NIST Special Publication
NIST SP 800-63-4 ipd

Digital Identity Guidelines

Initial 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.ipd

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

NIST Special Publication
NIST SP 800-63-4 ipd

Digital Identity Guidelines

Initial 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.ipd

December 2022

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 (2022) Digital Identity Guidelines. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-63-4 ipd. https://doi.org/10.6028/NIST.SP.800-63-4.ipd

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

December 16, 2022 - ~~March 24~~ April 14, 2023

Submit Comments

mailto:dig-comments@nist.gov

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

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.

Abstract

These guidelines provide technical requirements for federal agencies implementing digital identity services and are not intended to constrain the development or use of standards outside of this purpose. The 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. This publication will supersede NIST Special Publication 800-63-3.

Keywords

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

Note to Reviewers

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 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.

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

1. **Advance Equity:** This draft seeks to expand upon the risk management content of previous revisions and specifically mandates that agencies account for impacts to individuals and communities in addition to impacts to the organization. It also elevates risks to mission delivery – including challenges to providing services to all people who are eligible for and entitled to them – within the risk management process and when implementing digital identity systems. Additionally, the guidance now mandates continuous evaluation of potential impacts across demographics, provides biometric performance requirements, and additional parameters for the responsible use of biometric-based technologies, such as those that utilize face recognition.
2. **Emphasize Optionality and Choice for Consumers:** In the interest of promoting and investigating additional scalable, equitable, and convenient identify verification options, including those that do and do not leverage face recognition technologies, this draft expands the list of acceptable identity proofing alternatives to provide new mechanisms to securely deliver services to individuals with differing means, motivations, and backgrounds. The revision also emphasizes the need for digital identity services to support multiple authenticator options to address diverse consumer needs and secure account recovery.
3. **Deter Fraud and Advanced Threats:** This draft enhances fraud prevention measures from the third revision by updating risk and threat models to account for new attacks, providing new options for phishing resistant authentication, and introducing requirements to prevent automated attacks against enrollment processes. It also opens the door to new technology such as mobile driver's licenses and verifiable credentials.
4. **Address Implementation Lessons Learned:** This draft addresses areas where implementation experience has indicated that additional clarity or detail was required to effectively operationalize the guidelines. This includes re-working the federation assurance levels, providing greater detail on Trusted Referees, clarifying guidelines on identity attribute validation sources, and improving address confirmation requirements.

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

Identity Proofing and Enrollment

- NIST sees a need for inclusion of an unattended, fully remote Identity Assurance Level (IAL) 2 identity proofing workflow that provides security and convenience, but does not require face recognition. Accordingly, NIST seeks input on the following questions:

  - What technologies or methods can be applied to develop a remote, unattended IAL2 identity proofing process that demonstrably mitigates the same risks as the current IAL2 process?
  - Are these technologies supported by existing or emerging technical standards?
  - Do these technologies have established metrics and testing methodologies to allow for assessment of performance and understanding of impacts across user populations (e.g., bias in artificial intelligence)?
- What methods exist for integrating digital evidence (e.g., Mobile Driver's Licenses, Verifiable Credentials) into identity proofing at various identity assurance levels?
- What are the impacts, benefits, and risks of specifying a set of requirements for CSPs to establish and maintain fraud detection, response, and notification capabilities?
  - Are there existing fraud checks (e.g., date of death) or fraud prevention techniques (e.g., device fingerprinting) that should be incorporated as baseline normative requirements? If so, at what assurance levels could these be applied?
  - How might emerging methods such as fraud analytics and risk scoring be further researched, standardized, measured, and integrated into the guidance in the future?
  - What accompanying privacy and equity considerations should be addressed alongside these methods?
- Are current testing programs for liveness detection and presentation attack detection sufficient for evaluating the performance of implementations and technologies?
- What impacts would the proposed biometric performance requirements for identity proofing have on real-world implementations of biometric technologies?

Risk Management

- What additional guidance or direction can be provided to integrate digital identity risk with enterprise risk management?
- How might equity, privacy, and usability impacts be integrated into the assurance level selection process and digital identity risk management model?
- How might risk analytics and fraud mitigation techniques be integrated into the selection of different identity assurance levels? How can we qualify or quantify their ability to mitigate overall identity risk?

Authentication and Lifecycle Management

- Are emerging authentication models and techniques – such as FIDO passkey, Verifiable Credentials, and mobile driver's licenses – sufficiently addressed and accommodated, as appropriate, by the guidelines? What are the potential associated security, privacy, and usability benefits and risks?
- Are the controls for phishing resistance as defined in the guidelines for AAL2 and AAL3 authentication clear and sufficient?
- How are session management thresholds and reauthentication requirements implemented by agencies and organizations? Should NIST provide thresholds or leave session lengths to agencies based on applications, users, and mission needs?
- What impacts would the proposed biometric performance requirements for this volume have on real-world implementations of biometric technologies?

Federation and Assertions

- What additional privacy considerations (e.g., revocation of consent, limitations of use) may be required to account for the use of identity and provisioning APIs that had not previously been discussed in the guidelines?
- Is the updated text and introduction of "bound authenticators" sufficiently clear to allow for practical implementations of federation assurance level (FAL) 3 transactions? What complications or challenges are anticipated based on the updated guidance?

General

- Is there an element of this guidance that you think is missing or could be expanded?
- Is any language in the guidance confusing or hard to understand? Should we add definitions or additional context to any language?
- Does the guidance sufficiently address privacy?
- Does the guidance sufficiently address equity?
  - What equity assessment methods, impact evaluation models, or metrics could we reference to better support organizations in preventing or detecting disparate impacts that could arise as a result of identity verification technologies or processes?
- What specific implementation guidance, reference architectures, metrics, or other supporting resources may 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 impact 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 of the NIST SP 800-63-4 suite. NIST requests that all comments be submitted by 11:59pm Eastern Time on March 24, 2023. Please submit your comments to dig-comments@nist.gov. NIST will review all comments and make them available at the NIST Identity and Access Management website. Commenters are encouraged to use the comment template provided on the NIST Computer Security Resource Center website.

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

List of Tables

List of Figures

Acknowledgments

The authors would like to thank their fellow collaborators on the current revision of this special publication, Christine Abruzzi, James L. Fenton, and Justin P. Richer, as well as Kerrianne Buchanan for her contributions and review. The authors would like to also 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. Purpose

This section is informative.

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.

2. Introduction

This section is informative.

As the line between the virtual world and physical world blurs, and as digital and internet-enabled technologies continue to proliferate and connect, it is imperative that developers and consumers alike understand this changing hybrid ecosystem - including its associated opportunities and risks. Engagement across this ecosystem is often determined by an individual's ability and willingness to establish a digital identity - the unique representation of a person engaged in an online transaction.

A digital identity is always unique in the context of a digital service but does not always uniquely identify a person in all contexts. Further, while a digital identity may relay unique and specific meaning within the context of a digital service, the real-life identity of the individual behind the digital identity may not be known. For the purpose of this publication, a "person" refers to natural persons only (i.e., not all legal persons.)

Establishing 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 digital transaction. However, this process can present challenges. As in relationships and transactions in the physical world, 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, digital services must be designed with equity and flexibility in mind to ensure broad and enduring participation.

Risks associated with digital identity stretch beyond the potential impacts to enterprises and should be incorporated into enterprise decision-making. This publication endeavors to more robustly and explicitly account for risks to individuals, communities, and other organizations. Specifically, while using this guidance, organizations should consider how decisions related to digital identity that prioritize organizational cybersecurity objectives might affect or need to accommodate other objectives, such as those related to privacy, equity, usability, and other indicators of mission and business performance that center the experiences of the individuals interacting with programs and services. By taking a human-centered 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 effective and culturally appropriate redress options.

These guidelines lay out a model for federal programs and other organizations to assess and manage risks associated with digital identity systems, including the processes, policies, data, people, and technologies that support digital identity management. The model is supported by a series of processes: identity proofing, authentication, and federation. The identity proofing process establishes that a subject is a specific physical person. The digital authentication process determines the validity of one or

more authenticators used to claim a digital identity and establishes confidence that a subject attempting to access a digital service: (1) is in control of the technologies being used for authentication, and (2) is the same subject that previously accessed the service. Finally, the federation process allows for identity information to be shared in support of authentication across systems.

The composition, model, 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, and equitable services to diverse user communities. This revision addresses these challenges while facilitating the new models and architectures for identity services that have developed by clarifying requirements based on the function 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) and it describes the risk management processes that organizations should follow for implementing digital identity services and that supplement the *NIST Risk Management Framework* [NISTRMF] and its component special publications. The publication expands upon the NIST RMF by outlining how equity and usability considerations should be incorporated into digital identity risk management processes and it highlights the importance of considering impacts, not only on the enterprise operations and assets, but also on individuals, other organizations, and, more broadly, society. Further, while digital authentication supports privacy protection by mitigating risks of unauthorized access to individuals' information, given that identity proofing, authentication, authorization, and federation often involve the processing of individuals' information, these functions can also create privacy risks. These guidelines, therefore, include privacy requirements and considerations to help mitigate potential associated privacy risks.

Finally, while this publication provides organizations with technical requirements and recommendations for establishing, maintaining, and authenticating the digital identity of subjects in order to access digital systems over a network, additional support options outside the purview of information technology teams may need to be provided to address barriers and adverse impacts, foster equity, and successfully deliver on mission objectives.

2.1. Scope & Applicability

Not all digital services require identity proofing or authentication; however, this guidance applies to all online transactions for which some level of digital identity is required, regardless of the constituency (e.g., citizens, business partners, and government entities).

These guidelines primarily focus on organizational services that interact with external users, such as citizens 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, 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.

Transactions not covered by this guidance include those associated with national security systems as defined in 44 U.S.C. § 3542(b)(2). 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.

Additionally, these technical guidelines do not address the identity of subjects for physical access (e.g., to buildings), though some identities used for online transactions may also be used for physical access. Additionally, this revision of these guidelines does not explicitly address device identity, often referred to as machine-to-machine (such as router-to-router) authentication or interconnected devices, commonly referred to as the internet of things (IoT), although these guidelines are written to refer to generic subjects wherever possible to leave open the possibility for applicability to devices. Furthermore, these guidelines do not address authorization of access to Application Programming Interfaces (APIs) on behalf of subjects.

2.2. How to Use this Suite of SPs

These guidelines support the mitigation of the negative impacts induced by a digital identity error by separating the individual elements of digital identity into discrete, component parts. For non-federated systems, agencies will select two components, referred to as *Identity Assurance Level (IAL)* and *Authentication Assurance Level (AAL)*. For federated systems, a third component, *Federation Assurance Level (FAL)*, is included. Sec. 5, Digital Identity Risk Management provides details on the risk assessment process and how the results of the risk assessment, with additional context, inform organizational selection of IAL, AAL, and FAL combinations based on risk and mission.

By conducting appropriate risk management for business, security, and privacy, side-by-side with mission needs, organizations will select IAL, AAL, and FAL as distinct options. Specifically, organizations are required to individually select levels corresponding to each function being performed. While many systems could have the same numerical level for each IAL, AAL, and FAL, this is not a requirement and organizations should not assume they will be the same in any given system or application.

