DISCUSSION DRAFT

NIST PRIVACY FRAMEWORK: AN ENTERPRISE RISK MANAGEMENT TOOL

April 30, 2019

National Institute of Standards and Technology
U.S. Department of Commerce

Note to Reviewers

This document is provided for discussion purposes to promote the development of the NIST Privacy Framework: An Enterprise Risk Management Tool (Privacy Framework). NIST will use feedback on this discussion draft to develop a preliminary draft of the Privacy Framework.

Structure: Based on stakeholder feedback, this discussion draft is aligned with the structure of the Framework for Improving Critical Infrastructure Cybersecurity (Cybersecurity Framework) to support compatibility between the two frameworks. Feedback also supported use of additional organizing constructs referenced in NIST's Request for Information, such as privacy principles (e.g., the Fair Information Practice Principles), the information life cycle, and the NIST privacy engineering objectives (i.e., predictability, manageability, disassociability) or other constructs.[1] NIST welcomes feedback on how well these concepts have been integrated, as well as whether the Privacy Framework could be effectively implemented independently or in conjunction with the Cybersecurity Framework.

Privacy Risk Management: Based on feedback indicating a lack of a consistent or widespread understanding of privacy risks and privacy risk management, this discussion draft provides guidance on these topics in section 1.2 and Appendix D. NIST welcomes feedback on whether this guidance will be useful to organizations.

Core: This discussion draft provides a proposed Core, including functions, categories, and subcategories. NIST welcomes feedback on the Core, particularly regarding (i) gaps in, clarifications to, or usefulness of the categories and subcategories, (ii) organization of the functions, categories, and subcategories, and (iii) areas that need further development and may be more appropriate for the Roadmap section in Appendix F.

Informative References: This discussion draft defines informative references as specific sections of standards, guidelines, and practices that can be mapped to the Core subcategories and support achievement of the subcategory outcomes. In an effort to increase contributions of informative references and simplify updating, NIST is providing a mapping of the Core to relevant NIST guidance as a separate, companion document to this discussion draft. In addition, NIST will develop a process for accepting external informative references. NIST welcomes feedback regarding this approach to informative references.

Overall Discussion Draft: In general, NIST is interested in whether the Privacy Framework as proposed in this discussion draft could be readily usable as part of an enterprise's broader risk management processes and scalable to organizations of various sizes—and if not, how it could be improved to suit a greater range of organizations. Although these notes highlight key areas of interest, all feedback is welcome.

Please send feedback on this discussion draft to privacyframework@nist.gov.

[1] See Federal Register Notice 83 FR 56824, *Developing a Privacy Framework* at https://www.federalregister.gov/documents/2018/11/14/2018-24714/developing-a-privacy-framework.

Table of Contents

List of Figures

List of Tables

Executive Summary

The Executive Summary will be included in the preliminary draft.

Acknowledgements

Acknowledgements will be included in the preliminary draft.

1.0 Privacy Framework Introduction

For more than two decades in the United States and across the world, the Internet and associated information technologies have driven unprecedented innovation, economic value, and access to social services. Many of these benefits are fueled by data about individuals that flow through a complex ecosystem—so complex that individuals may not be able to understand or attend to the potential consequences for privacy resulting from their interactions with systems, products, and services. Similarly, organizations may not fully realize the consequences either. Failure to manage privacy risks can have direct adverse consequences for people at both the individual and societal level, with secondary effects on organizations and economic growth. Finding ways to continue to derive benefits from data while simultaneously protecting individuals' privacy is challenging, and not well-suited to one-size-fits-all solutions.

Approaches to Privacy are challenging because it is an all-encompassing concept. It is a condition or state that safeguards important values such as human autonomy and dignity, yet the means for achieving it vary. For example, in some circumstances it can be achieved through obscurity, in other circumstances through individuals' control of various facets of their identities (e.g. body, data, reputation).[2] Moreover, human autonomy and dignity are not fixed, quantifiable constructs; they are mediated through cultural diversity and individual differences. This broad and shifting nature of privacy makes it difficult to communicate clearly about privacy risks within and between organizations and with individuals. What has been missing is a shared lexicon and practical structure that is flexible enough to address diverse privacy needs.

To enable innovation and increase trust in systems, products and services, NIST has developed the voluntary NIST Privacy Framework: An Enterprise Risk Management Tool (Privacy Framework) to help organizations consider:

- How their systems, products, and services affect individuals; and
- How to integrate privacy practices into their organizational processes that result in effective solutions to mitigate these impacts and protect individuals' privacy.

The Privacy Framework has been developed to improve privacy risk management for organizations delivering or using *data processing* systems, products, or services in any sector of the economy or society, regardless of their focus or size. The common taxonomy that it provides is neither country- nor region-specific. Organizations outside the United States may also use the Privacy Framework to strengthen their own privacy efforts, and ideally, it can contribute to developing a common language for international cooperation on privacy.

[2] For more information, see Daniel Solove, *Understanding Privacy*, Harvard University Press, 2010; and Evan Selinger and Woodrow Hartzog, "Obscurity and Privacy," *Routledge Companion to Philosophy of Technology*, 2014, at https://ssrn.com/abstract=2439866.

1.1 Overview of the Privacy Framework

The Privacy Framework is composed of three parts: the Core, the Profiles, and the Implementation Tiers. Each component reinforces privacy risk management through the connection between business/mission drivers and privacy protection activities. These components are explained in more detail in section 2.0, but as an overview:

- The *Core* is a set of privacy protection activities and desired outcomes that allows for communicating prioritized privacy protection activities and outcomes across the organization from the executive level to the implementation/operations level. The Core consists of five concurrent and continuous functions—*Identify, Protect, Control, Inform,* and *Respond*. Together these functions provide a high-level, strategic view of the life cycle of an organization's management of privacy risk. The Core then identifies underlying key categories and subcategories—which are discrete outcomes—for each function.

- A *Profile* represents the privacy outcomes the organization aims to achieve. To develop a Profile, an organization can review all of the functions, categories, and subcategories to determine which are most important to achieving the desired privacy outcomes, based on business/mission drivers, types of data processing, and individuals' privacy needs. The organization can create or add functions, categories, and subcategories as needed. Profiles can be used to identify opportunities for improving privacy posture by comparing a "Current" Profile (the "as is" state) with a "Target" Profile (the "to be" state). Profiles can be used to conduct self-assessments and to communicate within an organization or between organizations about how privacy risks are being managed.

- *Implementation Tiers* ("Tiers") provide context on how an organization views privacy risk and whether it has adequate processes and resources in place to manage that risk. Tiers reflect a progression from informal, reactive responses to approaches that are agile and risk-informed. When selecting Tiers, an organization should consider its current risk management practices; its data processing systems, products, or services; legal and regulatory requirements; business/mission objectives; organizational privacy values and individuals' privacy needs; and organizational constraints.

1.2 Privacy Risk Management

Risk management is the ongoing set of processes for identifying, assessing, and responding to risk. To manage risk, organizations should seek to understand the likelihood that an event will occur and the potential impacts. While some organizations have a robust grasp of the underlying processes and resources needed to identify, assess, and respond to privacy risks, a common understanding of many aspects of privacy risk management is still not widespread.[3] To promote broader understanding, this section covers concepts and considerations that organizations may use to develop, improve, or communicate about their privacy risk management. Appendix D provides additional guidance on key privacy risk management practices.

[3] See *Summary Analysis of the Responses to the NIST Privacy Framework Request for Information* at p. 7 https://www.nist.gov/sites/default/files/documents/2019/02/27/rfi_response_analysis_privacyframework_2.27.19.pdf

1.2.1 Cybersecurity and Privacy Risk Management

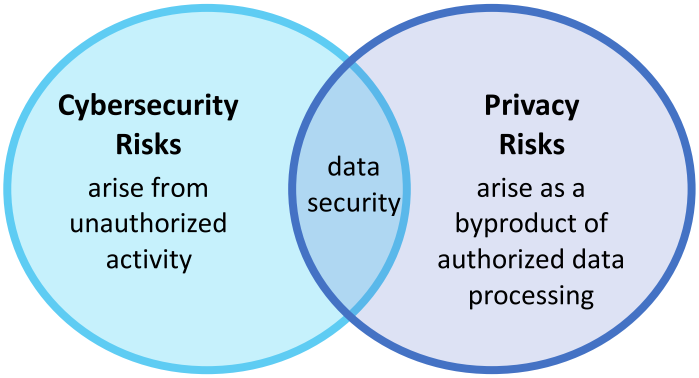
Cybersecurity Risks
arise from unauthorized activity
data security
Privacy Risks
arise as a byproduct of authorized data processing

Figure 1: Cybersecurity and Privacy Risk Relationship

Since its release in 2014, the Framework for Improving Critical Infrastructure Cybersecurity (Cybersecurity Framework) has helped organizations to communicate and manage cybersecurity risk.[4] While managing cybersecurity risk contributes to managing privacy risk, it is not sufficient as privacy risks can also arise outside the scope of cybersecurity risks. **Figure 1** illustrates how NIST considers the overlap and differences between cybersecurity and privacy risks. Cybersecurity risks arise from unauthorized activity related to the loss of confidentiality, integrity, or availability of a system or information asset. However, privacy risks arise as a byproduct of intentional (i.e., authorized) data processing occurring in systems, products, and services that help organizations to achieve their mission/business objectives. This data processing can lead to unintended problems or adverse consequences for individuals.[5] An example is the concerns that certain communities had about the installation of "smart meters" as part of the Smart Grid, a nationwide technological effort to increase energy efficiency.[6] The ability of these meters to collect, record, and distribute highly granular information about household electrical use could provide insight into people's behavior inside their homes.[7] The meters were operating as intended, but the data processing could lead to unintended consequences that people might feel surveilled.

Data Processing
An operation or set of operations performed upon data across the full data life cycle, including but not limited to the collection, retention, logging, generation, transformation, use, disclosure, transfer, and disposal of data.

Individuals can also experience problems or adverse consequences if the data being processed is subject to a loss of confidentiality, integrity, and availability. Figure 1 shows this data security issue as an overlap between managing cybersecurity and privacy risks.

Thus, *privacy risk* can be understood as the likelihood that individuals will experience problems resulting from data processing,

[4] See *Framework for Improving Critical Infrastructure Cybersecurity* at https://doi.org/10.6028/NIST.CSWP.04162018.

