and assessed?

- How can the effects of privacy policy approaches on privacy events, both domestically and internationally, be evaluated?

## 3.3 Develop Methods and Methodologies to Incorporate Privacy Preferences, Requirements, and Controls into Systems

Incorporating privacy preferences into systems requires an interdisciplinary approach to organizing the technical and managerial effort required to balance privacy goals with stakeholder needs, expectations, and constraints. When systems process personal information, whether by collecting, analyzing, generating, disclosing, retaining, or otherwise using the information, they can impact the privacy of individuals or groups. System designers and engineers need to account for all the stakeholders in the overall development and deployment of the solution. However, designing for privacy does not today have parity with other disciplines when it comes to engineering solutions that capture the appropriate protections and stakeholder interests for privacy. Designing for privacy must connect privacy goals with organizational and system requirements, constraints and controls in a way that effectively bridges these goals with technical development.

System designers often lack appropriate foundational constructs, effective tools, and knowledge for designing systems that incorporate effective privacy requirements and controls. Even when designers consider privacy at the beginning of the design process[14], challenges remain for understanding and assessing the risks that a system might pose to privacy, for identifying and expressing privacy requirements for a system, and for designing controls that can achieve those goals. It is difficult for privacy models to capture quantifiable risk, especially across different groups and with predictive validity that captures changes in data collection and integration and evolving societal preferences, expectations and harms. For this reason, systems development requires an interdisciplinary team including engineers, data scientists, compliance experts, and others who will interact with and/or be accountable for different parts of the system.

Beyond risk identification and management, system designers also need consistent privacy objectives oriented around engineering processes to allow them to develop system-level requirements and capabilities to specify and enforce privacy policies and legal requirements. System owners are often faced with conflicts among various organizational objectives such as efficiency, cost, functionality, mission, and system quality attributes (e.g., security, and safety to both the individual and others on the system, privacy, etc.) that force them to make tradeoffs[15]. Research is needed to find approaches that will minimize such tradeoffs and allow engineers to identify solutions that maximize both privacy and other objectives to the greatest extent possible.

Furthermore, research should be aimed at developing tools to help system designers and participants in the ecosystem choose, test, and validate among different privacy controls, as well as developing approaches for integrating multiple privacy-preserving mechanisms to protect established privacy guarantees. For example, organizations might choose among various privacy preserving technologies, which are applied differently to achieve privacy goals, based on the specific situation.

Utilizing existing frameworks and tools for privacy engineering and risk management, research can advance around technical controls and how system designers and participants can most effectively

[14] NIST. *Privacy Risk Assessment Methodology (PRAM).* https://csrc.nist.gov/glossary/term/privacy_risk_assessment_methodology

[15] NIST. (2017, January). *An Introduction to Privacy Engineering and Risk Management in Federal Systems*. https://csrc.nist.gov/pubs/ir/8062/final

apply them in systems, but challenges persist in managing risks across domains, including cybersecurity and AI risk, where overlaps and tradeoffs may exist in a given context. Novel frameworks and tools such as the NIST Privacy Framework[16] and Privacy Risk Assessment Methodology are critical to supporting privacy risk management and privacy engineering in a quickly evolving environment. Making combined progress on implementation of frameworks, risk models, and technical controls will improve the capability to assess privacy risk in specific systems and compare the effectiveness of different privacy controls. Ultimately these techniques should make it possible to transform measurements into end-to-end determinations as to how processing of personal information affects privacy.

Many systems prioritize the prevention of privacy events, amongst other harms, but recovery and remediation techniques must be established to address a privacy harm, should one occur. Existing recovery mechanisms are limited and are inconsistent in their efficacy. The difficulty of recovery magnifies the importance of privacy risks and increases the impact of the information asymmetry between data subjects and data collectors/users. By understanding how data actions operate (collections, flows, uses, disclosures of certain information, etc.) and how they may result in harm, better approaches for recovery might be devised.

Research is needed to measure the efficacy of existing technical, economic, and legal redress mechanisms (e.g., credit freezes and monitoring, privacy-protection insurance, liability regimes for privacy compromises, criminal and civil restitution mechanism against criminal actors, etc.), and to evaluate the consequences of a lack of redress. New approaches for recovering from privacy incidents need to be developed that are fast, predictable, and easy to implement. For example, research is needed to develop approaches for more quickly recovering from data breaches and problematic releases. Remediation techniques might also provide the capabilities to correct or delete erroneous data about individuals, exclude improperly used data, and effect a change in the processing systems that caused the privacy events. For instance, there is a growing need for considering protections such as the "right to be forgotten"[17] and redress from increasingly powerful AI systems; machine unlearning is emerging as a technical approach to address such issues. Research is also needed to develop new and effective techniques to effect redress, such as rendering the data useless, as well as mechanisms to delete or "forget" information.

Research is needed to develop methodologies to connect or integrate evolving privacy goals to system design to foster privacy by design and default approaches to system development and deployment. In security, the threats are constantly changing. Consequently, practitioners approach security objectives methodically and use risk-based processes to account for changing threat environments. Likewise, it is important to define and mature consistent privacy-related objectives and processes that allow for the interchange between privacy goals and the evolving technology environment.

Key research questions include:

- How can privacy impacts to individuals and groups be accurately measured in context?
- What kinds of system properties can be associated with privacy to support the implementation of privacy principles and policies?

