NIST Special Publication
NIST SP 800-63-4 2pd

Digital Identity Guidelines

Second Public Draft

David Temoshok
Diana Proud-Madruga
Yee-Yin Choong
Ryan Galluzzo
Sarbari Gupta
Connie LaSalle
Naomi Lefkovitz
Andrew Regenscheid

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-63-4.2pd

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Special Publication
NIST SP 800-63-4 2pd

Digital Identity Guidelines

Second Public Draft

David Temoshok
Ryan Galluzzo
Connie LaSalle
Naomi Lefkovitz
Applied Cybersecurity Division
Information Technology Laboratory

Andrew Regenscheid
Computer Security Division
Information Technology Laboratory

Yee-Yin Choong
Information Access Division
Information Technology Laboratory

Diana Proud-Madruga
Sarbari Gupta
Electrosoft

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-63-4.2pd

August 2024

U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at https://csrc.nist.gov/publications.

Authority

This publication has been developed by NIST in accordance with its statutory responsibilities under the Federal Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 et seq., Public Law (P.L.) 113-283. NIST is responsible for developing information security standards and guidelines, including minimum requirements for federal information systems, but such standards and guidelines shall not apply to national security systems without the express approval of appropriate federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130.

Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other federal official. This publication may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

NIST Technical Series Policies

Copyright, Fair Use, and Licensing Statements
NIST Technical Series Publication Identifier Syntax

Publication History

Approved by the NIST Editorial Review Board on YYYY-MM-DD [will be added upon final publication]

How to Cite this NIST Technical Series Publication

Temoshok D, Proud-Madruga D, Choong YY, Galluzzo R, Gupta S, LaSalle C, Lefkovitz N, Regenscheid A (2024) Digital Identity Guidelines. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-63-4 2pd. https://doi.org/10.6028/NIST.SP.800-63-4.2pd

Author ORCID iDs

David Temoshok: 0000-0001-6195-0331
Diana Proud-Madruga: 0000-0002-8972-7809
Yee-Yin Choong: 0000-0002-3889-6047
Ryan Galluzzo: 0000-0003-0304-4239
Sarbari Gupta: 0000-0003-1101-0856
Connie LaSalle: 0000-0001-6031-7550
Naomi Lefkovitz: 0000-0003-3777-3106
Andrew Regenscheid: 0000-0002-3930-527X

Public Comment Period

August 21, 2024 - October 7, 2024

Submit Comments

mailto:dig-comments@nist.gov

Additional Information

Additional information about this publication is available at https://csrc.nist.gov/pubs/sp/800/63/4/2pd, including related content, potential updates, and document history.

All comments are subject to release under the Freedom of Information Act (FOIA).

Abstract

These guidelines cover identity proofing and authentication of users (such as employees, contractors, or private individuals) interacting with government information systems over networks. They define technical requirements in each of the areas of identity proofing, registration, authenticators, management processes, authentication protocols, federation, and related assertions. They also offer technical recommendations and other informative text intended as helpful suggestions. The guidelines are not intended to constrain the development or use of standards outside of this purpose. This publication supersedes NIST Special Publication (SP) 800-63-3.

Keywords

authentication; authentication assurance; authenticator; assertions; credential service provider; digital authentication; digital credentials; identity proofing; federation; passwords; PKI.

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems. The Special Publication 800-series reports on ITL's research, guidelines, and outreach efforts in information system security, and its collaborative activities with industry, government, and academic organizations.

Note to Reviewers

In December 2022, NIST released the Initial Public Draft (IPD) of SP 800-63, Revision 4. Over the course of a 119-day public comment period, the authors received exceptional feedback from a broad community of interested entities and individuals. The input from nearly 4,000 specific comments has helped advance the improvement of these Digital Identity Guidelines in a manner that supports NIST's critical goals of providing foundational risk management processes and requirements that enable the implementation of secure, private, equitable, and accessible identity systems. Based on this initial wave of feedback, several substantive changes have been made across all of the volumes. These changes include but are not limited to the following:

1. Updated text and context setting for risk management. Specifically, the authors have modified the process defined in the IPD to include a context-setting step of defining and understanding the online service that the organization is offering and intending to potentially protect with identity systems.
2. Added recommended continuous evaluation metrics. The continuous improvement section introduced by the IPD has been expanded to include a set of recommended metrics for holistically evaluating identity solution performance. These are recommended due to the complexities of data streams and variances in solution deployments.
3. Expanded fraud requirements and recommendations. Programmatic fraud management requirements for credential service providers and relying parties now address issues and challenges that may result from the implementation of fraud checks.
4. Restructured the identity proofing controls. There is a new taxonomy and structure for the requirements at each assurance level based on the means of providing the proofing: Remote Unattended, Remote Attended (e.g., video session), Onsite Unattended (e.g., kiosk), and Onsite Attended (e.g., in-person).
5. Integrated syncable authenticators. In April 2024, NIST published interim guidance for syncable authenticators. This guidance has been integrated into SP 800-63B as normative text and is provided for public feedback as part of the Revision 4 volume set.
6. Added user-controlled wallets to the federation model. Digital wallets and credentials (called "attribute bundles" in SP 800-63C) are seeing increased attention and adoption. At their core, they function like a federated IdP, generating signed assertions about a subject. Specific requirements for this presentation and the emerging context are presented in SP 800-63C-4.

The rapid proliferation of online services over the past few years has heightened the need for reliable, equitable, secure, and privacy-protective digital identity solutions.

Revision 4 of NIST Special Publication SP 800-63, *Digital Identity Guidelines*, intends to respond to the changing digital landscape that has emerged since the last major revision of this suite was published in 2017, including the real-world implications of online risks. The guidelines present the process and technical requirements for meeting digital identity management assurance levels for identity proofing, authentication, and federation, including requirements for security and privacy as well as considerations for fostering equity and the usability of digital identity solutions and technology.

Based on the feedback provided in response to the June 2020 Pre-Draft Call for Comments, research into real-world implementations of the guidelines, market innovation, and the current threat environment, this draft seeks to:

- Address comments received in response to the IPD of Revision 4 of SP 800-63
- Clarify the text to address the questions and issues raised in the public comments
- Update all four volumes of SP 800-63 based on current technology and market developments, the changing digital identity threat landscape, and organizational needs for digital identity solutions to address online security, privacy, usability, and equity

NIST is specifically interested in comments and recommendations on the following topics:

1. Risk Management and Identity Models
    - Is the "user controlled" wallet model sufficiently described to allow entities to understand its alignment to real-world implementations of wallet-based solutions such as mobile driver's licenses and verifiable credentials?
    - Is the updated risk management process sufficiently well-defined to support an effective, repeatable, real-world process for organizations seeking to implement digital identity system solutions to protect online services and systems?
2. Identity Proofing and Enrollment
    - Is the updated structure of the requirements around defined types of proofing sufficiently clear? Are the types sufficiently described?
    - Are there additional fraud program requirements that need to be introduced as a common baseline for CSPs and other organizations?
    - Are the fraud requirements sufficiently described to allow for appropriate balancing of fraud, privacy, and usability trade-offs?
    - Are the added identity evidence validation and authenticity requirements and performance metrics realistic and achievable with existing technology capabilities?

3. Authentication and Authenticator Management

    - Are the syncable authenticator requirements sufficiently defined to allow for reasonable risk-based acceptance of syncable authenticators for public and enterprise-facing uses?
    - Are there additional recommended controls that should be applied? Are there specific implementation recommendations or considerations that should be captured?
    - Are wallet-based authentication mechanisms and "attribute bundles" sufficiently described as authenticators? Are there additional requirements that need to be added or clarified?

4. Federation and Assertions

    - Is the concept of user-controlled wallets and attribute bundles sufficiently and clearly described to support real-world implementations? Are there additional requirements or considerations that should be added to improve the security, usability, and privacy of these technologies?

5. General

    - What specific implementation guidance, reference architectures, metrics, or other supporting resources could enable more rapid adoption and implementation of this and future iterations of the Digital Identity Guidelines?
    - What applied research and measurement efforts would provide the greatest impacts on the identity market and advancement of these guidelines?

Reviewers are encouraged to comment and suggest changes to the text of all four draft volumes of the SP 800-63-4 suite. NIST requests that all comments be submitted by 11:59pm Eastern Time on October 7th, 2024. Please submit your comments to dig-comments@nist.gov. NIST will review all comments and make them available on the NIST Identity and Access Management website. Commenters are encouraged to use the comment template provided on the NIST Computer Security Resource Center website for responses to these notes to reviewers and for specific comments on the text of the four-volume suite.

Call for Patent Claims

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this Information Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be directly stated in this ITL Publication or by reference to another publication. This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to this ITL draft publication and of any relevant unexpired U.S. or foreign patents.

ITL may require from the patent holder, or a party authorized to make assurances on its behalf, in written or electronic form, either:

a) assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or

b) assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this ITL draft publication either:

   i. under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or

   ii. without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure that the commitments in the assurance are binding on the transferee, and that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

Such statements should be addressed to: mailto:dig-comments@nist.gov.

Table of Contents

312 List of Tables

317 List of Figures

Preface

This publication and its companion volumes, [SP800-63A], [SP800-63B], and [SP800-63C], provide technical guidelines to organizations for the implementation of digital identity services.

Acknowledgments

The authors would like to thank their fellow collaborators — Christine Abruzzi, James L. Fenton, and Justin P. Richer — on the current revision of this Special Publication, as well as Kerrianne Buchanan for her contributions and review. The authors would also like to acknowledge the past contributions of Donna F. Dodson, Elaine M. Newton, Ray A. Perlner, W. Timothy Polk, Emad A. Nabbus, Paul A. Grassi, Michael E. Garcia, Kaitlin Boeckl, Joni Brennan, Ellen Nadeau, Ben Piccareta, and Danna Gabel O'Rourke.

1. Introduction

This section is informative.

The rapid proliferation of online services over the past few years has heightened the need for reliable, equitable, secure, and privacy-protective digital identity solutions. A digital identity is always unique in the context of an online service. However, a person may have multiple digital identities and while a digital identity may relay a unique and specific meaning within the context of an online service, the real-life identity of the individual behind the digital identity may not be known. When confidence in a person's real-life identity is not required to provide access to an online service, organizations may use anonymous or pseudonymous accounts. In all other use cases, a digital identity is intended to demonstrate trust between the holder of the digital identity and the person, organization, or system on the other side of the online service. However, this process can present challenges. There are multiple opportunities for mistakes, miscommunication, impersonation, and other attacks that fraudulently claim another person's digital identity. Additionally, given the broad range of individual needs, constraints, capacities, and preferences, online services must be designed with equity, usability, and flexibility to ensure broad and enduring participation and access to digital devices and services.

Digital identity risks are dynamic and exist along a continuum; consequently, organizations' digital identity risk management approach should seek to manage risks using outcome-based approaches that are designed to meet the organization's unique needs. This guidance defines specific assurance levels which operate as baseline control sets designed to provide a common point for organizations seeking to address identity-related risks. Assurance levels provide multiple benefits, including a starting point for agencies in their risk management journey and a common structure for supporting interoperability between different entities. It is, however, impractical to create assurance levels that can comprehensively address the entire spectrum of risks, threats, or considerations and organization will face when deploying an identity solution. For this reason, these guidelines promote a risk-oriented approach to digital identity solution implementation rather than a compliance-oriented approach, and organizations are encouraged to tailor their control implementations based on the processes defined in these guidelines.

Additionally, risks associated with digital identity stretch beyond the potential impacts to the organization providing online services. These guidelines endeavor to account for risks to individuals, communities, and other organizations more robustly and explicitly. Organizations should consider how digital identity decisions that prioritize security might affect, or need to accommodate, the individuals who interact with the organization's programs and services. Privacy, equity, and usability for individuals should be considered along with security. Additionally, organizations should consider their digital identity approach alongside other mechanisms for identity management,

such as those used in call centers and in-person interactions. By taking a human-centric and continuously informed approach to mission delivery, organizations have an opportunity to incrementally build trust with the variety of populations they serve, improve customer satisfaction, identify issues more quickly, and provide individuals with culturally appropriate and effective redress options.

The composition, models, and availability of identity services has significantly changed since the first version of SP 800-63 was released, as have the considerations and challenges of deploying secure, private, usable, and equitable services to diverse user communities. This revision addresses these challenges by clarifying requirements based on the function that an entity may serve under the overall digital identity model.

Additionally, this publication provides instruction for credential service providers (CSPs), verifiers, and relying parties (RPs), that supplement the *NIST Risk Management Framework* [NISTRMF] and its component special publications. It describes the risk management processes that organizations should follow to implement digital identity services and expands upon the NIST RMF by outlining how equity and usability considerations should be incorporated. It also highlights the importance of considering impacts, not only on enterprise operations and assets, but also on individuals, other organizations, and — more broadly — society. Furthermore, digital identity management processes for identity proofing, authentication, and federation typically involve processing personal information, which can present privacy risks. Therefore, these guidelines include privacy requirements and considerations to help mitigate potential associated privacy risks.

Finally, while these guidelines provide organizations with technical requirements and recommendations for establishing, maintaining, and authenticating the digital identity of subjects who access digital systems over a network, additional support options outside of the purview of information technology teams may be needed to address barriers and adverse impacts, foster equity, and successfully deliver on mission objectives.

1.1. Scope and Applicability

This guidance applies to all online services for which some level of digital identity is required, regardless of the constituency (e.g., residents, business partners, and government entities). For this publication, “person” refers only to natural persons.

These guidelines primarily focus on organizational services that interact with external users, such as residents accessing public benefits or private-sector partners accessing collaboration spaces. However, it also applies to federal systems accessed by employees and contractors. The *Personal Identity Verification (PIV) of Federal Employees and Contractors* standard [FIPS201] and its corresponding set of Special Publications and organization-specific instructions extend these guidelines for the federal enterprise, by providing additional technical controls and processes for issuing and managing

Personal Identity Verification (PIV) Cards, binding additional authenticators as derived PIV credentials, and using federation architectures and protocols with PIV systems.

Online services not covered by this guidance include those associated with national security systems as defined in 44 U.S.C. § 3552(b)(6). Private-sector organizations and state, local, and tribal governments whose digital processes require varying levels of digital identity assurance may consider the use of these standards where appropriate.

These guidelines address logical access to online systems, services, and applications. They do not specifically address physical access control processes. However, the processes specified in these guidelines can be applied to physical access use cases where appropriate. Additionally, these guidelines do not explicitly address some subjects including, but not limited to, machine-to-machine authentication, interconnected devices (e.g., Internet of Things (IoT) devices), or access to Application Programming Interfaces (APIs) on behalf of subjects.

1.2. How to Use This Suite of SPs

These guidelines support the mitigation of the negative impacts of errors that occur during the identity system functions of identity proofing, authentication, and federation. Sec. 3, Digital Identity Risk Management, provides details on the risk assessment process and how the results of the risk assessment and additional context inform the selection of controls to secure the identity proofing, authentication, and federation processes. Controls are selected by determining the assurance level required to mitigate each applicable type of digital identity error for a particular service based on risk and mission.

Specifically, organizations are required to individually select assurance levels[1] that correspond to each function being performed:

- Identity Assurance Level (IAL) refers to the identity proofing process.
- Authentication Assurance Level (AAL) refers to the authentication process.
- Federation Assurance Level (FAL) refers to the federation process when the RP is connected to a CSP or an IdP through a federated protocol.

SP 800-63 is organized as the following suite of volumes:

- SP 800-63 *Digital Identity Guidelines* provides the digital identity models, risk assessment methodology, and process for selecting assurance levels for identity proofing, authentication, and federation. *SP 800-63 contains both normative and informative material.*
- [SP800-63A]: provides requirements for identity proofing and the enrollment of applicants, either remotely or in-person, that wish to gain access to resources at each of the three IALs. It details the responsibilities of CSPs with respect to

[1] When described generically or bundled, these guidelines will refer to IAL, AAL, and FAL as ***xAL***.

establishing and maintaining subscriber accounts and binding CSP issued or subscriber-provided authenticators to the subscriber account. *SP 800-63A contains both normative and informative material.*

- [SP800-63B] provides requirements for authentication processes, including choices of authenticators, that may be used at each of the three AALs. It also provides recommendations on events that may occur during the lifetime of authenticators, including invalidation in the event of loss or theft. *SP 800-63B contains both normative and informative material.*
- [SP800-63C] provides requirements on the use of federated identity architectures and assertions to convey the results of authentication processes and relevant identity information to an agency application. This volume offers privacy-enhancing techniques for sharing information about a valid, authenticated subject, and describes methods that allow for strong multi-factor authentication (MFA) while the subject remains pseudonymous to the online service. *SP 800-63C contains both normative and informative material.*

1.3. Enterprise Risk Management Requirements and Considerations

Effective enterprise risk management is multidisciplinary by design and involves the consideration of diverse sets of factors and equities. In a digital identity risk management context, these factors include, but are not limited to, information security, privacy, equity, and usability. It is important for risk management efforts to weigh these factors as they relate to enterprise assets and operations, individuals, other organizations, and society.

During the process of analyzing factors relevant to digital identity, organizations may determine that measures outside of those specified in this publication are appropriate in certain contexts (e.g., where privacy or other legal requirements exist or where the output of a risk assessment leads the organization to determine that additional measures or alternative procedural safeguards are appropriate). Organizations, including federal agencies, may employ compensating or supplemental controls that are not specified in this publication. They may also consider partitioning the functionality of an online service to allow less sensitive functions to be available at a lower level of assurance in order to improve equity and access without compromising security.

The considerations detailed below support enterprise risk management efforts and encourage informed, inclusive, and human-centered service delivery. While this list of considerations is not exhaustive, it highlights a set of cross-cutting factors that are likely to impact decision-making associated with digital identity management.

1.3.1. Security, Fraud, and Threat Prevention

It is increasingly important for organizations to assess and manage digital identity security risks, such as unauthorized access due to impersonation. As organizations

consult this guidance, they should consider potential impacts to the confidentiality, integrity, and availability of information and information systems that they manage and that their service providers and business partners manage on behalf of the individuals and communities that they serve.

Federal agencies implementing these guidelines are required to meet statutory responsibilities, including those under the *Federal Information Security Modernization Act (FISMA) of 2014* [FISMA] and related NIST standards and guidelines. NIST recommends that non-federal organizations implementing these guidelines follow comparable standards (e.g., ISO 27001) to ensure the secure operation of their digital systems.

FISMA requires federal agencies to implement appropriate controls to protect federal information and information systems from unauthorized access, use, disclosure, disruption, or modification. The NIST RMF [NISTRMF] provides a process that integrates security, privacy, and cyber supply-chain risk management activities into the system development life cycle. It is expected that federal agencies and organizations that provide services under these guidelines have already implemented the controls and processes required under FISMA and associated NIST risk management processes and publications.

The controls and requirements encompassed by the identity, authentication, and Federation Assurance Levels under these guidelines augment, but do not replace or alter, the information and information system controls determined under FISMA and the RMF.