[5] NIST has created an illustrative problem set with problems that can range, for example, from embarrassment to discrimination, to economic loss and physical harm), see *NIST Privacy Risk Assessment Methodology* at https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/resources. Other organizations may have created additional problem sets, or may refer to them as adverse consequences or harms.

[6] See e.g., NISTIR 7628 Revision 1 Volume 1, *Guidelines for Smart Grid Cybersecurity: Volume 1 – Smart Grid Cybersecurity Strategy, Architecture, and High-Level Requirements* at p. 26 https://doi.org/10.6028/NIST.IR.7628r1.

[7] See NIST Internal Report (NISTIR) 8062, *An Introduction to Privacy Engineering and Risk Management in Federal Systems* at p. 2 https://doi.org/10.6028/NIST.IR.8062. For additional types of privacy risks arising from authorized data processing, see Appendix E of NISTIR 8062.

and the impact should they occur.[8] Organizations typically determine the acceptable level of risk for achieving their organizational objectives and can express this as their risk tolerance. Privacy risk adds a layer of complexity to the determination of risk tolerance because it manifests as an externality—individuals, not organizations, experience the direct impact of the problems. Privacy risk management should help organizations to internalize consideration of these impacts to individuals, and appropriately account for them in their determination of risk tolerance. With an understanding of risk tolerance, organizations can better prioritize privacy activities, enabling organizations to make informed decisions about budgets and other resource allocations.

Given the relationship between privacy and cybersecurity, organizations may opt to use the Cybersecurity Framework and the Privacy Framework together. **Figure 2** illustrates how the five functions in those frameworks relate to each other.[9] While not exclusive:

- Identify, Protect, and Respond can be used to manage privacy risks whether they arise from loss of data security or more directly from authorized data processing;
- Detect and Recover can be used to manage privacy risks arising from loss of data security; and
- Control and Inform can be used to manage privacy risks arising directly from authorized data processing.

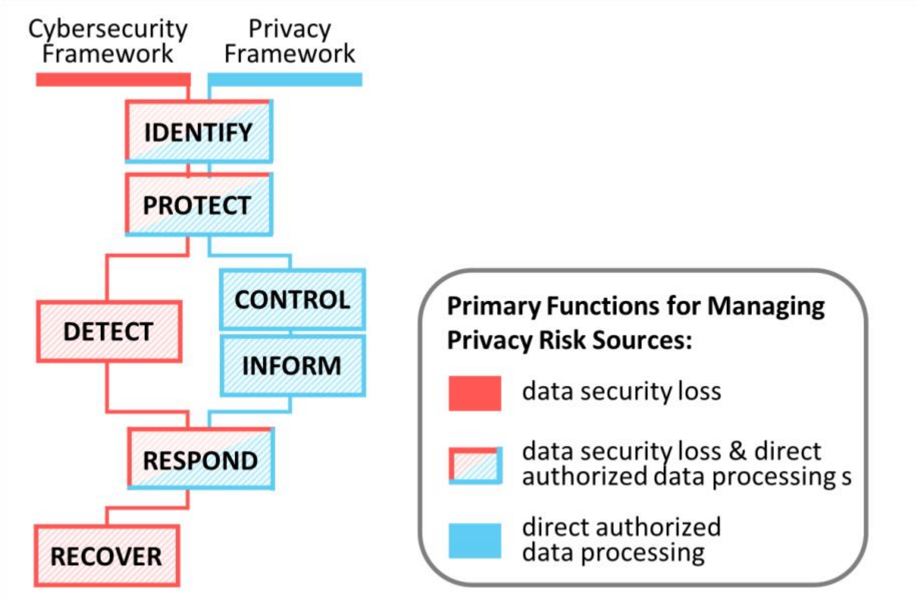
Cybersecurity Framework
Privacy Framework
IDENTIFY
PROTECT
DETECT
CONTROL
INFORM
RESPOND
RECOVER
Primary Functions for Managing Privacy Risk Sources:
data security loss
data security loss & direct authorized data processing s
direct authorized data processing

Figure 2: Cybersecurity Framework and Privacy Framework Functions Relationship

1.2.2 Relationship Between Privacy Risk Management and Risk Assessment

Privacy risk management is a cross-organizational set of processes that helps organizations to understand how their systems, products, and services may create problems for individuals and how to develop effective solutions to manage such risks. *Privacy risk assessment* is a sub-process for identifying, evaluating, prioritizing, and responding to specific privacy risks engendered by systems, products, or

[8] Id at p. 21

[9] Although the Privacy Framework can be used independently, modeling the structural design of the Core, Profiles, and Tiers after the Cybersecurity Framework allows the two frameworks to be used together more readily.

services. In general, privacy risk assessments should produce the information that can help organizations to weigh the benefits of the data processing against the risks and to determine the appropriate response (see Appendix D for more guidance on the operational aspects of privacy risk assessment). Organizations may choose to respond to privacy risk in different ways, depending on the potential impact to individuals (and secondarily, organizations). Approaches include:

- Mitigating the risk (e.g., organizations may be able to apply technical and/or policy measures to the systems, products, or services that minimize the risk to an acceptable degree);
- Transferring or sharing the risk (e.g., contracts are a means of sharing or transferring risk to other organizations, privacy notices and consent mechanisms are a means of sharing risk with individuals);
- Avoiding the risk (e.g., organizations may determine that the risks outweigh the benefits, and forego or terminate the data processing); or
- Accepting the risk (e.g., organizations may determine that problems for individuals are minimal or unlikely to occur, therefore the benefits outweigh the risks, and it is not necessary to invest resources in mitigation).

Privacy risk assessments are particularly important because as noted above privacy is a condition that safeguards multiple values. The methods for safeguarding these values may differ, and moreover, may be in tension with each other. For instance, if the organization is trying to achieve privacy through obscurity, this may lead to implementing measures such as data deletion schedules or privacy-enhancing cryptographic techniques that hide data even from the organization. If the organization is focused on control, obscurity measures could conflict with control measures. For example, if an individual requests access to data, the organization may not be able to produce the data if they have been deleted or encrypted in ways the organization cannot access. Privacy risk assessments can help an organization understand in a given context, the values to protect, the methods to employ, and how to balance implementation of different types of measures.

Lastly, privacy risk assessments help organizations distinguish between privacy risk and compliance risk. Identifying if data processing could create problems for individuals, even when the organization may be fully compliant with applicable laws or regulations, can help organizations make ethical decisions and avoid losses of trust that damage their reputations or slow adoption or cause abandonment of products and services.

1.3 Document Overview

The remainder of this document contains the following sections and appendices:

- Section 2.0 describes the Privacy Framework components: the Core, the Profiles, and the Implementation Tiers.
- Section 3.0 presents examples of how the Privacy Framework can be used.
- Appendix A presents the Privacy Framework Core in a tabular format: the functions, categories, and subcategories.
- Appendix B contains a glossary of selected terms.
- Appendix C lists acronyms used in this document.
- Appendix D considers key practices that contribute to successful privacy risk management.
- Appendix E defines the Implementation Tiers.

- Appendix F provides a placeholder for a companion roadmap covering NIST's next steps and identifying key areas where the relevant practices are not well enough understood to enable organizations to achieve a privacy outcome.

2.0 Privacy Framework Basics

The Privacy Framework provides a common language for understanding, managing, and communicating privacy risk with internal and external stakeholders. It can be used to help identify and prioritize actions for reducing privacy risk, and it is a tool for aligning policy, business, and technological approaches to managing that risk. Different types of entities—including sector-specific organizations—can use the Privacy Framework for different purposes, including the creation of common Profiles.

2.1 Core

The Core provides a set of activities to achieve specific privacy outcomes. The Core is not a checklist of actions to perform. It presents key privacy outcomes that are helpful in managing privacy risk. The Core comprises three elements: functions, categories, and subcategories, depicted in **Figure 3.**

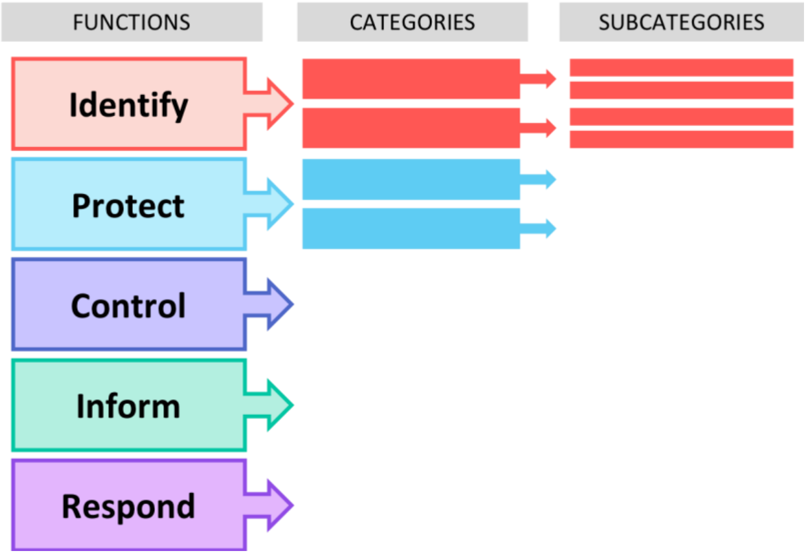

Figure 3: Privacy Framework Core Structure

The Core elements work together:

- *Functions* organize basic privacy activities at their highest level. These functions are Identify, Protect, Control, Inform, and Respond. They aid an organization in expressing its management of privacy risk by understanding and managing data processing, enabling risk management decisions, determining how to interact with individuals, and improving by learning from previous activities.

- *Categories* are the subdivisions of a function into groups of privacy outcomes closely tied to programmatic needs and particular activities. Examples include "Protected Processing," "Inventory and Mapping," and "Risk Assessment."

- *Subcategories* further divide a category into specific outcomes of technical and/or management activities. They provide a set of results that, while not exhaustive, help support achievement of the outcomes in each category. Examples include "Systems/products/services that process data, or with which individuals are interacting, are inventoried," "Data are processed to limit the identification of individuals," and "Individuals are informed when data are corrected or deleted."