[16] NIST. *Privacy Framework*. https://www.nist.gov/privacy-framework
[17] EUR-Lex. *General Data Protection Regulation*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02016R0679-20160504

- How should privacy properties be characterized, and how can they be assessed qualified or quantified?
- What privacy design patterns and use cases describe common solutions that would assist system designers, particularly in emerging areas such as immersive technologies?
- How can privacy-enhancing cryptographic and non-cryptographic technologies and techniques be developed to scale, as well as be integrated into the functional requirements and standards that are already widely adopted in systems?
- What metrics can be used to assess the effectiveness of privacy controls?
- How can privacy risk be considered and controlled in concert with system and data utility needs?
- What metrics and measurements can measure both privacy and system utility and support the development of systems that can maximize both?
- What technological or socio-technical mechanisms would effectively remediate a privacy harm?
- How can the effectiveness of remediation and recovery mechanisms be evaluated in terms of their financial, psychological, and societal impact?
- What effect does the existence of remediation and recovery mechanisms have on the likelihood of privacy events?
- What are methods to systematically identify and assess the interaction of privacy impacts with other individuals, organizations, and society goals such as convenience, cost savings, or public health and safety?
- What methods can organizations use to evaluate the adequacy of their privacy policies in preventing privacy events?
- How can cryptographic techniques be used to fundamentally enable privacy autonomy in various scenarios?

## 3.4 Increase Transparency of Data Collection, Sharing, Use, and Retention

Individuals face considerable burdens in understanding today's complex and dynamic information ecosystem. While some information is collected from individuals in a relatively transparent fashion, a great deal of information may be collected without an individual's full knowledge and by data collectors with whom the individual has no relationship. The growing use of sensors in both the home and in public space for public safety, transportation, and environmental purposes has also resulted in the collection of vast amounts of data on individuals. Similarly, increasingly powerful AI models such as LLMs are being built by using huge amounts of data available on the internet. Because much of this data collection and use is invisible to individuals, they often are unaware of when data about them is collected or for what purposes it will be used. In addition, individuals often do not understand the extent to which data about them is shared with third parties. The consequences as a result of lack of awareness lead to the individual being unable to make informed decisions about the tradeoffs involved in sharing personal information in exchange for some personal or social benefit.

Research designed to increase transparency of data collection and use would enable individuals to better evaluate the privacy implications and potential benefits of their activities and would permit data collectors/users to develop data practices that respect and protect individuals' privacy preferences. Increased transparency of data collection and use will also enable privacy technologists to develop solutions that better address the needs of individuals and data collectors/users, and it will provide regulators with improved visibility into data collection and use activities.

Today, many data collectors disclose their data practices through privacy policies in accordance with laws and regulations. Public posting of privacy policies promotes data collectors' accountability for their practices; however, privacy policies are often difficult to locate, overloaded with jargon, and ambiguous or open-ended in their meaning, rendering them confusing and even incomprehensible. The burden on individuals to read and understand these policies is further compounded in the mobile context where, because of the small size of the device, a privacy policy may be spread out over many separate screens. Some data collectors/users have begun to experiment with innovative approaches such as "just-in-time" disclosure that provides small, understandable amounts of information at relevant points in a transaction. AI techniques are being explored to devise privacy assistants that can guide data subjects to manage the privacy of their data. However, more research is needed to determine how traditional and newer transparency mechanisms can be improved and to identify other promising methods of disclosure.

Data has also become very durable. Because electronic storage is inexpensive and takes up very little space, data collectors are not only collecting greater amounts of information than they have in the past, but they are also storing that information for longer periods. Accordingly, developing effective means for informing individuals about prospective uses of their information is critical in achieving information symmetry between people and data collectors/users.

In addition, there has been insufficient effort to develop means to increase consumer awareness and understanding of today's systems, business practices, and information flows. Greater understanding regarding specific business models, the tools available to individuals to control the collection and use of their data, and the benefits and privacy implications of various data uses would alleviate much of the existing information disparity between people and data collectors/users. Education and literacy around privacy are critical in today's digital age, where personal data is constantly being collected, shared, and analyzed. As individuals increasingly rely on digital platforms and technologies, it is essential that they understand how their information is being used and the potential risks involved. Privacy education empowers people to make informed decisions about their data, safeguarding their personal rights and autonomy.

Key research questions include:

- What type(s) of experimental studies and field trials should be used to discover information asymmetry that impact privacy?
- What approaches, tools, or automated systems need to be built to effectively and efficiently measure and report information flows? Is it possible to measure such flows without inherently producing more privacy risk?
- What techniques could be effective in informing individuals about the information practices and risks of data collectors/users, and in informing data collectors/users about the desires and privacy preferences of individuals?
- How can the format and lexicon for describing data practices (e.g., notice and choice) across industries be standardized, considering the inevitability of changes in technology over time? What other measures could improve individuals' ability to compare data practices across the range of data collectors/users, thereby encouraging competition on privacy issues?
- What might be the appropriate level of transparency and choice for prospective changes to data-handling practices? How can the impact of these changes be measured?
- How can individuals be meaningfully provided with notice about the practices of data collectors that collect and use data without directly interacting with individuals?

How can privacy policies be improved to ensure reader comprehension, including examination of the efficacy of disclosure attributes such as text, font, and icons or graphics?

- How can data collectors/users provide meaningful notice of their data practices on mobile and similar devices? How can effective "just-in-time" disclosures be constructed?
- In what situations is the traditional notice-and-choice approach ineffective without other types of protections?
- How should the effectiveness of transparency mechanisms be evaluated? To what extent do design choices used for transparency mechanisms impact user consent? How can such impacts be measured?
- What are new methods for increasing public understanding of data collection, retention, and sharing practices reflecting differences across industries and sectors?
- How can data collectors and data users ensure transparency to individuals about the collection and long-term reuse of their data beyond its originally stated purpose and incorporate their privacy preferences and adhering to standards?

