NIST Special Publication 800-63-3

Digital Identity Guidelines

Paul A. Grassi
Michael E. Garcia
James L. Fenton

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-63-3

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

NIST Special Publication 800-63-3

Digital Identity Guidelines

Paul A. Grassi
Michael E. Garcia
Applied Cybersecurity Division
Information Technology Laboratory

James L. Fenton
Altmode Networks
Los Altos, Calif.

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-63-3

June 2017
INCLUDES UPDATES AS OF 03-02-2020; PAGE X

U.S. Department of Commerce
Wilbur L. Ross, Jr., Secretary

National Institute of Standards and Technology
Kent Rochford, Acting NIST Director and Under Secretary of Commerce for Standards and Technology

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.

National Institute of Standards and Technology Special Publication 800-63-3
Natl. Inst. Stand. Technol. Spec. Publ. 800-63-3, 75 pages (June 2017)
CODEN: NSPUE2

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SP.800-63-3

Comments on this publication may be submitted to:

National Institute of Standards and Technology
Attn: Applied Cybersecurity Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 2000) Gaithersburg, MD 20899-2000
Email: dig-comments@nist.gov

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 IT systems over open 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 supersedes NIST Special Publication 800-63-2.

Keywords

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

Acknowledgments

The authors gratefully acknowledge Kaitlin Boeckl for her artistic graphics contributions to all volumes in the SP 800-63 suite and the contributions of our many reviewers, including Joni Brennan from the Digital ID & Authentication Council of Canada (DIACC), Ellen Nadeau and Ben Piccarreta from NIST, and Danna Gabel O'Rourke from Deloitte & Touche LLP.

In addition, the authors would like to acknowledge the thought leadership and innovation of the original authors: Donna F. Dodson, Elaine M. Newton, Ray A. Perlner, W. Timothy Polk, Sarbari Gupta, and Emad A. Nabbus. Without their tireless efforts, we would not have had the incredible baseline from which to evolve SP 800-63 to the document it is today.

Requirements Notation and Conventions

The terms "SHALL" and "SHALL NOT" indicate requirements to be followed strictly in order to conform to the publication and from which no deviation is permitted.

The terms "SHOULD" and "SHOULD NOT" indicate that among several possibilities one is recommended as particularly suitable, without mentioning or excluding others, or that a certain course of action is preferred but not necessarily required, or that (in the negative form) a certain possibility or course of action is discouraged but not prohibited.

The terms "MAY" and "NEED NOT" indicate a course of action permissible within the limits of the publication.

The terms "CAN" and "CANNOT" indicate a possibility and capability, whether material, physical or causal or, in the negative, the absence of that possibility or capability.

Executive Summary

This section is informative.

Digital identity is the online persona of a subject, and a single definition is widely debated internationally. The term persona is apropos as a subject can represent themselves online in many ways. An individual may have a digital identity for email, and another for personal finances. A personal laptop can be someone's streaming music server yet also be a worker-bot in a distributed network of computers performing complex genome calculations. Without context, it is difficult to land on a single definition that satisfies all.

Digital identity as a legal identity further complicates the definition and ability to use digital identities across a range of social and economic use cases. Digital identity is hard. Proving someone is who they say they are — especially remotely, via a digital service — is fraught with opportunities for an attacker to successfully impersonate someone. As correctly captured by Peter Steiner in The New Yorker, "On the internet, nobody knows you're a dog." These guidelines provide mitigations to the vulnerabilities inherent online, while recognizing and encouraging that when accessing some low-risk digital services, "being a dog" is just fine; while other, high-risk services need a level of confidence that the digital identity accessing the service is the legitimate proxy to the real-life subject.

For these guidelines, digital identity is the unique representation of a subject engaged in an online transaction. A digital identity is always unique in the context of a digital service, but does not necessarily need to uniquely identify the subject in all contexts. In other words, accessing a digital service may not mean that the subject's real-life identity is known.

Identity proofing establishes that a subject is who they claim to be. Digital authentication establishes that a subject attempting to access a digital service is in control of one or more valid authenticators associated with that subject's digital identity. For services in which return visits are applicable, successfully authenticating provides reasonable risk-based assurances that the subject accessing the service today is the same as that which accessed the service previously. Digital identity presents a technical challenge because this process often involves proofing individuals over an open network, and always involves the authentication of individual subjects over an open network to access digital government services. The processes and technologies to establish and use digital identities offer multiple opportunities for impersonation and other attacks.