The five Core functions, defined below, are not intended to form a serial path or lead to a static desired end state. Rather, the functions should be performed concurrently and continuously to form or enhance an operational culture that addresses the dynamic nature of privacy risk. See Appendix A for the complete Core.

- *Identify* – Develop the organizational understanding to manage privacy risk for individuals arising from data processing or their interactions with systems, products, or services.

  The activities in the Identify function will be foundational for effective use of the Privacy Framework. Understanding the business context, including the circumstances under which data are processed, the privacy interests of individuals directly or indirectly served or affected by the organization, and legal/regulatory requirements will enable an organization to focus and prioritize its efforts, consistent with its risk management strategy and business needs. Examples of categories include: Inventory and Mapping, Business Environment, Governance, and Risk Assessment.

- *Protect* – Develop and implement appropriate data processing safeguards.

  The Protect function not only encapsulates data security, the primary overlap between privacy and cybersecurity, but also includes practices that enable authorized data processing to be conducted in a protected state. Examples of categories include: Identity Management, Authentication, and Access Control; Awareness and Training; Data Security; and Protected Processing.

- *Control* – Develop and implement appropriate activities to enable organizations or individuals to manage data with sufficient granularity to manage privacy risks.

  The Control function considers data management from both the standpoint of the organization and the individual. Examples of categories include: Policies, Processes, and Procedures; and Data Management.

- *Inform* – Develop and implement appropriate activities to enable organizations and individuals to have a reliable understanding about how data are processed.

  The Inform function recognizes that both organizations and individuals need to know how data are processed in order to manage privacy risk effectively. Examples of categories include: Transparency Processes and Procedures, and Data Processing Awareness.

- *Respond* – Develop and implement appropriate activities to take action regarding a privacy breach or event.

  The Respond function supports the ability to provide redress for individuals who have experienced a privacy breach or privacy event and to help organizations use lessons learned to improve their privacy practices. Examples of categories include: Mitigation and Redress.

2.2 Profile

The Profile is the alignment of the functions, categories, and subcategories with the business requirements, risk tolerance, privacy values, and resources of the organization. Under the Privacy Framework's risk-based approach, organizations may not need to achieve every outcome or activity reflected in the Core. When developing a Profile, an organization may select or tailor the Privacy Framework’s functions, categories, and subcategories to its specific needs. An organization or industry sector also may develop its own additional functions, categories, and subcategories to account for unique organizational risks. An organization determines these needs by considering organizational or industry sector goals, legal/regulatory requirements and industry best practices, the organization’s risk management priorities, and the privacy needs of individuals who are part of—or directly or indirectly served or affected by—an organization’s systems, products, or services.

Profiles can be used to describe the current state or the desired target state of specific privacy activities. A Current Profile indicates privacy outcomes that an organization is currently achieving, while a Target Profile indicates the outcomes needed to achieve the desired privacy risk management goals. The differences between the two Profiles enable an organization to identify gaps, develop an action plan for improvement, and gauge the resources that would be needed (e.g., staffing, funding) to achieve privacy goals. This forms the basis of an organization's plan for reducing privacy risk in a cost-effective, prioritized manner. Profiles also can aid in communicating risk within and between organizations by helping organizations understand and compare the current or desired state of privacy outcomes.

This Privacy Framework does not prescribe Profile templates, to allow for flexibility in implementation. An organization may choose to have multiple Profiles for specific organizational components, systems, products, or services, or categories of individuals (e.g., employees, customers).

2.3 Implementation Tiers

Tiers are meant to support organizational decision-making about how to manage privacy risk by taking into account the nature of the privacy risks engendered by the organization's systems, products, or services and the adequacy of the processes and resources the organization has in place to manage such risks. When selecting Tiers, an organization should consider its current risk management practices; its data processing systems, products, or services; legal and regulatory requirements; business/mission objectives; organizational privacy values and individuals' privacy needs; and organizational constraints.

There are four distinct tiers: Partial (Tier 1), Risk Informed (Tier 2), Repeatable (Tier 3), and Adaptive (Tier 4). Tiers do not represent maturity levels, although organizations identified as Tier 1 are encouraged to consider moving toward Tier 2. Some organizations may never need to achieve Tier 3 or 4 or may only focus on certain areas of these tiers. Progression to higher Tiers is appropriate when the nature of the privacy risks requires more multi-faceted risk management processes and resources.

Successful implementation of the Privacy Framework is based upon achieving the outcomes described in the organization's Target Profile(s) and not upon Tier determination. Still, Tier selection naturally affects Profiles, and helps to set the overall tone for how privacy risk will be managed within the organization. This should influence the prioritization of elements included in a Target Profile, and should influence assessments of progress in addressing gaps. The definitions of the Tiers are set forth in Appendix E.

3.0 How to Use the Privacy Framework

When used as a risk management tool, the Privacy Framework can assist an organization in its efforts to optimize beneficial uses of data and the development of innovative systems, products, and services while minimizing adverse consequences for individuals. The Privacy Framework can help organizations answer the fundamental question, "How are we considering the impacts to individuals as we develop our systems, products, and services?" As a result, the Privacy Framework can serve as the foundation for a new privacy program or a mechanism for improving an existing program. In either case, it is designed to complement existing business and system development operations, to provide a means of expressing privacy requirements to business partners and customers, and to support the identification of gaps in an organization's privacy practices.

To account for the unique needs of an organization, there are a wide variety of ways to use the Privacy Framework. The decision about how to apply it is left to the implementing organization. For example, one organization may choose to use the Implementation Tiers to articulate its envisioned privacy risk management processes. Another organization may already have robust privacy risk management processes, but may use the Core's five functions to analyze and articulate any gaps. Alternatively, an organization seeking to establish a privacy program can use the Core categories and subcategories as a reference. The variety of ways in which the Privacy Framework can be used by organizations should discourage the notion of "compliance with the Privacy Framework" as a uniform or externally referenceable concept.

The following subsections present different ways in which organizations can use the Privacy Framework.

3.1 Informative References

The Privacy Framework is technology neutral, but it supports technological innovation because any organization or industry sector can map the outcome-based subcategories in the Core to standards, guidelines, and practices that evolve with technology and related business needs. By relying on consensus-based standards, guidelines, and practices, the tools and methods available to achieve positive privacy outcomes can scale across borders, accommodate the global nature of privacy risks, and evolve with technological advances and business requirements. The use of existing and emerging standards will enable economies of scale and drive the development of systems, products, and services that meet identified market needs while being mindful of the privacy needs of individuals.

Mapping subcategories to specific sections of standards, guidelines, and practices supports the achievement of the outcomes associated with each subcategory. The subcategories also can be used to identify where additional or revised standards, guidelines, and practices would help an organization to address emerging needs. An organization implementing a given subcategory, or developing a new subcategory, might discover that there are insufficient informative references for a related activity. To address that need, the organization might collaborate with technology leaders and/or standards bodies to draft, develop, and coordinate standards, guidelines, or practices.

NIST has developed a mapping of the Core subcategories to relevant NIST guidance, as well as a process for organizations or industry sectors to submit additional informative references and mappings for publication on NIST's website at https://www.nist.gov/privacy-framework. These resources can support organizations' application of the Privacy Framework and achievement of better privacy practices.

3.2 Strengthening Accountability

Accountability is generally considered a key privacy principle, although conceptually it is not unique to privacy.[10] Accountability occurs throughout an organization, and it can be expressed at varying degrees of abstraction, for example as a cultural value, as governance policies and procedures, or as traceability relationships between privacy requirements and controls. Privacy risk management can be a means of supporting accountability at all organizational levels as it connects senior executives, who can communicate the organization's privacy values and risk tolerance, to those at the business/process

[10] See, e.g., Organisation for Economic Co-operation and Development (OECD), *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data* at https://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm; International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC), ISO/IEC 29100, *Information technology – Security techniques – Privacy framework* at https://standards.iso.org/ittf/PubliclyAvailableStandards/c045123_ISO_IEC_29100_2011.zip; Alliance of Automobile Manufacturers, Inc. and Association of Global Automakers, Inc., *Consumer Privacy Protection Principles: Privacy Principles for Vehicle Technologies and Services* at https://autoalliance.org/wp-content/uploads/2017/01/Consumer_Privacy_Principlesfor_VehicleTechnologies_Services-03-21-19.pdf.

manager level, who can collaborate on the development and implementation of governance policies and procedures that support the organizational privacy values. These policies and procedures can then be communicated to those at the implementation/operations level, who collaborate on defining the privacy requirements that support the expression of the policies and procedures in the organization's systems, products, and services. Personnel at the implementation/operations level also select, implement, and assess controls as the technical and policy measures that meet the privacy requirements, and report upward on progress, gaps and deficiencies, and changing privacy risks so that those at the business/process manager level and the senior executives can better understand and respond appropriately. **Figure 4** provides a graphical representation of this iterative cycle and how elements of the Privacy Framework can be incorporated to facilitate the process. In this way, organizations can use the Privacy Framework as a tool to support accountability. They can also use the Privacy Framework in conjunction with other frameworks and guidance that provide additional practices to achieve accountability within and between organizations (see section 3.6 on Communicating Privacy Requirements with Stakeholders).[11]