## 3.5 Ensure That Information Flows and Use are Consistent with Privacy Rules

Protecting privacy requires that both individuals and organizations understand the rules that govern flows and use of personal data. They need to have confidence that those rules are observed in practice. Research is needed to advance technologies that can ensure that personal data are linked with the rules appropriate for the context in which they are collected and that operations applied to those data are governed by those rules.[18] Research is also needed to determine whether privacy rules for the output data could be derived from rules associated with the inputs, the processing, and the permissible use (context) for the outputs. Attaining such capabilities could require new computational models and languages for addressing these concerns.

For example, techniques are needed that allow data to be reliably tagged and processed in a way that preserves the context under which they were collected and are maintained. "Context" is a broad concept that might include a person's consent and preferences, legal or regulatory requirements, geographical location, or data sharing agreements. Such tags could capture the acceptable data uses signaled by the individuals and allow data collectors to communicate the request to subsequent users to honor both the person's permissions and the specific requirements for individual data to the extent possible. More broadly, these techniques should facilitate people's expression of privacy preferences and their implementation.

Effective approaches are needed to translate requirements that allow for automated enforcement by machines and can be understood by people. Ways to associate these rules with code to make them machine readable are also needed, so that other code can verify that rules are being faithfully enforced and so that the resulting data can be associated with the rules under which they were collected and processed. Together with active engagement of stakeholders, appropriate processes, and governance, these approaches can help create *accountable* systems where violations of privacy policy can be detected and made known to affected persons.

Improved technology for managing data use would make it possible for data-processing and storage organizations to determine, rapidly and reliably, if their handling of private information meets legal,

[18] PCAST. (2014, May). *Big Data and Privacy: A Technological Perspective*. https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/PCAST/pcast_big_data_and_privacy_-_may_2014.pdf

regulatory, and ethical standards. Such technology would have the additional benefits of facilitating compliance from start to finish, including monitoring the system during its life and identifying potential privacy compliance issues including violations.

These approaches will help ensure that the responsibility for using personal data in accordance with the person's preferences will rest with data collectors, processors, aggregators and service providers who will balance these preferences with organizational objectives, considering the legal and statutory allowances and constraints. These approaches will also help support social norms and deter inappropriate data actions.

Accompanying research and technical solutions for privacy should focus on specification and enforcement of privacy policies, and regulations that govern data use. Data and data systems have evolved rapidly in recent decades, but the laws have been slower to catch up allowing inconsistent application of privacy laws and policies.

Key research questions include:

- What models and methods are needed to specify fine-grained contextually based privacy access and data use policies and what adaptable and sustainable in addition to effective and scalable to reflect that these mechanisms need to be able to adapt and be sustainable through the continual advances of the technology policies, and law?
- What are usable methods for specifying and managing information-flow based controls?
- How can hardware or software methods for establishing trustworthy execution environments support secure management of information flows and compliance with privacy policies?
- How can methods for tracking, assuring, and archiving the provenance of data and software components be used to assure privacy compliance?
- How can data provenance be implemented in a way that does not violate privacy itself?
- What program analysis methods can be developed for various kinds of information flow properties and privacy policy languages that are meaningful to legal experts, yet have precise semantics that system developers can use to restrict and provide accountability for how their code operates on personal information of users?
- What are effective methods for understanding the flow of personal data through systems of computer programs?
- In what ways can privacy rules for the results of data processing be derived from privacy rules of the inputs, processing, and context?
- How can the change in value or sensitivity of data, as they are combined with other information, be accounted for and properly acted upon by information processing systems?
- How can access control systems that incorporate usage-based and purpose-based constraints be adapted to the range of privacy issues now faced by system designers?
- What are effective information disclosure controls, methods for de-identifying data, and means for assessing these de-identification methods?
- How can anonymous and pseudonymous computing, computing with obscured or encrypted data, and management of multiple identities be made efficient and practical?
- How can existing Internet infrastructure and protocols be redesigned to better support privacy (i.e., support anonymous, censorship-resistant, and metadata-hiding communications)? Can privacy be built into core Internet services without adversely affecting cybersecurity?
- How should data use, and privacy policies be updated on a national level that will help organizations balance data use with individual privacy?

- How can AI technologies be leveraged to automate understanding and analysis of legal or regulatory requirements, as well as machine-readable policies to support privacy policy enforcement or compliance?
- What approaches are needed to facilitate scalable, effective, and privacy-preserving cross-border information flow?

## 3.6 Reduce Privacy Risks of Data Analytics and AI

Rapid advances in computing have led to the fast development of advanced data analytics and AI techniques. These analytics and AI algorithms can be used to analyze and predict human behavior and performance, lead to data-driven scientific discoveries, detect fraud, or perform other important functions. They are rapidly being used in both public and private sectors. The accelerated development of generative AI in recent years further adds to the AI capabilities to synthetize new content based on AI models trained on data at scale. These predictive and generative AI algorithms can benefit or harm individuals by categorizing a person in ways that enhance or limit their options and opportunities.