These technical guidelines supersede NIST Special Publication SP 800-63-2. Agencies use these guidelines as part of the risk assessment and implementation of their digital service(s). These guidelines provide mitigations of an authentication error's negative impacts by separating the individual elements of identity assurance into discrete, component parts. For non-federated systems, agencies will select two components, referred to as *Identity Assurance Level (IAL)* and *Authenticator Assurance Level (AAL)*. For federated systems, agencies will select a third component, *Federation Assurance Level (FAL)*.

These guidelines retire the concept of a level of assurance (LOA) as a single ordinal that drives implementation-specific requirements. Rather, by combining appropriate business and privacy risk management side-by-side with mission need, agencies will select IAL, AAL, and FAL as distinct options. While many systems will have the same numerical level for each of IAL, AAL, and FAL, this is not a requirement and agencies should not assume they will be the same in any given system.

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 strength of an assertion in a federated environment, used to communicate authentication and attribute information (if applicable) to a relying party (RP).

The separation of these categories provides agencies flexibility in choosing identity solutions and increases the ability to include privacy-enhancing techniques as fundamental elements of identity systems at any assurance level. For example, these guidelines support scenarios that will allow pseudonymous interactions even when strong, multi-factor authenticators are used. In addition, these guidelines encourage minimizing the dissemination of identifying information by requiring federated identity providers (IdPs) to support a range of options for querying data, such as asserting whether an individual is older than a certain age rather than querying the entire date of birth. While many agency use cases will require individuals to be fully identified, these guidelines encourage pseudonymous access to government digital services wherever possible and, even where full identification is necessary, limiting the amount of personal information collected as much as possible.

In today's environment, an organization's identity solution need not be a monolith, where one system or vendor provides all functionality. The market for identity services is componentized, allowing organizations and agencies to employ standards-based, pluggable identity solutions based on mission need. As such, SP 800-63 has been split into a suite of documents. The suite as a whole is referred to as "the guidelines," with the individual documents referred to as "volumes." RPs are required to use SP 800-63; the remaining volumes may be used independently or in an integrated fashion, depending on the component service(s) an agency requires.

Each volume has adopted verbs that are internationally recognized in standards organizations as normative and requirements-based. When used in a normative statement in these guidelines, they are CAPITALIZED for ease of identification. For example, SHALL is used to denote a mandatory requirement, while SHOULD refers to a technique, technology, or process that is recommended but not mandatory. For more details on the definitions of these terms see the Requirements Notation and Conventions at the beginning of each document.

These documents may inform — but do not restrict or constrain — the development or use of standards for application outside the federal government, such as e-commerce transactions.

These guidelines are organized as follows:

SP 800-63 Digital Identity Guidelines (This document)

SP 800-63 provides 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.*

SP 800-63A Enrollment and Identity Proofing

NIST SP 800-63-A addresses how applicants can prove their identities and become enrolled as valid subscribers within an identity system. It provides requirements by which applicants can both identity proof and enroll at one of three different levels of risk mitigation in both remote and physically-present scenarios. *SP 800-63A contains both normative and informative material.*

SP 800-63A sets requirements to achieve a given IAL. The three IALs reflect the options agencies may select from based on their risk profile and the potential harm caused by an attacker making a successful false claim of an identity. The IALs are as follows:

IAL1: There is no requirement to link the applicant to a specific real-life identity. Any attributes provided in conjunction with the authentication process are self-asserted or should be treated as such (including attributes a Credential Service Provider, or CSP, asserts to an RP).

IAL2: Evidence supports the real-world existence of the claimed identity and verifies that the applicant is appropriately associated with this real-world identity. IAL2 introduces the need for either remote or physically-present identity proofing. Attributes can be asserted by CSPs to RPs in support of pseudonymous identity with verified attributes.

IAL3: Physical presence is required for identity proofing. Identifying attributes must be verified by an authorized and trained representative of the CSP. As with IAL2, attributes can be asserted by CSPs to RPs in support of pseudonymous identity with verified attributes.