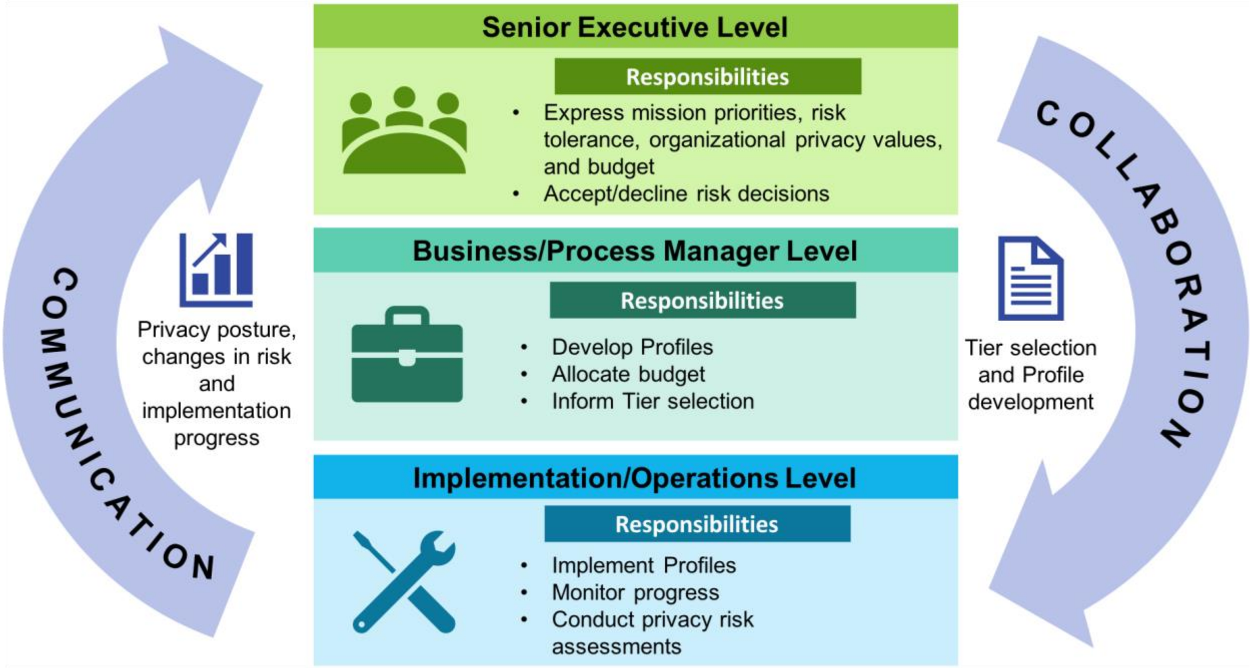
Senior Executive Level
Responsibilities
• Express mission priorities, risk tolerance, organizational privacy values, and budget
• Accept/decline risk decisions
Business/Process Manager Level
Responsibilities
• Develop Profiles
• Allocate budget
• Inform Tier selection
Implementation/Operations Level
Responsibilities
• Implement Profiles
• Monitor progress
• Conduct privacy risk assessments
COMMUNICATION
Privacy posture, changes in risk and implementation progress
COLLABORATION
Tier selection and Profile development

Figure 4: Notional Information and Decision Flows within an Organization

3.3 Basic Review of Privacy Practices

The Privacy Framework can be used to compare an organization's current privacy activities with those outlined in the Core. Through the creation of a Current Profile, organizations can examine the extent to which they are achieving the outcomes described in the Core categories and subcategories, aligned with the five high-level functions: Identify, Protect, Control, Inform, and Respond. An organization may find that it is already achieving the desired outcomes, thus managing privacy commensurate with the known

[11] See, e.g., NIST SP 800-37 Rev. 2, *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy* at https://doi.org/10.6028/NIST.SP.800-37r2; and Organization for the Advancement of Structured Information Standards (OASIS), *Privacy Management Reference Model and Methodology (PMRM) Version 1.0* at https://docs.oasis-open.org/pmrm/PMRM/v1.0/PMRM-v1.0.pdf.

risk. Alternatively, an organization may determine that it has areas for improvement. The organization can use that information to develop an action plan to strengthen existing privacy practices and reduce privacy risk. For example, an organization may be fulfilling its legal obligations with respect to processing data, but it may lack a robust privacy risk assessment process that would help it achieve better privacy outcomes for the individual. In building the Target Profile, the organization that is currently approaching its privacy program from a predominantly regulatory and compliance standpoint will be able to communicate the need for conducting privacy risk assessments to help it achieve more fine-tuned privacy benefits for individuals. The organization can use this information to reprioritize resources or adjust approaches to improve outcomes for both the organization and individuals.

While they do not replace a risk management process, the five high-level functions provide a concise way for senior executives and others to distill the fundamental concepts of privacy risk so that they can assess how identified risks are managed and how their organization stacks up at a high level against existing privacy standards, guidelines, and practices.

3.4 Establishing or Improving a Privacy Program

Using a straightforward model of “ready, set, go” phases, the Privacy Framework can support the creation of a new privacy program or improvement of an existing program. These phases should be repeated as necessary to continuously improve privacy.

A Simplified Method for Establishing or Improving Privacy Programs

Ready: use the Identify function to get “ready.”

Set: “set” an action plan based on the differences between Current and Target Profile(s).

Go: “go” forward with implementing the action plan.

Ready

Effective privacy risk management requires an organization to understand its business or mission environment; its legal environment; its enterprise risk tolerance; the privacy risks engendered by its systems, products, or services; and its role or relationship to other organizations in the ecosystem. An organization can use the Identify function to “get ready” by reviewing the categories and subcategories, and beginning to develop its Current Profile and Target Profile.[12]

An organization conducts privacy risk assessments pursuant to the Risk Assessment category of the Identify function. These assessments could be guided by the organization’s overall risk management process or previous risk assessment activities. It is important that an organization identifies emerging privacy risks to gain a better understanding of the impacts of its systems, products, or services on individuals. See Appendix D for more information on privacy risk assessments.

Set

The organization completes its Current Profile by indicating which category and subcategory outcomes from the remaining functions are being achieved. If an outcome is partially achieved, noting this fact will help support subsequent steps by providing baseline information. Informed by its privacy risk assessment, the organization creates its Target Profile focused on the assessment of the Core categories and subcategories describing the organization’s desired privacy outcomes. An organization also may

[12] For additional guidance, see the “Prepare” step, Section 3.1, NIST Special Publication (SP) 800-37 Revision 2, *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy* at https://doi.org/10.6028/NIST.SP.800-37r2.

develop its own additional functions, categories and subcategories to account for unique organizational risks. It may also consider influences and requirements of external stakeholders such as business customers and partners when creating a Target Profile. An organization can develop multiple Profiles to support its different business lines or processes, which may have different business needs and associated risk tolerances.

The organization compares the Current Profile and the Target Profile to determine gaps. Next, it creates a prioritized action plan to address gaps—reflecting mission drivers, costs and benefits, and risks—to achieve the outcomes in the Target Profile. An organization using the Cybersecurity Framework and the Privacy Framework together may develop integrated action plans. The organization then determines resources, including funding and workforce, necessary to address the gaps, which can inform the selection of a target Tier. Using Profiles in this manner encourages the organization to make informed decisions about privacy activities, supports risk management, and enables the organization to perform cost-effective, targeted improvements.

Go

With the action plan "set," the organization prioritizes which actions to take to address any gaps, and then adjusts its current privacy practices in order to achieve the Target Profile.[13] For further guidance, informative references that support outcome achievement for the categories and subcategories are available at https://www.nist.gov/privacy-framework. The organization should determine which standards, guidelines, and practices, including those that are sector specific, work best for its needs.

An organization can cycle through the phases non-sequentially as needed to continuously assess and improve its privacy posture. For instance, an organization may find that more frequent repetition of the Ready phase improves the quality of risk assessments. Furthermore, an organization may monitor progress through iterative updates to the Current Profile or the Target Profile to adjust to changing risks, subsequently comparing the Current Profile to the Target Profile. An organizations may also use this process to align its privacy program with its desired Tiers.

3.5 Application in the System Development Life Cycle

The Privacy Framework can be applied throughout the system development life cycle (SDLC) phases of plan, design, build/buy, deploy, operate, and decommission. The plan phase of the SDLC begins the cycle of any system and lays the groundwork for everything that follows. Overarching privacy considerations should be declared and described as clearly as possible. The plan should recognize that those considerations and requirements are likely to evolve during the remainder of the life cycle. A key milestone of the design phase is validating that the system privacy requirements match the needs and risk tolerance of the organization as they were expressed in a Profile. The desired privacy outcomes prioritized in a Target Profile should be incorporated when a) developing the system during the build phase and b) purchasing or outsourcing the system during the buy phase. That same Target Profile serves as a list of system privacy features that should be assessed when deploying the system to verify that all features are implemented. The privacy outcomes determined by using the Privacy Framework should then serve as a basis for ongoing operation of the system. This includes occasional reassessment, capturing results in a Current Profile, to verify that privacy requirements are still fulfilled.

[13] NIST SP 800-37 provides additional guidance on steps to execute on the action plan, including control selection, implementation, and assessment to close any gaps.

Privacy risk assessments typically focus on the information life cycle, the stages through which information passes, often characterized as creation or collection, processing, dissemination, use, storage, and disposition, to include destruction and deletion. Aligning the SDLC and the information lifecycle by identifying and understanding how data are processed during all stages of the SDLC helps organizations to better manage privacy risks and informs the selection and implementation of privacy controls throughout the SDLC.

3.6 Communicating Privacy Requirements with Stakeholders

The Privacy Framework provides a common language to communicate requirements among interdependent stakeholders. For example:

- An organization may use a Target Profile to express privacy risk management requirements to an external service provider (e.g., a cloud provider to which it is exporting data).
- An organization may express its privacy posture through a Current Profile to report results or to compare with acquisition requirements.
- An industry sector may establish a Target Profile that can be used among its constituents as an initial baseline Profile to build their tailored Target Profiles.
- An organization can better manage privacy risk affecting stakeholders by assessing their positions in the data processing ecosystem and the broader digital economy using Tiers.

Communication is especially important among stakeholders in supply chains. Supply chains are complex, globally distributed, and interconnected sets of resources and processes between multiple levels of organizations. Supply chains begin with the sourcing of products and services and extend from the design, development, manufacturing, processing, handling, and delivery of products and services to the end user. Given these complex and interconnected relationships, supply chain risk management (SCRM) is a critical organizational function.[14] SCRM practices should address the management of privacy risk associated with external parties, including both the effect an organization has on external parties and the effect external parties have on an organization. Such practices include identifying, assessing, and mitigating privacy risks arising from the processing of data, as well as from systems, products, and services that inherently lack the capabilities to mitigate privacy risks. Example activities may include:

- Determining privacy requirements for suppliers,
- Enacting privacy requirements through formal agreement (e.g., contracts),
- Communicating to suppliers how those privacy requirements will be verified and validated,
- Verifying that privacy requirements are met through a variety of assessment methodologies, and
- Governing and managing the above activities.

[14] Communicating Privacy Requirements with Stakeholders (section 3.6) and Buying Decisions (section 3.7) address only two uses of the Privacy Framework for SCRM and are not intended to address SCRM comprehensively.