Increasingly, analytical and AI algorithms, including LLMs and foundation models, are being developed to process or be trained on large-scale multi-modal data from many sources and data scraped from the web. Resulting models can leak private data that are retained from training. These algorithms are increasingly being embedded or implemented in many systems or applications to be used for making critical decisions based on the results of the algorithmic determination. The analytical and AI algorithms are used for prioritizing, classifying, filtering, and predicting so as to gain deeper insights from data or "generate" new content based on data that are used for these algorithmic processes of analytics, and/or training and inference. Their use can raise significant privacy issues when, for example, the information used by algorithms is inappropriate or inaccurate, when data used for analytics or to train AI models leads to incorrect decisions, when an individual's autonomy is directly related to algorithmic scoring or outcomes, or when the use of predictive or generative algorithms chill desirable behavior or encourages other privacy harms. In addition to privacy concerns, emerging generative AI algorithms are known to generate non-factual content (sometimes described as "hallucinating") that pose significant potential risks from deepfakes and disinformation that further exacerbate privacy concerns.

There are gaps in public knowledge about the range of increasingly powerful data-intensive analytical and AI algorithms that are in use, what they are used for, and their susceptibility to error and misuse. If not carefully designed, such algorithms have been shown to have disparate impact on different socio-demographic groups even when the algorithm does not explicitly use those attributes, and can result in different levels of privacy protection or risks for different groups represented in the data fed to the algorithms. Even when privacy-preserving technologies are employed, such algorithms have been shown to potentially amplify bias and exacerbate disparate impacts on different sub-groups represented in data. Use of such algorithms for employment, housing, policing, and other critical areas can have implications for federal equal opportunity laws and demand greater algorithmic transparency from the perspective of both privacy and fairness. Indeed, the lack of transparency around companies providing consumer data for credit and other eligibility determinations led to the adoption of the Fair Credit Reporting Act, passed in 1970. However, it is difficult, and sometimes infeasible, for those using these algorithms to know if they are producing a disparate impact. Outcomes-based studies have identified these issues in recent years, but such studies take substantial time and effort and may not be feasible when an algorithm is re-trained on a weekly or daily basis, or fine-tuning is done in foundation AI models for different downstream tasks.

Many anticipated uses of analytical or AI algorithms require that the outcomes of the algorithms be explainable for reasons of accountability, transparency, and auditing. In some cases, it may be appropriate (or legally required) for individuals to be able to control whether certain types of data are used in decision-making. For instance, the Equal Credit Opportunity Act of 1974 prohibits credit discrimination on the basis of race, color, religion, national origin, sex, marital status, age, or because a person receives public assistance. However, many analytical algorithms in use today provide little clarity in these areas. Data protection and privacy regulations, such as the California Consumer Privacy Act (CCPA) and the General Data Protection Regulation (GDPR), often include various rights, such as the right to be forgotten, that generate challenges for these algorithms by requiring complex machine unlearning techniques, further adding to the privacy challenges.

Research is needed to understand the current and planned usage of these algorithms, as well as to develop methods to increase transparency and improve accountability when these algorithms are employed. Improved capabilities are also needed to understand people's concerns, the type and extent of control that is feasible including through various rights allowed by data protection and privacy laws and regulations, and how to present information to both application developers and end users as new applications of predictive or generative algorithms arise. Techniques are also needed to detect, correct, and redress errors or harm that these algorithms might cause.

Key research questions include:

- What novel privacy-preserving techniques can be employed to ensure the end-to-end privacy safeguards of various analytics or AI/ML solutions for different application contexts?
- In what ways do analytical and AI algorithms and systems that act upon the results of the algorithms adversely affect the privacy of individuals or groups of people?
- What types of privacy concerns do individuals have with respect to analytical and predictive algorithms, and what information do they need to initially and ongoingly address these concerns?
- How can the provenance, accuracy, and quality of data used in making a decision or a prediction about an individual or groups be assessed for privacy issues, especially in small sub-populations?
- What are the privacy impacts on individuals or groups when analytical and AI algorithms use erroneous or inaccurate data? How can various rights such as the right to be forgotten or right to redress be supported by such analytical or AI techniques?
- What are the impacts of analytical or AI algorithms on individuals' autonomy and agency (i.e., the ability to make independent and free choices) and how do such impacts cause harm to people's privacy?
- How can foundation models be developed and continually updated in a privacy-preserving manner? How do we guarantee that privacy protection guarantees achieved during the training of foundation models are maintained when they are fine-tuned for different tasks?
- What techniques are needed to ensure that both privacy and fairness issues are addressed when employing analytical and AI algorithms?
- What privacy-by-design approaches are needed to ensure privacy protection is addressed for analytical and AI algorithms along with other contending issues such as security, bias, economics/incentives to achieve practical benefits while minimizing potential harms?
- What privacy auditing techniques are needed for analytical and AI algorithms, and how can systematic red-teaming approaches be established to ensure proper understanding of privacy risks and compliance?

- How can new technologies and algorithms, and combinations of technologies and algorithms, provide practical and theoretical privacy-preserving data analytics or machine learning?
- What effect does the use of remediation and recovery have on the investment in more robust privacy technologies including privacy-preserving analytics and ML?
- How could privacy-protecting and privacy-recovery technologies be integrated into algorithms and systems to create more effective and efficient solutions?
- What foundational approaches are needed to understand or anticipate and address new privacy issues introduced by emerging technologies such as artificial general intelligence, cognitive machine learning, and quantum machine learning?