SP 800-63B Authentication and Lifecycle Management

For services in which return visits are applicable, a successful authentication provides reasonable risk-based assurances that the subscriber accessing the service today is the same as that which accessed the service previously. The robustness of this confidence is described by an AAL categorization. NIST SP 800-63B addresses how an individual can securely authenticate to a CSP to access a digital service or set of digital services. *SP 800-63B contains both normative and informative material.*

The three AALs define the subsets of options agencies can select based on their risk profile and the potential harm caused by an attacker taking control of an authenticator and accessing agencies' systems. The AALs are as follows:

AAL1: AAL1 provides some assurance that the claimant controls an authenticator bound to the subscriber's account. AAL1 requires either single-factor or multi-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(s) bound to the subscriber's account. Proof of possession and control of two distinct authentication factors is required through secure authentication protocol(s). Approved cryptographic techniques are required at AAL2 and above.

AAL3: AAL3 provides very high confidence that the claimant controls authenticator(s) bound to the subscriber's account. Authentication at AAL3 is based on proof of possession of a key through a cryptographic protocol. AAL3 authentication SHALL use a hardware-based authenticator and an authenticator that provides verifier impersonation resistance; the same device MAY fulfill both these requirements. In order to authenticate at AAL3, claimants SHALL prove possession and control of two distinct authentication factors through secure authentication protocol(s). Approved cryptographic techniques are required.

SP 800-63C Federation and Assertions

NIST SP 800-63C provides requirements when using federated identity architectures and assertions to convey the results of authentication processes and relevant identity information to an agency application. In addition, 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.*

The three FALs reflect the options agencies can select based on their risk profile and the potential harm caused by an attacker taking control of federated transactions. The FALs are as follows:

FAL1: Allows for the subscriber to enable the RP to receive a bearer assertion. The assertion is signed by the IdP using approved cryptography.

FAL2: Adds the requirement that the assertion be encrypted using approved cryptography such that the RP is the only party that can decrypt it.

FAL3: Requires the subscriber to present proof of possession of a cryptographic key referenced in the assertion in addition to the assertion artifact itself. The assertion is signed by the IdP and encrypted to the RP using approved cryptography.

These guidelines are agnostic to the vast array of identity service architectures that agencies can develop or acquire, and are meant to be applicable regardless of the approach an agency selects. However, agencies are encouraged to use federation where possible, and the ability to mix and match IAL, AAL, and FAL is simplified when federated architectures are used. Furthermore, federation is a keystone in the ability to enhance the privacy of the federal government's constituents as they access valuable government digital services.

Table of Contents

List of Appendices

List of Figures

List of Tables

Errata

This table contains changes that have been incorporated into Special Publication 800-63-3. Errata updates can include corrections, clarifications, or other minor changes in the publication that are either editorial or substantive in nature.

Date	Type	Change	Location
2017-12-01	Editorial	Removed the term 'cryptographic' from the AAL3 description.	Executive Summary
	Editorial	Updated reference to Risk Management Framework	§5
	Editorial	Fixed verbiage in xAL flowcharts	Figures 6-1, 6-2, and 6-3
	Editorial	Added NISTIR 8062 as a reference	§8.1
	Editorial	Added definitions for disassociability, manageability, processing, and predictability	Appendix A
2020-03-02	Editorial	Fixed wording of FAL3 definition	§5.2
	Substantive	Clarified flowcharts for xAL selection	Figures 6-1, 6-2, and 6-3
	Substantive	Added definition for Authorization Component	Appendix A
	Editorial	Removed extraneous definition of Protected Session	Appendix A

1 Purpose

This section is informative.

This recommendation and its companion volumes, Special Publication (SP) 800-63A, SP 800-63B, and SP 800-63C, provide technical guidelines to agencies for the implementation of digital authentication.

2 Introduction

This section is informative.

Digital identity is the unique representation of a subject engaged in an online transaction. A digital identity is always unique in the context of a digital service, but does not necessarily need to uniquely identify the subject in all contexts. In other words, accessing a digital service may not mean that the subject's real-life identity is known. Identity proofing establishes that a subject is who they claim to be. Digital authentication is 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. Successful authentication provides reasonable risk-based assurances that the subject accessing the service today is the same as that which previously accessed the service. Digital identity presents a technical challenge because this process often involves proofing individuals over an open network, and typically involves the authentication of individual subjects over an open network to access digital government services. There are multiple opportunities for impersonation and other attacks that fraudulently claim another subject's digital identity.

This recommendation provides agencies with technical guidelines for digital authentication of subjects to federal systems over a network. This recommendation also provides guidelines for credential service providers (CSPs), verifiers, and relying parties (RPs).