As depicted in **Figure 5,** SCRM in the data processing ecosystem encompasses a range of entities and roles that may have multi-directional relationships with each other and individuals. Figure 5 displays entities as having distinct roles, but organizations may have multiple roles. For example, an organization may be both a service provider to other organizations and provide commercial products or services to individuals.

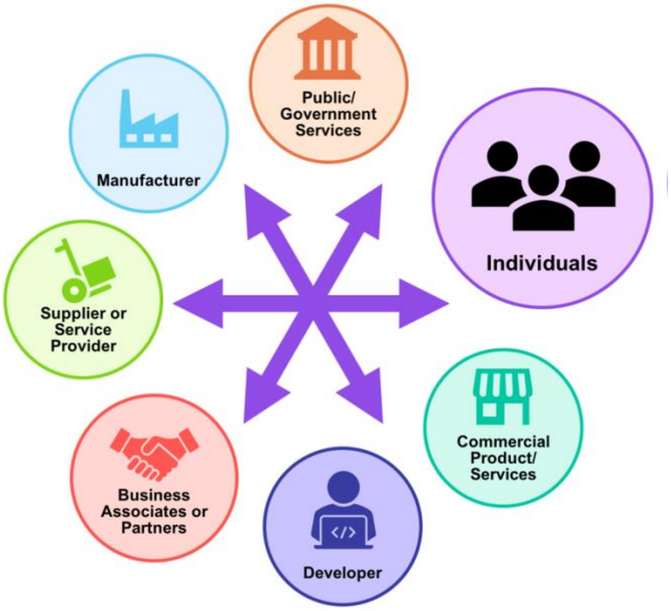
Manufacturer
Public/ Government Services
Individuals
Supplier or Service Provider
Commercial Product/ Services
Business Associates or Partners
Developer

Figure 5: Ecosystem Relationships

The parties described in Figure 5 comprise a data processing ecosystem. These relationships highlight the crucial role of SCRM in addressing privacy risk in the broader digital economy. These relationships, the systems, products, and services they provide, and the risks they present should be identified and factored into the data processing capabilities of organizations, as well as their response protocols.

Whether considering individual subcategories of the Core or the comprehensive considerations of a Profile, the Privacy Framework offers organizations and their partners a method to help ensure the system, product, or service meets critical privacy outcomes. By first selecting outcomes that are relevant to the context, the organization then can evaluate partners' systems, products, or services against this outcome. For example, if a device is being purchased for environmental monitoring of a forest, *manageability* may be important to support capabilities for minimizing the processing of data about people using the forest and should drive a manufacturer evaluation against applicable subcategories (see e.g., CT.DM-P1 in Appendix A: system or device configurations permit selective collection or disclosure of data elements to allow for implementation of privacy principles such as data minimization).

3.7 Buying Decisions

Since either a Current or Target Profile can be used to generate a prioritized list of organizational privacy requirements, these Profiles can also be used to inform decisions about buying products and services. In circumstances where it may not be possible to impose a set of privacy requirements on the supplier, the objective should be to make the best buying decision among multiple suppliers, given a carefully determined list of privacy requirements. Often, this means some degree of trade-off, comparing multiple products or services with known gaps to the Profile. If the system, product, or service purchased did not meet all of the objectives described in the Profile, the organization can address the residual risk through mitigation measures or other management actions.

Appendix A: Privacy Framework Core

This appendix presents the Core: a table of functions, categories, and subcategories that describe specific privacy activities that can support managing privacy risks when systems, products, and services are processing data or interacting with individuals. This presentation format does not suggest a specific implementation order—implementation may be non-sequential and iterative, depending on the SDLC stage or status of the privacy program—or imply a degree of importance between the categories and subcategories. The Core is not exhaustive; it is extensible, allowing organizations, sectors, and other entities to adapt or add additional functions, categories and subcategories to their Profile(s).

For ease of use, each component of the Core is given a unique identifier. Functions and categories each have a unique alphabetic identifier, as shown in **Table 1**. Subcategories within each category are referenced numerically; the unique identifier for each subcategory is included in **Table 2**.

Additional supporting material, including informative references, relating to the Privacy Framework can be found on the NIST website at https://www.nist.gov/privacy-framework.

Table 1: Privacy Framework Function and Category Unique Identifiers

	Function Unique Identifier	Function	Category Unique Identifier	Category
🟥	**ID**	Identify-P	ID.IM-P	Inventory and Mapping
			ID.BE-P	Business Environment
			ID.GV-P	Governance
			ID.RA-P	Risk Assessment
			ID.RM-P	Risk Management Strategy
			ID.SC-P	Supply Chain Risk Management
🟦	**PR**	Protect-P	PR.AC-P	Identity Management, Authentication, and Access Control
			PR.AT-P	Awareness and Training
			PR.DS-P	Data Security
			PR.DP-P	Data Protection Processes and Procedures
			PR.MA-P	Maintenance
			PR.PT-P	Protective Technology
			PR.PP-P	Protected Processing
🟪	**CT**	Control-P	CT.PO-P	Data Management Processes and Procedures
			CT.DM-P	Data Management
🟩	**IN**	Inform-P	IN.TP-P	Transparency Processes and Procedures
			IN.AW-P	Data Processing Awareness

	RS	Respond-P	RS.RP-P	Response Planning
			RS.CO-P	Communications
			RS.AN-P	Analysis
			RS.MI-P	Mitigation
			RS.IM-P	Improvements
			RS.RE-P	Redress

Table 2: Privacy Framework Core

	Function	Category	Subcategory
🟥	**IDENTIFY-P (ID)**	**Inventory and Mapping (ID.IM-P):** Data processing and individuals' interactions with systems, products, or services are understood and inform the management of privacy risk.	**ID.IM-P1:** Systems/products/services that process data, or with which individuals are interacting, are inventoried.
			ID.IM-P2: The owners or operators of systems/products/services that process data, or with which individuals are interacting, are identified.
			ID-IM-P3: Data elements that systems/products/services are processing are inventoried.
			ID.IM-P4: Data actions are identified.
			ID.IM-P5: The data processing environment is identified (e.g., internal, cloud).
			ID.IM-P6: Data processing is mapped, illustrating the processing of data elements by system components and their owner/operators, and interactions of individuals and organizations with the systems/products/services.
		Business Environment (ID.BE-P): The organization's mission, objectives, stakeholders, and activities are understood and prioritized; this information is used to inform privacy roles, responsibilities, and risk management decisions.	**ID.BE-P1:** The organization's role in the supply chain is identified and communicated.
			ID.BE-P2: Priorities for organizational mission, objectives, and activities are established and communicated.
			ID.BE-P3: Systems/products/services that support organizational priorities are identified and key functional requirements communicated.
		Governance (ID.GV-P): The policies, procedures, and processes to manage and monitor the organization's regulatory, legal, risk, environmental, and operational requirements are understood and inform the management of privacy risk.	**ID.GV-P1:** Organizational privacy policies are established and communicated.
			ID.GV-P2: Processes to instill organizational privacy values within system/product/service development operations are in place.
			ID.GV-P3: Privacy roles and responsibilities for the entire workforce are established.

	Function	Category	Subcategory
			ID.GV-P4: Privacy roles and responsibilities are coordinated and aligned with third-party stakeholders (e.g., suppliers, customers, partners).
			ID.GV-P5: Legal, regulatory, and contractual requirements regarding privacy are understood and managed.
			ID.GV-P6: Governance and risk management processes address privacy risks.
		Risk Assessment (ID.RA-P): The organization understands the privacy risks to individuals and how such privacy risks may create secondary impacts on organizational operations (including mission, functions, reputation, or workforce culture).	**ID.RA-P1:** The purposes for the data actions are identified.
			ID.RA-P2: Contextual factors related to the systems/products/services and the data actions are identified (e.g., individuals' privacy interests and perceptions, demographics, data sensitivity).
			ID.RA-P3: Potential problematic data actions and associated problems are identified.
			ID.RA-P4: Problematic data actions, likelihoods, and impacts are used to determine and prioritize risk.
			ID.RA-P5: Risk responses are identified and prioritized.
			ID.RA-P6: Risk is re-evaluated as data processing or individuals' interactions with systems/products/services change.
		Risk Management Strategy (ID.RM-P): The organization's priorities, constraints, risk tolerances, and assumptions are established and used to support operational risk decisions.	**ID.RM-P1:** Risk management processes are established, managed, and agreed to by organizational stakeholders.
			ID.RM-P2: Organizational risk tolerance is determined and clearly expressed.
			ID.RM-P3: The organization's determination of risk tolerance is informed by its role in the ecosystem.
		Supply Chain Risk Management (ID.SC-P): The organization's priorities, constraints, risk tolerances, and assumptions are established	**ID.SC-P1:** Supply chain risk management processes are identified, established, assessed, managed, and agreed to by organizational stakeholders.

	Function	Category	Subcategory
		and used to support risk decisions associated with managing privacy supply chain risk. The organization has established and implemented the processes to identify, assess, and manage privacy supply chain risks.	**ID.SC-P2:** Service providers/suppliers/third-party partners of data processing systems, products, and services are identified, prioritized, and assessed using a supply chain risk assessment process.
			ID.SC-P3: Contracts with service providers/suppliers/third-party partners are used to implement appropriate measures designed to meet the objectives of an organization's privacy program and supply chain risk management plan.
			ID.SC-P4: Service providers/suppliers/third-party partners are routinely assessed using audits, test results, or other forms of evaluations to confirm they are meeting their contractual obligations.
			ID.SC-P5: Response planning and testing are conducted with service providers/suppliers/third-party providers.
	PROTECT-P (PR)	**Identity Management, Authentication, and Access Control (PR.AC-P):** Access to data and devices is limited to authorized individuals, processes, and devices, and is managed consistent with the assessed risk of unauthorized access.	**PR.AC-P1:** Identities and credentials are issued, managed, verified, revoked, and audited for authorized individuals, processes, and devices.
			PR.AC-P2: Physical access to data and devices is managed.
			PR.AC-P3: Remote access is managed.
			PR.AC-P4: Access permissions and authorizations are managed, incorporating the principles of least privilege and separation of duties.
			PR.AC-P5: Network integrity is protected (e.g., network segregation, network segmentation).
			PR.AC-P6: Individuals and devices are proofed and bound to credentials, and authenticated commensurate with the risk of the transaction (e.g., individuals' security and privacy risks and other organizational risks).
			PR.AC-P7: Attribute references are used instead of attribute values.
			PR.AT-P1: All users are informed and trained.