It is increasingly important for organizations to assess and manage identity-related fraud risks associated with identity proofing and authentication processes. As organizations consult this guidance, they should consider the evolving threat environment, the availability of innovative anti-fraud measures in the digital identity market, and the potential impact of identity-related fraud. This is particularly important with respect to public-facing online services where the impact of identity-related fraud on e-government service delivery, public trust, and agency reputation can be substantial. This version enhances measures to combat identity theft and identity-related fraud by repurposing IAL1 as a new assurance level, updating authentication risk and threat models to account for new attacks, providing new options for phishing resistant authentication, introducing requirements to prevent automated attacks against enrollment processes, and preparing for new technologies (e.g., mobile driver's licenses and verifiable credentials) that can leverage strong identity proofing and authentication.

1.3.2. Privacy

When designing, engineering, and managing digital identity systems, it is imperative to consider the potential of that system to create privacy-related problems for individuals when processing (e.g., collection, storage, use, and destruction) personally identifiable information (PII) and the potential impacts of problematic data actions. If a breach of

PII or a release of sensitive information occurs, organizations need to ensure that the privacy notices describe, in plain language, what information was improperly released and, if known, how the information was exploited.

Organizations need to demonstrate how organizational privacy policies and system privacy requirements have been implemented in their systems. These guidelines recommend that organizations employ the full set of legal and regulatory mandates that may affect their users and technology providers including:

- The *NIST Privacy Framework* [NISTPF], which enables privacy engineering practices that support privacy by design concepts and helps organizations protect individuals' privacy.
- The *[PrivacyAct] of 1974, 2020 Edition* which established a set of fair information practices for the collection, maintenance, use, and disclosure of information about individuals that is maintained by federal agencies in systems of records.
- *OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002* [M-03-22], which describes the Privacy Impact Assessments that are supported by the privacy risk assessments that are required for PII processing or storing.
- [SP800-53] *Security and Privacy Controls for Information Systems and Organizations* , which lists privacy controls that can be implemented to mitigate the risks identified in the privacy risk and impact assessments.
- [SP800-122] *Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)*, which assists federal agencies in understanding what PII is, the relationship between protecting the confidentiality of PII, privacy, and the Fair Information Practices, and safeguards for protecting PII.

Furthermore, each volume of SP 800-63, ([SP800-63A], [SP800-63B], and [SP800-63C]) contains a specific section providing detailed privacy guidance and considerations for the implementation of the processes, controls, and requirements presented in that volume as well as normative requirements on data collection, retention, and minimization.

1.3.3. Equity

Equity has been defined as "the consistent and systematic fair, just, and impartial treatment of all individuals, including individuals who belong to underserved communities that have been denied such treatment" [EO13985]. Incorporating equity considerations when designing or operating a digital identity service helps ensure a person's ability to engage in an online service, such as accessing a critical service like healthcare. Accessing online services is often dependent on a person's ability to present a digital identity and use the required technologies successfully and safely. Many populations are either unable to successfully present a digital identity or face a higher

degree of burden in navigating online services than their more privileged peers. In a public service context, this poses a direct risk to successful mission delivery. In a broader societal context, challenges related to digital access can exacerbate existing inequities and continue systemic cycles of exclusion for historically marginalized and underserved groups.

To support the continuous evaluation and improvement program described in Sec. 3, it is important to maintain awareness of existing inequities faced by served populations and potential new inequities or disparities between populations that could be caused or exacerbated by the design or operation of digital identity systems. This can help identify the opportunities, processes, business partners, and multi-channel identity proofing and service delivery methods that best support the needs of those populations while also managing privacy, security, and fraud risks.

Further, section 508 of the Rehabilitation Act of 1973 (2011) [Section508] was enacted to eliminate barriers in information technology and require federal agencies to make electronic and information technologies accessible to people with disabilities. While these guidelines do not directly assert requirements from [Section508], federal agencies and their identity service providers are expected to design online services and systems with the experiences of people with disabilities in mind to ensure that accessibility is prioritized.

1.3.4. Usability

Usability refers to the extent to which a system, product, or service can be used to achieve goals with effectiveness, efficiency, and satisfaction in a specified context of use. Usability also supports major objectives such as equity, service delivery, and security. Like equity, usability requires an understanding of the people who interact with a digital identity system or process, as well as their unique goals and context of use.

Readers of this guidance should take a holistic approach to considering the interactions that each user will engage in throughout the process of enrolling in and authenticating to a service. Throughout the design and development of a digital identity system or process, it is important to conduct usability evaluations with demographically representative users, from all communities served and perform realistic scenarios and tasks in appropriate contexts of use. Additionally, following usability guidelines and considerations can help organizations meet customer experience goals articulated in federal policy [EO14058]. Digital identity management processes should be designed and implemented so that it is easy for users to do the right thing, hard to do the wrong thing, and easy to recover when the wrong thing happens.

1.4. Notations

This guideline uses the following typographical conventions in text:

- Specific terms in **CAPITALS** represent normative requirements. When these same terms are not in **CAPITALS**, the term does not represent a normative requirement.
  - The terms "**SHALL**" and "**SHALL NOT**" indicate requirements to be followed strictly in order to conform to the publication and from which no deviation is permitted.
  - The terms "**SHOULD**" and "**SHOULD NOT**" indicate that among several possibilities, one is recommended as particularly suitable without mentioning or excluding others, that a certain course of action is preferred but not necessarily required, or that (in the negative form) a certain possibility or course of action is discouraged but not prohibited.
  - The terms "**MAY**" and "**NEED NOT**" indicate a course of action that is permissible within the limits of the publication.
  - The terms "**CAN**" and "**CANNOT**" indicate a material, physical, or causal possibility and capability or — in the negative — the absence of that possibility or capability.

1.5. Document Structure

This document is organized as follows. Each section is labeled as either normative (i.e., mandatory for compliance) or informative (i.e., not mandatory).

- Section 1 provides an introduction to the document. This section is *informative*.
- Section 2 describes a general model for digital identity. This section is *informative*.
- Section 3 describes the digital identity risk model. This section is *normative*.
- The References section contains a list of publications that are cited in this document. This section is *informative*.
- Appendix A contains a selected list of abbreviations used in this document. This appendix is *informative*.
- Appendix B contains a glossary of selected terms used in this document. This appendix is *informative*.
- Appendix C contains a summarized list of changes in this document's history. This appendix is *informative*.

2. Digital Identity Model

This section is informative.

2.1. Overview

The SP 800-63 guidelines use digital identity models that reflect technologies and architectures that already currently available in the market. These models have a variety of entities and functions and vary in complexity. Simple models group functions, such as creating subscriber accounts and providing attributes, under a single entity. More complex models separate these functions among a larger number of entities. The entities, and their associated functions, found in digital identity models include:

Subject: In these guidelines, a subject is a person and is represented by one of three roles, depending on where they are in the digital identity process.

- Applicant — The subject to be identity-proofed and enrolled.
- Subscriber — - The subject who has successfully completed the identity proofing and enrollment process or authentication (i.e., when the subject is in an active on-line session).
- Claimant — The subject "making a claim" to be eligible for authentication.

Service provider: Service providers can perform any combination of functions involved in granting access to and delivering online services, such as a credential service provider, relyin party, verifier, and Identity provider.

Credential service provider (CSP): CSP functions include identity proofing applicants to the identity service and registering authenticators to subscriber accounts. A *subscriber account* is the CSP's established record of the subscriber, the subscriber's attributes, and associated authenticators. CSP functions may be performed by an independent third party.

Relying party (RP): RP functions rely on the information in the subscriber account from the CSP, typically to process a digital transaction or grant access to information or a system. When using federation, the RP accesses the information in the subscriber account through assertions from an identity provider.

Verifier: The function of a verifier is to verify the claimant's identity by verifying the claimant's possession and control of one or more authenticators using an authentication protocol. To do this, the verifier needs to confirm the binding of the authenticators with the subscriber account and check that the subscriber account is active.

Identity provider (IdP): When using federation, the IdP manages the subscriber's primary authenticators and issues assertions derived from the subscriber account.

2.2. Identity Proofing and Enrollment

Normative requirements can be found in [SP800-63A], *Identity Proofing and Enrollment*.

[SP800-63A] provides general information and normative requirements for the identity proofing and enrollment processes as well as requirements that are specific to IALs.

Figure 1 shows a sample of interactions for identity proofing and enrollment.

To start, an *applicant* opts to enroll with a CSP by requesting access. The CSP or the entity fulfilling CSP functions requests identity evidence and attributes, which the applicant provides. If the applicant is successfully identity-proofed, they are enrolled in the identity service as a *subscriber* of that CSP. A unique subscriber account is then created and one or more authenticators are registered to the subscriber account.

Subscribers have a responsibility to maintain control of their authenticators (e.g., guard against theft) and comply with CSP policies to remain in good standing with the CSP.

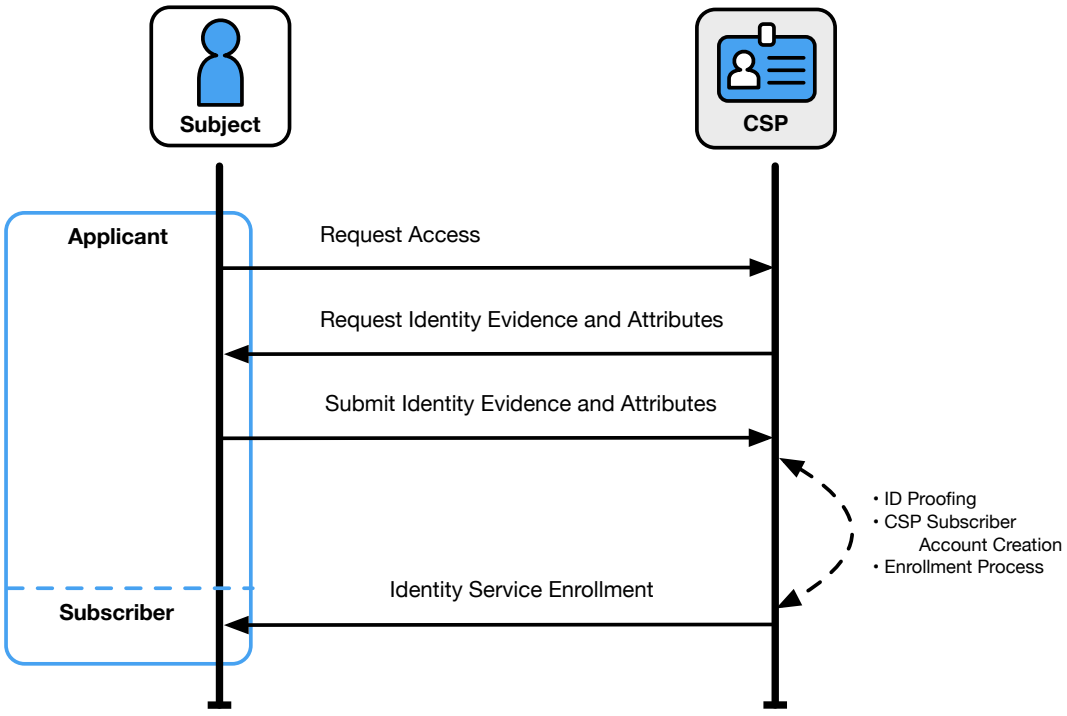

Fig. 1. Sample Identity Proofing and Enrollment Digital Identity Model

2.2.1. Subscriber Accounts

At the time of enrollment, the CSP establishes a subscriber account to uniquely identify each subscriber and record any authenticators registered (bound) to that subscriber account. The CSP may:

- Issue and register one or more authenticators to the subscriber at the time of enrollment,
- Register authenticators provided by the subscriber to the subscriber account,
- Register additional authenticators to the subscriber account at a later time as needed, or
- Provision the subscriber account to one or more general-purpose or subscriber-controlled wallets, for use in a federated protocol system.

See Sec. 5 of [SP800-63A], *Subscriber-Accounts*, for more information and normative requirements.

2.3. Authentication and Authenticator Management

Normative requirements can be found in [SP800-63B], *Authentication and Authenticator Management*.

2.3.1. Authenticators

[SP800-63B] provides normative descriptions of permitted authenticator types, their characteristics (e.g., phishing resistance), and authentication processes appropriate for each AAL.

This guidance defines three types of authentication factors used for authentication:

- Something you know (e.g., a password)
- Something you have (e.g., a device containing a cryptographic key)
- Something you are (e.g., a fingerprint or other biometric characteristic data)

Single-factor authentication requires only one of the above factors, most often "something you know". Multiple instances of the same factor still constitute single-factor authentication. For example, a user-generated PIN and a password do not constitute two factors as they are both "something you know." Multi-factor authentication (MFA) refers to the use of more than one distinct factor.

This guidance specifies that authenticators always contain or comprise a secret. The secrets contained in an authenticator are based on either key pairs (i.e., asymmetric cryptographic keys) or shared secrets (including symmetric cryptographic keys, seeds for generating one-time passwords (OTP), and passwords). Asymmetric key pairs are comprised of a public key and a related private key. The private key is stored on the

authenticator and is only available for use by the claimant who possesses and controls the authenticators. A verifier that has the subscriber's public key (e.g., through a public key certificate) can use an authentication protocol to verify that the claimant is a subscriber who has possession and control of the associated private key contained in the authenticator. Symmetric keys are generally chosen at random, complex and long enough to thwart network-based guessing attacks, and stored in hardware or software that the subscriber controls. Passwords typically have fewer characters and less complexity than cryptographic keys resulting in increased vulnerabilities that require additional defenses to mitigate.

Passwords used as activation factors for multi-factor authenticators are referred to as *activation secrets*. An activation secret is used to decrypt a stored key used for authentication or is compared against a locally held and stored verifier to provide access to the authentication key. In either of these cases, the activation secret remains within the authenticator and its associated user endpoint. An example of an activation secret would be the PIN used to activate a PIV card.

Biometric characteristics are unique, personal attributes that can be used to verify the identity of a person who is physically present at the point of authentication. This includes, but is not limited to, facial features, fingerprints, and iris patterns. While biometric characteristics cannot be used for single-factor authentication, they can be used as an authentication factor for multi-factor authentication when used in combination with a physical authenticator (i.e., something you have).

Some authentication methods used for in-person interactions do not apply directly to digital authentication. For example, a physical driver's license is something you have and may be useful when authenticating to a human (e.g., a security guard), but it is not an authenticator for online services.

Some commonly used authentication methods do not contain or comprise secrets and are therefore not acceptable for use under these guidelines. For example:

- Knowledge-based authentication, where the claimant is prompted to answer questions that are presumably known only by the claimant, does not constitute an acceptable secret for digital authentication.
- A biometric characteristic does not constitute a secret and cannot be used as a single-factor authenticator.

2.3.2. Authentication Process

The authentication process enables an RP to trust that a claimant is who they say they are. Some approaches are described in [SP800-63B], *Authentication and Authenticator Management*. The sample authentication process in Fig. 2 shows interactions between the RP, a claimant, and a verifier/CSP. The verifier is a functional role and is frequently implemented in combination with the CSP, the RP, or both (as shown in Fig. 4).

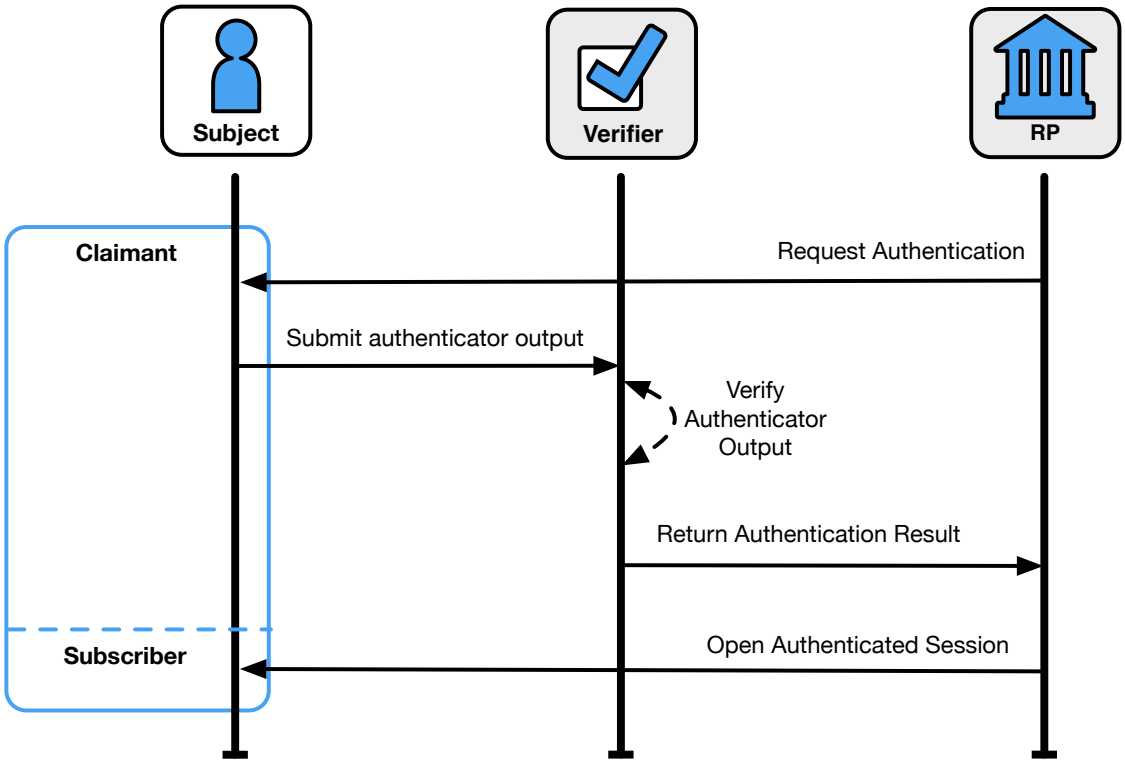

Fig. 2. Sample Authentication Process

A successful authentication process demonstrates that the claimant has possession and control of one or more valid authenticators that are bound to the subscriber's identity. In general, this is done using an authentication protocol that involves an interaction between the verifier and the claimant. The exact nature of the interaction is extremely important in determining the overall security of the system. Well-designed protocols can protect the integrity and confidentiality of communication between the claimant and the verifier both during and after the authentication and can help limit the damage done by an attacker masquerading as a legitimate verifier.

Additionally, mechanisms located at the verifier can mitigate online guessing attacks against lower entropy secrets (e.g., passwords and PINs) by limiting the rate at which an attacker can make authentication attempts, or otherwise delaying incorrect attempts. Generally, this is done by keeping track of and limiting the number of unsuccessful attempts, since the premise of an online guessing attack is that most attempts will fail.

2.4. Federation and Assertions

Normative requirements can be found in [SP800-63C], *Federation and Assertions*.

Section III of OMB [M-19-17] *Enabling Mission Delivery through Improved Identity, Credential, and Access Management* directs agencies to support cross-government

identity federation and interoperability. The term *federation* can be applied to several different approaches that involve the sharing of information between different trust domains. These approaches differ based on the kind of information that is being shared between the domains. These guidelines address the federation processes that allow for the conveyance of identity and authentication information based on trust agreements across a set of networked systems through federation assertions.

There are many benefits to using federated architectures including, but not limited to:

- Enhanced user experience (e.g., a subject can be identity proofed once but their subscriber account used at multiple RPs).
- Cost reduction to both the subscriber (e.g., reduction in authenticators) and the organization (e.g., reduction in information technology infrastructure and a streamlined architecture).
- Minimizing data in RPs that do not need to collect, store, or dispose of personal information.
- Minimizing data exposed to RPs by using pseudonymous identifiers and derived attribute values instead of copying account values to each application.
- Mission enablement, since organizations will need to focus fewer resources on complex identity management processes.