The components of identity assurance detailed in these guidelines are as follows:

- **IAL** refers to the identity proofing process.
- **AAL** refers to the authentication process.
- **FAL** refers to the federation process, when the RP is connected through a federated protocol.

Note: When described generically or bundled, these guidelines will refer to IAL, AAL, and FAL as ***xAL***.

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

SP 800-63 *Digital Identity Guidelines*: Provides the risk assessment methodology and an overview of general identity frameworks, using authenticators, credentials, and assertions together in a digital system, and a risk-based process of selecting assurance levels. *SP 800-63 contains both normative and informative material.*

[SP800-63A]: Provides requirements for enrollment and identity proofing of applicants, either remotely or in person, that wish to gain access to resources at each of the three identity assurance levels (IALs). It details the responsibilities of Credential Service Providers (CSPs) with respect to establishing and maintaining subscriber accounts and binding authenticators (either CSP-issued or subscriber-provided) to the subscriber account. *SP 800-63A contains both normative and informative material.*

[SP800-63B]: Provides recommendations on types of authentication processes, including choices of authenticators, that may be used at each of the three authentication assurance levels (AALs). It also provides recommendations on the lifecycle 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. Further, this volume offers privacy-enhancing techniques to share information about a valid, authenticated subject, and describes methods that allow for strong multi-factor authentication (MFA) while the subject remains pseudonymous to the digital service. *SP 800-63C contains both normative and informative material.*

2.3. Enterprise Risk Management Requirements and Considerations

Effective enterprise risk management is multidisciplinary by default and involves the consideration of a diverse set 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 not only to enterprise assets and operations but also to individuals, other organizations, and society more broadly.

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, for instance where privacy or other legal requirements exist or where the output of a risk assessment leads the organization to determine that additional measures or other process safeguards are appropriate. Organizations, including federal agencies, may

employ compensating or supplemental controls not specified in this publication. They may also consider partitioning the functionality of a digital service to allow less sensitive functions to be available at a lower level of assurance.

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

2.3.1. Security

It is increasingly important for enterprise organizations to assess and manage digital identity security risks, such as unauthorized access, availability issues, impersonation, and other types of fraudulent claims, as well as institute strong identity governance practices. 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 need to adhere to their 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 equivalent standards 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 as determined under FISMA and the RMF.

2.3.2. 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 PII — a problematic data action — and the potential impact of the problematic data action should it occur. Additionally, by focusing on the privacy

engineering objectives of predictability, manageability, and disassociability, organizations can determine the types of capabilities a given system may need to be able to demonstrate how organizational privacy policies and system privacy requirements have been implemented.

The *Privacy Act of 1974, 2010 Edition*, [PrivacyAct] 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.

When designing and implementing digital identity management processes and systems, privacy risk assessments are required for PII processing under these guidelines. Such privacy risk assessments can be used to support Privacy Impact Assessments under *OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002* [M-03-22] as well as to select controls from NIST Special Publication 800-53, *Security and Privacy Controls for Information Systems and Organizations* [SP800-53]. Further, each volume of 800-63 (63A, 63B, and 63C) contains a specific section providing detailed privacy requirements and considerations for the implementation of the processes, controls, and requirements presented in that volume.

2.3.3. Equity

As defined in Executive Order 13985, *Advancing Racial Equity and Support for Underserved Communities Through the Federal Government* [EO13985], equity refers to 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.

A person's ability to engage in an online transaction, such as accessing a critical service like healthcare, is often dependent on their ability to successfully and safely present a digital identity. Given the broad disparities that exist in the U.S. society and globally, many people are either unable to successfully present a digital identity, or they face a higher degree of burden in navigating online services than their more privileged peers, leaving them locked out of critical services or broader participation in the online world. 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.

Readers of this guidance are encouraged to consider existing inequities faced by the populations they serve to identify opportunities to design or operate digital identity systems and processes in ways that best support their needs. Readers are also encouraged

to consider any potential or actual impact to the experiences and outcomes of these populations, including disparities between populations, caused by the design or operation of digital identity systems.

For federal agencies implementing these guidelines, EO 13985 directs federal agencies to identify underserved communities for the programs and services that they provide and to determine and address any systemic barriers to underserved communities to provide equitable access to those programs and services. In alignment with the direction set by EO 13985, federal agencies should determine potential barriers communities and individuals may face to enrollment in and access to online benefits and services. They should also identify whether programmatic changes may be necessary to advance equity.

2.3.4. Usability

Usability refers to the extent to which a system, product, or service can be used by specified users to achieve specified goals with effectiveness, efficiency, and satisfaction in a specified context of use.

Similar to equity, usability requires an understanding of the people interacting with a digital identity system or process, as well as their unique goals and context of use. To provide an effective, efficient, and satisfactory experience, 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 lifecycle of a digital identity system or process, it is important to conduct usability evaluation with representative users performing realistic scenarios and tasks in appropriate context of use.

Digital identity management processes should be designed and implemented so 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.

3. Definitions and Abbreviations

See Appendix A for a complete set of definitions and abbreviations.

4. Digital Identity Model

This section is informative.

4.1. Overview

The SP 800-63 guidelines use digital identity models that reflect technologies and architectures 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 (represented by one of three roles):

- Applicant — the subject to be identity proofed
- Subscriber — the subject that has successfully completed the identity proofing process or has successfully completed authentication
- Claimant — the subject to be authenticated

Credential Service Provider (CSP): A trusted entity whose functions include identity proofing applicants to the identity service and the registration of authenticators to subscriber accounts. A *subscriber account* is the CSP's established record of the subscriber, the subscriber's attributes, and associated authenticators. A CSP may be an independent third party.

Relying Party (RP): An entity that relies upon the information in the subscriber account, or an identity provider (IdP) assertion when using federation, typically to process a transaction or grant access to information or a system.

Verifier: An entity whose function 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): An entity in a federated model that performs both the CSP and Verifier functions. The IdP is responsible for authenticating the subscriber and issuing assertions to communicate with one or more RPs.

The entities and interactions that comprise the non-federated digital identity model are illustrated in Figure 1. The federated digital identity model is illustrated in Figure 2.

Figure 1 shows an example of a common sequence of interactions in the non-federated model. Other sequences could also achieve the same functional requirements. The usual sequence of interactions for identity proofing and enrollment activities is as follows:

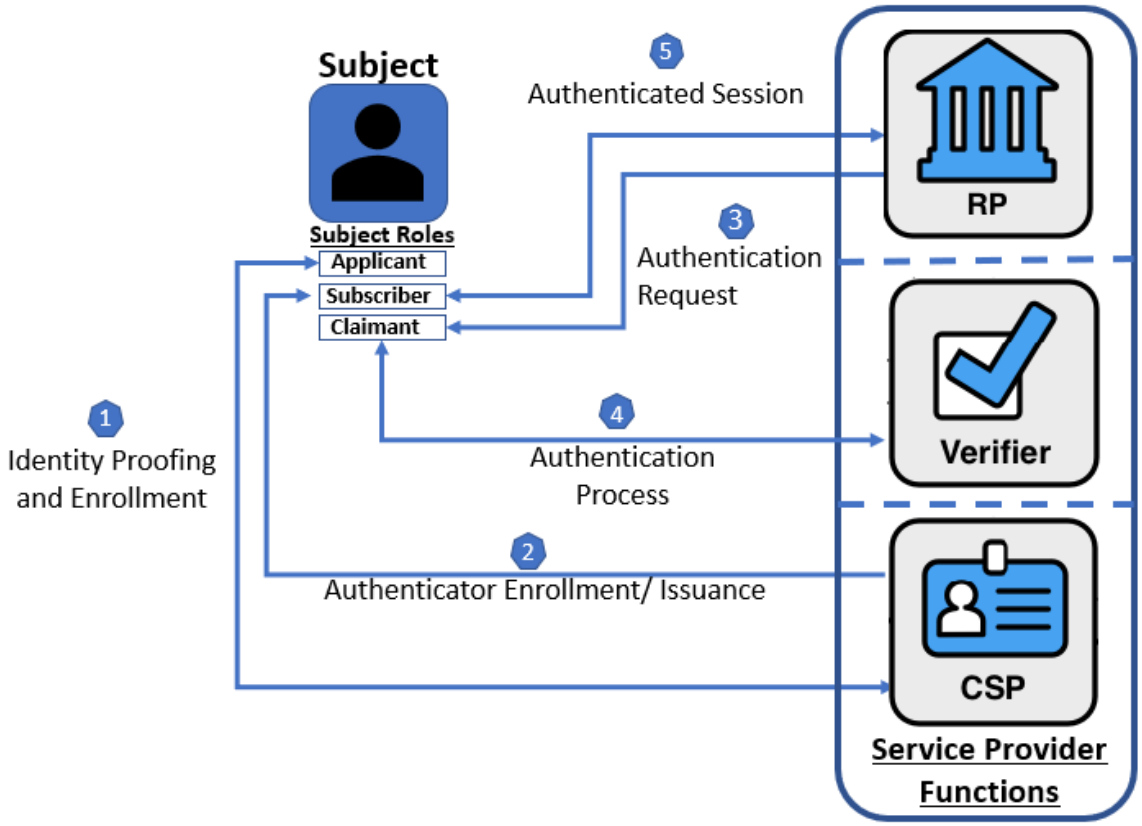

Figure 1. Non-Federated Digital Identity Model Example

- Step 1: An applicant applies to a CSP through an enrollment process. The CSP identity proofs that applicant.
- Step 2: Upon successful 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.

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 RP requests authentication from the claimant.
- 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 provide 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.

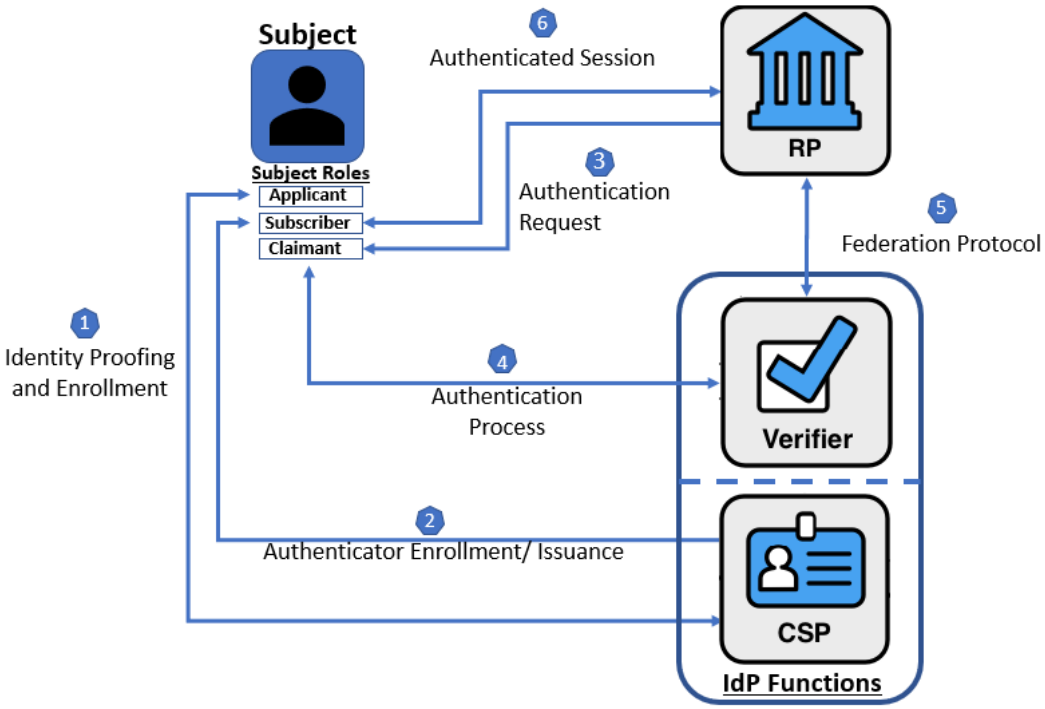

Figure 2. Federated Digital Identity Model Example

Figure 2 shows an example of those same common interactions in a federated model.