	Function	Category	Subcategory
		Awareness and Training (PR.AT-P): The organization's personnel and partners are provided privacy awareness education and are trained to perform their privacy-related duties and responsibilities consistent with related policies, procedures, and agreements.	**PR.AT-P2:** Privileged users understand their roles and responsibilities.
			PR.AT-P3: Third-party stakeholders (e.g., service providers, customers, partners) understand their roles and responsibilities.
			PR.AT-P4: Senior executives understand their roles and responsibilities.
			PR.AT-P5: Privacy personnel understand their roles and responsibilities.
		Data Security (PR.DS-P): Data are managed consistent with the organization's risk strategy to protect individuals' privacy and maintain data confidentiality, integrity, and availability.	**PR.DS-P1:** Data-at-rest is protected.
			PR.DS-P2: Data-in-transit is protected.
			PR.DS-P3: Systems/products/services and associated data are formally managed throughout removal, transfers, and disposition.
			PR.DS-P4: Adequate capacity to ensure availability is maintained.
			PR.DS-P5: Protections against data leaks are implemented.
			PR.DS-P6: Integrity checking mechanisms are used to verify software, firmware, and information integrity.
			PR.DS-P7: The development and testing environment(s) are separate from the production environment.
			PR.DS-P8: Integrity checking mechanisms are used to verify hardware integrity.
		Data Protection Processes and Procedures (PR.DP-P): Security and privacy policies (that address purpose, scope, roles, responsibilities, management commitment, and coordination among organizational entities), processes, and procedures are maintained and used to manage the protection of data.	**PR.DP-P1:** A baseline configuration of security and privacy controls is created and maintained.
			PR.DP-P2: A system development life cycle to manage systems and an information life cycle to manage data are aligned and implemented.
			PR.DP-P3: Configuration change control processes are in place.
			PR.DP-P4: Backups of information are conducted, maintained, and tested.
			PR.DP-P5: Policy and regulations regarding the physical operating environment for organizational assets are met.

	Function	Category	Subcategory
			PR.DP-P6: Data are destroyed according to policy.
			PR.DP-P7: Protection processes are improved.
			PR.DP-P8: Effectiveness of protection technologies is shared.
			PR.DP-P9: Response plans (Incident Response and Business Continuity) and recovery plans (Incident Recovery and Disaster Recovery) are in place and managed.
			PR.DP-P10: Response and recovery plans are tested.
			PR.DP-P11: Privacy procedures are included in human resources practices (e.g., deprovisioning, personnel screening).
			PR.DP-P12: A vulnerability management plan is developed and implemented.
		Maintenance (PR.MA-P): System maintenance and repairs are performed consistent with policies and procedures.	**PR.MA-P1:** Maintenance and repair of organizational assets are performed and logged, with approved and controlled tools.
			PR.MA-P2: Remote maintenance of organizational assets is approved, logged, and performed in a manner that prevents unauthorized access.
		Protective Technology (PR.PT-P): Technical security solutions are managed to ensure the security and resilience of systems/products/services and associated data, consistent with related policies, procedures, and agreements.	**PR.PT-P1:** Audit/log records are determined, documented, implemented, and reviewed in accordance with policy and incorporating the principle of data minimization.
			PR.PT-P2: Removable media is protected and its use restricted according to policy.
			PR.PT-P3: The principle of least functionality is incorporated by configuring systems to provide only essential capabilities.
			PR.PT-P4: Communications and control networks are protected.
			PR.PT-P5: Mechanisms (e.g., failsafe, load balancing, hot swap) are implemented to achieve resilience requirements in normal and adverse situations.
		Protected Processing (PR.PP-P): Technical data processing solutions increase disassociability consistent with related policies, procedures,	**PR.PP-P1:** Data are processed in an unobservable or unlinkable manner.
			PR.PP-P2: Data are processed to limit the identification of individuals.

	Function	Category	Subcategory
		and agreements and the organization’s risk strategy to protect individuals’ privacy.	**PR.PP-P3:** Data are processed to restrict the formulation of inferences about individuals’ behavior or activities.
			PR.PP-P4: Data are processed through a distributed system architecture.
			PR.PP-P5: Data are processed on local devices.
	CONTROL-P (CT)	**Data Management Processes and Procedures (CT.PO-P):** Policies (that address purpose, scope, roles, responsibilities, management commitment, and coordination among organizational entities), processes, and procedures are maintained and used to manage data consistent with the organization’s risk strategy to protect individuals’ privacy.	**CT.PO-P1:** Policies and procedures for authorizing data processing and maintaining authorizations are in place.
			CT.PO-P2: Processes for enabling data review, transmission/disclosure, alteration, or deletion are in place.
			CT.PO-P3: Processes and procedures for enabling individuals’ data processing preferences and requests (e.g., individual participation) are in place.
		Data Management (CT.DM-P): Data are managed consistent with the organization's risk strategy to protect individuals’ privacy and increase manageability.	**CT.DM-P1:** System or device configurations permit selective collection or disclosure of data elements to allow for implementation of privacy principles (e.g., data minimization).
			CT.DM-P2: Individuals’ authorization for the data action is obtained.
			CT.DM-P3: Data elements can be accessed for review.
			CT.DM-P4: Data elements can be accessed for transmission or disclosure.
			CT.DM-P5: Data elements can be accessed for alteration.
			CT.DM-P6: Data elements can be accessed for deletion.
			CT.DM-P7: Metadata containing processing permissions and related data values are transmitted with data elements.
			CT.DM-P8: Processing permissions are transmitted using standardized formats.
	INFORM-P (IN)	**Transparency Processes and Procedures (IN.TP-P):** Policies (that address purpose, scope, roles, responsibilities, management	**IN.TP-P1:** Transparency procedures and mechanisms (e.g., internal or public reports) for data processing practices are in place.

	Function	Category	Subcategory
		commitment, and coordination among organizational entities), processes, and procedures are maintained and used to increase transparency of the organization's data processing practices.	**IN.TP-P2:** Processes for communicating data processing purposes are in place.
		Data Processing Awareness (IN.AW-P): Individuals and organizations have an awareness of data processing practices, and processes and procedures are used and maintained to increase predictability consistent with the organization's risk strategy to protect individuals' privacy.	**IN.AW-P1:** Records of data disclosures are maintained and can be shared.
			IN.AW-P2: Individuals are informed about data processing practices.
			IN.AW-P3: System/product/service design enhances data processing visibility.
			IN.AW-P4: Data sources are informed of data deletion and correction.
			IN.AW-P5: Individuals are informed when data are corrected or deleted.
			IN.AW-P6: Data provenance is maintained and can be shared.
			IN.AW-P7: Data analytic inputs and outputs are understood and evaluated for bias.
	RESPOND-P (RS)	**Response Planning (RS.RP-P):** Response processes and procedures are executed and maintained to ensure response to privacy breaches and events.	**RS.RP-P1:** Response plan is executed during or after a privacy breach or event.
		Communications (RS.CO-P): Response activities are coordinated with internal and external stakeholders (e.g., external support from law enforcement agencies).	**RS.CO-P1:** Personnel know their roles and order of operations when a response is needed.
			RS.CO-P2: Privacy breaches and events are reported consistent with established criteria.
			RS.CO-P3: Information is shared consistent with response plans.
			RS.CO-P4: Coordination with stakeholders occurs consistent with response plans.
			RS.CO-P5: Data for voluntary information sharing is restricted to what is necessary for understanding the privacy breach or event.

	Function	Category	Subcategory
			RS.CO-P6: Impacted individuals are notified about a privacy breach or event.
		Analysis (RS.AN-P): Analysis is conducted to ensure effective response to privacy breaches and events.	**RS.AN-P1:** Notifications from detection systems or processes are investigated.
			RS.AN-P2: The impact of the privacy breach or event on individuals, the organization, and the ecosystem is understood.
			RS.AN-P3: Forensics are performed.
			RS.AN-P4: Privacy breaches and events are categorized consistent with response plan.
			RS.AN-P5: Processes are established to receive, analyze, and respond to problematic data actions disclosed to the organization from internal and external sources (e.g., internal testing, privacy researchers).
		Mitigation (RS.MI-P): Activities are performed to prevent expansion of, mitigate, and resolve privacy breaches and events.	**RS.MI-P1:** Privacy breaches and events are contained.
			RS.MI-P2: Privacy breaches and events are mitigated.
			RS.MI-P3: Newly identified problematic data actions are mitigated or documented as accepted risks.
		Improvements (RS.IM-P): Organizational privacy practices are improved by incorporating lessons learned from privacy breaches and events.	**RS.IM-P1:** Policies and processes incorporate lessons learned.
		Redress (RS.RE-P): Organizational response activities include processes or mechanisms to address impacts to individuals that arise from data processing.	**RS.RE-P1:** Processes for receiving and responding to complaints, concerns, and questions from individuals about organizational privacy practices are in place.
			RS.RE-P2: Individuals are provided with mitigation mechanisms.

Appendix B: Glossary

This appendix defines selected terms used for the purposes of this publication.

Availability [NIST SP 800-37]	Ensuring timely and reliable access to and use of information.
Category	The subdivision of a function into groups of privacy outcomes, closely tied to programmatic needs and particular activities. Examples of categories include "Protected Processing," "Inventory and Mapping," and "Risk Assessment."
Confidentiality [NIST SP 800-37]	Preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information.
Control (function)	Develop and implement appropriate activities to enable organizations or individuals to manage data with sufficient granularity to manage privacy risks.
Core	A set of privacy protection activities, desired outcomes, and applicable references. The Framework Core comprises three types of elements: functions, categories, and subcategories.
Data	A representation of information with the potential for adverse consequences for individuals when processed.
Data Action [NISTIR 8062, Adapted]	A system/product/service operation that processes data.
Data Elements	The smallest named item of data that conveys meaningful information.
Data Processing [NISTIR 8062, Adapted]	An operation or set of operations performed upon data across the full data life cycle, including but not limited to the collection, retention, logging, generation, transformation, use, disclosure, transfer, and disposal of data.
Disassociability [NISTIR 8062, Adapted]	Enabling the processing of data or events without association to individuals or devices beyond the operational requirements of the system.
Function	One of the main components of the Privacy Framework. Functions provide the highest level of structure for organizing basic privacy activities into categories and subcategories. The five functions are Identify, Protect, Control, Inform, and Respond.
Identify (function)	Develop the organizational understanding to manage privacy risk for individuals arising from data processing or their interactions with systems, products, or services.
Implementation Tier	The degree to which an organization's current or target risk management practices demonstrate an understanding of privacy risk and how systematic the practices are.
Inform (function)	Develop and implement appropriate activities to enable organizations and individuals to have a reliable understanding about how data are processed.
Integrity [NIST SP 800-37]	Guarding against improper information modification or destruction, and includes ensuring information non-repudiation and authenticity.