## 4. Executing the National Privacy Research Strategy

This strategy presents privacy research priorities based on a joint assessment by federal agencies participating in the federal Networking and Information Technology Research and Development (NITRD) Program. As a strategic plan, this document provides guidance to the Executive Branch, policymakers, researchers, and the public in determining how to direct resources into R&D activities that have the greatest potential to generate the greatest impact. The strategy is not intended to provide a detailed roadmap of national privacy research activities. It is each agency's responsibility to incorporate these research priorities into its research plans and programs, drawing on its individual strengths and in the context of its mission.

The execution of the National Privacy Research Strategy vests in the federal agencies, which develop and execute R&D activities based on their missions and capabilities, provide leadership across sectors to focus on critical national research and development needs, and advance fundamental research in federal laboratories and through funding at academic institutions and private research firms. The NITRD Program coordinates federal research investments in various areas of IT through its interagency working groups. In particular, the NITRD Program will ensure that federal privacy research is well coordinated by helping agencies understand each other's activities, by supporting agencies in minimizing duplication and gaps, promoting and sharing best practices, maximizing impact, by supporting multi-agency collaboration, and by considering how to align the overall NITRD privacy research portfolio with this strategy.

The Strategy also provides research entities outside of government – academic and private – with a set of focused open questions. Privacy research funded under this strategy, both government-led and academic or private sector-led, can have a broad range of effects. Research on current practices in the information ecosystem can inform the public debate on privacy issues and provide useful information to policymakers and leaders in academia and the private sector. Research that creates new privacy theory and models creates intellectual frameworks that can help individuals understand privacy, guide the creation of privacy-preserving technologies, and serve as the basis for further theoretical development. Work on new privacy-preserving approaches creates foundational theory, prototypes of socio-technical privacy solutions, and products that can be used to help society to realize the benefits of increasingly data centric digital environments without sacrificing privacy or endangering individuals or the public.

Among the first steps in executing the strategy should be a comprehensive review of literature and studies across sectors to assess existing knowledge relevant to the research priorities defined in this plan. Identifying and connecting the variety of foundational, use-inspired, and translational research in

privacy in the many sectors and domains where such work is conducted would be a valuable contribution of the NPRS.

As part of the national strategy, funding agencies are strongly encouraged to create opportunities for researchers to meet with potential users of the research and the public throughout the research process to ensure that research remains aligned with real-world needs and requirements. These opportunities can include "matchmaking" events for researchers to discuss their work, and for potential users to discuss their needs and requirements, ensuring ongoing relationships between researchers, potential customers, and the public, creating opportunities for testing prototypes on real data, and providing governmental assistance for pilot studies and field testing. Funders should encourage those submitting proposals to have clear plans for technology transfer at the successful conclusion of a research project.

Funding agencies should also explicitly account for the multidisciplinary nature of privacy and enable research that requires joint contributions from two or more disciplines. Various existing funding programs such as NSF's Secure and Trustworthy Cyberspace (SaTC) and the recent Privacy-preserving Data Sharing in Practice (PDaSP) are examples of federal initiatives that have been and continue to foster foundational, use-inspired and translational research. Such initiatives should be further expanded and strengthened to establish a robust funding ecosystem to achieve the research goals of this strategy.

While many privacy-preserving techniques and solutions are developed for a specific application, they can frequently be applied in other areas or generalized to broader classes of problems. NITRD agencies are therefore strongly encouraged to create or support the creation of catalogs, or other shared mechanisms, of privacy-preserving solutions so that such solutions can be shared among agencies and with the public. To help ensure that new and better methods and tools are adopted, the government may need to create incentives or requirements for adoption.

To enable and support the research goals in this strategy, organizations must understand the context for research versus real world implementation and deployment. These contexts are vastly different and should be treated with care so as to not prevent or severely delay research. Different business functions may also have different goals, but also different privacy concerns. The challenge is to develop widely understood, consistent but context-based privacy policies and practices to enable research and successful implementation of technology solutions in the most efficient and effective manner possible.

Finally, while this strategy does not directly address many specific privacy policies and practices, no research or any mission can ignore privacy concerns and be considered successful. Leadership and stakeholders at all levels should understand privacy concerns in context and not relegate these decisions to the sole domain of policy and legal experts. Federal agencies should establish efficient and inclusive governance where privacy protections are the norm to foster developing and successfully transitioning new technologies into use with privacy by design from the beginning. Support for privacy by design includes adequately resourcing the teams conducting this work proportionately to growing privacy requirements.[19]

[19]In the federal space, existing privacy requirements include the Privacy Act of 1974, the E-government Act of 2002, Federal Information Security Modernization Act (FISMA), NIST 800-Rev 5. With the growth of AI, additional privacy reviews must be completed to ensure protection from privacy risks related to AI, pursuant to Executive Order 14110 and the Office of Management and Budget Memoranda 24-10 and 24-18.

## Appendix A: National Privacy Research Strategy Background

Efforts by the federal government to protect the privacy of individuals are numerous including, for example, the strict confidentiality provisions of the 1929 Census Act which made a disclosure of private information by an agent of the Census Bureau a felony, punishable with up to two years of imprisonment. The federal government enacted the Privacy Act of 1974 (the Privacy Act") to engender public trust in personal information collection, handling, and use. Continuing in this vein, in 2012, the document *Consumer Data Privacy in a Networked World: a Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy*[20] articulates policy on consumer privacy, and the subsequent discussion draft of a legislative proposal[21] suggests a path forward to address privacy challenges in today's information technology-driven world.