These guidelines describe the risk management processes for selecting appropriate digital identity services and the details for implementing identity assurance, authenticator assurance, and federation assurance levels based on risk. Risk assessment guidance in these guidelines supplements the *NIST Risk Management Framework* [NIST RMF] and its component special publications. This guideline does not establish additional risk management processes for agencies. Rather, requirements contained herein provide specific guidance related to digital identity risk while executing all relevant RMF lifecycle phases.

Digital authentication supports privacy protection by mitigating risks of unauthorized access to individuals' information. At the same time, because identity proofing, authentication, authorization, and federation 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.

These guidelines support the mitigation of the negative impacts induced by an authentication error by separating the individual elements of identity assurance into discrete, component parts. For non-federated systems, agencies will select two components, referred to as *Identity Assurance Level (IAL)* and *Authenticator Assurance Level (AAL)*. For federated systems, a third component, *Federation Assurance Level (FAL)*, is included. Section 5, Digital Identity Risk Management provides details on the risk assessment process. Section 6, Selecting Assurance Levels combines the results of the risk assessment with additional context to support agency selection of the appropriate IAL, AAL, and FAL combinations based on risk.

These guidelines do not consider nor result in a composite level of assurance (LOA) in the context of a single ordinal that drives implementation-specific requirements. Rather, by combining appropriate risk management for business, security, and privacy side-by-side with mission need, agencies will select IAL, AAL, and FAL as distinct options. Specifically, this document does not recognize the four LOA model previously used by federal agencies and described in OMB M-04-04, instead requiring agencies to individually select levels corresponding to each function being performed. While many systems will have the same numerical level for each IAL, AAL, and FAL, this is not a requirement, and agencies 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 assertion protocol used in a federated environment to communicate authentication and attribute information (if applicable) to an RP.

As such, SP 800-63 is organized as a suite of volumes as follows:

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.*

SP 800-63A *Enrollment and Identity Proofing*: Addresses how applicants can prove their identities and become enrolled as valid subjects within an identity system. It provides requirements for processes by which applicants can both proof and enroll at one of three different levels of risk mitigation in both remote and physically-present scenarios. *SP 800-63A contains both normative and informative material.*

SP 800-63B *Authentication and Lifecycle Management*: Addresses how an individual can securely authenticate to a CSP to access a digital service or set of digital services. This volume also describes the process of binding an authenticator to an identity. *SP 800-63B contains both normative and informative material.*

SP 800-63C *Federation and Assertions*: 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. Furthermore, 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.*

NIST anticipates that individual volumes in these guidelines will be revised asynchronously. At any time, the most recent revision of each should be used (e.g., if at a time in the future SP 800-63A-1 and SP 800-63B-2 are the most recent revisions of each volume, they should be used together even though the revision numbers do not match). To minimize the risk of compatibility

errors, a reference to the base document (i.e., SP 800-63 rather than SP 800-63-3) always refers to the current version of the document.

The following table states which sections of this volume are normative and which are informative:

Table 2-1 Normative and Informative Sections of SP 800-63-3

Section Name	Normative/Informative
1. Purpose	Informative
2. Introduction	Informative
3. Definitions and Abbreviations	Informative
4. Digital Identity Model	Informative
5. Digital Identity Risk Management	Normative
6. Selecting Assurance Levels	Normative
7. Federation Considerations	Informative
8. References	Informative

2.1 Applicability

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

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 assurance may consider the use of these standards where appropriate.

These guidelines primarily focus on agency services that interact with the non-federal workforce, such as citizens accessing benefits or private sector partners accessing information sharing collaboration spaces. However, it also applies to internal agency systems accessed by employees and contractors. These users are expected to hold a valid government-issued credential, primarily the Personal Identity Verification (PIV) card or a derived PIV. Therefore SP 800-63A and SP 800-63B are secondary to the requirements of FIPS 201 and its corresponding set of special publications and agency-specific instructions. However, SP 800-63C and the risk-based selection of an appropriate FAL applies, regardless of the credential type the internal user holds. FAL

selection provides agencies guidance and flexibility in how to PIV-enable their applications based on system risk.

2.2 Considerations, Other Requirements, and Flexibilities

Agencies may employ other risk mitigation measures and compensating controls not specified herein. Agencies need to ensure that any mitigations and compensating controls do not degrade the selected assurance level's intended security and privacy protections. Agencies may consider partitioning the functionality of a digital service to allow less sensitive functions to be available at a lower level of authentication and identity assurance.