Manageability [NISTIR 8062, Adapted]	Providing the capability for granular administration of data, including alteration, deletion, and selective disclosure.
Metadata [NIST SP 800-53, Adapted]	Information describing the characteristics of data including, for example, structural metadata describing data structures (i.e., data format, syntax, semantics) and descriptive metadata describing data contents.
Predictability [NISTIR 8062, Adapted]	Enabling reliable assumptions by individuals, owners, and operators about data and its processing by a system.
Privacy Breach	The loss of control, compromise, unauthorized disclosure, unauthorized acquisition, or any similar occurrence where (1) a person other than an authorized user accesses or potentially accesses data or (2) an authorized user accesses data for an other than authorized purpose.
Privacy Control [NIST SP 800-37, Adapted]	The administrative, technical, and physical safeguards employed within an organization to satisfy privacy requirements.
Privacy Event	The occurrence of problematic data actions.
Privacy Requirement	A specification for system/product/service functionality to meet stakeholders' desired privacy outcomes.
Privacy Risk	The likelihood that individuals will experience problems resulting from data processing, and the impact should they occur.
Privacy Risk Assessment	A privacy risk management sub-process for identifying, evaluating, prioritizing, and responding to specific risks arising from data processing.
Privacy Risk Management	A cross-organizational set of processes for identifying, assessing, and responding to privacy risks.
Problematic Data Action [NISTIR 8062]	A data action that can cause an adverse effect, or problem, for individuals.
Profile	A representation of the outcomes based on business/mission objectives, types of data processing, and individuals' privacy needs that an organization has selected from the Privacy Framework categories and subcategories.
Protect (function)	Develop and implement appropriate data processing safeguards.
Provenance [NISTIR 8112, Adapted]	Metadata pertaining to the origination or source of specified data.
Respond (function)	Develop and implement appropriate activities to take action regarding a privacy breach or event.
Risk [NIST SP 800-30]	A measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence.
Subcategory	The further divisions of a category into specific outcomes of technical and/or management activities. Examples of subcategories include "Systems/products/services that process data, or with which individuals are interacting, are inventoried," "Data are processed to limit the identification of individuals," and "Individuals are informed when data are corrected or deleted."

Appendix C: Acronyms

This appendix defines selected acronyms used in the publication.

IAB	Interactive Advertising Bureau
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization
IT	Information Technology
NIST	National Institute of Standards and Technology
NISTIR	National Institute of Standards and Technology Internal Report
OASIS	Organization for the Advancement of Structured Information Standards
OECD	Organisation for Economic Co-operation and Development
OMB	Office of Management and Budget
PMRM	Privacy Management Reference Model and Methodology
PRAM	Privacy Risk Assessment Methodology
SCRM	Supply Chain Risk Management
SDLC	System Development Life Cycle
SP	Special Publication

Appendix D: Privacy Risk Management Practices

Section 1.2 introduces a number of considerations around privacy risk management, including the relationship between cybersecurity and privacy risk and the role of privacy risk assessment. This appendix considers some of the key practices that contribute to successful privacy risk management, including organizing preparatory resources, determining privacy capabilities, defining privacy requirements, conducting privacy risk assessments, creating privacy requirements traceability, and monitoring for changing privacy risks. Category and subcategory references are included to facilitate use of the Core to support these practices; these references appear in parentheticals.

Organizing Preparatory Resources

The right resources facilitate informed decision-making about privacy risks at all levels of an organization. As a practical matter, the responsibility for the development of various resources may lie with different components of the organization. Therefore, a component of the organization depending on certain resources may find that they do not exist, or may not adequately address privacy. In these circumstances, the dependent component can consider the purpose of the resource and either seek the information through other sources, or make the best decision it can with the available information. In short, good resources are helpful, but any deficiencies should not prevent organizational components from making the best risk decisions they can within their capabilities.

The following resources, while not exhaustive, build a foundation for better decision-making.

- **Risk management role assignments** (ID.GV-P3, ID.GV-P4)

  Enabling cross-organizational understanding of who has responsibility for different risk management tasks in the organization supports better coordination and accountability for decision-making. In addition, a broad range of perspectives can improve the process of identifying, assessing, and responding to privacy risks. A diverse and cross-functional team can help to identify a more comprehensive range of risks to individuals' privacy, and to select a wider set of mitigations. Determining which roles to include in the risk management discussions depends on organizational context and makeup, although collaboration between an organization's privacy and cybersecurity programs will be important. If one individual is being assigned to multiple roles, managing potential conflicts of interest should be considered.

- **Organizational risk management strategy** (ID.RM)

  An organization's risk management strategy helps to align the organization's mission and values with organizational risk assumptions and constraints. Limitations on resources to achieve mission/business objectives and to manage risk will likely require trade-offs. Enabling personnel involved in the risk management process to better understand the organization's risk tolerance should help to guide decisions about how to allocate resources and improve decisions around risk response.

- **Key stakeholders** (ID.GV-P4, ID.SC)

  Privacy stakeholders are those who have an interest or concern in the privacy outcomes of the system, product, or service. For example, legal concerns likely focus on whether the system, product, or service is operating in a way that would cause the organization to be out of compliance with privacy laws or regulations or its business agreements. Business owners that want to maximize usage may be concerned about loss of trust in the system, product, or service due to poor privacy. Individuals whose data are being processed or who are interacting with the system, product, or service will be interested in not experiencing problems or adverse consequences. Understanding the stakeholders and the types of privacy outcomes they are interested in will facilitate system/product/service design that appropriately addresses stakeholders' needs.

- **Organizational-level privacy requirements** (ID.GV)

  Organizational-level privacy requirements are a means of expressing the legal obligations, privacy values, and policies to which the organization intends to adhere. Understanding these requirements is key to ensuring that the system/product/service design complies with its obligations. Organizational-level privacy requirements may be derived from a variety of sources, including:

  - Legal environment (e.g., laws, regulations, contracts)
  - Organizational policies or cultural values
  - Relevant global standards
  - Privacy principles

- **System/product/service design artifacts** (ID.BE-P3)

  Design artifacts may take many forms such as system design architectures or data flow diagrams. These artifacts help an organization build systems, products, and services that meet an organization's mission/business priorities and objectives. Therefore, they can help privacy programs understand how systems, products, and services need to function so that controls or measures that help to mitigate privacy risk can be selected and implemented in ways that maintain functionality while protecting privacy.

- **Data maps** (ID.IM)

  Data maps illustrate data processing and individuals' interactions with systems, products, and services. A comprehensive data map shows the data processing environment and includes the components through which data are being processed or with which individuals are interacting, the owners or operators of the components, and discrete data actions and the specific data elements being processed. A data map can be overlaid on existing system/product/service design artifacts for convenience and ease of communication between organizational components. As discussed below, a data map is an important artifact in privacy risk assessment.

Determining Privacy Capabilities

Privacy capabilities can be used to describe the system, product, or service property or feature that achieves the desired privacy outcome (e.g., "the service enables data minimization.") Security system engineers use the security objectives confidentiality, integrity, and availability along with organizational-level security requirements to consider the security capabilities for a system, product, or service. As set forth in **Table 3**, NIST has developed an additional set of privacy engineering objectives to support the determination of privacy capabilities. An organization may also use the privacy engineering objectives as a high-level prioritization tool. Systems, products, or services that are low in predictability, manageability, or disassociability may be a signal of increased privacy risk, and therefore merit a more comprehensive privacy risk assessment.

In determining privacy capabilities, an organization may consider which of the privacy engineering and security objectives are most important with respect to its mission/business needs, risk tolerance, and organizational-level privacy requirements (see Organizing Preparatory Resources above). Not all of the objectives may be equally important, or trade-offs may be necessary among them. Although the privacy capabilities inform the privacy risk assessment by supporting risk prioritization decisions, the privacy capabilities may also be informed by the risk assessment and adjusted to support the management of specific privacy risks or address changes in the environment, including design changes to the system, product, or service.

Table 3: Privacy Engineering and Security Objectives[15]

	Objective	Definition	Principal Related Functions from the Privacy Framework Core
Privacy Engineering Objectives	Predictability	Enabling reliable assumptions by individuals, owners, and operators about data and its processing by a system	Identify, Protect, Control, Inform, Respond
	Manageability	Providing the capability for granular administration of data, including alteration, deletion, and selective disclosure	Identify, Control, Respond
	Disassociability	Enabling the processing of data or events without association to individuals or devices beyond the operational requirements of the system	Identify, Protect, Respond
Security Objectives	Confidentiality	Preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information	Identify, Protect, Respond
	Integrity	Guarding against improper information modification or destruction; includes ensuring information non-repudiation and authenticity	Identify, Protect, Respond
	Availability	Ensuring timely and reliable access to and use of information	Identify, Protect, Respond

Defining Privacy Requirements

Privacy requirements specify the way the system, product, or service needs to function to meet stakeholders' desired privacy outcomes (e.g., "the application is configured to allow users to select specific data elements"). To define privacy requirements, consider organizational-level privacy requirements (see Organizing Preparatory Resources above) and the outputs of a privacy risk assessment. This process helps an organization to answer two questions: 1) what a system, product, or service *can* do with data processing and interactions with individuals, and 2) what it *should* do. Then an organization can allocate resources to design a system, product, or service in a way that achieves the defined requirements. Ultimately, this can lead to the development of systems, products, and services that are more mindful of individuals' privacy, and are based on informed risk decisions.

[15] The privacy engineering objectives are adapted from NISTIR 8062, *An Introduction to Privacy Engineering and Risk Management in Federal Systems* at https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8062.pdf. The security objectives are from NIST SP 800-37 Revision 2, *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy* at https://doi.org/10.6028/NIST.SP.800-37r2.