The technological challenges and opportunities in protecting privacy have received increased attention as well. The President's Council of Advisors on Science and Technology (PCAST) 2015,[22] 2013,[23] and 2010[24] reviews of the NITRD Program[25] have identified challenges to personal privacy in the digital era as a significant impairment undermining societal benefits from large-scale deployments of networking and IT systems. Underscoring the impairment of societal benefits, a national survey[26] sponsored by the National Telecommunications and Information Administration (NTIA) revealed that 45% of online households have been deterred from participating in online activities such as conducting financial transactions, buying goods or services, or expressing opinions on controversial issues via the internet, due to concerns about online privacy and security.

Consequently, PCAST has called upon federal research agencies to create a multi-agency initiative focused on developing scientific and engineering foundations for protecting privacy, which could then be the basis for new technologies and solutions in this space.

In 2014, the National Coordination Office (NCO) for the NITRD Program surveyed federal agencies to assess the size and scope of federally funded privacy research activities. It identified investments of

---

[20] The White House. (2012, February). *Consumer Data Privacy in a Networked World: a Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy.* https://obamawhitehouse.archives.gov/sites/default/files/privacy-final.pdf

[21] The White House. (2015, February). *Administration Discussion Draft: Consumer Privacy Bill of Rights Act.* https://obamawhitehouse.archives.gov/sites/default/files/omb/legislative/letters/cpbr-act-of-2015-discussion-draft.pdf

[22] PCAST. (2015, August). *Report to the President and Congress: Ensuring Leadership in Federally Funded Research and Development in Information Technology.* https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/PCAST/nitrd_report_aug_2015.pdf

[23] PCAST. (2013, January). *Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology*. https://www.nitrd.gov/pubs/PCAST-NITRD-report-2013.pdf

[24] PCAST. (2010, December). *Designing a Digital Future: Federally Funded Research and Development Networking and Information Technology*. https://www.nitrd.gov/pubs/PCAST-NITRD-report-2010.pdf

[25] Networking and Information Technology Research and Development (NITRD) Program provides a framework in which many US Government agencies come together to coordinate networking and information technology research and development efforts. More information is available at http://www.nitrd.gov

[26] NTIA. (2016, May). *Lack of Trust in Internet Privacy and Security May Deter Economic and Other Online Activities*. https://www.ntia.doc.gov/blog/2016/lack-trust-internet-privacy-and-security-may-deter-economic-and-other-online-activities

approximately $80 million/year in R&D activities across a broad spectrum of topics and interests related to privacy. The resulting document, *Report on Privacy Research within NITRD*,[27] provides a summary of the survey. The review showed that there are many innovative research projects within NITRD that are classified by their agencies as relevant to a broad range of privacy challenges. At the same time, the survey demonstrated the need for an interagency research framework that will help maximize research impact and ensure the coordination of R&D investments in this area.

Consequently, NITRD began examining both governmental and societal needs in privacy-enhancing technologies and began defining a framework for research to guide federal R&D investments in this area. In September 2014, the NITRD Cyber Security and Information Assurance Research and Development Senior Steering Group (CSIA R&D SSG) convened a task group of representatives from various agencies, including Air Force Office of Scientific Research (AFOSR), Census Bureau, Defense Advanced Research Projects Agency (DARPA), Department of Homeland Security (DHS), Department of Energy (DOE), Federal Bureau of Investigation (FBI), Federal Trade Commission (FTC), Intelligence Advanced Research Projects Activity (IARPA), Office of the Director of National Intelligence (ODNI), Office of Naval Research (ONR), Office of the Secretary of Defense (OSD), Office of Science and Technology Policy (OSTP), National Institute of Standards and Technology (NIST), National Institutes of Health (NIH), National Security Agency (NSA), and the National Science Foundation (NSF). CSIA R&D SSG tasked the group with developing a strategy to establish objectives and prioritization guidance for federally funded privacy research, providing a framework for coordinating R&D in privacy-enhancing technologies, and encouraging multi-disciplinary research that recognizes the responsibilities of the Government and the needs of society, as well as enhances opportunities for innovation in the digital realm.

The task group reviewed agency needs and existing research activities related to privacy. The group also obtained public input in three ways: (1) by issuing a Request For Information published in the Federal Register in September 2014, (2) by hosting a National Privacy Research Strategy Workshop in Arlington, Virginia in February 2015, and (3) by reviewing the report *Towards a Privacy Research Roadmap for the Computing Community* prepared by the Computing Community Consortium in May 2015. Details of these engagements are available on the NITRD website.[28]

[27] National Coordination Office for NITRD. (2014, April). *Report on Privacy Research within NITRD," National Coordination Office for NITRD*. https://www.nitrd.gov/Pubs/Report_on_Privacy_Research_within_NITRD.pdf

[28] NITRD. *National Privacy Research Strategy*. https://www.nitrd.gov/coordination-areas/privacy-rd/national-privacy-research-strategy/

## Appendix B: Legal and Policy Context for Privacy

The U.S. privacy regulatory structure encompasses three basic areas: regulation of commercial entities, government delivery of services, and national security and law enforcement. Each of these areas has a long history of law and policymaking aimed at protecting individual privacy from intrusions by private and governmental actors. These existing laws and policy approaches have begun the work of developing a conceptual basis for privacy, articulating basic expectations and values, and establishing principles such as use limitation and access.