While the federation process is generally the preferred approach to authentication when the RP and IdP are not administered together under a common security domain, federation can also be applied within a single security domain for a variety of benefits including centralized account management and technical integration.

The SP 800-63 guidelines are agnostic to the identity proofing, authentication, and federation architectures that an organization selects, and they allow organizations to deploy a digital identity scheme according to their own requirements. However, there are scenarios that an organization may encounter that make federation potentially more efficient and effective than establishing identity services that are local to the organization or individual applications. The following lists detailed potential scenarios in which the organization may consider federation to be a viable option:

- Potential users already have an authenticator at or above the required AAL.
- Multiple types of authenticators are required to cover all possible user communities.
- An organization does not have the necessary infrastructure to support management of subscriber accounts (e.g., account recovery, authenticator issuance, help desk).
- There is a desire to allow primary authenticators to be added and upgraded over time without changing the RP's implementation.

- There are different environments to be supported, since federation protocols are network-based and allow for implementation on a wide variety of platforms and languages.
- Potential users come from multiple communities, each with its own existing identity infrastructure.
- The organization needs the ability to centrally manage account lifecycles, including account revocation and the binding of new authenticators.

An organization may want to consider accepting federated identity attributes if any of the following apply:

- Pseudonymity is required, necessary, feasible, or important to stakeholders accessing the service.
- Access to the service requires a defined list of attributes.
- Access to the service requires at least one derived attribute value.
- The organization is not the authoritative source or issuing source for required attributes.
- Attributes are only required temporarily during use (e.g., to make an access decision), and the organization does not need to retain the data.

2.5. Examples of Digital Identity Models

The entities and interactions that comprise the non-federated digital identity model are illustrated in Fig. 3. The general-purpose federated digital identity model is illustrated in Fig. 4, and a federated digital identity model with a subscriber-controlled wallet is illustrated in Fig. 5.

Figure 3 shows an example of a common sequence of interactions in the non-federated model. Other sequences could also achieve the same functional requirements. One common sequence of interactions for identity proofing and enrollment activities is as follows:

- Step 1: An applicant applies to a CSP through an identity proofing and enrollment process. The CSP identity proofs that applicant.
- Step 2: Upon successful identity proofing, the applicant is enrolled in the identity service as a subscriber.
  - A subscriber account and corresponding authenticators are established between the CSP and the subscriber. The CSP maintains the subscriber account, its status, and the enrollment data. The subscriber maintains their authenticators.

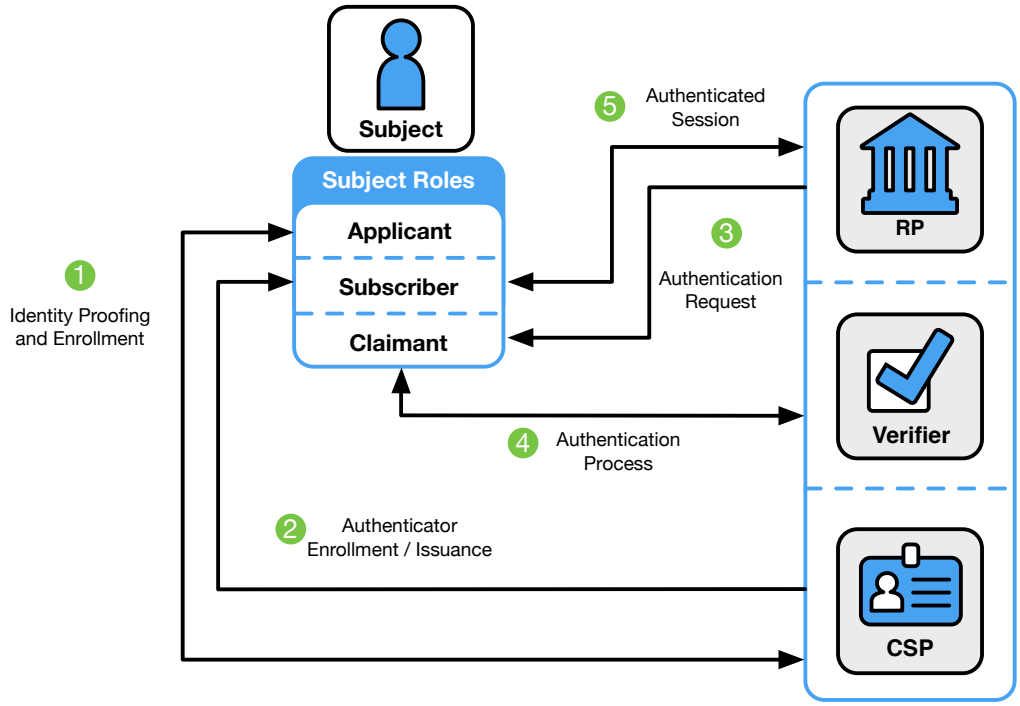

Fig. 3. Non-Federated Digital Identity Model Example

Steps 3 through 5 may immediately follow steps 1 and 2 or they may be done at a later time. The usual sequence of interactions involved in using one or more authenticators to perform digital authentication in the non-federated model is as follows:

- Step 3: The claimant initiates an online interaction with the RP and the RP requests that the claimant authenticate.
- Step 4: The claimant proves possession and control of the authenticators to the verifier through an authentication process:
  - The verifier interacts with the CSP to verify the binding of the claimant's identity to their authenticators in the subscriber account and to optionally obtain additional subscriber attributes.
  - The CSP or verifier functions of the service provider give information about the subscriber. The RP requests the attributes it requires from the CSP. The RP optionally uses this information to make authorization decisions.
- Step 5: An authenticated session is established between the subscriber and the RP.

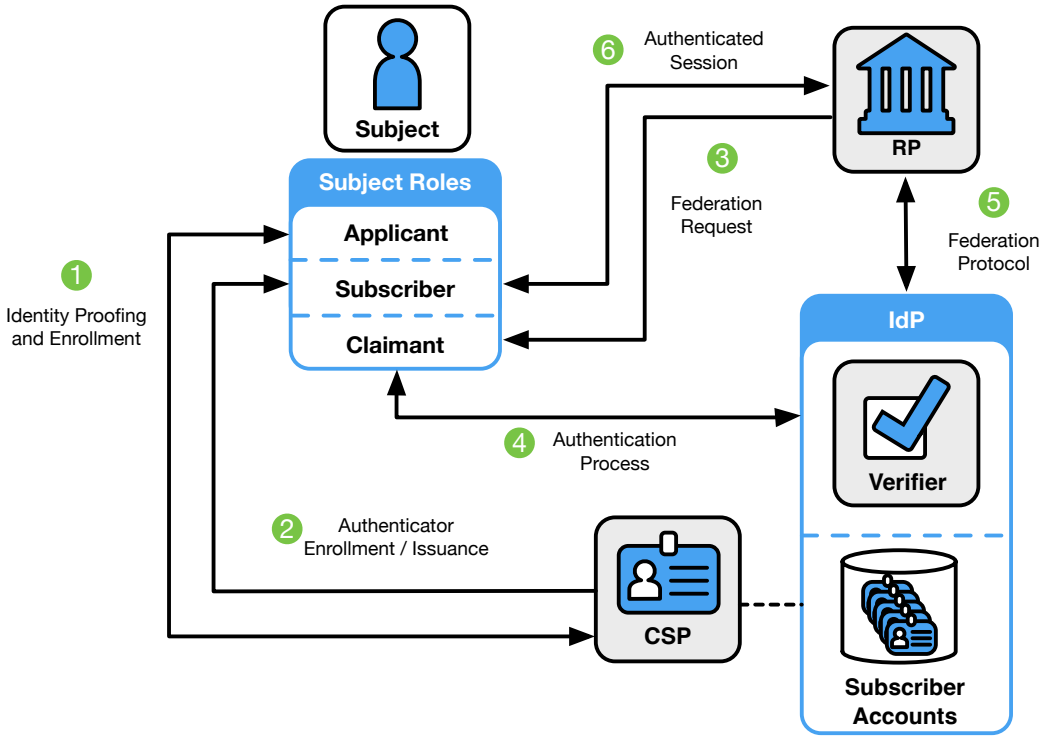

Fig. 4. Federated Digital Identity Model Example

Figure 4 shows an example of those same common interactions in a federated model.

- Step 1: An applicant applies to a CSP through an identity proofing and enrollment process. The CSP identity proofs that applicant.
- Step 2: Upon successful identity proofing, the applicant is enrolled in the identity service as a subscriber.
  - A subscriber account and corresponding authenticators are established between the CSP and the subscriber.
  - Unlike in Fig. 3, the IdP is provisioned either directly by the CSP or indirectly through access to attributes of the subscriber account. The CSP maintains the subscriber account, its status, and the enrollment data collected in accordance with the record retention and disposal requirements described in Sec. 3.1.1 of [SP800-63A]. The subscriber maintains their authenticators. The IdP maintains its view of the subscriber account, any federated identifiers assigned to the subscriber account, and authorizations to RPs.

The usual sequence of interactions involved in using one or more authenticators in the federated model to perform digital authentication is as follows:

- Step 3: The RP requests that the claimant authenticate. This triggers a request for federated authentication to the IdP.
- Step 4: The claimant proves possession and control of the authenticators to the verifier function of the IdP through an authentication process.
  - Within the IdP, the verifier and CSP functions interact to verify the binding of the claimant's authenticators with those bound to the claimed subscriber account and optionally to obtain additional subscriber attributes.
- Step 5: The RP and the IdP communicate through a federation protocol. The IdP provides an assertion and optionally additional attributes to the RP through a federation protocol. The RP verifies the assertion to establish confidence in the identity and attributes of a subscriber for an online service at the RP. RPs may use a subscriber's federated identity (pseudonymous or non-pseudonymous), IAL, AAL, FAL, and other factors to make authorization decisions.
- Step 6: An authenticated session is established between the subscriber and the RP.

In the two cases described in Fig. 3 and Fig. 4, the verifier does not always need to communicate in real time with the CSP to complete the authentication activity (e.g., digital certificates can be used). Therefore, the line between the verifier and the CSP represents a logical link between the two entities. In some implementations, the verifier, RP, and CSP functions may be distributed and separated. However, if these functions reside on the same platform, the interactions between the functions are signals between applications or application modules that run on the same system rather than using network protocols.

Figure 5 shows an example of the interactions in a federated digital identity model in which the subscriber controls a device with software (i.e., a digital wallet) that acts as the IdP. In the terminology of the "three-party model", the CSP is the issuer, the IdP is the holder, and the RP is the verifier. In this model, it is common for the RP to establish a trust agreement with the CSP through the use of a federation authority as defined in [SP800-63C]. This arrangement allows the RP to accept assertions from the subscriber-controlled wallet without needing a direct trust relationship with the wallet.

- Step 1: An applicant applies to a CSP identity proofing and enrollment process.
- Step 2: Upon successful identity proofing, the applicant goes through an onboarding process and is enrolled in the identity service as a subscriber.
- Step 3: The subscriber-controlled wallet is onboarded by the CSP.
  - The subscriber authenticates to the CSP's onboarding function.

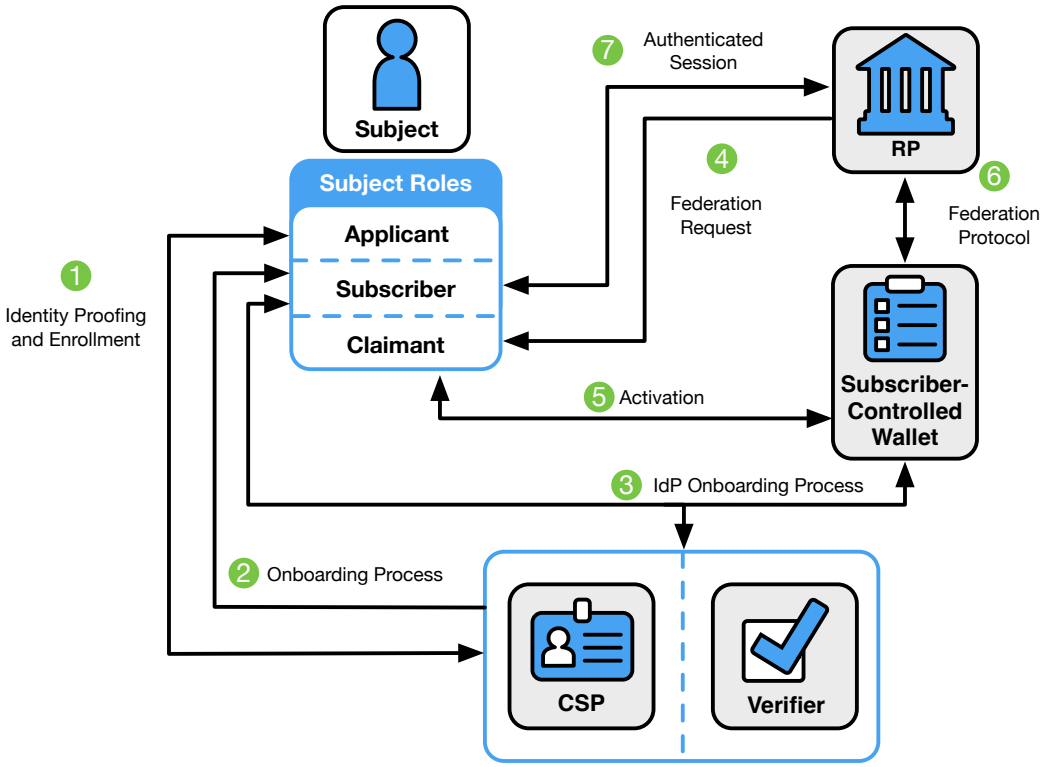

Fig. 5. Federated Digital Identity Model with Subscriber-Controlled Wallet Example

- The subscriber activates the subscriber-controlled wallet using an activation factor.
- The wallet sends a request to the CSP, including proof of a key held by the wallet.
- The CSP creates an attribute bundle that contains a reference for the key of the wallet and any additional attributes.

The usual sequence of interactions involved in providing an assertion to the RP from a subscriber-controlled wallet is as follows:

- Step 4: The RP requests that the claimant authenticate. This triggers a request for federated authentication to the wallet.
- Step 5: The claimant proves possession and control of the subscriber-controlled wallet.
  - The subscriber activates the wallet using an activation factor.

911 – The wallet prepares an assertion including the attribute bundle provided by
912 the CSP for the subscriber account.

913 • Step 6: The RP and the wallet communicate through a federation protocol. The
914 wallet provides an assertion and optionally additional attributes to the RP through
915 a federation protocol. The RP verifies the assertion to establish confidence in the
916 identity and attributes of a subscriber for an online service at the RP. RPs may use
917 a subscriber's federated identity (pseudonymous or non-pseudonymous), IAL, AAL,
918 FAL, and other factors to make authorization decisions.

919 • Step 7: An authenticated session is established between the subscriber and the RP.

> Note: Other protocols and specifications often refer to attribute bundles as *credentials*. These guidelines use the term *credentials* for a different concept. To avoid a conflict, the term *attribute bundle* is used within these guidelines. Normative requirements for attribute bundles can be found including Sec. 3.11.1 of [SP800-63C].

920

3. Digital Identity Risk Management

This section is normative.

This section provides details on the methodology for assessing digital identity risks associated with online services and the residual risks to users of the online service, communities impacted by the service, the service provider organization, and its mission and business partners. It offers guidance on selecting usable, equitable, and privacy-enhancing security and anti-fraud controls that mitigate those risks. Additionally, it emphasizes the importance of continuously evaluating the performance of the selected controls.

The Digital Identity Risk Management (DIRM) process focuses on the identification and management of risks according to two dimensions: (1) risks to the online service that might be addressed by an identity system; and (2) risks from the identity system to be implemented.

The first dimension of risk informs initial assurance level selections and seeks to identify the risks associated with a compromise of the online service that might be addressed by an identity system. For example:

- Identity proofing: What negative impacts would reasonably be expected if an imposter were to gain access to a service or receive a credential using the identity of a legitimate user (e.g., an attacker successfully impersonates someone)?
- Authentication: What negative impacts would reasonably be expected if a false claimant accessed an account that was not rightfully theirs (e.g., an attacker who compromises or steals an authenticator)?
- Federation: What negative impacts would reasonably be expected if the wrong subject successfully accessed an online service, system, or data (e.g., compromising or replaying an assertion)?

All three types of errors can result in the wrong subject successfully accessing an online service, system, or data.

If it is determined that there are risks associated with a compromise of the online service that could be addressed by an identity system, an initial assurance level is selected and the second dimension of risk is then considered. The second dimension of risk seeks to identify the risks posed by the identity system and informs the tailoring process. Tailoring provides a process to modify an initially assessed assurance level, implement compensating or supplemental controls, or modify selected controls based on ongoing detailed risk assessments.

For example, assuming that aspects of the identity system are not sufficiently privacy-enhancing, usable, equitable, or able or necessary to address specific real-world threats:

- Identity proofing: What is the impact of not successfully identity proofing and enrolling a legitimate subject due to barriers faced by the subject throughout the process of identity proofing, including biases? What is the impact of falling victim to a breach of information that was excessively collected and retained to support identity proofing processes? What is the impact if the initial IAL does not completely address specific threats, threat actors, and fraud?
- Authentication: What is the impact of failing to authenticate the correct subject due to barriers faced by the subject in presenting their authenticator, including biases or usability issues? What is the impact if the initial AAL does not completely address targeted account takeover models or specific authenticator types fail to mitigate anticipated attacks?
- Federation: What is the impact of releasing subscriber attributes to the wrong online service or system?

The outcomes of the DIRM process depend on the role that an entity plays within the digital identity model.

1. For **relying parties**, the intent of this process is to determine the assurance levels and any tailoring required to protect online services and the applications, transactions, and systems that comprise or are impacted by those services. This directly contributes to the selection, development, and procurement of CSP services. Federal RPs **SHALL** implement the DIRM process for all online services.
2. For **credential service providers and identity providers**, the intent of this process is to design service offerings that meet the requirements of the defined assurance levels, continuously guard against compromises to the identity system, and meet the needs of RPs. Whenever a service offering deviates from normative guidance, those deviations must be clearly communicated to the RPs that utilize the service. All CSPs **SHALL** implement the DIRM process for the services they offer and **SHALL** make a Digital Identity Acceptance Statement (DIAS) for each offering available to all current or potential RPs. CSPs **MAY** base their assessment on anticipated or representative digital identity services they wish to support. In creating this risk assessment, CSPs **SHOULD** seek input from real-world RPs on their user populations and their anticipated context.

This process augments the risk management processes required by the Federal Information Security Modernization Act [FISMA]. The results of the DIRM impact assessment for the online service may be different from the FISMA impact level for the underlying application or system. Identity process failures may result in different levels of impact for various user groups. For example, the overall assessed FISMA impact level for a payment system may result in a 'FISMA Moderate' impact category due to sensitive

financial data processed by the system. However, for individuals who are making guest payments where no persistent account is established, the authentication and proofing impact levels may be lower as associated data may not be retained or made accessible. Agency authorizing officials **SHOULD** require documentation demonstrating adherence to the DIRM process as a part of the Authority to Operate (ATO) for the underlying information system that supports an online service. Agency authorizing officials **SHOULD** require documentation from CSPs demonstrating adherence to the DIRM as part of procurement or ATO processes for integration with CSPs.