- Step 1: An applicant applies to an IdP through an enrollment process. Using its CSP function, the IdP identity proofs the applicant.
- Step 2: Upon successful proofing, the applicant is enrolled in the identity service as a subscriber.
  - A subscriber account and corresponding authenticators are established between the IdP and the subscriber. The IdP maintains the subscriber account, its status, and the enrollment data collected for the lifetime of the subscriber account (at a minimum). The subscriber maintains their authenticators.

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 authentication from the claimant. The IdP provides an assertion and optionally additional attributes to the RP through a federation protocol.
- 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: All communication, including assertions, between the RP and the IdP happens through federation protocols.
- Step 6: The IdP provides the RP with the authentication status of the subscriber and relevant attributes and an authenticated session is established between the subscriber and the RP.

For both models, the verifier does not always need to communicate in real time with the CSP to complete the authentication activity (e.g., some uses of digital certificates). 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 running on the same system rather than using network protocols.

In all cases, the RP should request the attributes it requires from a CSP or IdP before authenticating the claimant.

The following sections provide more detailed digital identity models for identity proofing, authentication, and federation.

4.2. Enrollment and Identity Proofing

The previous section introduced the entities and interactions in the conceptual digital identity model. This section provides additional details regarding the participants' relationships and responsibilities with respect to identity proofing and enrollment processes.

[SP800-63A], *Enrollment and Identity Proofing* provides general information and normative requirements for the identity proofing and enrollment processes as well as requirements specific to identity assurance levels (IALs). In addition to a "no identity proofing" level, IAL0, this document defines three IALs that indicate the relative strength of an identity proofing process.

An individual, referred to as an *applicant* at this stage, opts to enroll with a CSP. If the applicant is successfully proofed, the individual is then enrolled in the identity service as a *subscriber* of that CSP.

The CSP then establishes a subscriber account to uniquely identify each subscriber and record any authenticators registered (bound) to that subscriber account. The CSP may:

- issue one or more authenticators to the subscriber at the time of enrollment,
- bind authenticators provided by the subscriber, and/or
- bind authenticators to the subscriber account at a later time as needed.

CSPs generally maintain subscriber accounts according to a documented lifecycle, which defines specific events, activities, and changes that affect the status of a subscriber account. CSPs generally limit the lifetime of a subscriber account and any associated authenticators in order to ensure some level of accuracy and currency of attributes associated with a subscriber. When there is a status change or when the authenticators near expiration and any renewal requirements are met, they may be renewed and/or re-issued. Alternately, the authenticators may be invalidated and destroyed according to the CSPs written policy and procedures.

Subscribers have a duty to maintain control of their authenticators and comply with CSP policies in order to remain in good standing with the CSP.

In order to request issuance of a new authenticator, typically the subscriber authenticates to the CSP using their existing, unexpired authenticators. If the subscriber fails to request authenticator re-issuance prior to their expiration or revocation, they may be required to repeat the identity proofing (either complete or abbreviated) and enrollment processes in order to obtain a new authenticator.

Figure 3 shows a sample of interactions for identity proofing and enrollment.

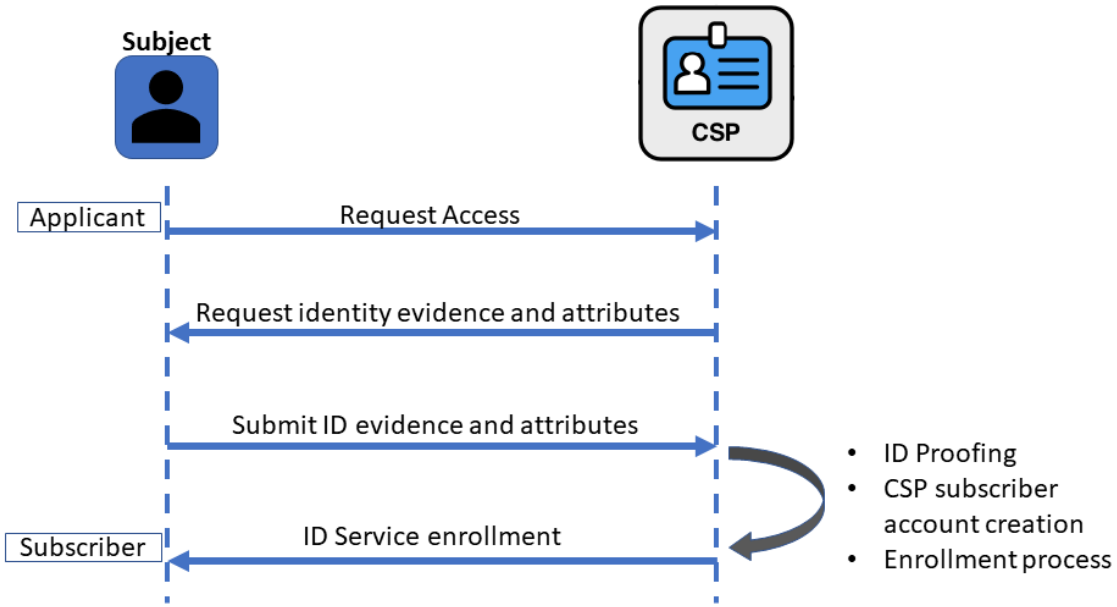

Figure 3. Sample Identity Proofing and Enrollment Digital Identity Model

4.3. Authentication and Lifecycle Management

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

4.3.1. Authenticators

The classic paradigm for authentication systems identifies three factors as the cornerstones of authentication:

- Something you know (e.g., a password)
- Something you have (e.g., an ID badge or 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. For the purposes of these guidelines, using two factors is adequate to meet the highest security requirements. Other types of information, such as location data or device identity, may also be used by a verifier to evaluate the risk in a claimed identity but they are not considered authentication factors.

In digital authentication, the claimant possesses and controls one or more authenticators. The authenticators will have been bound with the subscriber account. The authenticators contain secrets the claimant can use to prove they are a legitimate subscriber. The claimant authenticates to a system or application over a network by demonstrating they have possession and control of the authenticator. Once authenticated, the claimant is referred to as a subscriber.

The secrets contained in an authenticator are based on either key pairs (asymmetric cryptographic keys) or shared secrets (including symmetric cryptographic keys and memorized secrets). 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, for example through a public key certificate, can use an authentication protocol to verify the claimant has possession and control of the associated private key contained in the authenticators and, therefore, is a subscriber.

As mentioned above, shared secrets stored on an authenticator may be either symmetric keys or memorized secrets (e.g., passwords and PINs). While both keys and memorized secrets can be used in similar protocols, one important difference between the two is how they relate to the claimant. Symmetric keys are generally chosen at random and are complex and long enough to thwart network-based guessing attacks, and stored in hardware or software that the subscriber controls. Memorized secrets typically have fewer

characters and less complexity than cryptographic keys to facilitate memorization and ease of entry. The result is that memorized secrets have increased vulnerabilities that require additional defenses to mitigate.

There is another type of memorized secret used as an activation factor for a multi-factor authenticator. These 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 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.

As used in these guidelines, authenticators always contain or comprise a secret; however, 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 also does not constitute a secret and can not be used as a single-factor authenticator.

A digital authentication system may incorporate multiple factors in one of two ways:

1. The system may be implemented so that multiple factors are presented to the verifier, or
2. Some factors may be used to protect a secret that will be presented to the verifier.

For example, item 1 can be satisfied by pairing a memorized secret (something you know) with an out-of-band device (something you have). Both authenticator outputs are presented to the verifier to authenticate the claimant. For item 2, the authenticator and authenticator secret could be a piece of hardware that contains a cryptographic key (something you have) that is controlled by the claimant where access is protected with a fingerprint (something you are). When used with the biometric factor, the cryptographic key produces an output that is used to authenticate the claimant.

As noted above, biometrics do not constitute acceptable secrets for digital authentication and, therefore, cannot be used for single-factor authentication. However, biometrics authentication can be used as an authentication factor for multi-factor authentication when used in combination with a possession-based authenticator. 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 verification. This includes, but is not limited to, facial features, fingerprints, iris patterns, and voiceprints.

4.3.2. Subscriber Accounts

As described in the preceding sections, a subscriber account binds one or more authenticators to the subscriber via an identifier as part of the registration process. A subscriber account is created, stored, and maintained by the CSP. The subscriber account records all identity attributes validated during the identity proofing process.

4.3.3. Authentication Process

The authentication process enables an RP to trust that a claimant is who they say they are. Figure 4 shows a sample authentication process. Other approaches are described in [SP800-63B], *Authentication and Lifecycle Management*. This sample authentication process 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, as shown in Fig. 4, the RP, or both.

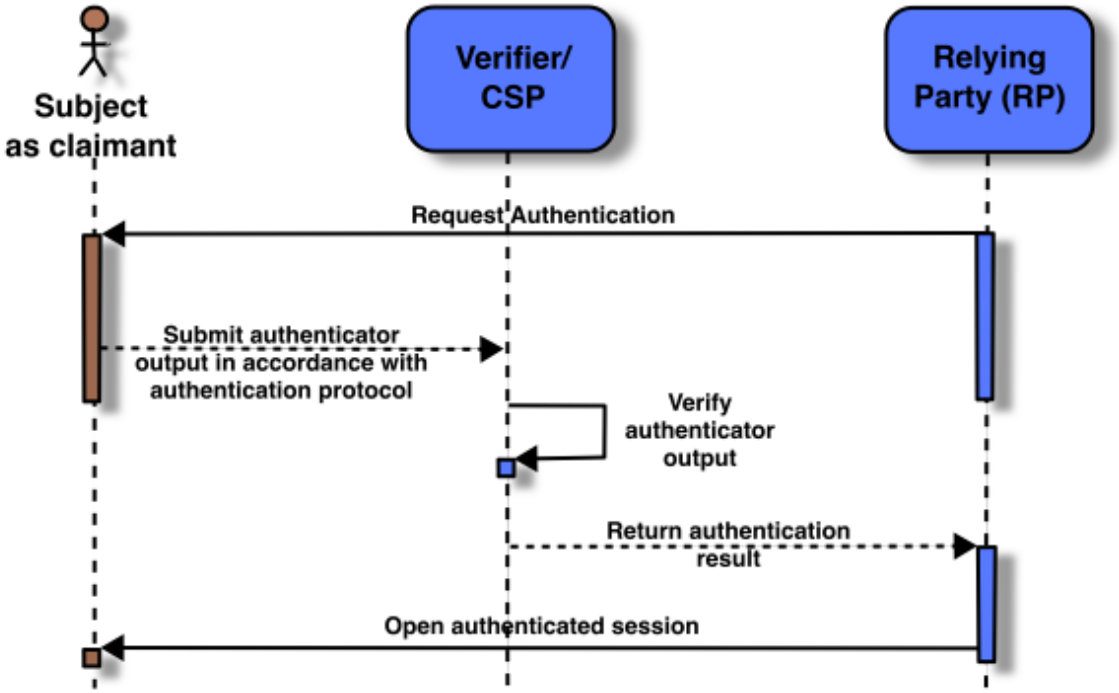

Figure 4. Sample Authentication Process

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 involving 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 that can be done by an attacker masquerading as a legitimate verifier.