When considering privacy in the public sector, the Fair Information Practice Principles (FIPPs)[29] have shaped federal laws, regulations, and guidance. The Privacy Act is the foundation for privacy protection at the federal level, and there are similar statutes among the states. The Privacy Act establishes obligations for federal agencies to limit information collection and maintain accurate information about systems of records, about conditions for disclosure, about provisions for individuals' access to their information, as well as requirements for how data can be shared among separate systems of records. The Privacy Act is often augmented at the agency level through additional statutes or regulations that specifically protect materials such as tax information, census filings, student information, and other kinds of information and, in the process, reflecting various FIPPs principles such as use limitation, purpose specification, security safeguards, and accountability. "Appendix-J" of NIST Special Publication 800-53 Revision 4, Security and Privacy Controls for Federal Information Systems and Organizations,[30] describes 25 different privacy controls that have been implemented by the federal government, providing the agencies with supplemental guidance and the appropriate legislative justification for each one. Based on the FIPPs and reflecting best practices, the privacy control catalog is intended to complement and augment federal information security programs and reflects the ever-increasing importance of the intersection of privacy and information security programs.

Regulation of commercial actors has become an area of tremendous importance in the U.S. privacy structure as advances in IT have led to novel commercial uses of personal information across a variety of industries. Whereas the privacy laws of many other nations protect all personal data broadly, the U.S. consumer data protection structure has no comprehensive statutory protection specifically addressing privacy across all sectors. Instead, the U.S. approach is sectoral, with most data privacy statutes only applying to specific sectors such as health care, education, communications, and financial services. The sectoral approach permits controls tailored to particular context but can also leave gaps. For instance, between 1974 and 2004, the United States passed legislation providing significant privacy protections

[29] First presented in the "Records, Computers and the Rights of Citizens," Report of the Secretary's Advisory Committee on Automated Personal Data Systems, U.S. Department of Health, Education & Welfare, July 1973, https://www.justice.gov/opcl/docs/rec-com-rights.pdf. The principles were subsequently tailored by policy documents, such as by the "Privacy Policy Guidance Memorandum (2008), Department of Homeland Security, http://www.dhs.gov/xlibrary/assets/privacy/privacy_policyguide_2008-01.pdf, by the "National Strategy for Trusted Identities in Cyberspace: Enhancing Online Choice, Efficiency, Security, and Privacy," The White House, April 2011, https://obamawhitehouse.archives.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf, and by the "Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy," The White House, February 2012, https://obamawhitehouse.archives.gov/sites/default/files/privacy-final.pdf.

[30] NIST. (2013, April). *Security and Privacy Controls for Federal Information Systems and Organizations," NIST Special Publication 800-53 Revision 4*. http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf.

for consumer information in government databanks (1974),[31] educational records (1974),[32] financial records (1978),[33] cable television records (1984),[34] e-mail (1986),[35] video rental records (1988),[36] unwanted phone calls (1991),[37] driver's license records (1994),[38] healthcare records (1996),[39] telecommunications data (1996),[40] information collected from children online (2001),[41] and satellite television records (2004).[42] In each of these cases, Congress protected information that was collected during the course of obtaining services commonly used by citizens. In addition, the Federal Trade Commission (FTC) can take action against companies engaged in "unfair or deceptive" privacy practices where they, for example, make false or misleading claims about privacy or data security or fail to employ reasonable security measures and, as a result, cause or are likely to cause substantial consumer injury.

In the United States, self-regulation has played an important role in helping to police commercial markets. Self-regulation through trade associations and certification programs can frequently prove to be more capable of adapting more quickly and in a more tailored fashion than government regulation. Self-regulation is a market-based solution that can quickly reward players who deliver products and policies responsive to consumer needs and desires. In addition, self-regulation can handle a variety of tasks, creating rules, playing a role in enforcement, and/or being involved in adjudication. The notice-and-choice model, based on the right to know about what data is collected and to consent (or withhold consent) from its collection and use, encourages companies to develop privacy policies describing their information collection and use practices so that individuals can make informed choices. Some critics claim, however, that self-regulation, and in particular the notice-and-choice model on which it relies, has failed to provide meaningful protection. Instead of providing transparency and empowering individuals with market choices, critics argue that this model has led to long, incomprehensible privacy policies that individuals do not read and have difficulty understanding and are often substantially more expansive than the actual and expected use of the data. In extreme cases, notice-and-choice has allowed players to engage in aggressive data sharing practices as long as the practices are documented, and the consumers give their consent.

In keeping with its mission of promoting free market competition while preventing "deceptive or unfair practices," the FTC has established itself as a backstop in the self-regulatory scheme. If a company deceives consumers about its compliance with a self-regulatory scheme, the FTC can take action alleging a deceptive practice under the FTC Act. State attorneys general have similar consumer protection authorities and play an important role in collaboration with the FTC.

[31] The Privacy Act of 1974, 5 U.S.C. § 552a.
[32] The Family Educational Rights and Privacy Act (FERPA), 20 U.S.C. § 1232g.
[33] The Right to Financial Privacy Act of 1978, 12 U.S.C. 3401.
[34] The Cable Communications Policy Act of 1984, 47 U.S.C. ch. 5, subch. V–A.
[35] The Electronic Communications Privacy Act of 1986 (ECPA), 18 U.S.C. § 2510-22.
[36] The Video Privacy Protection Act of 1988, 18 U.S.C. § 2710.
[37] The Telephone Consumer Protection Act of 1991 (TCPA), 47 U.S.C. § 227.
[38] The Driver's Privacy Protection Act of 1994, 18 U.S.C. § 2725.
[39] The Health Insurance Portability and Accountability Act of 1996 (HIPAA), Pub.L. 104–191.
[40] The Telecommunications Act of 1996, 47 U.S.C. § 222.
[41] The Children's Online Privacy Protection Act of 1998 (COPPA), 15 U.S.C. § 6501–6506.
[42] Carriage of local television signals by satellite carriers, 47 U.S.C. § 338.