There are 5 steps in the DIRM process:

1. **Define the online service**: As a starting point, the organization documents a description of the online service in terms of its functional scope, the user groups it is intended to serve, the types of online transactions available to each user group, and the underlying data that the online service processes through its interfaces. If the online service is one element of a broader business process, its role is documented, as are the impacts of any data collected and processed by the online service. Additionally, an organization needs to determine the entities that will be impacted by the online service and the broader business process of which it is a part. The outcome is a description of the online service, its users, and the entities that may be impacted by its functionality.
2. **Conduct initial impact assessment**: In this step, organizations evaluate their user population and assess the impacts of a compromise of the online service that might be addressed by an identity system (i.e., identity proofing, authentication, or federation). Each function of the online service is assessed against a defined set of harms and impact categories. Each user group of the online service is considered separately based on the transactions available to that user group (i.e., the permissions that the group is granted relative to the data and functions of the online service). The outcome of this step is a documented set of impact categories and associated impact levels (i.e., Low, Moderate, or High) for each user group of the online service.
3. **Select initial assurance levels**: In this step, the impact categories and impact levels are evaluated to determine the initial assurance levels to protect the online service from unauthorized access and fraud. Using the assurance levels, the organization identifies the baseline controls for the IAL, AAL, and FAL for each user group based on the requirements from companion volumes [SP800-63A], [SP800-63B], and [SP800-63C], respectively. The outcome of this step is an identified initial IAL, AAL, and FAL, as applicable, for each user group.
4. **Tailor and document assurance level determinations**: In this step, detailed assessments are conducted or leveraged to determine the potential impact of the initially selected assurance levels and their associated controls on privacy, equity, usability, and resistance to the current threat environment. Tailoring may

result in a modification of the initially assessed assurance level, the identification of compensating or supplemental controls, or both. All assessments and final decisions are documented and justified. The outcome is a DIAS (see Sec. 3.4.4) with a defined and implementable set of assurance levels and a final set of controls for the online service.

5. **Continuously evaluate and improve**: In this step, information on the performance of the identity management approach is gathered and evaluated. This evaluation considers a diverse set of factors, including business impacts, effects on fraud rates, and impacts on user communities. This information is crucial in determining if the selected assurance level and controls meet mission, business, security, and — where applicable — program integrity needs. It also helps monitor for unintended harms that impact privacy and equitable access. Opportunities for improvement should also be considered by closely monitoring the evolving threat landscape and investigating new technologies and methodologies that can counter those threats or improve usability, equity, or privacy. The outcomes of this step are performance metrics, documented and transparent processes for evaluation and redress, and ongoing improvements to the identity management approach.

Figure 6 illustrates the major actions and outcomes for each step of the DIRM process flow. While presented as a "stepwise" approach, there can be many points in the process that require divergence from the sequential order, including the need for iterative cycles between initial task execution and revisiting tasks. For example, the introduction of new regulations or requirements while an assessment is ongoing may require organizations to revisit a step in the process. Additionally, new functionality, changes in data usage, and changes to the threat environment may require an organization to revisit steps in the Digital Identity Risk Management process at any point, including potentially modifying the assurance level and/or the related controls of the online service.

Organizations **SHOULD** adapt and modify this overall approach to meet organizational processes, governance, and enterprise risk management practices. At a minimum, organizations **SHALL** execute and document each step, consult with a representative sample of the online service's user population to inform the design and performance evaluation of the identity management approach, and complete and document the normative mandates and outcomes of each step regardless of operational approach or enabling tools.

3.1. Define the Online Service

The purpose of defining the online service is to establish a common understanding of the context and circumstances that influence the organization's risk management decisions. The context-rich information ascertained during this step is intended to inform subsequent steps of the DIRM process. The role of the online service is contextualized as part of the broader business environment and associated processes, resulting in

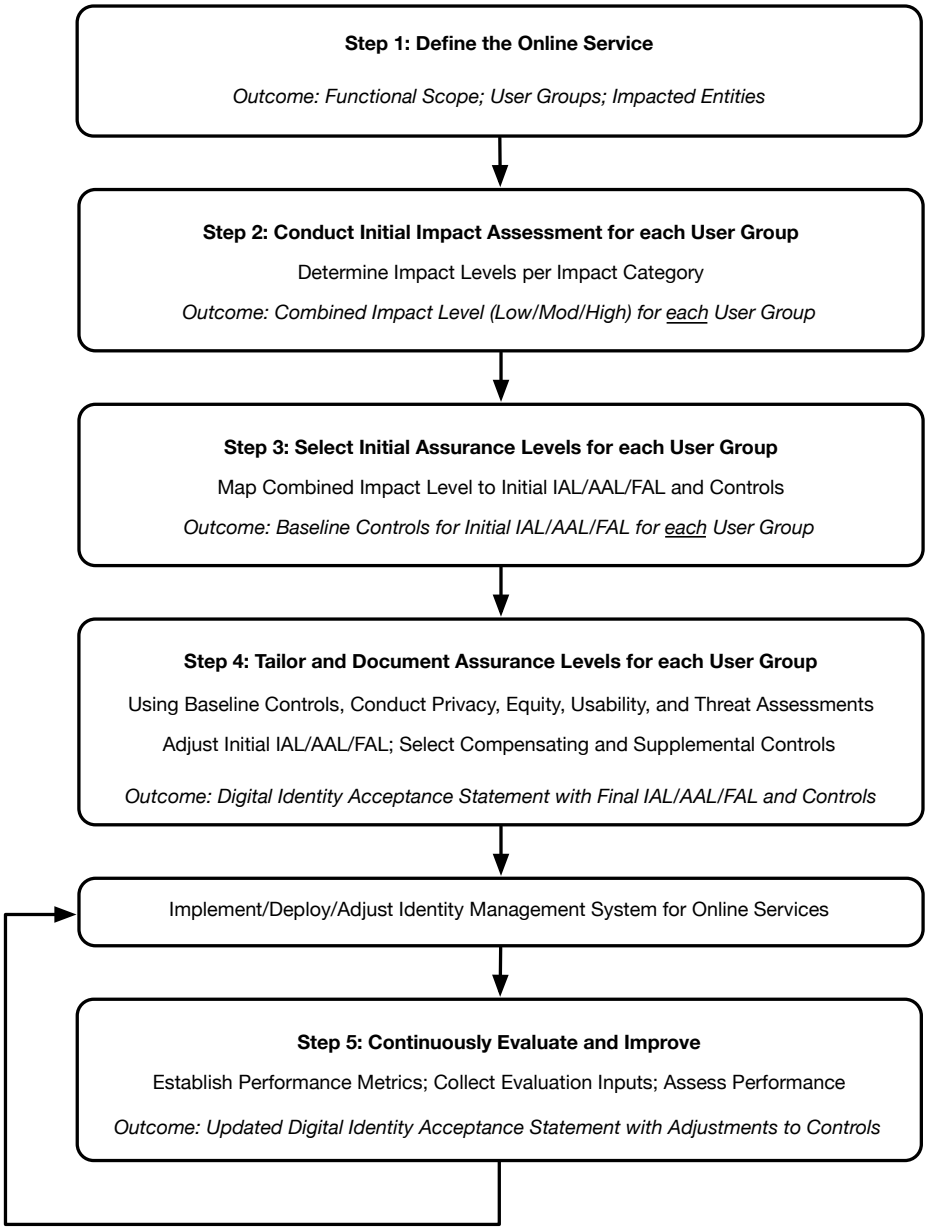

Fig. 6. High-level diagram of the Digital Identity Risk Management Process Flow

a documented description of the online service functionality, user groups and their expectations, data processed and other pertinent details.

RPs **SHALL** develop a description of the online service that includes, at minimum:

- Organizational mission and business objectives supported by the online service
- Mission and business partner dependencies associated with the online service
- Legal, regulatory, and contractual requirements, including privacy and civil liberties obligations that apply to the online service
- Functionality of the online service and the underlying data that it is expected to process
- User groups that need to have access to the online service as well as the types of online transactions and privileges available to each user group
- User expectations for the online service, including functionality, features, identity verification and authentication options, accessibility and language requirements, and culturally responsive communication alternatives
- The results of any pre-existing DIRA assessments (as an input) and the current state of any pre-existing identity technologies (i.e., proofing, authentication, or federation)
- Across all users served, the estimated availability of forms of identity evidence to support the identity proofing process for services that require identity proofing.

Additionally, an organization needs to determine the entities that will be impacted by the online service and the broader business process of which it is a part. It is imperative to consider the unexpected and undesirable impacts on different entities, populations, or demographic groups that result from an unauthorized user gaining access to the online service due to a failure of the digital identity system. For example, if an attacker obtained unauthorized access to an application that controls a power plant, the actions taken by the bad actor could have devastating environmental impacts on the local populations that live near the facility as well as cause power outages for the localities served by the plant.

It is important to differentiate between *user groups* and *impacted entities* as described in this document. The online service will allow access to a set of users who may be partitioned into a few user groups based on the kind of functionality that is offered to that user group. For example, an income tax filing and review online service may have the following user groups: (i) citizens who need to check on the status of their personal tax returns; (2) tax preparers who file tax returns on behalf of their clients; and (3) system administrators who assign privileges to different groups of users or create new user groups as needed. In contrast, impacted entities include all populations impacted by the online service and its functionality. For example, an online service that allows

remote access to control, operate and monitor a water treatment facility may have the following types of impacted entities: (1) populations that drink the water from that water treatment facility; (2) technicians who control and operate the water treatment facility; (3) the organization that owns and operates the facility; and (4) auditors and other officials who provide oversight of the facility and its compliance with applicable regulations.

Accordingly, impact assessments **SHALL** include individuals who use the online application as well as the organization itself. Additionally, organizations **SHOULD** identify other entities (e.g., mission partners, communities, and those identified in [SP800-30]) that need to be specifically included based on mission and business needs. At a minimum, agencies **SHALL** document all impacted when conducting their impact assessments.

The output of this step is a documented description of the online service including a list of entities that are impacted by the functionality provided by the online service. This information will serve as a basis and establish the context for effectively applying the impact assessments as detailed in the following sections.

3.2. Conduct Initial Impact Assessment

This step of the DIRM process addresses the first dimension of risk (i.e., risks to the identity system) and seeks to identify the risks to the online service that might be addressed by an identity system.

The purpose of the initial impact assessment is to identify the potential adverse impacts of failures in identity proofing, authentication, and federation that are specific to an online service, yielding an initial set of assurance levels. RPs **SHOULD** consider historical data and results from user focus groups when performing this step. The impact assessment **SHALL** include:

- Identifying a set of impact categories and the potential harms for each impact category,
- Identifying the levels of impact, and
- Assessing the level of impact for each user group.

The level of impact for each user group identified in Sec. 3.1 **SHALL** be considered separately based on the transactions available to that user group. Assessing the user groups separately allows organizations maximum flexibility in selecting and implementing an identity approach and assurance levels that are appropriate for each user group.

The output of this assessment is a defined impact level (i.e., Low, Moderate, or High) for each user group. This serves as the primary input to the initial assurance level selection. The effort focuses on defining and documenting the impact assessment to promote consistent application across an organization.

3.2.1. Identify Impact Categories and Potential Harms

Initial assurance levels for online services **SHALL** be determined by assessing the potential impact of — at a minimum — each of the following categories:

- Degradation of mission delivery
- Damage to trust, standing or reputation
- Unauthorized access to information
- Financial loss or financial liability
- Loss of life or danger to human safety, human health, or environmental health

Organizations **SHOULD** include additional impact categories, as appropriate, based on their mission and business objectives. Each impact category **SHALL** be documented and consistently applied when implementing the DIRM process across different online services offered by the organization.

Harms refer to any adverse effects that would be experienced by an entity. They provide a means to effectively understand the impact categories and how they may apply to specific entities impacted by the online service. For each impact category, agencies **SHALL** consider potential harms for each of the impacted entities identified in Sec. 3.1.

Examples of harms associated with each category include, but are not limited to:

- Degradation of mission delivery:
  - Harms to individuals may include the inability to access government services or benefits for which they are eligible.
  - Harms to the organization (including the organization offering the online service as well as organizations supported by the online service) may include an inability to perform current mission/business functions in a sufficiently timely manner, with sufficient confidence and/or correctness, or within planned resource constraints or an inability or limited ability to perform mission/business functions in the future.
- Damage to trust, standing or reputation:
  - Harms to individuals may include damage to image or reputation as a result of impersonation.
  - Harms to the organization may include damage to reputation resulting in damage to existing trust relationships, image, or reputation or the inability to forge future, potential trust relationships.
- Unauthorized access to information:

  - Harms to individuals may include breach of PII or other sensitive information, which may result in secondary harms such as financial loss, loss of life, physical or psychological injury, impersonation, identity theft, or persistent inconvenience.
  - Harms to the organization may include exfiltration, deletion, degradation, or exposure of intellectual property or unauthorized disclosure of other information assets such as classified materials or controlled unclassified information (CUI).
- Financial loss or liability:
  - Harms to individuals may include debts incurred or assets lost as a result of fraud or other harm, damage to or loss of credit, actual or potential employment, or sources of income, loss of accessible affordable housing and/or other financial loss.
  - Harms to the organization may include costs related to fraud or other criminal activity, loss of assets, devaluation, or loss of business.
- Loss of life or danger to human safety, human health, or environmental health:
  - Harms to individuals may include death; damage to or loss of physical, mental, or emotional well-being; or impact to environmental health that could result in uninhabitability of the local environment and require some level of intervention to address potential or actual damage.
  - Harms to the organization may include damage to or loss of the organization's workforce or the impact of unsafe conditions that render the organization unable to operate or reduce its capacity to operate.

The outcome of this activity will be a list of impact categories and harms that will be used to assess impacts on entities identified in Sec. 3.1.

3.2.2. Identify Potential Impact Levels

Initial assurance levels for digital transactions are determined by assessing the potential level of impact caused by a compromise of the online service that might be addressed by an identity system for each of the impact categories selected for consideration by the organization (from Sec. 3.2.1). Impact levels can be assigned using one of the following potential impact values:

- **Low**: Could be expected to have a limited adverse effect
- **Moderate**: Could be expected to have a serious adverse effect
- **High**: Could be expected to have a severe or catastrophic adverse effect

In this step, the impact of access by an unauthorized individual **SHALL** be considered for each user group, each impact category, and each of the impacted entities. Examples of potential impacts in each of the categories are provided below. However, to provide a more objective basis for impact level assignments, organizations **SHOULD** develop thresholds and examples for the impact levels for each impact category. Where this is done, particularly with specifically defined quantifiable values, these thresholds **SHALL** be documented and used consistently in the DIRM assessments across an organization to allow for a common understanding of risks.

- **Degradation of mission delivery:**
  - **Low**: Expected to result in limited mission capability degradation such that the organization is still able to perform its primary functions but with noticeably reduced effectiveness.
  - **Moderate**: Expected to result in serious mission capability degradation such that the organization is still able to perform its primary functions but with significantly reduced effectiveness.
  - **High**: Expected to result in severe or catastrophic mission capability degradation or loss over a duration such that the organization is unable to perform one or more of its primary functions.
- **Damage to trust, standing or reputation:**
  - **Low**: Expected to result in limited, short-term inconvenience, distress, or embarrassment to any party.
  - **Moderate**: Expected to result in serious short-term or limited long-term inconvenience, distress, or damage to the standing or reputation of any party.
  - **High**: Expected to result in severe or serious long-term inconvenience, distress, or damage to the standing or reputation of any party; ordinarily reserved for situations with particularly severe effects or that potentially affect many individuals.
- **Unauthorized access to information:**
  - **Low**: Expected to have a limited adverse effect on organizational operations, organizational assets, or individuals as defined in [FIPS199].
  - **Moderate**: Expected to have a serious adverse effect on organizational operations, organizational assets, or individuals as defined in [FIPS199].
  - **High**: Expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, or individuals as defined in [FIPS199].

- **Financial loss or financial liability:**
  - **Low**: Expected to result in limited financial loss or liability to any party.
  - **Moderate**: Expected to result in a serious financial loss or liability to any party.
  - **High**: Expected to result in severe or catastrophic financial loss or liability to any party.
- **Loss of life or danger to human safety, human health, or environmental health:**
  - **Low**: Expected to result in minor injury or an acute health issue that resolves itself and does not require medical attention, including mental health treatment; or an impact to environmental health that requires at most some limited intervention to prevent further or reverse existing damage.
  - **Moderate**: Expected to result in moderate risk of minor injury or limited risk of injury that requires medical attention, including mental health treatment; an impact to environmental health that results in a period of uninhabitability and requires intervention to prevent further or reverse existing damage; or the compounding impacts of multiple low-impact events.
  - **High**: Expected to result in serious injury, trauma, or death; impacts to environmental health that results in long-term or permanent uninhabitability and require significant intervention to prevent further or reverse existing damage, if possible; or the compounding impacts of multiple moderate impact events.

This guidance provides three impact levels. However, agencies **MAY** define more granular impact levels and develop their own methodologies for their initial impact assessment activities.

3.2.3. Impact Analysis

The impact analysis considers the level of impact (i.e., Low, Moderate or High) of compromises of the online service that might be addressed by the identity system functions (i.e., identity proofing, authentication, and federation). The impact analysis considers the following dimensions:

- User groups Sec. 3.1
- Impacted entities Sec. 3.1
- Impact categories Sec. 3.2.1
- Impact levels Sec. 3.2.2

If there is no harm or impact for a given impact category for any entity, the impact level can be marked as None.

For each user group, the impact analysis **SHALL** consider the level of impact for each impact category for each type of impacted entity. Because different sets of transactions are available to each user group, it is important to consider each user group separately for this analysis.

For example, for an online service that allows for the control, operation and monitoring of a water treatment facility, each group of users (e.g., technicians who control and operate the facility, auditors and monitoring officials, system administrators, etc.) is considered separately based on the transactions available to that user group through the online service. In other words, the impact analysis tries to determine if a bad actor obtained unauthorized access to the online service as a member of a user group and performed some nefarious actions and the level of impact (i.e., Low, Moderate or High) on various impacted entities (e.g., citizens who drink the water, the organization that owns the facility, auditors, monitoring officials, etc.) for each of the impact categories being considered.

The impact analysis **SHALL** be performed for each user group that has access to the online service. For each impact category, the impact level is estimated for each impacted entity as a result of a compromise of the online service caused by failures in the identity management functions.

The output of this impact analysis is a set of impact levels for each user group that SHALL be documented in a suitable format for further analysis in accordance with the next subsection below.

3.2.4. Determine Combined Impact Level for Each User Group

The impact assessment level results for each user group generated from the previous step are combined to establish a single impact level for that user group. This single impact level represents the risks to impacted entities that result from a compromise of identity proofing, authentication, and/or federation functions for that user group.

Organizations can apply a variety of methods for this combinatorial analysis to determine the effective impact level for each user group. Some options include:

- Using a high-water mark approach across the various impact categories and impacted entities to derive the effective impact level
- Assigning different weights to different impact categories and/or impacted entities and taking an average to derive the effective impact level
- Some other combinatorial logic that aligns with the organization's mission and priorities

Organizations **SHALL** document the approach they use to combine their impact assessment into an overall impact score for each of their defined user groups and **SHALL** apply it consistently across all its online services. At the conclusion of the combinatorial analysis, organizations **SHALL** document the impact for each user group.