Additionally, mechanisms located at the verifier can mitigate online guessing attacks against lower entropy secrets — like 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.

4.4. Federation and Assertions

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

In general usage, the term *federation* can be applied to a number of different approaches involving the sharing of information between different trust domains. These approaches differ based on the kind of information that is being shared between the domains. Some common examples include:

- sharing identifiers (e.g., using a driver's license number or an email address),
- sharing authenticators (e.g., using a PKI authenticator for multiple applications),
- sharing identity assertions (e.g., a federation protocol like OpenID Connect or SAML),
- sharing account attributes (e.g., a provisioning protocol like SCIM), and
- sharing authorization decisions (e.g., a policy protocol like XACML).

The SP 800-63 guidelines are agnostic to the identity proofing, authentication, and federation architectures 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 local to the organization or individual applications. The following lists detail scenarios where the organization may consider federation a viable option. These lists are provided for consideration and are not intended to be comprehensive.

An organization should consider accepting federated identity assertions if any of the following apply:

1. Potential users already have an authenticator at or above the required AAL.
2. Multiple types of authenticators are required to cover all possible user communities.
3. An organization does not have the necessary infrastructure to support management of subscriber accounts (e.g., account recovery, authenticator issuance, help desk).

4. There is a desire to allow primary authenticators to be added and upgraded over time without changing the RP's implementation.
5. There are different environments to be supported, as federation protocols are network-based and allow for implementation on a wide variety of platforms and languages.
6. Potential users come from multiple communities, each with its own existing identity infrastructure.
7. The ability to centrally manage account lifecycles, including account revocation and binding of new authenticators is important.

An organization should consider accepting federated identity attributes if any of the following apply:

1. Pseudonymity is required, necessary, feasible, or important to stakeholders accessing the service.
2. Access to the service requires a partial attribute list.
3. Access to the service requires at least one derived attribute value.
4. The organization is not the authoritative source or issuing source for required attributes.
5. Attributes are only required temporarily during use (such as to make an access decision), and the organization does not need to retain the data.

4.4.1. Federation Benefits

Federated architectures have many significant benefits, including, but not limited to:

- Enhanced user experience: For example, an individual can be identity proofed once and reuse the subscriber account at multiple RPs.
- Cost reduction to both the user (reduction in authenticators) and the organization (reduction in information technology infrastructure).
- Minimizing data in applications as organizations do not need to collect, store, or dispose of personal information.
- Minimizing data exposed to applications, using pseudonymous identifiers and derived attribute values instead of copying account values to each application.
- Mission enablement: Organizations can focus on their mission without worrying about expending resources on identity management.

The following sections discuss the components of a federated identity architecture should an organization elect this type of model.

4.4.2. Federation Protocols and Assertions

Federation protocols allow for the conveyance of assertions, authentication attributes, and subscriber attributes across networked systems. In a federation scenario, as shown in Figure 2, the CSP provides a service known as an identity provider, or IdP. The IdP acts as a verifier for authenticators issued by the CSP. Using federation protocols, the IdP sends a message, called an assertion, about this authentication event to the RP. Assertions are verifiable statements from an IdP to an RP that represent an authentication event for a subscriber. The RP receives and uses the assertion provided by the IdP, but the RP does not verify authenticators directly.

Federation is generally used when the RP and the IdP are not a single entity or are not under common administration, though this technology can be applied within a single security domain for a variety of reasons. The RP uses the information in the assertion to identify the subscriber and make authorization decisions about their access to resources controlled by the RP.

Examples of assertions include:

- Security Assertion Markup Language (SAML) assertions are specified using a mark-up language intended for describing security assertions. They can be used by a verifier to make a statement to an RP about the identity of a claimant. SAML assertions may optionally be digitally signed.
- OpenID Connect claims are specified using JavaScript Object Notation (JSON) for describing security, and optionally, user claims. JSON user information claims may optionally be digitally signed.
- Kerberos tickets allow a ticket-granting authority to issue session keys to two authenticated parties using symmetric or asymmetric key establishment schemes.

4.4.3. Relying Parties

An RP relies on results of an authentication protocol to establish confidence in the identity or attributes of a subscriber for the purpose of conducting an online transaction. RPs may use a subscriber's federated identity (pseudonymous or non-pseudonymous), IAL, AAL, FAL, and other factors to make authorization decisions.

When using federation, the verifier is not a function of the RP. A federated RP receives an assertion from the IdP, which provides the verifier function, and the RP ensures that the assertion came from an IdP that is trusted by the RP. The RP also processes any additional information in the assertion, such as personal attributes or expiration times. The RP is the final arbiter concerning whether a specific assertion presented by a verifier meets the RP's established criteria for system access regardless of IAL, AAL, or FAL.

5. Digital Identity Risk Management

This section is normative.

This section provides details on the methodology for assessing digital identity risks for each xAL. This process augments the risk management processes for information and information system risk under NIST guidance for implementing Federal Information Security Modernization Act [FISMA] requirements.

There are 4 steps to the digital identity risk management process:

1. **Conduct Initial Impact Assessment**: In this step, organizations evaluate their user population and assess the impact of a failure of each function in the identity system (i.e., proofing, authentication, and federation) for their protected application or service against a defined set of impact categories. The outcome of this step is a documented set of impact categories and associated impact levels.
2. **Select Initial Assurance Levels**: In this step, the impact categories and impact levels are evaluated to determine the appropriate assurance levels to protect the application. The outcome of this step is an identified initial level for each applicable xAL.
3. **Tailor and Document Assurance Level Determinations**: In this step, detailed privacy, equity, usability, and threat assessments are conducted to determine the potential impact of the initially selected assurance level on the specific user population and threat environment of the application. The initial assurance level is tailored, compensating or supplemental controls are identified, and all decisions are documented. The outcome is a Digital Identity Acceptance Statement (see Sec. 5.3.4) with a defined implementable assurance level.
4. **Continuously Evaluate & Improve**: In this step, information is collected on performance of the identity system across a diverse set of factors based on organization needs and evolving threat vectors. This information is used to determine if the selected assurance level and controls are meeting mission, business, and security needs and to monitor for unintended harms that may have emerged. The outcomes of this step are performance metrics, documented and transparent processes for evaluation and redress, and ongoing improvements to the identity system as needed.

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 at any point require an organization to revisit steps in the digital identity risk management process.

Organizations **SHOULD** adapt and modify this overall approach to meet organizational processes, governance, and integration with enterprise risk management practices. At a minimum, organizations **SHALL** ensure that each step is executed and the normative mandates and outcomes of each step are completed and documented regardless of operational approach and enabling tools.

5.1. Conduct Initial Impact Assessment

The purpose of the initial impact analysis is to identify the potential adverse impacts of failures in identity proofing, authentication, and federation specific to an RP application or service, yielding an initial set of assurance levels. Assessing these areas separately allows organizations maximum flexibility in developing or acquiring a digital identity service that best enables them to successfully deliver on mission objectives.

The impact assessment includes:

- Identifying impacted entities,
- Identifying a set of impact categories for which harms will be assessed,
- Identifying potential harms for each of the impact categories,
- Identifying the levels of impact those potential harms would inflict should failures occur, and
- Assessing the impact of each type of failure (proofing, authentication, and federation) and the resulting impact level to all affected entities.

The output of this assessment is a defined impact level — High, Moderate, or Low — for each possible type of failure. This serves as the primary input to the initial assurance level selection.

5.1.1. Identify Impacted Entities

When assessing impacts, an organization needs to determine the entities that will be impacted by the application or transaction under consideration. As mentioned earlier in this guideline, it is imperative to consider the impact on different entities resulting from a failure of the digital identity system. Of particular importance is ensuring that the potential impacts to individuals are considered alongside those of the enterprise.

Accordingly, impact assessments **SHALL** include individuals using the system or application in addition to the organization itself. Additionally, organizations **SHOULD** identify other entities, such as 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 entities to which impacts will be assessed when conducting their impact analysis.

The outcome of this activity is a list of entities subject to the application or transaction under consideration for whom impacts will be assessed.

5.1.2. Identify Impact Categories and Potential Harms

Initial assurance levels for digital transactions **SHALL** be determined by assessing the potential impact of, at a minimum, each of the following categories:

- Damage to mission delivery
- Damage to trust or reputation
- Loss of sensitive information
- Damage to or loss of economic stability
- Loss of life or damage to safety, health, or environmental stability
- Noncompliance with laws, regulations, and/or contractual obligations

Organizations **SHOULD** include additional impact categories as appropriate based on their mission. Each impact category **SHALL** be documented and consistently applied across different applications assessed by the organization.

Harms are any adverse effects that would be experienced by an entity. They provide a means to more effectively understand the impact categories and how they may apply to specific entities associated with that application. Agencies **SHOULD** consider specific harms for each of the defined impact categories to better inform their impact analysis. Identification of harms for each category **SHALL** be done for each of the entities identified during "entity identification" process.

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

Damage to mission delivery:

- Harms to individuals may include the inability to access government services or benefits for which they are eligible.
- Harms to the organization may include an inability to perform current mission/business functions in a sufficiently timely manner, with sufficient confidence and/or correctness, within planned resource constraints, or an inability, or limited ability, to perform mission/business functions in the future.

Damage to trust or reputation:

- Harms to individuals may include impersonation or damage to image or reputation.
- Harms to the organization may include damage to trust relationships, image, or reputation including future, potential trust relationships.

Loss of sensitive information:

- Harms to individuals includes loss of PII or other sensitive information, which may result in secondary harms such as loss of economic stability, loss of life, physical or psychological injury, impersonation, identity theft, or persistent inconvenience.

- Harms to the organization may include loss or degradation of intellectual property or other information assets such as classified materials or controlled unclassified information (CUI).

Damage to or loss of economic stability:

- 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, and/or other financial loss.
- Harms to the organization may include costs incurred related to fraud or other criminal activity, loss of assets, devaluation, or loss of business.

Loss of life or damage to safety, health, or environmental stability:

- Harms to individuals may include death, damage to or loss of physical, mental, or emotional well-being, damage to the environment, or loss of accessible, affordable housing.
- Harms to the organization may include damage to or loss of the organization's workforce or the impact of unsafe conditions rendering the organization unable to operate or operating at reduced capacity.