Agencies may determine based on their risk analysis that additional measures are appropriate in certain contexts. In particular, privacy requirements and legal risks may lead agencies to determine that additional authentication measures or other process safeguards are appropriate. When developing digital authentication processes and systems, agencies should consult *OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002* [M-03-22]. See the *Use of Electronic Signatures in Federal Organization Transactions* [ESIG] for additional information on legal risks, especially those related to the need to 1) satisfy legal standards of proof and 2) prevent repudiation.

Additionally, federal agencies implementing these guidelines should adhere to their 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 [FISMA], and related NIST standards and guidelines. FISMA directs federal agencies to develop, document, and implement agency-wide programs to provide security for the information and systems that support the agency's operations and assets. This includes the security authorization and accreditation (SA&A) of IT systems that support digital authentication. NIST recommends that non-federal entities implementing these guidelines follow equivalent standards to ensure the secure operations of their digital systems.

2.3 A Few Limitations

These technical guidelines do not address the authentication of subjects for physical access (e.g., to buildings), though some authenticators used for digital access may also be used for physical access authentication. 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). That said, these guidelines are written to refer to generic subjects wherever possible to leave open the possibility for applicability to devices. Also excluded are specific requirements for issuing authenticators to devices when they are used in authentication protocols with people.

2.4 How to Use this Suite of SPs

The business model, marketplace, and composition of how identity services are delivered has drastically changed since the first version of SP 800-63 was released. Notably, CSPs can be componentized and comprised of multiple independently-operated and owned business entities. Furthermore, there may be a significant security benefit to using strong authenticators even if no

identity proofing is required. Therefore, in this revision, a suite of SPs under the 800-63 moniker has been created to facilitate these new models and make it easy to access the specific requirements for the function an entity may serve under the overall digital identity model.

2.5 Change History

2.5.1 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.

2.5.2 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 Section 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 Section 5 were minor explanations and clarifications.

2.5.3 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.
- Terminology changes, including the use of authenticator in place of token to avoid conflicting use of the word token in assertion technologies.
- Updates to authentication and assertion requirements to reflect advances in both security technology and threats.
- Requirements on the storage of long-term secrets by verifiers.
- Restructured identity proofing model.
- Updated requirements regarding remote identity proofing.
- Clarification on the use of independent channels and devices as "something you have".

- **Removal** of pre-registered knowledge tokens (authenticators), with the recognition that they are special cases of (often very weak) passwords.
- Requirements regarding account recovery in the event of loss or theft of an authenticator.
- **Removal** of email as a valid channel for out-of-band authenticators.
- Expanded discussion of re-authentication and session management.
- Expanded discussion of identity federation; restructuring of assertions in the context of federation.

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 digital identity model used in these guidelines reflects technologies and architectures currently available in the market. More complex models that separate functions — such as issuing credentials and providing attributes — among a larger number of parties are also available and may have advantages in some application classes. While a simpler model is used in this document, it does not preclude agencies from separating these functions. Additionally, certain enrollment, identity proofing, and issuance processes performed by the CSP are sometimes delegated to an entity known as either the registration authority (RA) or identity manager (IM). A close relationship between the RA and CSP is typical, and the nature of this relationship may differ among RAs, IMs, and CSPs. The type of relationship and its requirements is outside of the scope of this document. Accordingly, the term CSP will be inclusive of RA and IM functions. Finally, a CSP may provide other services in addition to digital identity services. In these situations, the requirements specified throughout these guidelines only apply to the CSP function(s), not the additional services.

Digital identity is the unique representation of a subject engaged in an online transaction. The process used to verify a subject's association with their real-world identity is called *identity proofing*. In these guidelines, the party to be proofed is called an *applicant*. When the applicant successfully completes the proofing process, they are referred to as a *subscriber*.

The strength of identity proofing is described by an ordinal measurement called the IAL. At IAL1, identity proofing is not required, therefore any attribute information provided by the applicant is self-asserted, or should be treated as self-asserted and not verified (even if provided by a CSP to an RP). IAL2 and IAL3 require identity proofing, and the RP may request the CSP assert information about the subscriber, such as verified attribute values, verified attribute references, or pseudonymous identifiers. This information assists the RP in making authorization decisions. An RP may decide that it requires IAL2 or IAL3, but may only need specific attributes, resulting in the subject retaining some degree of pseudonymity. This privacy-enhancing approach is a benefit of separating the strength of the proofing process from that of the authentication process. An RP may also employ a federated identity approach where the RP outsources all identity proofing, attribute collection, and attribute storage to a CSP.