The outcome of this step is an effective impact level for each user group due to a compromise of the identity management system functions (i.e., identity proofing, authentication, federation).

3.3. Select Initial Assurance Levels and Baseline Controls

The initial impact analysis of the last step yields an effective impact level (i.e., Low, Moderate, or High) that serves as a primary input to the process of selecting the initial assurance levels for identity proofing, authentication, and federation for each user group.

The purpose of the initial assurance level is to identify baseline digital identity controls (including process and technology elements) for each identity management function, from the requirements and guidelines in the companion volumes [SP800-63A], [SP800-63B], and [SP800-63C].

The initial set of digital identity controls and processes selected will be assessed and tailored in Step 4 based on potential risks generated by the identity management system.

3.3.1. Assurance Levels

Depending on the functionality and deployed architecture of the online service, it may require the support of one or more of the identity management functions (i.e., identity proofing, authentication, and federation). The strength of these functions is described in terms of assurance levels. The RP **SHALL** identify the types of assurance levels that apply to their online service from the following:

- **IAL**: The robustness of the identity proofing process to determine the identity of an individual. The IAL is selected to mitigate risks that result from potential identity proofing failures.
- **AAL**: The robustness of the authentication process itself, and the binding between an authenticator and a specific individual's identifier. The AAL is selected to mitigate risks that result from potential authentication failures.
- **FAL**: The robustness of the federation process used to communicate authentication and attribute information to an RP from an IdP. The FAL is selected to mitigate risks that result from potential federation failures.

3.3.2. Assurance Level Descriptions

A summary of each of the xALs is provided below. While high-level descriptions of the assurance levels are provided in this subsection, readers of this guidance are encouraged to refer to companion volumes [SP800-63A], [SP800-63B], and [SP800-63C] for normative guidelines and requirements for each assurance level.

3.3.2.1. Identity Assurance Level

- **IAL1**: Supports the real-world existence of the claimed identity. Core attributes are obtained from identity evidence or asserted by the applicant. All core attributes are validated against authoritative or credible sources and steps are taken to link the attributes to the person undergoing the identity proofing process.
- **IAL2**: IAL2 adds rigor by requiring the collection of additional evidence and a more rigorous process for validating the evidence and verifying the identity.
- **IAL3**: IAL3 adds the requirement for a trained CSP representative (i.e., proofing agent) to interact directly with the applicant as part of an on-site attended identity proofing session as well as the collection of at least one biometric.

Table 1. IAL Summary

IAL	Control Objectives
IAL1	Limit highly scalable attacks; provide protection against synthetic identity. Provide protections against attacks using compromised PII.
IAL2	Limit scaled and targeted attacks. Provide protections against basic evidence falsification and evidence theft. Provide protections against basic social engineering.
IAL3	Limit sophisticated attacks. Provide protections against advanced evidence falsification, theft, and repudiation. Provide protection against advanced social engineering attacks.

3.3.2.2. Authentication Assurance Level

- **AAL1**: AAL1 provides a basic level of confidence that the claimant controls an authenticator bound to the subscriber account being authenticated. AAL1 requires only single-factor authentication using a wide range of available authentication technologies. However, it is recommended that online services assessed at AAL1 offer multi-factor authentication options. Successful authentication requires that the claimant prove possession and control of the authenticator.

- **AAL2**: AAL2 provides high confidence that the claimant controls one or more authenticators bound to the subscriber account being authenticated. Proof of possession and control of two distinct authentication factors is required. A phishing-resistant authentication option must be offered for online services assessed at AAL2.
- **AAL3**: AAL3 provides very high confidence that the claimant controls one or more authenticators bound to the subscriber account being authenticated. Authentication at AAL3 is based on the proof of possession of a key through the use of a public-key cryptographic protocol. AAL3 authentication requires a hardware-based authenticator with a non-exportable private key and a phishing-resistant authenticator; the same device may fulfill both requirements. In order to authenticate at AAL3, claimants are required to prove possession and control of two distinct authentication factors.

Table 2. AAL Summary

AAL	Control Objectives
AAL1	Provide minimal protections against attacks. Deter password focused attacks.
AAL2	Support multifactor authentication. Offer phishing-resistant options.
AAL3	Provide phishing resistance and verifier compromise protections.

3.3.2.3. Federation Assurance Level

- **FAL1**: FAL1 allows a subscriber to authenticate to the RP using an assertion from an IdP in a federation protocol. FAL1 provides assurance that the assertion came from a specific IdP and was intended for a specific RP.
- **FAL2**: FAL2 additionally requires that the trust agreement between the IdP and RP be established prior to the federation transaction, and that the RP have robust protections against injection of assertions from attackers.
- **FAL3**: FAL3 additionally requires the subscriber to authenticate directly to the RP with a bound authenticator and present the assertion from the IdP. Additionally, the IdP and RP establish their identities and cryptographic key material with each other through a highly trusted process that is often manual.

Table 3. FAL Summary

FAL	Control Objectives
FAL1	Provide protections against forged assertions.
FAL2	Provide protections against forged assertions and injection attacks.
FAL3	Provide protection against IdP compromise.

3.3.3. Initial Assurance Level Selection

The overall impact level for each user group is used as the basis for the selection of the initial assurance level and related technical and process controls for the digital identity functions for the organization's online service under assessment. These initial assurance levels and control selections are primarily based on the impacts arising from failures within the digital identity functions that allow an unauthorized entity to gain access to the online service. The initial assurance levels and controls will be further assessed and tailored, as appropriate, in the next step of the DIRM process.

Organizations **SHALL** develop and document a process and governance model for selecting initial assurance levels and controls based on the potential impact of failures in the digital identity approach. This section provides guidance on the major elements to include in that process.

While online service providers must assess and determine the xALs that are appropriate for protecting their applications, the selection of these assurance levels does not mean that the online service provider must implement the controls independently. Based on the identity model that the online service provider chooses to implement, some or all of the assurance levels may be implemented by an external entity such as a third-party CSP or IdP.

3.3.3.1. Selecting Initial IAL

Before selecting an initial assurance level, RPs must determine if identity proofing is needed for the users of their online services. Identity proofing is not required if the online service does not require any personal information to execute digital transactions. If personal information is needed, the RP needs to determine if validated attributes are required or if self-asserted attributes are acceptable. The system may also be able to operate without identity proofing if the potential harms from accepting self-asserted attributes are insignificant. In such cases, the identity proofing processes described in [SP800-63A] are not applicable to the system.

If the online service does require identity proofing, an initial IAL is selected through a simple mapping process, as follows:

- Low impact: IAL1
- Moderate impact: IAL2
- High impact: IAL3

The organization **SHALL** document whether identity proofing is required for their application and, if it is, **SHALL** select an initial IAL for each user group based on the effective impact level determination from Sec. 3.2.4.

The IAL reflects the level of assurance that an applicant holds the claimed real-life identity. The initial selection assumes that higher potential impacts of failures in the identity proofing process should be mitigated by higher assurance processes.

3.3.3.2. Selecting Initial AAL

Not all online services require authentication. Online services that offer access to public information and do not utilize subscriber accounts do not necessarily need to implement authentication mechanisms. However, authentication is needed for online services that do offer access to personal information, protected information, or subscriber accounts. In addition to the impact assessments mandated by these guidelines, when making decisions regarding the application of authentication assurance levels and authentication mechanisms, it is important that organizations consider legal, regulatory, or policy requirements that govern online services. For example, [EO13681] states “that all organizations making personal data accessible to citizens through digital applications require the use of multiple factors of authentication,” which requires a minimum selection of AAL2 for applications meeting those criteria.

If the online service requires an authenticator to be implemented, an initial AAL is selected through a simple mapping process, as follows:

- Low impact: AAL1
- Moderate impact: AAL2
- High impact: AAL3

The organization **SHALL** document whether authentication is needed for their online service and, if it is, **SHALL** select an initial AAL for each user group based on the effective impact level determination from Sec. 3.2.4.

The AAL reflects the level of assurance that the claimant is the same individual to whom the credential or authenticator was issued. The initial selection assumes that higher potential impacts of failures in the authentication process should be mitigated by higher assurance processes.

3.3.3.3. Selecting Initial FAL

Identity federation brings many benefits including a convenient user experience that avoids redundant, costly, and often time-consuming identity processes. The benefits of federation through a general-purpose IdP model or a subscriber-controlled wallet model are covered in Sec. 5 of [SP800-63C]. However, not all online services will be able to make use of federation, whether for risk-based reasons or due to legal or regulatory requirements. Consistent with [M-19-17], federal agencies that operate online services **SHOULD** implement federation as an option for user access.

If the online service implements identity federation, an initial FAL is selected through a simple mapping process, as follows:

- Low impact: FAL1
- Moderate impact: FAL2
- High impact: FAL3

The organization **SHALL** document whether federation will be used for their online service and, if it is, **SHALL** select an initial FAL for each user group based on the effective impact level determination from Sec. 3.2.4.

The FAL reflects the level of assurance in identity assertions that convey the results of authentication processes and relevant identity information to RP online services. The preliminary selection assumes that higher potential impacts of failures in federated identity architectures should be mitigated by higher assurance processes.

3.3.4. Identify Baseline Controls

The selection of the initial assurance levels for each of the applicable identity functions (i.e., IAL, AAL, and FAL) serves as the basis for the selection of the baseline digital identity controls from the guidelines in companion volumes [SP800-63A], [SP800-63B], and [SP800-63C]. As described in Sec. 3.4, the baseline controls include technology and process controls that will be assessed against additional potential impacts.

The output of this step **SHALL** include the relevant xALs and controls for each user group, as follows:

- Initial IAL and related technology and process controls from [SP800-63A]
- Initial AAL and related technology and process controls from [SP800-63B]
- Initial FAL and related technology and process controls from [SP800-63C]

3.4. Tailor and Document Assurance Levels

The second dimension of risk addressed by the Digital Identity Risk Management process focuses on risks from the identity management system. These risks inform the tailoring process and seeks to identify the risks and unintended consequences that result from the initial selection of xALs and the related technical and process controls in Sec. 3.3.4.

Tailoring provides a process to modify an initially assessed assurance level and implement compensating or supplemental controls based on ongoing detailed risk assessments. It provides a pathway for flexibility and enables organizations to achieve risk management objectives that align with their specific context, users, and threat environment. This process focuses on assessing for unintended risks and equity, privacy, and usability impacts, and specific environmental threats. It does not prioritize any specific risk area or outcomes for agencies. Making decisions that balance different types of risks to meet organizational outcomes remains the responsibility of organizations. Organizations **SHOULD** employ tailoring with the objective of aligning of digital identity controls to their specific context, users, and threat environment.

Within the tailoring step, organizations **SHALL** focus on impacts to mission delivery due to the implementation of identity management controls that result in disproportionate impact on marginalized or historically underserved populations. Organizations **SHALL** consider not only the possibility of certain intended subjects failing to access the online service, but also the burdens, frustrations, and frictions experienced as a result of the identity management controls.

As a part of the tailoring process, organizations **SHALL** review the impact assessment documentation and practice statements[2] from CSPs and IdPs that they use or intend to use. However, organizations **SHALL** also conduct their own analysis to ensure that the organization's specific mission and the communities being served by the online service are given due consideration for tailoring purposes. As a result the organization may require their chosen CSP to strengthen or provide optionality in the implementation of certain controls to address risks and unintended impacts to the organization's mission and the communities served.

To promote interoperability and consistency across organizations, third-party CSPs **SHOULD** implement their (assessed or tailored) xALs consistent with the normative guidance in this document. However, these guidelines provide flexibility to allow organizations to tailor the initial xALs and related controls to meet specific mission needs, address unique risk appetites, and provide secure and accessible online services. In doing so, CSPs **MAY** offer and organizations **MAY** utilize tailored sets of controls that differ from the normative statements in this guidance.

[2] Further information on practice statements and their contents can be found in Section 3.1 of SP800-63A.

Therefore, organizations **SHALL** establish and document an xAL tailoring process. At a minimum this process:

- **SHALL** follow a documented governance approach to allow for decision-making.
- **SHALL** document all decisions in the tailoring process, including the assessed xALs, modified xALs, and supplemental and compensating controls in the Digital Identity Acceptance Statement (see Sec. 3.4.4).
- **SHALL** justify and document all risk-based decisions or modifications to the initially assessed xALs in the Digital Identity Acceptance Statement (see Sec. 3.4.4).
- **SHOULD** establish a cross-functional capability to support subject matter analysis of xAL selection impacts in the tailoring process (e.g., subject matter experts who can speak about risks and considerations related to privacy, usability, fraud and impersonation impacts, equity, and other germane areas).
- **SHOULD** be a continuous process that incorporates real-world operational data to evaluate the impacts of selected xAL controls.

The tailoring process promotes a structured means of balancing risks and impacts in the furtherance of protecting online services, systems, and data in a manner that enables mission success while supporting equity, privacy, and usability for individuals.

3.4.1. Assess Privacy, Equity, Usability and Threat Resistance

When selecting and tailoring assurance levels for specific online services, it is critical that insights and inputs to the process extend beyond the initial impact assessment in Sec. 3.2. When transitioning from the initial assurance level selection in Sec. 3.3.4 to the final xAL selection and implementation, organizations **SHALL** conduct detailed assessments of the controls defined for the initially selected xALs to identify potential impacts in the operational environment.

At a minimum, organizations **SHALL** assess the impacts and potential unintended consequences related to the following areas:

- **Privacy** – Identify unintended consequences to the privacy of individuals that will be subject to the controls at an assessed xAL and of individuals affected by organizational or third-party practices related to the establishment, management, or federation of a digital identity. Privacy assessments **SHOULD** leverage existing Privacy Threshold Assessments (PTAs) and Privacy Impact Assessments (PIAs) as inputs to the privacy assessment process. However, as the goal of the privacy assessment is to identify privacy risks that arise from the initial assurance level selection, additional assessments and evaluations that are specific to the baseline controls for the assurance levels may be required for the underlying information system.

- **Equity** – Determine whether implementation of the initial assurance levels may create, maintain, or exacerbate inequities across communities. Equity assessments **SHALL** evaluate impacts on the communities being served by considering factors such as: proficiency with and access to technology, the availability of end devices with required technical capabilities (e.g., cameras), shared computing or device scenarios, housing status, access to internet, internet speed, family income bracket, credit score, disability status, sex, skin tone, age, native language, English fluency, and education. The intent of this assessment is to mitigate potential impacts on marginalized and historically underserved groups and limit disproportionate impacts from the requirements of the identity management functions.
- **Usability** – Determine whether implementation of the initial assurance levels will result in challenges to the end-user experience. Usability assessments **SHALL** consider usability impacts that result from the identity management controls to ensure that they do not cause undue burdens, frustrations, or frictions for the communities served and that there are pathways to provide accessibility to users of all capabilities.
- **Threat Resistance** – Determine whether the defined assurance level and related controls will address specific threats to the online service based on the operational environment, its threat actors, and known tactics, techniques, and procedures (TTPs). Threat assessments **SHALL** consider specific and known threats, threat actors, and TTPs within the implementation environment for the identity management functions. For example, certain benefits programs may be more subject to familial threats or collusion. Supplemental controls **MAY** need to be implemented to address specific threats within communities served by the online service. Conversely, agencies **MAY** tailor their assessed xAL down or modify their baseline controls if their threat assessment indicates that a reduced threat posture is appropriate based on their environment.

Organizations **SHOULD** leverage consultation and feedback to ensure that the tailoring process addresses the constraints of the entities and communities served. Organizations **MAY** establish mechanisms through which civil society organizations that work with marginalized groups can provide input on the impacts felt or likely to be felt.

Additionally, organizations **SHOULD** conduct additional business-specific assessments as appropriate to fully represent mission- and domain-specific considerations not captured here. These assessments **SHALL** be extended to any compensating or supplemental controls as defined in Sec. 3.4.2 and Sec. 3.4.3.

The outcome of this step is a set of risk assessments for privacy, equity, usability, threat resistance and other dimensions that informs the tailoring of the initial assurance levels and the selection of compensating and supplemental controls.

3.4.2. Identify Compensating Controls

A compensating control is a management, operational, or technical control employed by an organization in lieu of a normative control in the defined xALs. They are intended to address the same risks as the baseline control is intended to address to the greatest degree practicable.

Organizations **MAY** choose to implement a compensating control when they are unable to implement a baseline control or when a risk assessment indicates that a compensating control sufficiently mitigates risk in alignment with organizational risk tolerance. This control **MAY** be a modification to the normative statements as defined in these guidelines, but **MAY** also be applied elsewhere in an application, digital transaction, or service lifecycle. For example:

- A federal agency could choose to use a federal background investigation and checks, as referenced by *Personal Identity Verification* [FIPS201], to compensate for the identity evidence validation with authoritative sources requirement under these guidelines.
- An organization could choose to implement stricter auditing and transactional review processes on a payment application where verification processes using weaker forms of identity evidence were accepted due to the lack of required evidence in the end-user population.

Where compensating controls are implemented, organizations **SHALL** document the compensating control, the rationale for the deviation, comparability of the chosen alternative, and resulting residual risk (if any). CSPs and IDPs who implement compensating controls **SHALL** communicate this information to all potential RPs prior to integration to allow the RP to assess and determine the acceptability of the compensating controls for their use cases.

The process of tailoring allows agencies and service providers to make risk-based decisions regarding how they implement their xALs and related controls. It also provides a mechanism for documenting and communicating decisions through the Digital Identity Acceptance Statement described in Sec. 3.4.4.

3.4.3. Identify Supplemental Controls

Supplemental controls are those that may be added to further strengthen the baseline controls specified for the organization's selected assurance levels. Organizations **SHOULD** identify and implement supplemental controls to address specific threats in the operational environment that may not be addressed by the baseline controls. For example:

- To complete the proofing process, an organization could choose to verify an end user against additional pieces of identity evidence, beyond what is required by the assurance level, due to a high prevalence of fraudulent attempts.

- An organization could restrict users to only phishing-resistant authentication at AAL2.
- An organization could choose to implement risk-scoring analytics, coupled with re-proofing mechanisms, to confirm a user's identity when their access attempts exhibit certain risk factors.

Any supplemental controls **SHALL** be assessed for impacts based on the same factors used to tailor the organization's assurance level and **SHALL** be documented.

3.4.4. Digital Identity Acceptance Statement (DIAS)

Organizations **SHALL** develop a Digital Identity Acceptance Statement (DIAS) to document the results of the Digital Identity Risk Management process for each online service managed by the organization. A CSP/IdP **SHALL** make their DIAS and practice statements available to RPs. RPs who intend to use a particular CSP/IdP **SHALL** review the latter's DIAS and practice statements and incorporate relevant information into the organization's DIAS for each online service.

The DIAS **SHALL** include, at a minimum:

- Initial impact assessment results,
- Initially assessed xALs,
- Tailored xAL and rationale, if the tailored xAL differs from the initially assessed xAL,
- All compensating controls with their comparability or residual risk, and
- All supplemental controls.

Federal agencies **SHOULD** include this information in the information system authorization package described in [NISTRMF].