Noncompliance with laws, regulations, and/or contractual obligations:

- Harms to individuals may include damage to or loss of economic stability, safety, privacy, civil liberties, equity, and/or usability due to violations of local, state, and federal laws, regulations, and/or contractual obligations.
- Harms to the organization may include financial costs, sanctions, liability, etc, due to noncompliance with applicable laws, regulations, contractual requirements, or other requirements in other binding agreements.

The outcome of this activity will be a list of impact categories and harms which will be used to assess impacts to identified entities.

5.1.3. Identify Potential Impact Levels

Initial assurance levels for digital transactions are determined by assessing the potential impact a failure would have on each of the categories from Sec. 5.1.2 using one of the following potential impact values:

1. Low potential impact: could be expected to have a limited adverse effect
2. Moderate potential impact: could be expected to have a serious adverse effect
3. High potential impact: could be expected to have a severe or catastrophic adverse effect

Note: If a failure in the identity system causes no measurable consequences for a category, there is no impact.

Each assurance level, IAL, AAL, and FAL (if accepting or asserting a federated identity) **SHALL** be evaluated separately. Ideally, any evaluation will include different viewpoints such as harm to individuals, the organization, other organizations, and the nation as applicable to successful delivery of the organization's mission. Examples of potential impacts in each of the categories include:

Damage to mission delivery:

- **Low**: at worst, slight outcome disparities exist between individuals that participate in federally funded programs and those that are eligible but unable to participate, or a limited adverse effect on organizational operations or assets, or public interests. Examples of limited adverse effects are: mission capability degradation to the extent and duration that the organization is able to perform its primary functions with noticeably reduced effectiveness, or minor damage to organizational assets or public interests.
- **Moderate**: at worst, outcome disparities are evident between individuals that participate in federally funded programs and those that are eligible but unable to participate, or a serious adverse effect on organizational operations or assets, or public interests. Examples of serious adverse effects are: significant mission capability degradation to the extent and duration that the organization is able to perform its primary functions with significantly reduced effectiveness; or significant damage to organizational assets or public interests.
- **High**: outcome disparities endure across communities, indicating a systemic pattern of exclusion, avoidance, or other barriers to participation in federally funded programs, or a severe or catastrophic adverse effect on organizational operations or assets, or public interests. Examples of severe or catastrophic effects are: severe mission capability degradation or loss to the extent and duration that the organization is unable to perform one or more of its primary functions; or major damage to organizational assets or public interests.

Damage to trust and reputation:

- **Low**: at worst, limited, short-term inconvenience, distress, or embarrassment to any party.
- **Moderate**: at worst, serious short-term or limited long-term inconvenience, distress, or damage to the standing or reputation of any party.
- **High**: severe or serious long-term inconvenience, distress, or damage to the standing or reputation of any party. This is ordinarily reserved for situations with particularly severe effects or which potentially affect many individuals.

- **Low**: at worst, a limited release of personal, U.S. government sensitive, or commercially sensitive information to unauthorized parties resulting in a loss of confidentiality with a low impact as defined in [FIPS199].
- **Moderate**: at worst, a release of personal, U.S. government sensitive, or commercially sensitive information to unauthorized parties resulting in loss of confidentiality with a moderate impact as defined in [FIPS199].
- **High**: a release of personal, U.S. government sensitive, or commercially sensitive information to unauthorized parties resulting in loss of confidentiality with a high impact as defined in [FIPS199].

Damage to or loss of economic stability:

- **Low**: at worst, an insignificant or inconsequential financial loss to any party.
- **Moderate**: at worst, a serious financial loss to any party.
- **High**: severe or catastrophic financial loss to any party.

Loss of life or damage to safety, health, or environmental stability:

- **Low**: at worst, minor injury or acute health issue that resolves itself and does not require medical, including mental health, treatment; limited risk of environmental impact in locality where program operations take place.
- **Moderate**: at worst, moderate risk of minor injury or limited risk of injury requiring medical, including mental health, treatment; or the compounding impact of multiple low impact events; moderate risk of environmental impact in locality where program operations take place.
- **High**: a risk of serious injury, trauma, or death; or the compounding impact of multiple moderate impact events; high risk of environmental impact in locality where program operations take place.

Noncompliance with laws, regulations, and/or contractual obligations:

- **Low**: at worst, a risk of civil or criminal violations of a nature that would not ordinarily be subject to enforcement efforts, or at worst, an insignificant or inconsequential organization liability.
- **Moderate**: at worst, a risk of civil or criminal violations that may be subject to enforcement efforts, or a serious organization liability.
- **High**: a risk of civil or criminal violations that are of special importance to enforcement programs, or severe or catastrophic organization liability.

5.1.4. Impact Analysis

The impact analysis helps determine the extent to which risk must be mitigated by the identity proofing, authentication, and federation processes. These determinations drive the relevant choices of applicable technologies and mitigation strategies, rather than the desire for any given technology driving risk determinations.

To determine the appropriate level of assurance of the user's asserted identity, organizations **SHALL** assess the potential risks and identify measures to minimize their impact. Organizations **SHALL** assess the risk of identity proofing, authentication, and federation failures separately to determine the required assurance level for each transaction. This process **SHALL** include consideration of potentially varying impacts of harms to different entities impacted by the digital identity system, as described in Sec. 5.1.1. Business processes, policies, and technologies may help reduce risk. Entities **SHOULD** consider the impact of specific modes of failures related to identity proofing, authentication, and federation this includes, but may not be limited to:

Identity Proofing:

- The impact of providing a service to the wrong subject (e.g., an attacker successfully proofs as someone else).
- The impact of not providing service to an eligible subject due to barriers, including biases, faced by the subject throughout the process of identity proofing.
- The impact of excessive information collection and retention to support identity proofing processes.

Authentication:

- The impact of authenticating the wrong subject (e.g., an attacker who compromises or steals an authenticator).
- The impact of failing to authenticate the correct subject due to barriers, including biases, faced by the subject in presenting their authenticator.

Federation:

- The impact of the wrong subject successfully accessing an application, system, or data (e.g., compromising or replaying an assertion).
- The impact of releasing subscriber attributes to the wrong application or system.

Using a worksheet similar to Table 1 can assist organizations with compiling the information gathered in order to complete the analysis. This kind of analysis would be done for each type of potential failure for identity proofing, authentication, and federation to determine the overall risks to entities interacting with the digital identity system.

Table 1. Impact Categories

Impact Categories	Harm to Individuals	Harm to the Organization	(Other harm categories)	Combined Impact Level
Damage to mission delivery	L / M / H	L / M / H	L / M / H	
Damage to trust or reputation	L / M / H	L / M / H	L / M / H	
Loss of sensitive information	L / M / H	L / M / H	L / M / H	
Damage to or loss of economic stability	L / M / H	L / M / H	L / M / H	
Loss of life or damage to safety, health, or environmental stability	L / M / H	L / M / H	L / M / H	
Noncompliance with laws, regulations, and/or contractual obligations	L / M / H	L / M / H	L / M / H	

The output of this step is a defined impact level for failures of identity proofing, authentication, and federation which serve as the primary input to the initial assurance level selection.

5.2. Select Initial Assurance Levels

The impact analysis serves as a primary input to the process of selecting initial assurance levels for identity proofing, authentication and federation. The assurance levels may differ across these areas based on the analysis of the potential impact of failures in each area. The purpose of these initial assurance levels is to identify baseline digital identity controls and processes, reflected in the requirements and guidelines in the companion volumes of [SP800-63A], [SP800-63B], and [SP800-63C], which will be assessed and tailored based on mission need, cybersecurity risk, and other potential impacts to the organization and users of the digital identity systems.

5.2.1. Assurance Levels

An organization RP **SHALL** select, based on cybersecurity risk and mission needs, the following individual initial assurance levels:

- **IAL**: The robustness of the identity proofing process to confidently determine the identity of an individual. IAL is selected to mitigate potential identity proofing failures.
- **AAL**: The robustness of the authentication process itself, and the binding between an authenticator and a specific individual's identifier. AAL is selected to mitigate potential authentication failures.
- **FAL**: The robustness of the federation process used to communicate authentication and attribute information (if applicable) to an RP from an IdP. FAL is optional as not all digital systems will leverage federated identity architectures. FAL is selected to mitigate potential federation failures.

5.2.2. xAL Descriptions

A summary of each of the identity, authenticator, and federation assurance levels is provided below.

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

5.2.2.1. Identity Assurance Level

IAL1: IAL1 requires validation of identifying attributes against authoritative or credible sources and use of basic processes to verify the claimed identity of the applicant.

IAL2: IAL2 requires identifying attributes to be supported by strong evidence and validated against authoritative or credible sources and use of processes to verify the claimed identity of the applicant.

IAL3: IAL3 requires identifying attributes to be verified by an authorized CSP representative through examination of physical documentation using an interactive process with a CSP representative.

5.2.2.2. Authentication Assurance Level

AAL1: AAL1 provides some assurance that the claimant controls an authenticator registered to the subscriber. AAL1 requires single-factor authentication using a wide range of available authentication technologies. Successful authentication requires that the claimant prove possession and control of the authenticator through a secure authentication protocol.

AAL2: AAL2 provides high confidence that the claimant controls authenticator registered to the subscriber. Proof of possession and control of two different authentication factors is

required through a secure authentication protocol. Approved cryptographic techniques are required at AAL2 and above.

AAL3: AAL3 provides very high confidence that the claimant controls authenticator registered to the subscriber. Authentication at AAL3 is based on proof of possession of a key through a cryptographic authentication protocol capable of resisting phishing attacks.

5.2.2.3. Federation Assurance Level

FAL1: FAL1 allows for the subscriber to log into the RP using an assertion from the IdP that can be verified by the RP as coming from the IdP and targeted for a specific RP. The assertion is protected from modification or construction by an attacker. The trust agreement and registration between the IdP and RP can happen dynamically.

FAL2: FAL2 adds the requirement that the assertion be robust against injection at the RP. One means of this is to have the assertion presented directly to the RP from the IdP instead of passing through an intermediary like a browser. The trust agreement between the IdP and RP cannot happen dynamically, but dynamic registration of the specific IdP and RP can occur at runtime.

FAL3: FAL3 adds the requirement that the subscriber authenticate directly to the RP using a bound authenticator along with presenting the authentication assertion. The presence of this additional authenticator provides a very high assurance to the RP that the party accessing the RP is the party identified in the assertion. The trust agreement and registration cannot be dynamic.

5.2.3. Initial Assurance Level Selection

The identification and assessment of the potential impacts of failures in identity proofing, authentication, and federation processes informs the organization's digital identity risk management process and the initial selection of assurance levels for those areas. These initial selections are primarily based on cybersecurity risk, but will be tailored, based on mission needs and other potential impacts to the organization, users, and mission partners.

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

5.2.3.1. Selecting Initial IAL

The IAL reflects the level of assurance that an applicant holds the claimed real-life identity. Organizations **SHALL** use a risk-based approach to select the most appropriate identity proofing requirements for their RP application. The impact analysis described in Sec. 5.3.1 informs the selection of the initial IAL selection. This initial selection **SHALL** be tailored, as described in Sec. 5.3, based on mission needs, risk tolerance, and potential impacts to privacy, equity, and usability, before making a final IAL determination.

The IAL selection does not mean the RP application owner will need to perform the proofing themselves since identity proofing is the function of the CSP.