A white paper entitled *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy was released in 2012*.[43] This paper described a four-point strategy for protecting consumer privacy: the creation of a Consumer Privacy Bill of Rights (CPBR); fostering multistakeholder processes to develop enforceable codes of conduct; strengthening FTC enforcement; and improving global interoperability. The CPBR laid out general principles, including respect for context and individual control, among others, that afforded companies discretion in how they were implemented. Legislation was recommended to codify the CPBR and implement a process for the creation of codes of conduct through voluntary participation in multistakeholder processes. The proposed legislation would set forth a process through which the FTC could grant safe harbor status to these codes. Finally, the white paper laid out the goals of increasing global interoperability of privacy enforcement. This framework was put forward in actionable form in the 2015 Consumer Privacy Bill of Rights Act Discussion Draft.[44] While this draft was not taken up by Congress, the 2012 Administration continued in its belief that it presents the best way forward to both protect and promote consumer privacy and trust while maintaining the flexibility needed to promote innovation and growth.

Another area that has been a significant focus of privacy law and policymaking in the United States is law enforcement and national security. Today, some law enforcement and intelligence agencies have the ability, subject to lawful due process and oversight, to collect, connect, and analyze a wide array of data that may facilitate the creation of a "virtual picture" of individuals to help with solving crimes, preventing attacks, and tracking terrorists.

Recognizing the potential privacy concerns that such activities can raise, these activities are bound by the rule of law, and, in many respects, subject to both Congressional and judicial oversight as well. The legal constraints include Constitutional protections such as the First Amendment's protection for freedom of speech, the Fourth Amendment's requirement of judicial review and prohibition on unreasonable searches and seizures, as well as the Fourteen Amendment's guarantee of due process of law for all. In addition, all federal executive agencies are also bound by statutory laws such as the Wiretap Act, the Electronic Communications Privacy Act (ECPA), the Privacy Act, and the Foreign Intelligence Surveillance Act (FISA) which control and condition government's particularized access to personal information.

Establishing an effective approach to privacy protection that allows individuals to realize the benefits of information technology without compromising their privacy has been difficult—in part, because of differences in individuals' understanding, attitudes, expectations, and behavior, as well as the rapid pace of change in technology. By focusing research efforts on these challenges and prioritizing the translation of research results into government policy and commercial imperatives, the NPRS aims to meet and overcome the challenges that have confronted policy- and lawmaking on privacy issues to date.

In recent years, the discussion of comprehensive privacy laws has been invigorated, in light of the introduction of the European Union's (EU) General Data Protection Regulation (GDPR) in 2018. The

[43] The White House. (2012, February). *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy*. https://obamawhitehouse.archives.gov/sites/default/files/privacy-final.pdf

[44] The White House. (2015, February). *Administration Discussion Draft: Consumer Privacy Bill of Rights Act.* https://obamawhitehouse.archives.gov/sites/default/files/omb/legislative/letters/cpbr-act-of-2015-discussion-draft.pdf

GDPR introduced significant changes to data protection laws in the enhancing and harmonizing privacy rights for individuals and imposing obligations on organizations handling personal data. Since its implementation, the GDPR has influenced discussions and actions related to privacy in the United States. GDPR's emphasis on individuals' rights to access, rectify, delete, and control their personal data has sparked discussions around similar rights for U.S. consumers. Some states have introduced or updated privacy laws, such as the California Consumer Privacy Act (CCPA) and subsequent amendments, offering enhanced rights to consumers regarding their personal data. The GDPR has also served as a model for some state-level privacy legislation in the U.S. For instance, the California Privacy Rights Act (CPRA) expanded upon the CCPA, introducing stricter data protection rules and establishing a dedicated enforcement agency. Finally, the GDPR has contributed to discussions and calls for a comprehensive federal privacy law in the United States.

## Appendix C: National Privacy Research Strategy Working Group (2016)

**Marjory Blumenthal**, Office of Science and Technology Policy

**Sean Brooks**, National Institute of Standards and Technology

**Chris Clifton**, National Science Foundation

**Milton Corn**, National Institutes of Health

**Lorrie Cranor**, Federal Trade Commission

**Ed Doray**, Department of Defense

**Simson Garfinkel**, National Institute of Standards and Technology

**Marc Groman**, Office of Management and Budget

**Karyn Higa-Smith**, Department of Homeland Security

**Christa Jones**, Census Bureau

**Anthony Kelly**, National Science Foundation

**Erin Kenneally**, Department of Homeland Security

**Eva Kleederman**, Office of the Director of National Intelligence

**Carl Landwehr**, National Security Agency

**John Launchbury**, Defense Advanced Research Projects Agency

**Naomi Lefkovitz**, National Institute of Standards and Technology

**Jessica Lyon**, Federal Trade Commission

**David Marcos**, National Security Agency

**Keith Marzullo**, National Coordination Office for NITRD

**Gregg Motta**, Federal Bureau of Investigation

**Tristan Nguyen**, Air Force Office of Scientific Research

**Eugene Sullivan**, National Security Agency

**Ralph Wachter**, National Science Foundation

**Heng Xu**, National Science Foundation

*Staff*

**Tomas Vagoun**, National Coordination Office for Networking and Information Technology Research and Development