3.5. Continuously Evaluate and Improve

Threat actors adapt; user capabilities, expectations, and needs shift; seasonal surges occur; and missions evolve. As such, risk assessments and identity solutions must be continuously improved. In addition to keeping pace with the threat and technology environment, continuous improvement is a critical tool for illustrating programmatic gaps that — if unaddressed — may hinder the implementation of identity management systems in a manner that balances risk management objectives. For instance, an organization may determine that a portion of the target population intended to be served by the online service does not have access to affordable high-speed internet services needed to support remote identity proofing. The organization could address this gap with a program that implements local proofing capabilities within the community or by offering appointments with proofing agents who will meet the individual at an

address that is more accessible and convenient, such as their local community center, closest post office, an affiliated business partner facility, or the individual's home.

To address the shifting environment in which they operate and more rapidly address service capability gaps, organizations **SHALL** implement a continuous evaluation and improvement program that leverages input from end users who have interacted with the identity management system as well as performance metrics for the online service. This program **SHALL** be documented, including the metrics that are collected, the sources of data required to enable performance evaluation, and the processes in place for taking timely actions based on the continuous improvement process. This program and its effectiveness **SHOULD** be assessed on a defined basis to ensure that outcomes are being achieved and that programs are addressing issues in a timely manner.

Additionally, organizations **SHALL** monitor the evolving threat landscape to stay informed of the latest threats and fraud tactics. Organizations **SHALL** regularly assess the effectiveness of current security measures and fraud detection capabilities against the latest threats and fraud tactics.

3.5.1. Evaluation Inputs

To fully understand the performance of their identity system, organizations will need to identify critical inputs to their evaluation process. At a minimum these **SHALL** include:

- Integrated CSP, IdP, and authenticator functions as well as validation, verification, and fraud management systems as appropriate.
- Customer feedback mechanisms such as complaint processes, help-desk statistics, and other user feedback (e.g., surveys, interviews, or focus groups)
- Threat analysis, threat reporting, and threat intelligence feeds as available to the organization.
- Fraud trends, fraud investigation results, and fraud metrics as available to the organization.
- The results of ongoing equity assessments, privacy assessments, and usability assessments.

Organizations **SHALL** document their metrics, reporting requirements, and data inputs for any CSP, IdP, or other integrated identity services to ensure that expectations are appropriately communicated to partners and vendors.

3.5.2. Performance Metrics

The exact metrics available to organizations will vary based on the technologies, architectures, and deployment patterns they follow. Additionally, what is available and what is useful may vary over time. Therefore, these guidelines do not attempt

to define a comprehensive set of metrics for all scenarios. Table 4 provides a set of recommended metrics that organizations **SHOULD** capture as part of their continuous evaluation program. However, organizations are not constrained by this table and **SHOULD** implement metrics that are not defined here based on their specific systems, technology, and program needs. In Table 4, all references to unique users include both legitimate users and imposters.

Table 4. Performance Metrics

Title	Description	Type
Pass Rate (Overall)	Percentage of unique users who successfully proof.	Proofing
Pass Rate (Per Proofing Type)	Percentage of unique users who successfully proof for each offered type (i.e., Remote Unattended, Remote Attended, Onsite Attended, Onsite Unattended).	Proofing
Fail Rate (Overall)	Percentage of unique users who start the identity proofing process but are unable to successfully complete all the steps.	Proofing
Estimated Adjusted Fail Rate	Percentage adjusted to account for digital transactions that are terminated based on suspected fraud.	Proofing
Fail Rate (Per Proofing Type)	Percentage of unique users who do not complete proofing due to a process failure for each offered type (i.e., Remote Unattended, Remote Attended, Onsite Attended, Onsite Unattended)	Proofing
Abandonment Rate (Overall)	Percentage of unique users who start the identity proofing process, but do not complete it without failing a process.	Proofing
Abandonment Rate (Per Proofing Type)	Percentage of unique users who start a specific type of identity proofing process, but do not complete it without failing a process.	Proofing
Failure Rates (Per Proofing Process Step)	Percentage of unique users who are unsuccessful at completing each identity proofing step in a CSP process.	Proofing
Completion (Times Per Proofing Type)	Average time that it takes a user to complete each defined proofing type offered as part of an identity service.	Proofing

Authenticator Type Usage	Percentage of subscribers who have an active authenticator by each type available.	Authentication
Authentication Failures	Percentage of authentication events that fail (not to include attempts that are successful after re-entry of an authenticator output).	Authentication
Account Recovery Attempts	The number of account or authenticator recovery processes initiated by subscribers	Authentication
Confirmed Fraud	Percentage of digital transactions that are confirmed to be fraudulent through investigation or self-reporting.	Fraud
Suspected Fraud	Percentage of digital transactions that are suspected of being fraudulent.	Fraud
Reported Fraud	Percentage of digital transactions reported to be fraudulent by users.	Fraud
Fraud (Per Proofing Type	Number of digital transactions that are suspected, confirmed, and reported by each available type of proofing.	Fraud
Fraud (Per Authentication Type)	Number of digital transactions suspected, confirmed, and reported by each available type of authentication	Fraud
Help Desk Calls	Number of calls received by the CSP or identity service.	Customer Support
Help Desk Calls (Per Type)	Number of calls received related to each offered service (e.g., proofing failures, authenticator resets, complaints)	Customer Support
Help Desk Resolution Times	Average length of time it takes to resolve a complaint or help desk ticket.	Customer Support
Customer Satisfaction Surveys	The results of customer feedback surveys conducted by CSPs, RP, or both.	User Experience
Redress requests	The number of redress requests received related to the identity management system.	User Experience
Redress resolution times	The average time it takes to resolve redress requests related to the identity management system.	User Experience

The data used to generate continuous evaluation metrics may not always reside with the identity program or the organizational entity responsible for identity management systems. The intent of these metrics is not to establish redundant processes but to integrate with existing data sources whenever possible to collect information that is critical to identity program evaluation. For example, customer service representative (CSR) teams may already have substantial information on customer requests, complaints,

or concerns. Identity management systems would be expected to coordinate with these teams to acquire the information needed to discern identity management system-related complaints or issues.

3.5.3. Measurement in Support of Equity Assessments and Outcomes

A primary purpose of continuous improvement is to improve equity and accessibility outcomes for different user populations. As a result, the metrics collected by organizations **SHOULD** be further evaluated to provide insights into the performance of their identity management systems for their supported communities and demographics. Where possible, these efforts **SHOULD** avoid the collection of additional personal information and instead use informed analysis of proxy data to help provide indicators of potential disparities. This can include comparing and filtering the metrics to identify deviations in performance across different user populations based on other available data such as zip code, geographic region, age, or sex.

Organizations are encouraged to consult the OMB Report *A Vision for Equitable Data: Recommendations from the Equitable Data Working Group* [EO13985-vision] for guidance on incorporating performance metrics into equity assessments across demographic groups and generating disaggregated statistical estimates to assess equitable performance outcomes.

3.6. Redress

An important part of designing services that support a wide range of populations is the inclusion of processes to adjudicate issues and provide redress[3] as warranted. Service failures, disputes, and other issues tend to arise as part of normal operations, and their impact can vary broadly, from minor inconveniences to major disruptions or damage. Barriers to access, as well as cybersecurity incidents and data breaches, have real-world consequences for affected individuals. Furthermore, the same issue experienced by one person or community as an inconvenience can have a disproportionately damaging impacts on other individuals and communities, particularly those that are currently experiencing other harms or barriers. Left unchecked, these issues can result in harms that exacerbate existing inequities and allow systemic cycles of exclusion to continue.

To enable equitable access to critical services while deterring identity-related fraud and cybersecurity threats, it is essential for organizations to plan for potential issues and to design redress approaches that aim to be fair, transparent, easy for legitimate claimants to navigate, and resistant to exploitation attempts.

Understanding when and how harms might be occurring is a critical first step for organizations to take informed action. Continuous evaluation and improvement programs can play a key role in identifying instances and patterns of potential harm.

[3] Redress generally refers to a remedy that is made after harm occurs.

Moreover, there may be business processes in place outside of those established to support identity management that can be leveraged as part of a comprehensive approach to issue adjudication and redress. Beyond these activities, additional practices can be implemented to ensure that users of identity management systems are able to voice their concerns and have a path to redress. Requirements for these practices include:

- RPs and CSPs **SHALL** enable people to convey grievances and seek redress through an issue handling process that is documented, accessible, trackable, and usable by all people, and whose instructions are easy to find on a public-facing website.
- RPs and CSPs **SHALL** institute a governance model, including documented roles and responsibilities, for implementing this issue handling process.
- The issue handling process **SHALL** be implemented as a dedicated function that includes:
  - Procedures for the impartial review of evidence pertinent to issues;
  - Procedures for requesting and collecting additional evidence that informs the issues; and
  - Procedures to expeditiously resolve issues and determine corrective action.
- RPs and CSPs **SHALL** make human support personnel available to intervene and override issue adjudication outputs generated by algorithmic support mechanisms, such as chatbots.
- RPs and CSPs **SHALL** educate support personnel on issue handling procedures for the digital identity management system, the avenues for redress, and the alternatives available to gain access to services.
- RPs and CSPs **SHALL** implement a process for personnel and technologies that provides support functions to report major barriers that end users face and commonly expressed grievances. This process **SHALL** enable tracing (e.g., who/what is reported) and tracking (e.g. progress/state of action taken).
- RPs and CSPs **SHALL** incorporate findings derived from the issue handling process into continuous evaluation and improvement activities.

Organizations are encouraged to consider these and other emerging redress practices. Prior to adopting any new redress practice, including supporting technology, organizations **SHOULD** test the practice with target populations to avoid the introduction of unintended consequences, particularly those that may counteract or contradict the goals associated with redress.

3.7. Cybersecurity, Fraud, and Identity Program Integrity

Identity solutions should not operate in a vacuum. Close coordination of identity functions with teams that are responsible for cybersecurity, privacy, threat intelligence, fraud detection, and program integrity can enable a more complete protection of business capabilities, while constantly improving identity solution capabilities. For example, payment fraud data collected by program integrity teams could provide indicators of compromised subscriber accounts and potential weaknesses in identity proofing implementations. Similarly, threat intelligence teams may learn of new TTPs that Could impact identity proofing, authentication, and federation processes. Organizations **SHALL** establish consistent mechanisms for the exchange of information between critical internal security and fraud stakeholders. Organizations **SHOULD** do the same for external stakeholders and identity services that are part of the protection plan for their online services.

When supporting identity service providers (e.g., CSPs) are external to an organization, the exchange of data related to security, fraud, and other RP functions may be complicated by regulation or policy. However, establishing the necessary mechanisms and guidelines to enable effective information-sharing **SHOULD** be considered in contractual and legal mechanisms. All data collected, transmitted, or shared **SHALL** be minimized and subject to a detailed privacy and legal assessment by the generating entity.

This section is meant to address coordination and integration with various organizational functional teams to achieve better outcomes for the identity functions. Ideally, such coordination is performed throughout the risk management process and operations lifecycle. Companion volumes [SP800-63A], [SP800-63B], and [SP800-63C] provide specific fraud mitigation requirements related to each of the identity functions.

3.8. Artificial Intelligence (AI) and Machine Learning (ML) in Identity Systems

Identity solutions have used and will continue to use AI and ML for multiple purposes, such as improving the performance of biometric matching systems, documenting authentication, detecting fraud, and even assisting users (e.g., chatbots). The potential applications of AI/ML are extensive. They also introduce distinct risks and potential issues, including disparate outcomes, biased outputs, and the exacerbation of existing inequities and access issues.

The following requirements apply to all uses of AI and ML regardless of how they are used in identity systems:

- All uses of AI and ML **SHALL** be documented and communicated to organizations that rely on these systems. The use of integrated technologies that leverage AI and ML by CSPs, IdPs, or verifiers **SHALL** be disclosed to all RPs that make access decisions based on information from these systems.

- All organizations that use AI and ML **SHALL** provide information to any entities that use their technology on the methods and techniques used for training their models, a description of the data sets used in training, information on the frequency of model updates, and the results of all testing completed on their algorithms.
- All organizations that use AI and ML systems or rely on services that use these systems **SHALL** implement *NIST AI Risk Management Framework* ([NISTAIRMF]) to evaluate the risks that may be introduced by the use of AI and ML. 4. All organizations that use AI and ML **SHALL** consult [SP1270], *Towards a Standard for Managing Bias in Artificial Intelligence*.

NIST continues to advance efforts to promote safe and trustworthy AI implementations through a number of venues. In particular, the U.S. AI Safety Institute, housed at NIST [US-AI-Safety-Inst], is creating a portfolio of safety-focused resources, guidance, and tools that can improve how organizations assess, deploy, and manage their AI systems. Organizations are encouraged to follow the U.S. AI Safety Institute's efforts and make use of their resources.

References

This section is informative.

[A-130] Office of Management and Budget (2016) Managing Information as a Strategic Resource. (The White House, Washington, DC), OMB Circular A-130, July 28, 2016. Available at https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/OMB/circulars/a130/a130revised.pdf

[EO13985] Biden J (2021) Advancing Racial Equity and Support for Underserved Communities Through the Federal Government. (The White House, Washington, DC), Executive Order 13985, January 25, 2021. https://www.federalregister.gov/documents/2021/01/25/2021-01753/advancing-racial-equity-and-support-for-underserved-communities-through-the-federal-government

[EO13985-vision] Office of Management and Budget (2022) A Vision for Equitable Data: Recommendations from the Equitable Data Working Group. (The White House, Washington, DC), OMB Report Pursuant to Executive Order 13985, April 22, 2022. https://www.whitehouse.gov/wp-content/uploads/2022/04/eo13985-vision-for-equitable-data.pdf

[EO14012] Biden J (2021) Restoring Faith in Our Legal Immigration Systems and Strengthening Integration and Inclusion Efforts for New Americans. (The White House, Washington, DC), Executive Order 14012, February 02, 2021. https://www.whitehouse.gov/briefing-room/presidential-actions/2021/02/02/executive-order-restoring-faith-in-our-legal-immigration-systems-and-strengthening-integration-and-inclusion-efforts-for-new-americans/

[EO14058] Biden J (2021) Transforming Federal Customer Experience and Service Delivery to Rebuild Trust in Government. (The White House, Washington, DC), Executive Order 14058, December 13, 2021. https://www.federalregister.gov/documents/2021/12/16/2021-27380/transforming-federal-customer-experience-and-service-delivery-to-rebuild-trust-in-government

[EO14091] Biden J (2023) Further Advancing Racial Equity and Support for Underserved Communities Through the Federal Government. (The White House, Washington, DC), Executive Order 14091, February 16, 2023. https://www.whitehouse.gov/briefing-room/presidential-actions/2023/02/16/executive-order-on-further-advancing-racial-equity-and-support-for-underserved-communities-through-the-federal-government/

[FIPS199] National Institute of Standards and Technology (2004) Standards for Security Categorization of Federal Information and Information Systems. (U.S. Department of Commerce, Washington, DC), Federal Information Processing Standards Publication (FIPS) 199. https://doi.org/10.6028/NIST.FIPS.199

[FIPS201] National Institute of Standards and Technology (2022) Personal Identity Verification (PIV) of Federal Employees and Contractors. (U.S. Department of Commerce, Washington, DC), Federal Information Processing Standards Publication (FIPS) 201-3. https://doi.org/10.6028/NIST.FIPS.201-3

[FISMA] Federal Information Security Modernization Act of 2014, Pub. L. 113-283, 128 Stat. 3073. https://www.govinfo.gov/app/details/PLAW-113publ283

[ISO/IEC9241-11] International Standards Organization (2018) *ISO/IEC 9241-11 Ergonomics of human-system interaction – Part 11: Usability: Definitions and concepts* (ISO, Geneva, Switzerland). Available at https://www.iso.org/standard/63500.html

[M-03-22] Office of Management and Budget (2003) OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002. (The White House, Washington, DC), OMB Memorandum M-03-22, September 26, 2003. Available at https://georgewbush-whitehouse.archives.gov/omb/memoranda/m03-22.html

[M-19-17] Office of Management and Budget (2019) Enabling Mission Delivery through Improved Identity, Credential, and Access Management. (The White House, Washington, DC), OMB Memorandum M-19-17, May 21, 2019. Available at https://www.whitehouse.gov/wp-content/uploads/2019/05/M-19-17.pdf

[NISTAIRMF] Tabassi E (2023) Artificial Intelligence Risk Management Framework (AI RMF 1.0). (National Institute of Standards and Technology (U.S.), Gaithersburg, MD), NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

[NISTIR8062] Brooks SW, Garcia ME, Lefkovitz NB, Lightman S, Nadeau EM (2017) An Introduction to Privacy Engineering and Risk Management in Federal Systems. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) 8062. https://doi.org/10.6028/NIST.IR.8062

[NISTRMF] Joint Task Force (2018) Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-37, Rev. 2. https://doi.org/10.6028/NIST.SP.800-37r2

[NISTPF] National Institute of Standards and Technology (2020) NIST Privacy Framework: A Tool for Improving Privacy Through Enterprise Risk Management, Version 1.0. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Cybersecurity White Paper (CSWP) NIST CSWP 10. https://doi.org/10.6028/NIST.CSWP.10

[PrivacyAct] Privacy Act of 1974, Pub. L. 93-579, 5 U.S.C. § 552a, 88 Stat. 1896 (1974). https://www.govinfo.gov/content/pkg/USCODE-2018-title5/pdf/USCODE-2018-title5-partI-chap5-subchapII-sec552a.pdf

[RFC5246] Rescorla E, Dierks T (2008) The Transport Layer Security (TLS) Protocol Version 1.2. (Internet Engineering Task Force (IETF)), IETF Request for Comments (RFC) 5246. https://doi.org/10.17487/RFC5246

[RFC5280] Cooper D, Santesson S, Farrell S, Boeyen S, Housley R, Polk W (2008) Internet X.509 Public Key Infrastructure Certification and Certificate Revocation List (CRL) Profile. (Internet Engineering Task Force (IETF)), IETF Request for Comments (RFC) 5280. https://doi.org/10.17487/RFC5280

[RFC9325] Sheffer Y, Saint-Andre P, Fossati T (2022) Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS). (Internet Engineering Task Force (IETF)), IETF Request for Comments (RFC) 9325. https://doi.org/10.17487/RFC9325

[Section508] Section 508 of the Rehabilitation Act of 1973 (2011), 29 U.S.C. § 794(d). https://www.govinfo.gov/content/pkg/USCODE-2011-title29/html/USCODE-2011-title29-chap16-subchapV-sec794d.htm

[SP800-30] Blank R, Gallagher P (2012) Guide for Conducting Risk Assessments. (National Institute of Standards and Technology, Gaithersburg, MD) NIST Special Publication (SP) 800-30 Revision 1. https://doi.org/10.6028/NIST.SP.800-30r1

[SP800-52] McKay K, Cooper D (2019) Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations. (National Institute of Standards and Technology), NIST Special Publication (SP) 800-52 Rev. 2. https://doi.org/10.6028/NIST.SP.800-52r2

[SP800-53] Joint Task Force (2020) Security and Privacy Controls for Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-53 Rev. 5, Includes updates as of December 10, 2020. https://doi.org/10.6028/NIST.SP.800-53r5