In these guidelines, the party to be authenticated is called a *claimant* and the party verifying that identity is called a *verifier*. When a claimant successfully demonstrates possession and control of one or more authenticators to a verifier through an authentication protocol, the verifier can verify that the claimant is a valid subscriber. The verifier passes on an assertion about the subscriber, who may be either pseudonymous or non-pseudonymous, to the RP. That assertion includes an identifier, and may include identity information about the subscriber, such as the name, or other attributes that were collected in the enrollment process (subject to the CSP's policies, the RP's

needs, and consent for disclosure of attributes given by the subject). Where the verifier is also the RP, the assertion may be implicit. The RP can use the authenticated information provided by the verifier to make authorization decisions.

Authentication establishes confidence that the claimant has possession of an authenticator(s) bound to the credential, and in some cases in the attribute values of the subscriber (e.g., if the subscriber is a U.S. citizen, is a student at a particular university, or is assigned a particular number or code by an agency or organization). Authentication does not determine the claimant's authorizations or access privileges; this is a separate decision, and is out of these guidelines' scope. RPs can use a subscriber's authenticated identity and attributes with other factors to make authorization decisions. Nothing in this document suite precludes RPs from requesting additional information from a subscriber that has successfully authenticated.

The strength of the authentication process is described by an ordinal measurement called the AAL. AAL1 requires single-factor authentication and is permitted with a variety of different authenticator types. At AAL2, authentication requires two authentication factors for additional security. Authentication at the highest level, AAL3, additionally requires the use of a hardware-based authenticator and verifier impersonation resistance.

The various entities and interactions that comprise the digital identity model used here are illustrated in Figure 4-1.

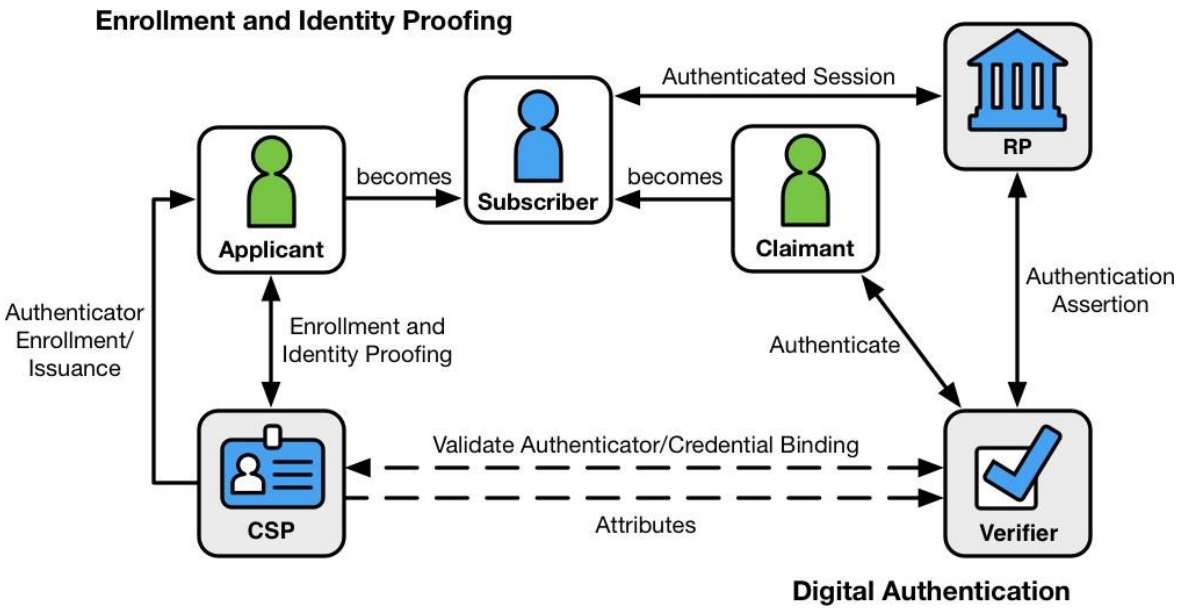

Figure 4-1 Digital Identity Model

The left side of the diagram shows the enrollment, credential issuance, lifecycle management activities, and various states of an identity proofing and authentication process. The usual sequence of interactions is as follows:

1. An applicant applies to a CSP through an enrollment process.
2. The CSP identity proofs that applicant. Upon successful proofing, the applicant becomes a subscriber.
3. Authenticator(s) and a corresponding credential are established between the CSP and the subscriber.
4. The CSP maintains the credential, its status, and the enrollment data collected for the lifetime of the credential (at a minimum). The subscriber maintains his or her authenticator(s).