Not all RP applications will require identity proofing. If the RP application does not require any personal information to execute any digital transactions, the system can operate without identity proofing users of the RP application. If personal information is needed, the RP needs to determine if validated and verified attributes are required or if self-asserted attributes are acceptable. If there are insignificant potential harms from accepting self-asserted attributes, the system may also be able to operate without identity proofing users. In such cases, the identity proofing processes described in [SP800-63A] are not applicable to the system.

If an organization determines that identity proofing is necessary, the initial IAL **SHALL** be assessed based on the potential impacts of identity proofing failures. As described in Sec. 5.1, potential impacts **SHALL** be considered from the perspective of the organization, individuals, other organizations, and the nation, for harms incurred through the use or operation of the RP application. While the organization may not be negatively impacted, the user could be significantly harmed, as could individuals whose privacy or other rights have been violated by the business practices of a service provider. Organizations **SHOULD** consider the worst-case when identifying the overall impact level of the RP application, but may use risk management processes to tailor their initial selection when there are differing impacts.

When assessing the overall impact level of the RP application, the organization **SHOULD** consider impacts to mission delivery separately from other impact categories. Potential failures in the identity proofing process that could lead to harms in mission delivery should be assessed by the organization to determine if the associated impacts would be mitigated or exacerbated by the implementation of more rigorous identity proofing processes. As such, the organization **MAY** exclude the mission delivery category when initially identifying the overall impact level of the RP application, as these impacts will need to be considered in the tailoring process.

The overall impact level assessed by the organization leads to a preliminary selection of the IAL from which further tailoring may be done:

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

The preliminary selection assumes that higher potential impacts of failures in the identity proofing process should be mitigated by higher assurance processes. While this is often the case, organizations should consider the specific failures, impact categories, and impacted entities identified as part of the impact analysis to determine if additional tailoring is warranted. For example, if a failure to enroll a legitimate applicant could lead

to excessive harm, organizations should assess whether lower-assurance identity proofing processes would be appropriate.

The result of this process, including any additional tailoring, is the initial assessment of the IAL, which will be assessed against additional potential impacts as described in Sec. 5.3.

5.2.3.2. Selecting Initial AAL

The AAL reflects the level of assurance from the authentication process that the claimant is who they claim to be. Organizations **SHALL** use a risk-based approach to select the most appropriate authentication requirements for their RP application. The impact analysis described in Sec. 5.1.3 informs the selection of the initial AAL selection. This initial selection **SHALL** be tailored, as described in Sec. 5.3, based on mission needs, risk tolerance, and potential impacts to privacy, equity, and usability, before making a final AAL determination.

The AAL selection does not mean the RP application owner will need to issue authenticators themselves.

The initial AAL **SHALL** be assessed based on the potential impacts of authentication failures. As described in Sec. 5.1, potential impacts **SHALL** be considered from the perspective of the organization, individuals, other organizations, and the nation, for harms incurred through the use or operation of the RP application, as the level of harm from a failure could vary significantly across these entities. Organizations **SHOULD** consider the worst-case when identifying the overall impact level of the RP application, but may use risk management processes to tailor their initial selection when there are differing impacts.

When assessing the overall impact level of the RP application, the organization **SHOULD** consider impacts to mission delivery separately from other impact categories. Potential failures in the authentication process that could lead to harms in mission delivery should be assessed by the organization to determine if the associated impacts would be mitigated or exacerbated by the implementation of more rigorous authentication controls. As such, the organization **MAY** exclude the mission delivery category when initially identifying the overall impact level of the RP application, as these impacts will need to be considered in the tailoring process.

The overall impact level assessed by the organization leads to a preliminary selection of the AAL from which further tailoring may be done:

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

The preliminary selection assumes that higher potential impacts of failures in the authentication process should be mitigated by higher assurance processes. While this is often the case, organizations should consider the specific failures, impact categories, and impacted entities identified as part of the impact analysis to determine if additional tailoring is warranted. Further, organizations should consider legal, regulatory, or policy requirements that govern digital services. For example, the terms of [EO13681] requiring "that all organizations making personal data accessible to citizens through digital applications require the use of multiple factors of authentication," which would drive the selection of AAL2 or AAL3.

The result of this process, including any additional tailoring, is the initial assessment of the AAL, which will be as assessed against additional potential impacts as described in Sec. 5.3.

5.2.3.3. Selecting Initial FAL

The FAL reflects the level of assurance in identity assertions that convey the results of authentication processes and relevant identity information to RP systems. Organizations **SHALL** use a risk-based approach to select the most appropriate federation requirements for their RP application. The impact analysis described in Sec. 5.3.1 informs the selection of the initial FAL selection. This initial selection **SHALL** be tailored, as described in Sec. 5.3, based on mission needs, risk tolerance, and potential impacts to privacy, equity, and usability, before making a final FAL determination.

The initial FAL **SHALL** be assessed based on the potential impacts of failures in the presentation or acceptance of assertions in federated identity architectures. Examples of compromise include use of assertion replay to impersonate a valid user or leakage of assertion information through the browser. As described in Sec. 5.1, potential impacts **SHALL** be considered from the perspective of the organization, individuals, other organizations, and the nation, for harms incurred through the use or operation of the RP application, as the level of harm from a failure could vary significantly across these entities. Organizations **SHOULD** consider the worst-case when identifying the overall impact level of the RP application, but may use risk management processes to tailor their initial selection when there are differing impacts.

When assessing the overall impact level of the RP application, the organization **SHOULD** consider impacts to mission delivery separately from other impact categories. Potential failures in federated architectures that could lead to harms in mission delivery **MAY** be assessed by the organization to determine if the associated impacts would be mitigated or exacerbated by the implementation of more rigorous controls by identity providers. As such, the organization may exclude the mission delivery impact category when initially identifying the overall impact level of the RP application, as these impacts will need to be considered in the tailoring process.

The overall impact level assessed by the organization leads to a preliminary selection of the FAL from which further tailoring may be done:

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

The preliminary selection assumes that higher potential impacts of failures in federated identity architectures should be mitigated by higher assurance processes. While this is often the case, organizations should consider the specific failures, impact categories, and impacted entities identified as part of the impact analysis to determine if additional tailoring is warranted.

The result of this process, including any additional tailoring, is the initial assessment of the FAL, which will be as assessed against additional potential impacts as described in Sec. 5.3.

5.3. Tailor and Document Assurance Levels

Tailoring provides a process to modify an initially assessed assurance level or implement compensating controls based on ongoing detailed impact and risk assessments. Organizations **SHOULD** implement the assessed assurance level as defined in these guidelines. However, these guidelines provide flexibility to allow for organizations to meet specific mission needs and address unique risk appetites and considerations. Therefore, organizations **SHALL** establish and document an xAL tailoring process. At a minimum this process:

1. **SHALL** include a structured governance approach to allow for decision-making and conflict resolution.
2. **SHALL** document all decisions in the tailoring process, including the assessed xALs, modified xALs, and compensating controls in the Digital Identity Acceptance Statement (see Sec. 5.3.4).
3. **SHALL** justify and document all risk-based decisions or modifications to the initially assessed xALs in the Digital Identity Acceptance Statement (see Sec. 5.3.4).
4. **SHOULD** establish a cross-functional capability to support subject matter analysis of xAL selection impacts in the tailoring process.
5. **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 to balance risks and impacts in the furtherance of protecting systems, data, and services in a manner that enables mission success while supporting equity, privacy, and usability for individuals.

5.3.1. Assess Privacy, Equity, Usability and Threats

When selecting and tailoring assurance levels for specific applications, it is critical that insights and inputs to the process extend beyond an initial, static impact assessment. When transitioning from an initial assurance level to the final xAL selection and implementation, organizations **SHALL** conduct detailed assessments of the controls defined at the assurance level to determine potential impacts in their operational environment. At a minimum, organizations **SHALL** assess impacts related to the following areas:

- **Privacy** – to determine 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.
- **Equity** – to determine whether implementation of controls may create or maintain inequities across demographics or user groups.
- **Usability** – to determine whether implementation of the selected controls will result in challenges to end-user experience.
- **Threat** – to determine whether the defined assurance level will address specific threats based on environment, threat actors, and known tactics, techniques, and procedures (TTPs).

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. 5.3.2 and Sec. 5.3.3.

5.3.2. Identify Compensating Controls

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

Organizations **SHOULD** implement their identity services per the requirements in these guidelines for their tailored assurance level. However, where organizations are unable to implement a specific control associated with their baseline or tailored assurance level, they **MAY** select to implement a compensating control. This control **MAY** be a modification to a digital identity process as defined in these guidelines, but **MAY** also be applied elsewhere in an application, 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 availability of evidence in the end-user population.

Where compensating controls are implemented, organizations **SHALL** demonstrate comparability of a chosen alternative or document residual risk incurred by deviating from normative requirements. Organizations **SHALL** implement procedures to document both the justification for any departure from normative requirements and detail the compensating controls employed. The inclusion of compensating controls does not imply that an organization must tailor to a lower xAL. The process of tailoring allows for agencies and service providers to make risk-based decisions in how they implement their xALs and provides a mechanism for documenting and communicating decisions through the Digital Identity Acceptance Statement described in Sec. 5.3.4.

5.3.3. Identify Supplemental Controls

Supplemental controls are those that may be added, in addition to those specified in the organizations tailored assurance level, in order to address specific threats or attacks. Organizations **SHOULD** identify and implement supplemental controls where they identify threats that may not be addressed in baseline controls. For example:

- 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 to complete the proofing process.
- 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.

Where organizations implement supplemental controls, these **SHALL** be assessed for impacts based on the same factors used to tailor the organization's assurance level. Supplemental controls **SHALL** be documented.

5.3.4. Document Results - The Digital Identity Acceptance Statement

The Digital Identity Acceptance Statement documents the results of the digital identity risk management process. This includes the Impact Assessment, Initial Assurance Level Selection, and Tailoring process.

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

1. Initial Impact Assessment Results
2. Initially assessed xAL,
3. Tailored xAL and rationale, if tailored xAL differs from initially assessed xAL,

4. All compensating controls and their comparability or residual risk associated with compensating controls
5. All supplemental controls

Federal agencies **SHOULD** include this information in the system authorization package described in [SP800-37].

5.4. Continuously Evaluate and Improve

Threat actors adapt, user expectations and needs shift, and missions evolve. As such, risk assessments and identity solutions are not to be set and forgotten. To maintain pace with the constantly shifting environment in which they operate, organizations **SHOULD** implement a continuous evaluation and improvement program that leverages input from people interacting with the identity system. These programs **SHOULD** consider feedback from application performance metrics, threat intelligence, fraud analytics, assessments of equity impacts, privacy impact analysis, and user inputs.

5.5. Cyber, Fraud, and Identity Program Integrity

Typically, identity solutions are the front door for a critical business or service function. Accordingly, they should not operate in a vacuum. Close coordination of identity functions with cybersecurity teams, threat intelligence teams, and program integrity teams 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 receive indication of new TTPs that may impact identity proofing, authentication, and federation processes. Organizations **SHOULD** establish consistent mechanisms for the exchange of information between critical security and fraud stakeholders.

Where supporting service providers, such as CSPs, are external, this may be complicated, but **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.