[SP800-57Part1] Barker EB (2020) Recommendation for Key Management: Part 1 – General. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-57 Part 1, Rev. 5. https://doi.org/10.6028/NIST.SP.800-57pt1r5

[SP800-63A] Temoshok D, Abruzzi C, Choong YY, Fenton JL, Galluzzo R, LaSalle C, Lefkovitz N, Regenscheid A (2024) Digital Identity Guidelines: Identity Proofing and Enrollment. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-63A-4 2pd. https://doi.org/10.6028/NIST.SP.800-63a-4.2pd

[SP800-63B] Temoshok D, Fenton JL, Choong YY, Lefkovitz N, Regenscheid A, Galluzzo R, Richer JP (2024) Digital Identity Guidelines: Authentication and Authenticator Management. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-63B-4 ipd. https://doi.org/10.6028/NIST.SP.800-63b-4.2pd

[SP800-63C] Temoshok D, Richer JP, Choong YY, Fenton JL, Lefkovitz N, Regenscheid A, Galluzzo R (2024) Digital Identity Guidelines: Federation and Assertions. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-63C-4 2pd. https://doi.org/10.6028/NIST.SP.800-63c-4.2pd

[SP800-122] McCallister E, Grance T, Scarfone KA (2010) Guide to Protecting the Confidentiality of Personally Identifiable Information (PII). (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-122. https://doi.org/10.6028/NIST.SP.800-122

[SP1270] Schwartz R, Vassilev A, Greene K, Perine L, Burt A, Hall P (2022) Towards a standard for identifying and managing bias in artificial intelligence. (National Institute of Standards and Technology (U.S.), Gaithersburg, MD), NIST SP 1270. https://doi.org/10.6028/NIST.SP.1270

[US-AI-Safety-Inst] U.S. Artificial Intelligence Safety Institute (2023) NIST. Available at https://www.nist.gov/aisi

Appendix A. List of Symbols, Abbreviations, and Acronyms

1:1 Comparison
One-to-One Comparison

ABAC
Attribute-Based Access Control

AAL
Authentication Assurance Level

CAPTCHA
Completely Automated Public Turing test to tell Computers and Humans Apart

CSP
Credential Service Provider

CSRF
Cross-Site Request Forgery

XSS
Cross-Site Scripting

DNS
Domain Name System

FACT Act
Fair and Accurate Credit Transaction Act of 2003

FAL
Federation Assurance Level

FEDRAMP
Federal Risk and Authorization Management Program

FMR
False Match Rate

FNMR
False Non-Match Rate

IAL
Identity Assurance Level

IdP
Identity Provider

JOSE
JSON Object Signing and Encryption

JWT
JSON Web Token

KBA
Knowledge-Based Authentication

KBV
Knowledge-Based Verification

KDC
Key Distribution Center

MAC
Message Authentication Code

MFA
Multi-Factor Authentication

NARA
National Archives and Records Administration

OTP
One-Time Password

PAD
Presentation Attack Detection

PIA
Privacy Impact Assessment

PII
Personally Identifiable Information

PIN
Personal Identification Number

PKI
Public Key Infrastructure

PSTN
Public Switched Telephone Network

RMF
Risk Management Framework

RP
Relying Party

SA&A
Security Authorization & Accreditation

SAML
Security Assertion Markup Language

SAOP
Senior Agency Official for Privacy

SSL
Secure Sockets Layer

SSO
Single Sign-On

SMS
Short Message Service

SORN
System of Records Notice

TEE
Trusted Execution Environment

TLS
Transport Layer Security

TPM
Trusted Platform Module

TTP
Tactics, Techniques, and Procedures

VOIP
Voice-Over-IP

XSS
Cross-Site Scripting

Appendix B. Glossary

This section is informative.

A wide variety of terms are used in the realm of digital identity. While many definitions are consistent with earlier versions of SP 800-63, some have changed in this revision. Many of these terms lack a single, consistent definition, warranting careful attention to how the terms are defined here.

account linking
The association of multiple *federated identifiers* with a single *RP subscriber account*, or the management of those associations.

account recovery
The ability to regain ownership of a *subscriber account* and its associated information and privileges.

account resolution
The association of an *RP subscriber account* with information already held by the *RP* prior to the *federation transaction* and outside of a *trust agreement*.

activation
The process of inputting an *activation factor* into a *multi-factor authenticator* to enable its use for *authentication*.

activation factor
An additional *authentication factor* that is used to enable successful *authentication* with a *multi-factor authenticator*.

activation secret
A *password* that is used locally as an *activation factor* for a *multi-factor authenticator*.

allowlist
A documented list of specific elements that are allowed, per policy decision. In *federation* contexts, this is most commonly used to refer to the list of *RPs* allowed to connect to an *IdP* without subscriber intervention. This concept has historically been known as a *whitelist*.

applicant
A *subject* undergoing the processes of *identity proofing* and *enrollment*.

applicant reference
A representative of the *applicant* who can vouch for the identity of the applicant, specific *attributes* related to the applicant, or conditions relative to the context of the individual (e.g., emergency status, homelessness).

approved cryptography
An encryption algorithm, *hash function*, random bit generator, or similar technique that is *Federal Information Processing Standard* (FIPS)-approved or NIST-recommended. Approved algorithms and techniques are either specified or adopted in a FIPS or NIST recommendation.

assertion
A statement from an *IdP* to an *RP* that contains information about an authentication event for a subscriber. Assertions can also contain identity *attributes* for the subscriber.

assertion reference
A data object, created in conjunction with an *assertion*, that is used by the *RP* to retrieve an assertion over an *authenticated* protected channel.

assertion presentation
The method by which an *assertion* is transmitted to the *RP*.

asymmetric keys
Two related keys, comprised of a *public key* and a *private key*, that are used to perform complementary operations such as encryption and decryption or signature *verification* and generation.

attestation
Information conveyed to the *CSP*, generally at the time that an *authenticator* is bound, describing the characteristics of a connected authenticator or the *endpoint* involved in an authentication operation.

attribute
A quality or characteristic ascribed to someone or something. An identity attribute is an attribute about the identity of a subscriber.

attribute bundle
A package of *attribute values* and *derived attribute values* from a *CSP*. The package has necessary cryptographic protection to allow *validation* of the bundle independent from interaction with the CSP or *IdP*. Attribute bundles are often used with subscriber-controlled wallets.

attribute provider
The provider of an *identity API* that provides access to a subscriber's attributes without necessarily asserting that the subscriber is present to the *RP*.

attribute validation
The process or act of confirming that a set of attributes are accurate and associated with a real-life identity. See *validation*.

attribute value
A complete statement that asserts an identity attribute of a subscriber, independent of format. For example, for the *attribute* "birthday," a value could be "12/1/1980" or "December 1, 1980."

audience restriction
The restriction of a message to a specific target audience to prevent a receiver from unknowingly *processing* a message intended for another recipient. In *federation protocols*, *assertions* are audience *restricted* to specific *RPs* to prevent an RP from accepting an assertion generated for a different RP.

authenticate
See *authentication*.

authenticated protected channel
An encrypted communication channel that uses *approved cryptography* where the connection initiator (client) has authenticated the recipient (server). Authenticated protected channels are encrypted to provide confidentiality and protection against active intermediaries and are frequently used in the user *authentication* process. *Transport Layer Security* (TLS) and Datagram Transport Layer Security (DTLS) [RFC9325] are examples of authenticated protected channels in which the certificate presented by the recipient is verified by the initiator. Unless otherwise specified, authenticated protected channels do not require the server to authenticate the client. Authentication of the server is often accomplished through a certificate chain that leads to a trusted root rather than individually with each server.

authenticated session
See *protected session*.

authentication
The process by which a *claimant* proves possession and control of one or more *authenticators* bound to a *subscriber account* to demonstrate that they are the subscriber associated with that account.

Authentication Assurance Level (AAL)
A category that describes the strength of the authentication process.

authentication factor
The three types of authentication factors are *something you know*, *something you have*, and *something you are*. Every *authenticator* has one or more authentication factors.

authentication intent
The process of confirming the *claimant's* intent to *authenticate* or reauthenticate by requiring user intervention in the authentication flow. Some *authenticators* (e.g., OTPs) establish authentication intent as part of their operation. Others require a specific step, such as pressing a button, to establish intent. Authentication intent is a countermeasure against use by malware at the *endpoint* as a proxy for authenticating an attacker without the subscriber's knowledge.

authentication protocol
A defined sequence of messages between a *claimant* and a *verifier* that demonstrates that the claimant has possession and control of one or more valid *authenticators* to establish their identity, and, optionally, demonstrates that the claimant is communicating with the intended verifier.

authentication secret
A generic term for any secret value that an attacker could use to impersonate the subscriber in an *authentication protocol*.

These are further divided into *short-term authentication secrets*, which are only useful to an attacker for a limited period of time, and *long-term authentication secrets*, which allow an attacker to impersonate the subscriber until they are manually reset. The *authenticator* secret is the canonical example of a long-term authentication secret, while the *authenticator output* — if it is different from the *authenticator secret* — is usually a short-term authentication secret.

authenticator
Something that the subscriber possesses and controls (e.g., a *cryptographic module* or *password*) and that is used to *authenticate* a *claimant's* identity. See *authenticator type* and *multi-factor authenticator*.

authenticator binding
The establishment of an association between a specific *authenticator* and a *subscriber account* that allows the *authenticator* to be used to *authenticate* for that subscriber account, possibly in conjunction with other authenticators.

authenticator output
The output value generated by an *authenticator*. The ability to generate valid authenticator outputs on demand proves that the *claimant* possesses and controls the authenticator. Protocol messages sent to the *verifier* depend on the authenticator output, but they may or may not explicitly contain it.

authenticator secret
The secret value contained within an *authenticator*.

authenticator type
A category of *authenticators* with common characteristics, such as the types of *authentication factors* they provide and the mechanisms by which they operate.

authenticity
The property that data originated from its purported source.

authoritative source
An entity that has access to or verified copies of accurate information from an *issuing source* such that a *CSP* has high confidence that the source can confirm the validity of the identity attributes or evidence supplied by an *applicant* during *identity proofing*. An issuing source may also be an authoritative source. Often, authoritative sources are determined by a policy decision of the agency or CSP before they can be used in the identity proofing *validation* phase.

authorize
A decision to grant access, typically automated by evaluating a *subject*'s *attributes*.

authorized party
In *federation*, the organization, person, or entity that is responsible for making decisions regarding the release of information within the *federation transaction*, most notably subscriber *attributes*. This is often the subscriber (when runtime decisions are used) or the party operating the *IdP* (when *allowlists* are used).

back-channel communication
Communication between two systems that relies on a direct connection without using redirects through an intermediary such as a browser.

bearer assertion
An *assertion* that can be presented on its own as proof of the identity of the presenter.

biometric reference
One or more stored *biometric samples*, templates, or models attributed to an individual and used as the object of biometric comparison in a database, such as a facial image stored digitally on a passport, fingerprint minutiae template on a National ID card or Gaussian Mixture Model for speaker recognition.

biometric sample
An analog or digital representation of biometric characteristics prior to biometric feature extraction, such as a record that contains a fingerprint image.

biometrics
Automated recognition of individuals based on their biological or behavioral characteristics. Biological characteristics include but are not limited to fingerprints, palm prints, facial features, iris and retina patterns, voiceprints, and vein patterns. Behavioral characteristics include but are not limited to keystrokes, angle of holding a smart phone, screen pressure, typing speed, mouse or mobile phone movements, and gyroscope position.

blocklist
A documented list of specific elements that are blocked, per policy decision. This concept has historically been known as a *blacklist*.

challenge-response protocol
An *authentication protocol* in which the *verifier* sends the *claimant* a challenge (e.g., a random value or *nonce*) that the claimant combines with a secret (e.g., by hashing the challenge and a *shared secret* together or by applying a *private-key* operation to the challenge) to generate a response that is sent to the verifier. The verifier can independently verify the response generated by the claimant (e.g., by re-computing the hash of the challenge and the shared secret and comparing to the response or performing a public-key operation on the response) and establish that the claimant possesses and controls the secret.

claimant
A *subject* whose identity is to be verified using one or more *authentication protocols*.

claimed address
The physical location asserted by a *subject* where they can be reached. It includes the individual's residential street address and may also include their mailing address.

claimed identity
An *applicant's* declaration of unvalidated and unverified personal *attributes*.

compensating controls
Alternative *controls* to the normative controls for the assessed and selected xALs of an organization based on that organization's mission, risk tolerance, business processes, and *risk assessments* and considerations for the privacy, *usability*, and *equity* of the populations served by the *online service*.

controls
Policies, procedures, guidelines, practices, or organizational structures that manage security, privacy, and other risks. See *supplemental controls* and *compensating controls*

core attributes
The set of identity *attributes* that the *CSP* has determined and documented to be required for *identity proofing*.

credential
An object or data structure that authoritatively binds an identity — via an *identifier* — and (optionally) additional *attributes*, to at least one *authenticator* possessed and controlled by a subscriber.

A credential is issued, stored, and maintained by the *CSP*. Copies of information from the credential can be possessed by the subscriber, typically in the form of one or more digital certificates that are often contained in an authenticator along with their associated *private keys*.

credential service provider (CSP)
A trusted entity whose functions include *identity proofing applicants* to the identity service and registering *authenticators* to *subscriber accounts*. A CSP may be an independent third party.

credible source
An entity that can provide or validate the accuracy of *identity evidence* and *attribute* information. A credible source has access to attribute information that was validated through an *identity proofing* process or that can be traced to an *authoritative source*, or it maintains identity attribute information obtained from multiple sources that is checked for data correlation for accuracy, consistency, and currency.

cross-site request forgery (CSRF)
An attack in which a subscriber who is currently *authenticated* to an *RP* and connected through a secure session browses an attacker's website, causing the subscriber to unknowingly invoke unwanted actions at the RP.

For example, if a bank website is vulnerable to a CSRF attack, it may be possible for a subscriber to unintentionally *authorize* a large money transfer by clicking on a malicious link in an email while a connection to the bank is open in another browser window.

cross-site scripting (XSS)
A vulnerability that allows attackers to inject malicious code into an otherwise benign website. These scripts acquire the permissions of scripts generated by the target website to compromise the confidentiality and integrity of data transfers between the website and clients. Websites are vulnerable if they display user-supplied data from requests or forms without sanitizing the data so that it is not executable.

cryptographic authenticator
An *authenticator* that proves possession of an *authentication secret* through direct communication with a *verifier* through a cryptographic *authentication protocol*.

cryptographic key
A value used to control cryptographic operations, such as decryption, encryption, signature generation, or signature *verification*. For the purposes of these guidelines, key requirements shall meet the minimum requirements stated in Table 2 of [SP800-57Part1]. See *asymmetric keys* or *symmetric keys*.

cryptographic module
A set of hardware, software, or firmware that implements approved security functions including cryptographic algorithms and key generation.

data integrity
The property that data has not been altered by an unauthorized entity.

derived attribute value
A statement that asserts a limited identity *attribute* of a subscriber without containing the attribute value from which it is derived, independent of format. For example, instead of requesting the attribute "birthday," a derived value could be "older than 18". Instead of requesting the attribute for "physical address," a derived value could be "currently residing in this district." Previous versions of these guidelines referred to this construct as an "attribute reference."

digital authentication
The process of establishing confidence in user identities that are digitally presented to a system. In previous editions of SP 800-63, this was referred to as electronic authentication.

digital identity
An *attribute* or set of attributes that uniquely describes a *subject* within a given context.

Digital Identity Acceptance Statement (DIAS)
Documents the results of the *digital identity risk management* process. This includes the impact assessment, initial assurance level selection, and *tailoring* process.

digital signature
An *asymmetric key* operation in which the *private key* is used to digitally sign data and the *public key* is used to verify the signature. Digital signatures provide *authenticity* protection, integrity protection, and *non-repudiation* support but not confidentiality or *replay attack* protection.

digital transaction
A discrete digital event between a user and a system that supports a business or programmatic purpose.

disassociability
Enabling the *processing* of PII or events without association to individuals or devices beyond the operational requirements of the system. [NISTIR8062]

electronic authentication (e-authentication)
See *digital authentication*.

endpoint
Any device that is used to access a *digital identity* on a *network*, such as laptops, desktops, mobile phones, tablets, servers, Internet of Things devices, and virtual environments.

enrollment
The process through which a *CSP*/*IdP* provides a successfully identity-proofed *applicant* with a *subscriber account* and binds *authenticators* to grant persistent access.

entropy
The amount of uncertainty that an attacker faces to determine the value of a secret. Entropy is usually stated in bits. A value with *n* bits of entropy has the same degree of uncertainty as a uniformly distributed *n*-bit random value.

equity
The consistent and systematic fair, just, and impartial treatment of all individuals, including individuals who belong to underserved communities that have been denied such treatment, such as Black, Latino, and Indigenous and Native American persons, Asian Americans and Pacific Islanders, and other persons of color; members of religious minorities; lesbian, gay, bisexual, transgender, and queer (LGBTQ+) persons; persons with disabilities; persons who live in rural areas; and persons otherwise adversely affected by persistent poverty or inequality. [EO13985]

factor
See *authentication factor*

Federal Information Processing Standard (FIPS)
Under the Information Technology Management Reform Act (Public Law 104-106), the Secretary of Commerce approves the standards and guidelines that the National Institute of Standards and Technology (NIST) develops for federal computer systems. NIST issues these standards and guidelines as Federal Information Processing Standards (FIPS) for government-wide use. NIST develops FIPS when there are compelling federal government requirements, such as for security and interoperability, and there are no acceptable industry standards or solutions. See background information for more details.

FIPS documents are available online on the FIPS home page: https://www.nist.gov/itl/fips.cfm

federated identifier
The combination of a *subject identifier* within an *assertion* and an *identifier* for the *IdP* that issued that assertion. When combined, these pieces of information uniquely identify the *subscriber* in the context of a *federation transaction*.

federation
A process that allows for the conveyance of identity and authentication information across a set of *networked* systems.

Federation Assurance Level (FAL)
A category that describes the process used in a *federation transaction* to communicate authentication events and subscriber *attributes* to an *RP*.

federation protocol
A technical protocol that is used in a *federation transaction* between *networked* systems.

federation proxy
A component that acts as a logical *RP* to a set of *IdPs* and a logical IdP to a set of RPs, bridging the two systems with a single component. These are sometimes referred to as "brokers."

federation transaction
A specific instance of *processing* an authentication using a *federation* process for a specific *subscriber* by conveying an *assertion* from an *IdP* to an *RP*.

front-channel communication
Communication between two systems that relies on passing messages through an intermediary, such as using redirects through the subscriber's browser.

hash function
A function that maps a bit string of arbitrary length to a fixed-length bit string. Approved hash functions satisfy the following properties:

1. One-way — It is computationally infeasible to find any input that maps to any pre-specified output.
2. Collision-resistant — It is computationally infeasible to find any two distinct inputs that map to the same output.

identifier
A data object that is associated with a single, unique entity (e.g., individual, device, or *session*) within a given context and is never assigned to any other entity within that context.

identity
See *digital identity*

identity API
A protected API accessed by an *RP* to access the *attributes* of a specific subscriber.