Other sequences are less common, but could also achieve the same functional requirements.

The right side of Figure 4-1 shows the entities and interactions involved in using an authenticator to perform digital authentication. A subscriber is referred to as a claimant when he or she needs to authenticate to a verifier. The interactions are as follows:

1. The claimant proves possession and control of the authenticator(s) to the verifier through an authentication protocol.
2. The verifier interacts with the CSP to validate the credential that binds the subscriber's identity to their authenticator and to optionally obtain claimant attributes.
3. The CSP or verifier provides an assertion about the subscriber to the RP, which may use the information in the assertion to make an authorization decision.
4. An authenticated session is established between the subscriber and the RP.

In all cases, the RP should request the attributes it requires from a CSP before authenticating the claimant. In addition, the claimant should be requested to consent to the release of those attributes prior to generation and release of an assertion.

In some cases, the verifier does not need to communicate in real time with the CSP to complete the authentication activity (e.g., some uses of digital certificates). Therefore, the dashed 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 as shown in Figure 4-1. However, if these functions reside on the same platform, the interactions between the components are local messages between applications running on the same system rather than protocols over shared, untrusted networks.

As noted above, a CSP maintains status information about the credentials it issues. CSPs will generally assign a finite lifetime when issuing credentials to limit the maintenance period. When the status changes, or when the credentials near expiration, credentials may be renewed or re-issued; or, the credential may be revoked and destroyed. Typically, the subscriber authenticates to the CSP using their existing, unexpired authenticator and credential in order to request issuance of a new authenticator and credential. If the subscriber fails to request authenticator and credential re-issuance prior to their expiration or revocation, they may be required to repeat the enrollment process to obtain a new authenticator and credential. Alternatively, the CSP may choose to accept a request during a grace period after expiration.

4.2 Enrollment and Identity Proofing

Normative requirements can be found in SP 800-63A, *Enrollment and Identity Proofing*.

The previous section introduced the participants in the conceptual digital identity model. This section provides additional details regarding the participants' relationships and responsibilities in enrollment and identity proofing.

An individual, referred to as an *applicant* at this stage, opts to be identity proofed by a CSP. If the applicant is successfully proofed, the individual is then termed a subscriber of that CSP.

The CSP establishes a mechanism to uniquely identify each subscriber, register the subscriber's credentials, and track the authenticators issued to that subscriber. The subscriber may be given authenticators at the time of enrollment, the CSP may bind authenticators the subscriber already has, or they may be generated later as needed. Subscribers have a duty to maintain control of their authenticators and comply with CSP policies in order to maintain active authenticators. The CSP maintains enrollment records for each subscriber to allow recovery of authenticators, for example, when they are lost or stolen.

4.3 Authentication and Lifecycle Management

Normative requirements can be found in SP 800-63B, *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 data).

MFA refers to the use of more than one of the above factors. The strength of authentication systems is largely determined by the number of factors incorporated by the system — the more factors employed, the more robust the authentication system. For the purposes of these guidelines, using two factors is adequate to meet the highest security requirements. As discussed in Section 5.1, other types of information, such as location data or device identity, may be used by an RP or 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 that have been registered with the CSP and are used to prove the claimant's identity. The authenticator(s) contains secrets the claimant can use to prove that he or she is a valid subscriber, the claimant authenticates to a system or application over a network by proving that he or she has possession and control of one or more authenticators.

The secrets contained in authenticators are based on either public key pairs (asymmetric keys) or shared secrets (symmetric keys). A public key and a related private key comprise a public key pair. The private key is stored on the authenticator and is used by the claimant to prove possession and control of the authenticator. A verifier, knowing the claimant's public key through some credential (typically a public key certificate), can use an authentication protocol to verify the claimant's identity by proving that the claimant has possession and control of the associated private key authenticator.

Shared secrets stored on authenticators may be either symmetric keys or memorized secrets (e.g., passwords and PINs), as opposed to the asymmetric keys described above, which subscribers need not share with the verifier. While both keys and passwords can be used in similar protocols, one important difference between the two is how they relate to the subscriber. While symmetric keys are generally stored in hardware or software that the subscriber controls, passwords are intended to be memorized by the subscriber. Since most users choose short passwords to facilitate memorization and ease of entry, passwords typically have fewer characters than cryptographic keys. Furthermore, whereas systems choose keys at random, users attempting to choose memorable passwords will often select from a very small subset of the possible passwords of a given length, and many will choose very similar values. As such, whereas cryptographic keys are typically long enough to make network-based guessing attacks untenable, user-chosen passwords may be vulnerable, especially if no defenses are in place.