Conducting Privacy Risk Assessments

Conducting a privacy risk assessment helps an organization to identify privacy risks engendered by the system, product, or service and prioritize them to be able to make informed decisions about how to respond to the risks (ID.RA-P, ID.RM-P). Methodologies for conducting privacy risk assessments may vary, but organizations should consider the following characteristics: [16]

- **Risk model** (ID.RA-P6)
  Risk models define the risk factors to be assessed and the relationships among those factors.[17] If an organization is not using a pre-defined risk model, the organization should clearly define which risk factors it will be assessing and the relationships among these factors. Although cybersecurity has a widely-used risk model based on the risk factors of threats, vulnerabilities, likelihood, and impact, there is not one commonly accepted privacy risk model. NIST has developed a privacy risk model based on the risk factors of problematic data actions, likelihood, and impact, each explained below.

  NIST Privacy Risk Factors:
  Likelihood | Problematic Data Action | Impact

  - A problematic data action is any action a system takes to process data that could result in a problem for individuals. Organizations consider the type of problems that are relevant to the population of individuals. Problems can take any form and may consider the experience of individuals singly or as a group.[18]
  - Likelihood is defined as a contextual analysis that a data action is likely to create a problem for a representative set of individuals. Context can include organizational factors (e.g., the public perception about participating organizations with respect to privacy), system factors (e.g., the nature and history of individuals' interactions with the system), or individual factors (e.g., demographics of the population set).[19]
  - Impact is an analysis of the costs should the problem occur. As noted in section 1.2, the experience of individuals is a type of externality for organizations. Moreover, individuals' experiences may be subjective. Thus, impact may be difficult to assess accurately. Organizations should consider the best means of internalizing impact to individuals in order to appropriately prioritize and respond to privacy risks.[20]

[16] NIST has developed a Privacy Risk Assessment Methodology (PRAM) that can help organizations identify, assess, and respond to privacy risks. It is comprised of a set of worksheets available at https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/resources.

[17] See NIST SP 800-30 Revision 1, *Guide for Conducting Risk Assessments* at p. 8 https://doi.org/10.6028/NIST.SP.800-30r1

[18] As part of its PRAM, NIST has created an illustrative catalog of problematic data actions and problems for consideration. See https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/resources. Other organizations may have created additional problem sets, or may refer to them as adverse consequences or harms.

[19] See NIST PRAM for more information about contextual factors. Id at Worksheet 2.

[20] The NIST PRAM uses indirect costs to an organization as a proxy for considering individual impact such as non-compliance costs, direct business costs, reputational costs, and internal culture costs. Id at Worksheet 3, Impact Tab.

- **Assessment approach**
  The assessment approach is the mechanism by which identified risks are prioritized. Assessment approaches can be categorized as quantitative, semi-quantitative, or qualitative.[21] [22]
- **Prioritizing risks** (ID.RA-P7)
  Given the applicable limits of an organization's resources, organizations prioritize the risks to facilitate communication about how to respond.[23]
- **Responding to risks**
  As described in section 1.2, responding to risk is usually categorized as mitigation, transfer/sharing, avoidance, or acceptance.[24]

Creating Privacy Requirements Traceability

Once the organization has determined which risks to mitigate, the organization can refine the privacy requirements and then select and implement controls (i.e., technical and/or policy safeguards) to meet the defined requirements.[25] An organization may use a variety of sources to select controls, such as NIST Special Publication 800-53, *Security and Privacy Controls for Information Systems and Organizations.*[26] After implementation, an organization iteratively assesses the controls for their effectiveness in meeting the privacy requirements and managing privacy risk. In this way, an organization creates traceability between the controls and the privacy requirements and demonstrate accountability between its systems, products, and services and its organizational privacy goals. (ID.RA-P7)

Monitoring Changing Privacy Risks

Privacy risk management is not a static process. An Organizations monitors how changes in its business environment and corresponding changes to its systems, products, and services may be affecting privacy risk, and iteratively use the practices in this appendix to adjust accordingly. (ID.RA-P8)

[21] See NIST SP 800-30, *Guide for Conducting Risk Assessments* at p. 14 https://doi.org/10.6028/NIST.SP.800-30r1
[22] The NIST PRAM uses a semi-quantitative approach based on a scale of 1-10.
[23] The NIST PRAM provides various prioritization representations, including a heat map. See https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/resources, Worksheet 3.
[24] The NIST PRAM provides a process for responding to prioritized privacy risks. Id at Worksheet 4.
[25] See NIST SP 800-37 Revision 2, *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy* at https://doi.org/10.6028/NIST.SP.800-37r2
[26] See NIST SP 800-53, *Security and Privacy Controls for Information Systems and Organizations*, as updated at https://csrc.nist.gov/publications/sp800-53

Appendix E: Implementation Tiers Definitions

The Tiers are defined through four areas summarized below:

Tier 1: Partial

- **Privacy Risk Management Process** – Organizational privacy risk management practices are not formalized, and risk is managed in an ad hoc and sometimes reactive manner. Prioritization of privacy activities may not be directly informed by organizational risk objectives, privacy risk assessments, or business/mission requirements.
- **Integrated Privacy Risk Management Program** – There is limited awareness of privacy risk at the organizational level. The organization implements privacy risk management on an irregular, case-by-case basis due to varied experience or information gained from outside sources. The organization may not have processes that enable the sharing of information about data processing and resulting privacy risks within the organization.
- **Ecosystem Relationships** – There is limited understanding of an organization's role in the larger ecosystem with respect to other entities (e.g., buyers, suppliers, service providers, business associates, partners). The organization does not have processes for identifying how privacy risks may proliferate throughout the ecosystem or for communicating privacy risks or requirements to other entities in the ecosystem.
- **Workforce** – Some personnel may have a limited understanding of privacy risks or privacy risk management processes, but have no specific privacy responsibilities. If available, privacy training is ad hoc and the content is not kept current with best practices.

Tier 2: Risk Informed

- **Privacy Risk Management Process** – Risk management practices are approved by management but may not be established as organizational-wide policy. Prioritization of privacy activities is directly informed by organizational risk objectives, privacy risk assessments, and business/mission requirements.
- **Integrated Privacy Risk Management Program** – There is an awareness of privacy risk at the organizational level, but an organization-wide approach to managing privacy risk has not been established. Information about data processing and resulting privacy risks is shared within the organization on an informal basis. Consideration of privacy in organizational objectives and programs may occur at some but not all levels of the organization. Privacy risk assessment occurs, but is not typically repeatable or reoccurring.
- **Ecosystem Relationships** – There is some understanding of an organization's role in the larger ecosystem with respect to other entities (e.g., buyers, suppliers, service providers, business associates, partners). The organization is aware of the privacy ecosystem risks associated with the products and services it provides and uses, but does not act consistently or formally upon those risks.
- **Workforce** – There are personnel with specific privacy responsibilities, but they may have non-privacy responsibilities as well. Privacy training is conducted regularly for privacy personnel, although there is no consistent process for updates on best practices.

Tier 3: Repeatable

- **Privacy Risk Management Process** – The organization's risk management practices are formally approved and expressed as policy. Organizational privacy practices are regularly updated based on the application of risk management processes to changes in business/mission requirements and a changing risk, policy, and technology landscape.

- **Integrated Privacy Risk Management Program** – There is an organization-wide approach to manage privacy risk. Risk-informed policies, processes, and procedures are defined, implemented as intended, and reviewed. Consistent methods are in place to respond effectively to changes in risk. The organization consistently and accurately monitors privacy risk. Senior privacy and non-privacy executives communicate regularly regarding privacy risk. Senior executives ensure consideration of privacy through all lines of operation in the organization.
- **Ecosystem Relationships** – The organization understands its role, dependencies, and dependents in the larger ecosystem and may contribute to the community's broader understanding of risks. The organization is aware of the privacy ecosystem risks associated with the products and services it provides and it uses. Additionally, it usually acts formally upon those risks, including mechanisms such as written agreements to communicate baseline requirements, governance structures, and policy implementation and monitoring.
- **Workforce** – Dedicated privacy personnel possess the knowledge and skills to perform their appointed roles and responsibilities. There is regular, up-to-date privacy training for all personnel.

Tier 4: Adaptive

- **Privacy Risk Management Process** – The organization adapts its privacy practices based on lessons learned from privacy breaches and events, and identification of new privacy risks. Through a process of continuous improvement incorporating advanced privacy technologies and practices, the organization actively adapts to a changing policy and technology landscape and responds in a timely and effective manner to evolving privacy risks.
- **Integrated Privacy Risk Management Program** – There is an organization-wide approach to managing privacy risk that uses risk-informed policies, processes, and procedures to address potential privacy events. The relationship between privacy risk and organizational objectives is clearly understood and considered when making decisions. Senior executives monitor privacy risk in the same context as cybersecurity risk, financial risk, and other organizational risks. The organizational budget is based on an understanding of the current and predicted risk environment and risk tolerance. Business units implement executive vision and analyze system-level risks in the context of the organizational risk tolerances. Privacy risk management is part of the organizational culture and evolves from lessons learned and continuous awareness of data processing and resulting privacy risks. The organization can quickly and efficiently account for changes to business/mission objectives in how risk is approached and communicated.
- **Ecosystem Relationships** – The organization understands its role, dependencies, and dependents in the larger ecosystem and contributes to the community's broader understanding of risks. The organization uses real-time or near-real-time information to understand and consistently act upon privacy ecosystem risks associated with the products and services it provides and it uses. Additionally, it communicates proactively, using formal (e.g., agreements) and informal mechanisms to develop and maintain strong ecosystem relationships.
- **Workforce** – The organization has specialized privacy skillsets throughout the organizational structure; personnel with diverse perspectives contribute to the management of privacy risks. There is regular, up-to-date, specialized privacy training for all personnel. Personnel at all levels understand the organizational privacy values and their role in maintaining them.

Appendix F: Roadmap

This appendix will provide a companion roadmap to the Privacy Framework covering next steps and identifying key areas where the relevant practices are not well enough understood to enable organizations to achieve a privacy outcome. These areas will be based on input and feedback received from stakeholders through the Privacy Framework development process.