References

This section is informative.

General References

[A-130] OMB Circular A-130, *Managing Federal Information as a Strategic Resource*, July 28, 2016, available at: https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/OMB/circulars/a130/a130revised.pdf.

[EO13681] Executive Order 13681, *Improving the Security of Consumer Financial Transactions*, October 17, 2014, available at: https://www.federalregister.gov/d/2014-25439.

[EO13985] Executive Order 13985, *Advancing Racial Equity and Support for Underserved Communities Through the Federal Government*, January 25, 2021, available at: https://www.federalregister.gov/documents/2021/01/25/2021-01753/advancing-racial-equity-and-support-for-underserved-communities-through-the-federal-government

[FISMA] *Federal Information Security Modernization Act of 2014, 44 U.S.C. § 3551 et seq., Public Law (P.L.) 113-283*, available at: https://www.congress.gov/bill/113th-congress/senate-bill/2521.

[M-03-22] OMB Memorandum M-03-22, *OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002*, September 26, 2003, available at: https://georgewbush-whitehouse.archives.gov/omb/memoranda/m03-22.html.

[NISTIR8062] NIST Internal Report 8062, *An Introduction to Privacy Engineering and Risk Management in Federal Systems*, January 2017, available at: https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8062.pdf.

[NISTRMF] *Risk Management Framework Overview*, available at https://csrc.nist.gov/groups/SMA/fisma/framework.html.

[NISTPF] *NIST Privacy Framework*, available at https://www.nist.gov/privacy-framework/privacy-framework.

[PrivacyAct] *The Privacy Act of 1974*, available at https://www.govinfo.gov/content/pkg/USCODE-2020-title5/pdf/USCODE-2020-title5-partI-chap5-subchapII-sec552a.pdf

[SORN] United States Office of Personnel Management (OPM), *System of Records Notice (SORN) Guide*, April 22, 2010, available at: https://www.opm.gov/information-management/privacy-policy/privacy-references/sornguide.pdf

Standards

[BCP195] Sheffer, Y., Holz, R., and P. Saint-Andre, *Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)*, BCP 195, RFC 7525,DOI 10.17487/RFC7525, May 2015, available at: https://doi.org/10.17487/RFC7525.

[ISO9241-11] International Standards Organization, ISO/IEC 9241-11 *Ergonomic requirements for office work with visual display terminals (VDTs) — Part 11: Guidance on usability*, March 1998, available at: https://www.iso.org/standard/16883.html.

[OIDC] Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and C. Mortimore, *OpenID Connect Core 1.0 incorporating errata set 1*, December, 2014. Available at: https://openid.net/specs/openid-connect-core-1_0.html.

[RFC5246] Dierks, T. and E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.2*, RFC 5246, DOI 10.17487/RFC5246, August 2008, https://www.rfc-editor.org/info/rfc5246.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*, RFC 5280, DOI 10.17487/RFC5280, May 2008, https://www.rfc-editor.org/info/rfc5280.

NIST Special Publications

NIST 800 Series Special Publications are available at: < https://csrc.nist.gov/publications/sp800l>. The following publications may be of particular interest to those implementing systems of applications requiring digital authentication.

[SP800-30] NIST Special Publication 800-30 Revision 1, *Guide for Conducting Risk Assessments*, September 2012, available at: https://doi.org/10.6028/NIST.SP.800-30r1.

[SP800-37] NIST Special Publication 800-37 Revision 2, *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy*, December 2018, available at: https://csrc.nist.gov/publications/detail/sp/800-37/rev-2/final.

[SP800-52] NIST Special Publication 800-52 Revision 2, *Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations*, August 2019, available at: https://csrc.nist.gov/publications/detail/sp/800-52/rev-2/final.

[SP800-53] NIST Special Publication 800-53 Revision 5, *Security and Privacy Controls for Information Systems and Organizations*, September 2020 (incudes updates as of Dec. 10, 2020), available at: https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final.

[SP800-53A] NIST Special Publication 800-53A Revision 5, *Assessing Security and Privacy Controls in Information Systems and Organizations*, January 2022, available at:

https://csrc.nist.gov/publications/detail/sp/800-53a/rev-5/final.

[SP800-57Part1] NIST Special Publication 800-57 Part 1, Revision 5, *Recommendation for Key Management, Part 1: General*, May 2020, https://dx.doi.org/10.6028/NIST.SP.800-57pt1r5.

[SP800-63A] NIST Special Publication 800-63B-4, *Digital Identity Guidelines: Enrollment and Identity Proofing*, December 2022, https://doi.org/10.6028/NIST.SP.800-63a-4.ipd.

[SP800-63B] NIST Special Publication 800-63B-4, *Digital Identity Guidelines: Authentication and Lifecycle Management*, November 2022, https://doi.org/10.6028/NIST.SP.800-63b-4.ipd.

[SP800-63C] NIST Special Publication 800-63C-4, *Digital Identity Guidelines: Assertions and Federation*, November 2022, https://doi.org/10.6028/NIST.SP.800-63c-4.ipd.

Federal Information Processing Standards

[FIPS199] Federal Information Processing Standard 199, *Standards for Security Categorization of Federal Information and Information Systems*, February 2004, available at: https://doi.org/10.6028/NIST.FIPS.199.

[FIPS201] Federal Information Processing Standard Publication 201-3, *Personal Identity Verification (PIV) of Federal Employees and Contractors*, January 2022, available at: https://csrc.nist.gov/publications/detail/fips/201/3/final.

Appendix A. Definitions and Abbreviations

This section is informative.

A.1. Definitions

A wide variety of terms is used in the realm of authentication. While many terms' 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.

Access
To make contact with one or more discrete functions of an online, digital service.

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*. Since all multi-factor authenticators are physical authenticators, activation factors are either *memorized secrets* or *biometric* factors.

Active Attack
An attack on the authentication protocol where the attacker transmits data to the claimant, Credential Service Provider (CSP), verifier, or Relying Party (RP). Examples of active attacks include attacker-in-the-middle (AitM), impersonation, and session hijacking.

Address of Record
The validated and verified location (physical or digital) where a subscriber can receive communications using approved mechanisms.

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 enrollment and identity proofing.

Approved Cryptography

Federal Information Processing Standard (FIPS)-approved or NIST recommended. An algorithm or technique that is either 1) specified in a FIPS or NIST Recommendation, or 2) adopted in a FIPS or NIST Recommendation.

Assertion
A statement from a verifier to an RP that contains information about a subscriber. Assertions may also contain verified attributes.

Assertion Reference
A data object, created in conjunction with an assertion, that identifies the verifier and includes a pointer to the full assertion held by the verifier.

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.

Attack
An unauthorized entity's attempt to fool a verifier or RP into believing that the unauthorized individual in question is the subscriber.

Attacker
A party, including an insider, who acts with malicious intent to compromise a system.

Attacker-in-the-Middle Attack (AitM)
An attack in which an attacker is positioned between two communicating parties in order to intercept and/or alter data traveling between them. In the context of authentication, the attacker would be positioned between claimant and verifier, between registrant and CSP during enrollment, or between subscriber and CSP during authenticator binding.

Attribute
A quality or characteristic ascribed to someone or something.

Attribute API
An API that provides *attribute values*, *derived attribute values*, and related information about one or more subscribers. Access to these APIs are often granted to RPs in the context of an *identity API* (for a single subscriber) or a *provisioning API* (for multiple subscribers). This is distinct from an *attribute verification API* which is used to verify attribute values for a CSP during the identity proofing process.

Attribute Bundle

A packaged set of attributes, usually contained within an assertion. Attribute bundles offer RPs a simple way to retrieve the most relevant attributes they need from IdPs. OpenID Connect scopes [OIDC] are an implementation of attribute bundles.

Attribute Provider
A service that provides a subscriber's attributes without asserting that the subscriber is present to the RP. An *Identity Provider (IdP)* is one type of attribute provider used in federated scenarios. Attribute providers often make these attributes available by means of an *attribute API*.

Attribute Value
A complete statement asserting a property of a subscriber, independent of format. For example, for the attribute "birthday," a value could be "12/1/1980" or "December 1, 1980."

Attribute Verification API
An API that provides verification of *attribute values* for use during an *identity proofing* process. This API accepts attribute values as input queries and returns whether or not the attribute values can be verified. This is distinct from an *attribute API* which is used to convey attributes to an 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 provide confidentiality and MitM protection and are frequently used in the user authentication process. Transport Layer Security (TLS) [BCP195] is an example of an authenticated protected channel where 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 leading to a trusted root rather than individually with each server.

Authentication
The process of determining the validity of one or more authenticators used to claim a digital identity. Authentication establishes that a subject attempting to access a digital service is in control of the technologies used to authenticate.

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 including a process requiring user intervention in the authentication flow. Some authenticators (e.g., OTP devices) 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 of 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 the claimant possesses and controls (typically a cryptographic module or password) that is used to authenticate the claimant's identity. In some previous editions of SP 800-63, this was referred to as a *token*.

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

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 are dependent upon 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. Some authenticator types provide one authentication factor, others provide two.

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 can confirm the validity of the identity 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 (allowing for standard protocol-level proxies), without using redirects through an intermediary such as a browser. This can be accomplished using HTTP requests and responses.

Bearer Assertion
The assertion a party presents as proof of identity, where possession of the assertion itself is sufficient proof of identity for the assertion bearer.

Binding
An association between a subscriber identity and an authenticator or given subscriber session.

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

Biometric Sample
An analog or digital representation of biometric characteristics prior to biometric feature extraction. An example is a record containing a fingerprint image.

Biometrics
Automated recognition of individuals based on their biological and behavioral characteristics.

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 where the verifier sends the claimant a challenge (usually a random value or nonce) that the claimant combines with a secret (such as 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 (such as 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.

For example, a person with a foreign passport living in the U.S. will need to give an address when going through the identity proofing process. This address would not be an "address of record" but a "claimed address."

Claimed Identity
An applicant's declaration of unvalidated and unverified personal attributes.

Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA)
An interactive feature added to web forms to distinguish whether a human or automated agent is using the form. Typically, it requires entering text corresponding to a distorted image or a sound stream.

Core Attributes
The set of identity attributes 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 or identifiers - 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 a one or more digital certificates that are often contained, along with their associated private keys, in an authenticator.

Credential Service Provider (CSP)
A trusted entity whose functions include identity proofing applicants to the identity service and the registration of authenticators to subscriber accounts. A CSP may be an independent third party.

Cross-site Request Forgery (CSRF)
An attack in which a subscriber currently authenticated to an RP and connected through a secure session browses to 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, merely by viewing a malicious link in a webmail message 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 and can therefore compromise the confidentiality and integrity of data transfers between the website and client. 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, via the endpoint, with a verifier.

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 NIST [SP800-57Part1].

See also Asymmetric Keys, Symmetric Key.

Cryptographic Module
A set of hardware, software, and/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 asserting a property of a subscriber without necessarily containing identity information, 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 presented digitally to a system. In previous editions of SP 800-63, this was referred to as *Electronic Authentication*.

Digital Signature
An asymmetric key operation where 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, but not confidentiality protection.

Disassociability
Per [NISTIR8062]: The processing of PII or events without association to individuals or devices beyond the operational requirements of the system.