Identity Assurance Level (IAL)
A category that conveys the degree of confidence that the *subject*'s *claimed identity* is their real identity.

identity evidence
Information or documentation that supports the real-world existence of the *claimed identity*. Identity evidence may be physical (e.g., a driver's license) or digital (e.g., a mobile driver's license or digital *assertion*). Evidence must support both *validation* (i.e., confirming *authenticity* and accuracy) and *verification* (i.e., confirming that the *applicant* is the true owner of the evidence).

identity proofing
The processes used to collect, validate, and verify information about a *subject* to establish assurance in the subject's *claimed identity*.

identity provider (IdP)
The party in a *federation transaction* that creates an *assertion* for the subscriber and transmits the assertion to the *RP*.

identity resolution
The process of collecting information about an *applicant* to uniquely distinguish an individual within the context of the population that the *CSP* serves.

identity verification
See *verification*

injection attack
An attack in which an attacker supplies untrusted input to a program. In the context of federation, the attacker presents an untrusted *assertion* or *assertion reference* to the *RP* in order to create an *authenticated session* with the RP.

issuing source
An authority responsible for the generation of data, digital evidence (i.e., *assertions*), or physical documents that can be used as *identity evidence*.

knowledge-based verification (KBV)
A process of validating knowledge of personal or private information associated with an individual for the purpose of verifying the *claimed identity* of an *applicant*. KBV does not include collecting personal *attributes* for the purposes of *identity resolution*.

legal person
An individual, organization, or company with legal rights.

login
Establishment of an *authenticated session* between a person and a system. Also known as *"sign in"*, *"log on"*, and *"sign on."*

manageability
Providing the capability for the granular administration of *personally identifiable information*, including alteration, deletion, and selective disclosure. [NISTIR8062]

memorized secret
See *password*.

message authentication code (MAC)
A cryptographic checksum on data that uses a *symmetric key* to detect both accidental and intentional modifications of the data. MACs provide *authenticity* and integrity protection, but not *non-repudiation* protection.

mobile code
Executable code that is normally transferred from its source to another computer system for execution. This transfer is often through the *network* (e.g., JavaScript embedded in a web page) but may transfer through physical media as well.

multi-factor authentication (MFA)
An authentication system that requires more than one distinct type of *authentication factor* for successful authentication. MFA can be performed using a *multi-factor authenticator* or by combining *single-factor* authenticators that provide different types of factors.

multi-factor authenticator
An *authenticator* that provides more than one distinct *authentication factor*, such as a cryptographic authentication device with an integrated biometric sensor that is required to activate the device.

natural person
A real-life human being, not synthetic or artificial.

network
An open communications medium, typically the Internet, used to transport messages between the *claimant* and other parties. Unless otherwise stated, no assumptions are made about the network's security; it is assumed to be open and subject to active (e.g., impersonation, *session* hijacking) and passive (e.g., eavesdropping) attacks at any point between the parties (e.g., claimant, *verifier*, *CSP*, *RP*).

nonce
A value used in security protocols that is never repeated with the same key. For example, nonces used as challenges in *challenge-response authentication protocols* must not be repeated until authentication keys are changed. Otherwise, there is a possibility of a *replay attack*. Using a nonce as a challenge is a different requirement than a random challenge, because a nonce is not necessarily unpredictable.

non-repudiation
The capability to protect against an individual falsely denying having performed a particular transaction.

offline attack
An attack in which the attacker obtains some data (typically by eavesdropping on an authentication transaction or by penetrating a system and stealing security files) that the attacker is able to analyze in a system of their own choosing.

one-to-one (1:1) comparison
The process in which a *biometric sample* from an individual is compared to a *biometric reference* to produce a comparison score.

online attack
An attack against an *authentication protocol* in which the attacker either assumes the role of a *claimant* with a genuine *verifier* or actively alters the authentication channel.

online guessing attack
An attack in which an attacker performs repeated logon trials by guessing possible values of the *authenticator* output.

online service
A service that is accessed remotely via a *network*, typically the internet.

pairwise pseudonymous identifier
A *pseudonymous identifier* generated by an IdP for use at a specific *RP*.

passphrase
A *password* that consists of a sequence of words or other text that a *claimant* uses to *authenticate* their identity. A passphrase is similar to a password in usage but is generally longer for added security.

password
A type of *authenticator* consisting of a character string that is intended to be memorized or memorable by the subscriber to permit the *claimant* to demonstrate *something they know* as part of an authentication process. Passwords are referred to as *memorized secrets* in the initial release of SP 800-63B.

personal identification number (PIN)
A *password* that typically consists of only decimal digits.

personal information
See *personally identifiable information*.

personally identifiable information (PII)
Information that can be used to distinguish or trace an individual's identity, either alone or when combined with other information that is linked or linkable to a specific individual. [A-130]

personally identifiable information processing
An operation or set of operations performed upon *personally identifiable information* that can include the collection, retention, logging, generation, transformation, use, disclosure, transfer, or disposal of personally identifiable information.

pharming
An attack in which an attacker corrupts an infrastructure service such as DNS (e.g., Domain Name System [DNS]) and causes the subscriber to be misdirected to a forged *verifier/RP*, which could cause the subscriber to reveal sensitive information, download harmful software, or contribute to a fraudulent act.

phishing
An attack in which the subscriber is lured (usually through an email) to interact with a counterfeit *verifier/RP* and tricked into revealing information that can be used to masquerade as that subscriber to the real verifier/RP.

phishing resistance
The ability of the *authentication protocol* to prevent the disclosure of *authentication secrets* and valid *authenticator* outputs to an impostor *verifier* without reliance on the vigilance of the *claimant*.

physical authenticator
An *authenticator* that the *claimant* proves possession of as part of an authentication process.

possession and control of an authenticator
The ability to activate and use the *authenticator* in an *authentication protocol*.

practice statement
A formal statement of the practices followed by the parties to an authentication process (e.g., *CSP* or *verifier*). It usually describes the parties' policies and practices and can become legally binding.

predictability
Enabling reliable assumptions by individuals, owners, and operators about PII and its *processing* by an information system. [NISTIR8062]

private key
In *asymmetric key* cryptography, the private key (i.e., a secret key) is a mathematical key used to create *digital signatures* and, depending on the algorithm, decrypt messages or files that are encrypted with the corresponding *public key*. In *symmetric key* cryptography, the same private key is used for both encryption and decryption.

processing
Operation or set of operations performed upon PII that can include, but is not limited to, the collection, retention, logging, generation, transformation, use, disclosure, transfer, and disposal of PII. [NISTIR8062]

presentation attack
Presentation to the biometric data capture subsystem with the goal of interfering with the operation of the biometric system.

presentation attack detection (PAD)
Automated determination of a *presentation attack*. A subset of presentation attack determination methods, referred to as *liveness detection*, involves the measurement and analysis of anatomical characteristics or voluntary or involuntary reactions, to determine if a *biometric sample* is being captured from a living *subject* that is present at the point of capture.

process assistant
An individual who provides support for the proofing process but does not support decision-making or risk-based evaluation (e.g., translation, transcription, or accessibility support).

proofing agent
An agent of the *CSP* who is trained to attend *identity proofing sessions* and can make limited risk-based decisions – such as physically inspecting *identity evidence* and making physical comparisons of the *applicant* to identity evidence.

Privacy Impact Assessment (PIA)
A method of analyzing how *personally identifiable information* (PII) is collected, used, shared, and maintained. PIAs are used to identify and mitigate privacy risks throughout the development lifecycle of a program or system. They also help ensure that handling information conforms to legal, regulatory, and policy requirements regarding privacy.

protected session
A *session* in which messages between two participants are encrypted and integrity is protected using a set of *shared secrets* called "session keys."

A protected session is said to be *authenticated* if — during the session — one participant proves possession of one or more *authenticators* in addition to the session keys, and if the other party can verify the identity associated with the authenticators. If both participants are authenticated, the protected session is said to be *mutually authenticated*.

Provisioning API
A protected API that allows an *RP* to access identity *attributes* for multiple subscribers for the purposes of provisioning and managing RP *subscriber accounts*.

pseudonym
A name other than a legal name.

pseudonymity
The use of a *pseudonym* to identify a *subject*.

pseudonymous identifier
A meaningless but unique *identifier* that does not allow the *RP* to infer anything regarding the subscriber but that does permit the RP to associate multiple interactions with a single subscriber.

public key
The public part of an *asymmetric key* pair that is used to verify signatures or encrypt data.

public key certificate
A digital document issued and digitally signed by the *private key* of a certificate authority that binds an *identifier* to a subscriber's *public key*. The certificate indicates that the subscriber identified in the certificate has sole control of and access to the private key. See also [RFC5280].

public key infrastructure (PKI)
A set of policies, processes, server platforms, software, and workstations used to administer certificates and public-_private key_ pairs, including the ability to issue, maintain, and revoke *public key certificates*.

reauthentication
The process of confirming the subscriber's continued presence and intent to be *authenticated* during an extended usage *session*.

registration
See *enrollment*.

relying party (RP)
An entity that relies upon a *verifier*'s *assertion* of a subscriber's identity, typically to process a transaction or grant access to information or a system.

remote
A process or transaction that is conducted through connected devices over a *network*, rather than in person.

replay attack
An attack in which the attacker is able to replay previously captured messages (between a legitimate *claimant* and a *verifier*) to masquerade as that claimant to the verifier or vice versa.

replay resistance
The property of an authentication process to resist *replay attacks*, typically by the use of an *authenticator* output that is valid only for a specific authentication.

resolution
See *identity resolution*.

restricted
An *authenticator* type, class, or instantiation that has additional risk of false acceptance associated with its use and is therefore subject to additional requirements.

risk assessment
The process of identifying, estimating, and prioritizing risks to organizational operations (i.e., mission, functions, image, or reputation), organizational assets, individuals, and other organizations that result from the operation of a system. A risk assessment is part of *risk management*, incorporates threat and vulnerability analyses, and considers mitigations provided by security *controls* that are planned or in-place. It is synonymous with "risk analysis."

risk management
The program and supporting processes that manage information security risk to organizational operations (including mission, functions, image, reputation), organizational assets, individuals, and other organizations and includes (i) establishing the context for risk-related activities, (ii) assessing risk, (iii) responding to risk once determined, and (iv) monitoring risk over time.

RP subscriber account
An account established and managed by the *RP* in a federated system based on the RP's view of the *subscriber account* from the *IdP*. An RP subscriber account is associated with one or more *federated identifiers* and allows the subscriber to access the account through a *federation transaction* with the IdP.

salt
A non-secret value used in a cryptographic process, usually to ensure that the results of computations for one instance cannot be reused by an attacker.

Secure Sockets Layer (SSL)
See *Transport Layer Security (TLS)*.

security domain
A set of systems under a common administrative and access control.

Senior Agency Official for Privacy (SAOP)
Person responsible for ensuring that an agency complies with privacy requirements and manages privacy risks. The SAOP is also responsible for ensuring that the agency considers the privacy impacts of all agency actions and policies that involve PII.

session
A persistent interaction between a subscriber and an *endpoint*, either an *RP* or a *CSP*. A session begins with an authentication event and ends with a session termination event. A session is bound by the use of a session secret that the subscriber's software (e.g., a browser, application, or OS) can present to the RP to prove association of the session with the authentication event.

session hijack attack
An attack in which the attacker is able to insert themselves between a *claimant* and a *verifier* subsequent to a successful authentication exchange between the latter two parties. The attacker is able to pose as a subscriber to the verifier or vice versa to control *session* data exchange. Sessions between the claimant and the *RP* can be similarly compromised.

shared secret
A secret used in authentication that is known to the subscriber and the verifier.

side-channel attack
An attack enabled by the leakage of information from a physical cryptosystem. Characteristics that could be exploited in a side-channel attack include timing, power consumption, and electromagnetic and acoustic emissions.

single-factor
A characteristic of an authentication system or an *authenticator* that requires only one *authentication factor* (i.e., something you know, something you have, or something you are) for successful authentication.

single sign-on (SSO)
An authentication process by which one account and its *authenticators* are used to access multiple applications in a seamless manner, generally implemented with a *federation protocol*.

social engineering
The act of deceiving an individual into revealing sensitive information, obtaining unauthorized access, or committing fraud by associating with the individual to gain confidence and trust.

subject
A person, organization, device, hardware, *network*, software, or service. In these guidelines, a subject is a *natural person*.

subscriber
An individual enrolled in the *CSP* identity service.

subscriber account
An account established by the *CSP* containing information and *authenticators* registered for each subscriber enrolled in the CSP identity service.

supplemental controls
Controls that may be added, in addition to those specified in the organization's tailored assurance level, in order to address specific threats or attacks.

symmetric key
A *cryptographic key* used to perform both the cryptographic operation and its inverse. (e.g., to encrypt and decrypt or create a *message authentication code* and to verify the code).

sync fabric
Any on-premises, cloud-based, or hybrid service used to store, transmit, or manage authentication keys generated by syncable *authenticators* that are not local to the user's device.

syncable authenticators
Software or hardware cryptographic *authenticators* that allow authentication keys to be cloned and exported to other storage to sync those keys to other authenticators (i.e., devices).

synthetic identity fraud
The use of a combination of *personally identifiable information* (PII) to fabricate a person or entity in order to commit a dishonest act for personal or financial gain.

system of record (SOR)
An SOR is a collection of records that contain information about individuals and are under the control of an agency. The records can be retrieved by the individual's name or by an identifying number, symbol, or other *identifier*.

System of Record Notice (SORN)
A notice that federal agencies publish in the Federal Register to describe their systems of records.

tailoring
The process by which xALs and specified *controls* are modified by: considerations for the impacts on privacy, *usability*, and *equity* on the user population, identifying and designating common controls, applying scoping considerations on the applicability and implementation of specified controls, selecting any *compensating controls*, assigning specific values to organization-defined security control parameters, supplementing xAL controls with additional controls or control enhancements, and providing additional specification information for control implementation.

token
See *authenticator*.

transaction
See *digital transaction*

Transport Layer Security (TLS)
An authentication and security protocol widely implemented in browsers and web servers. TLS is defined by [RFC5246]. TLS is similar to the older SSL protocol, and TLS 1.0 is effectively SSL version 3.1. SP 800-52, Guidelines for the Selection and Use of Transport Layer Security (TLS) Implementations [SP800-52], specifies how TLS is to be used in government applications.

trust agreement
A set of conditions under which a *CSP*, *IdP*, and *RP* are allowed to participate in a *federation transaction* for the purposes of establishing an authentication *session* between the subscriber and the RP.

trust anchor
A public or *symmetric key* that is trusted because it is built directly into hardware or software or securely provisioned via out-of-band means rather than because it is vouched for by another trusted entity (e.g., in a *public key* certificate). A trust anchor may have name or policy constraints that limit its scope.

trusted referee
An agent of the *CSP* who is trained to make risk-based decisions regarding an *applicant's identity proofing* case when that applicant is unable to meet the expected requirements of a defined IAL proofing process.

usability
The extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency, and satisfaction in a specified context of use. [ISO/IEC9241-11]

validation
The process or act of checking and confirming that the evidence and *attributes* supplied by an *applicant* are authentic, accurate and associated with a real-life identity. Specifically, evidence validation is the process or act of checking that the presented evidence is authentic, current, and issued from an acceptable source. See also *attribute validation*.

verification
The process or act of confirming that the *applicant* undergoing *identity proofing* holds the claimed real-life identity represented by the validated identity *attributes* and associated evidence. Synonymous with "identity verification."

verifier
An entity that verifies the *claimant's* identity by verifying the claimant's possession and control of one or more *authenticators* using an *authentication protocol*. To do this, the verifier needs to confirm the binding of the authenticators with the *subscriber account* and check that the subscriber account is active.

verifier impersonation
See *phishing*.

zeroize
Overwrite a memory location with data that consists entirely of bits with the value zero so that the data is destroyed and unrecoverable. This is often contrasted with deletion methods that merely destroy references to data within a file system rather than the data itself.

zero-knowledge password protocol
A password-based *authentication protocol* that allows a *claimant* to *authenticate* to a *verifier* without revealing the *password* to the verifier. Examples of such protocols are EKE, SPEKE and SRP.

Appendix C. Change Log

C.1. SP 800-63-1

NIST SP 800-63-1 updated NIST SP 800-63 to reflect current authenticator (then referred to as "token") technologies and restructured it to provide a better understanding of the digital identity architectural model used here. Additional (minimum) technical requirements were specified for the CSP, protocols used to transport authentication information, and assertions if implemented within the digital identity model.

C.2. SP 800-63-2

NIST SP 800-63-2 was a limited update of SP 800-63-1 and substantive changes were made only in Sec. 5, *Registration and Issuance Processes*. The substantive changes in the revised draft were intended to facilitate the use of professional credentials in the identity proofing process, and to reduce the need to send postal mail to an address of record to issue credentials for level 3 remote registration. Other changes to Sec. 5 were minor explanations and clarifications.

C.3. SP 800-63-3

NIST SP 800-63-3 is a substantial update and restructuring of SP 800-63-2. SP 800-63-3 introduces individual components of digital authentication assurance — AAL, IAL, and FAL — to support the growing need for independent treatment of authentication strength and confidence in an individual's claimed identity (e.g., in strong pseudonymous authentication). A risk assessment methodology and its application to IAL, AAL, and FAL has been included in this guideline. It also moves the whole of digital identity guidance covered under SP 800-63 from a single document describing authentication to a suite of four documents (to separately address the individual components mentioned above) of which SP 800-63-3 is the top-level document.

Other areas updated in 800-63-3 include:

- Renamed to *Digital Identity Guidelines* to properly represent the scope includes identity proofing and federation, and to support expanding the scope to include device identity, or machine-to-machine authentication in future revisions.
- Changed terminology, including the use of *authenticator* in place of *token* to avoid conflicting use of the word *token* in assertion technologies.
- Updated authentication and assertion requirements to reflect advances in both security technology and threats.
- Added requirements on the storage of long-term secrets by verifiers.
- Restructured identity proofing model.
- Updated requirements regarding remote identity proofing.

- Clarified the use of independent channels and devices as “something you have”.
- Removed pre-registered knowledge tokens (authenticators), with the recognition that they are special cases of (often very weak) passwords.
- Added requirements regarding account recovery in the event of loss or theft of an authenticator.
- Removed email as a valid channel for out-of-band authenticators.
- Expanded discussion of reauthentication and session management.
- Expanded discussion of identity federation; restructuring of assertions in the context of federation.

C.4. SP 800-63-4

NIST SP 800-63-4 has substantial updates and re-organization from SP 800-63-3. Updates to 800-63-4 include:

- Expanded security and privacy considerations and added equity and usability considerations.
- Updated digital identity models and added a user-controlled wallet federation model that addresses the increased attention and adoption of digital wallets and attribute bundles.
- Expanded digital identity risk management process to include definition of the protected online services, user groups, and impacted entities.
- A more descriptive introduction to establish the context of the DIRM process, the two dimensions of risk it addresses, and the intended outcomes. This context-setting step includes defining and understanding the online service that the organization is offering and intending to protect with identity systems.
- Expanded digital identity risk management process to include definition of the protected online services, user groups, and impacted entities.
- Updated digital identity risk management process for additional assessments for tailoring initial baseline control selections.
- Added performance metrics for the continuous evaluation of digital identity systems.
- Added a new subsection on redress processes and requirements.
- Added a new Artificial Intelligence subsection to address the use of Artificial Intelligence in digital identity services.