Eavesdropping Attack
An attack in which an attacker listens passively to the authentication protocol to capture information that can be used in a subsequent active attack to masquerade as the claimant.

Electronic Authentication (E-Authentication)
See *Digital Authentication*.

Enrollment
The process through which an applicant applies to become a subscriber of a CSP and the CSP validates the applicant's identity.

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

Equity
Per EO 13985, Equity refers to 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.

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 the conveyance of identity and authentication information across a set of networked systems.

Federation Assurance Level (FAL)
A category describing the assertion protocol used by the federation to communicate authentication and attribute information (if applicable) to an RP.

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 redirects through an intermediary such as a browser. This is normally accomplished by appending HTTP query parameters to URLs hosted by the receiver of the message.

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; and
2. Collision resistant - It is computationally infeasible to find any two distinct inputs that map to the same output.

Identity
An attribute or set of attributes that uniquely describe a subject within a given context.

Identity API
An *attribute API* accessed by an RP for accessing attributes of a specific subscriber. Access to the identity API is generally granted as part of a federation authentication process and limited to the information for a single, specific subscriber.

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

Identity Evidence
Information or documentation provided by the applicant to support the claimed identity. Identity evidence may be physical (e.g. a driver license) or digital (e.g. an assertion generated and issued by a CSP based on the applicant successfully authenticating to the CSP).

Identity Proofing
The process by which a CSP collects, validates, and verifies information about a person.

Identity Provider (IdP)
When using federation, this is the party that manages the subscriber's primary authenticators and issues assertions derived from the subscriber account.

Identity Resolution
The process of collecting information about an applicant in order to uniquely distinguish an individual within the context of the population the CSP serves.

Issuing Source
An authority responsible for the generation of data, digital evidence (such as assertions), or physical documents that can be used as identity evidence.

Kerberos
A widely used authentication protocol developed at MIT. In "classic" Kerberos, users share a secret password with a Key Distribution Center (KDC). The user (Alice) who wishes to communicate with another user (Bob) authenticates to the KDC and the KDC furnishes a "ticket" to use to authenticate with Bob.

See [SP800-63C] Sec. 11.2 for more information.

Knowledge-Based Verification (KBV)
Identity verification method based on knowledge of private information associated with the claimed identity. This is often referred to as knowledge-based authentication (KBA) or knowledge-based proofing (KBP).

Manageability
Per NISTIR 8062: Providing the capability for granular administration of personally identifiable information, including alteration, deletion, and selective disclosure.

Memorized Secret
A type of authenticator comprised of a character string intended to be memorized or memorable by the subscriber, permitting the subscriber to demonstrate *something they know* as part of an authentication process.

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

A characteristic of an authentication system or an authenticator that requires more than one distinct authentication factor for successful authentication. MFA can be performed using a single authenticator that provides more than one factor or by a combination of authenticators that provide different factors.

The three authentication factors are something you know, something you have, and something you are.

Multi-Factor Authentication (MFA)
An authentication system that requires more than one distinct authentication factor for successful authentication. Multi-factor authentication can be performed using a multi-factor authenticator or by a combination of authenticators that provide different factors.

The three authentication factors are *something you know*, *something you have*, and *something you are*.

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.

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, attacker-in-the-middle, session hijacking) and passive (e.g., eavesdropping) attack 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.

Offline Attack
An attack where 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 where 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.

Pairwise Pseudonymous Identifier
An opaque unguessable subscriber identifier generated by a CSP for use at a specific individual RP. This identifier is only known to and only used by one CSP-RP pair.

Passive Attack
An attack against an authentication protocol where the attacker intercepts data traveling along the network between the claimant and verifier, but does not alter the data (i.e., eavesdropping).

Passphrase
A passphrase is a memorized secret consisting 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
See *memorized secret*.

Personal Data
See *Personally Identifiable Information*.

Personal Identification Number (PIN)
A memorized secret typically consisting of only decimal digits.

Personal Information
See *Personally Identifiable Information*.

Personally Identifiable Information (PII)
As defined by OMB Circular A-130, PII is 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.

Personally Identifiable Information Processing
An operation or set of operations performed upon personally identifiable information that can include, but is not limited to, the collection, retention, logging, generation,

transformation, use, disclosure, transfer, and disposal of personally identifiable information.

Pharming
An attack in which an attacker corrupts an infrastructure service such as DNS (Domain Name System) causing 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.

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

Private Key
The secret part of an asymmetric key pair that is used to digitally sign or decrypt data.

Processing
Per [NISTIR8062]: 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.

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 measurement and analysis of anatomical characteristics or involuntary or voluntary reactions, in order to

2054 determine if a biometric sample is being captured from a living subject present at the
2055 point of capture.

2056 **Protected Session**
2057 A session wherein messages between two participants are encrypted and integrity is
2058 protected using a set of shared secrets called session keys.

2059 A protected session is said to be *authenticated* if, during the session, one participant
2060 proves possession of one or more authenticators in addition to the session keys, and if the
2061 other party can verify the identity associated with the authenticator(s). If both participants
2062 are authenticated, the protected session is said to be *mutually authenticated*.

2063 **Provisioning API**
2064 An *attribute API* that allows an RP to access to attributes for multiple subscribers for
2065 the purposes of provisioning RP subscriber accounts. Access to a provisioning API is
2066 generally granted to the RP outside of a specific federated authentication transaction.

2067 **Pseudonym**
2068 A name other than a legal name.

2069 **Pseudonymity**
2070 The use of a pseudonym to identify a subject.

2071 **Pseudonymous Identifier**
2072 A meaningless but unique number that does not allow the RP to infer anything regarding
2073 the subscriber but which does permit the RP to associate multiple interactions with the
2074 subscriber's claimed identity.

2075 **Public Key**
2076 The public part of an asymmetric key pair that is used to verify signatures or encrypt data.

2077 **Public Key Certificate**
2078 A digital document issued and digitally signed by the private key of a certificate authority
2079 that binds an identifier to a subscriber to a public key. The certificate indicates that the
2080 subscriber identified in the certificate has sole control and access to the private key. See
2081 also [RFC5280].

2082 **Public Key Infrastructure (PKI)**
2083 A set of policies, processes, server platforms, software, and workstations used for the
2084 purpose of administering certificates and public-private key pairs, including the ability to
2085 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
(*In the context of remote authentication or remote transaction*) An information exchange between network-connected devices where the information cannot be reliably protected end-to-end by a single organization's security controls.

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 use of an authenticator output that is valid only for a specific authentication.

Resolution
See *Identity Resolution*.

Restricted
An authenticator type, class, or instantiation having additional risk of false acceptance associated with its use that is therefore subject to additional requirements.

Risk Assessment
The process of identifying, estimating, and prioritizing risks to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, and other organizations, resulting from the operation of a system. It is part of risk management, incorporates threat and vulnerability analyses, and considers mitigations provided by security controls planned or in place. Synonymous with risk analysis.

Risk Management
The program and supporting processes to manage information security risk to organizational operations (including mission, functions, image, reputation), organizational

assets, individuals, 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.

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)*.

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 use of a session secret that the subscriber's software (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 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 (something you know, something you have, or something you are) for successful authentication.

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.

Software Statement

Software Statement
A list of attributes describing a piece of software that is cryptographically signed by an authority. Software statements are used most commonly with RPs in a federated scenario.

Special Publication (SP)
A type of publication issued by NIST. Specifically, the SP 800-series reports on the Information Technology Laboratory's research, guidelines, and outreach efforts in computer security, and its collaborative activities with industry, government, and academic organizations.

Subject
A person, organization, device, hardware, network, software, or service.

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.

Supervised Remote Identity Proofing
A remote identity proofing process that employs physical, technical and procedural measures that provide sufficient confidence that the remote session can be considered equivalent to a physical, in-person identity proofing process.

Symmetric Key
A cryptographic key used to perform both the cryptographic operation and its inverse. For example, to encrypt and decrypt or create a message authentication code and to verify the code.

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.

Token
See Authenticator.

Transaction
A discrete event between a user and a system that supports a business or programmatic purpose. A government digital system may have multiple categories or types of

transactions, which may require separate analysis within the overall digital identity risk assessment.

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. NIST 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 Anchor
A public or symmetric key that is trusted because it is directly built 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 limiting its scope.

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 presented evidence is authentic, current, and issued from an acceptable source; attribute validation is the process or act of confirming the a set of attributes are accurate and associated with a real-life identity.

Verification
The process or act of confirming that the applicant holds the claimed identity represented by the validated identity attributes and associated evidence. In NIST SP 800-63, the term "verification" is 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 consisting entirely of bits with the value zero so that the data is destroyed and not recoverable. This is often contrasted with deletion methods that merely destroy reference 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.

A.2. Abbreviations

Selected abbreviations in these guidelines are defined below.

ABAC
Attribute Based Access Control

AAL
Authentication Assurance Level

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

CSP
Credential Service Provider

CSRF
Cross-site Request Forgery

XSS
Cross-site Scripting

DNS
Domain Name System

EO
Executive Order

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

FIPS
Federal Information Processing Standard

FISMA
Federal Information Security Modernization Act

1:1 Comparison
One-to-one Comparison

IAL
Identity Assurance Level

IdP
Identity Provider

IoT
Internet of Things

ISO/IEC
International Organization for Standardization/International Electrotechnical Commission

JOSE
JSON Object Signing and Encryption

JSON
JavaScript Object Notation

JWT
JSON Web Token

KBA
Knowledge-Based Authentication

KBV
Knowledge-Based Verification

KDC
Key Distribution Center

LOA
Level of Assurance

MAC
Message Authentication Code

MFA
Multi-Factor Authentication

N/A
Not Applicable

NARA
National Archives and Records Administration

OMB
Office of Management and Budget

OTP
One-Time Password

PAD
Presentation Attack Detection

PIA
Privacy Impact Assessment

PII
Personally Identifiable Information

PIN
Personal Identification Number

PKI
Public Key Infrastructure

PL
Public Law

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

SMS
Short Message Service

SP
Special Publication

SORN
System of Records Notice

TEE
Trusted Execution Environment

TGS
Ticket Granting Server

TGT
Ticket Granting Ticket

TLS
Transport Layer Security

TPM
Trusted Platform Module

TTP
Tactics, Techniques, and Procedures

VOIP
Voice-over-IP

Appendix B. Change Log

B.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.

B.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.

B.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.

B.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:

- Section 2.3 expands security and privacy consideration content of previous revisions. It also adds equity and usability considerations.
- Section 4.1 includes updated non-federated and federated digital identity models and descriptions.
- Section 4.4 consolidates informative descriptions and considerations on the use of federated identity architectures and assertions into one section.
- Section 5 expands upon the risk management content of previous revisions and specifically mandates that organizations take into account impacts to individuals and communities in addition to impacts to the organization. It also elevates risks to mission delivery, including challenges to the provisioning of services to all people who are eligible for and entitled to them, within the risk management process and when implementing digital identity systems. The xAL selection flowcharts, previously found in 800-63-3, section 6, have been replaced with text that elaborates the risk management process along with a sample risk assessment matrix that supports xAL selection. Additionally, the guidelines now mandate continuous evaluation of potential impacts to individuals, communities, and organizations.