In this volume, authenticators always contain a secret. Some of the classic authentication factors 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 is not in itself an authenticator for digital authentication. Authentication factors classified as something you know are not necessarily secrets, either. Knowledge-based authentication, where the claimant is prompted to answer questions that are presumably known only by the claimant, also does not constitute an acceptable secret for digital authentication. A biometric also does not constitute a secret. Accordingly, these guidelines only allow the use of biometrics for authentication when strongly bound to a physical 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 (what you know) with an out-of-band device (what you have). Both authenticator outputs are presented to the verifier to authenticate the claimant. For item 2, consider a piece of hardware (the authenticator) that contains a cryptographic key (the authenticator secret) where access is protected with a fingerprint. When used with the biometric, the cryptographic key produces an output that is used to authenticate the claimant.

As noted above, biometrics, when employed as a single factor of authentication, do not constitute acceptable secrets for digital authentication — but they do have their place in the authentication of digital identities. 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. They include facial features, fingerprints, iris patterns, voiceprints, and many other characteristics. SP 800-63A, *Enrollment and Identity Proofing* recommends that biometrics be collected in the enrollment process to later help prevent a registered subscriber from repudiating the enrollment, and to help identify those who commit enrollment fraud.

4.3.2 Credentials

As described in the preceding sections, a credential binds an authenticator to the subscriber, via an identifier, as part of the issuance process. A credential is stored and maintained by the CSP, though the claimant may possess it. The claimant possesses an authenticator, but is not necessarily in possession of the credential. For example, database entries containing the user attributes are considered to be credentials for the purpose of this document but are possessed by the verifier. X.509 public key certificates are a classic example of credentials the claimant can, and often does, possess.

4.3.3 Authentication Process

The authentication process begins with the claimant demonstrating to the verifier possession and control of an authenticator that is bound to the asserted identity through an authentication protocol. Once possession and control have been demonstrated, the verifier verifies that the credential remains valid, usually by interacting with the CSP.

The exact nature of the interaction between the verifier and the claimant during the authentication protocol 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.

The verifier is a functional role, but is frequently implemented in combination with the CSP, the RP, or both. If the verifier is a separate entity from the CSP, it is often desirable to ensure that the verifier does not learn the subscriber's authenticator secret in the process of authentication, or at least to ensure that the verifier does not have unrestricted access to secrets stored by the CSP.

4.4 Federation and Assertions

Normative requirements can be found in SP 800-63C, *Federation and Assertions*.

Overall, SP 800-63 does not presuppose a federated identity architecture; rather, these guidelines are agnostic to the types of models that exist in the marketplace, allowing agencies to deploy a

digital authentication scheme according to their own requirements. However, identity federation is preferred over a number of siloed identity systems that each serve a single agency or RP.

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 issued credential at multiple RPs.
- Cost reduction to both the user (reduction in authenticators) and the agency (reduction in information technology infrastructure).
- Data minimization as agencies do not need to pay for collection, storage, disposal, and compliance activities related to storing personal information.
- Pseudonymous attribute assertions as agencies can request a minimized set of attributes, to include claims, to fulfill service delivery.
- Mission enablement as agencies can focus on mission, rather than the business of identity management.

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

4.4.1 Assertions

Upon completion of the authentication process, the verifier generates an assertion containing the result of the authentication and provides it to the RP. The assertion is used to communicate the result of the authentication process, and optionally information about the subscriber, from the verifier to the RP. Assertions may be communicated directly to the RP, or can be forwarded through the subscriber, which has further implications for system design.

An RP trusts an assertion based on the source, the time of creation, how long the assertion is valid from time of creation, and the corresponding trust framework that governs the policies and processes of CSPs and RPs. The verifier is responsible for providing a mechanism by which the integrity of the assertion can be confirmed.

The RP is responsible for authenticating the source (the verifier) and for confirming the integrity of the assertion. When the verifier passes the assertion through the subscriber, the verifier must protect the integrity of the assertion in such a way that it cannot be modified. However, if the verifier and the RP communicate directly, a protected session may be used to preserve the integrity of the assertion. When sending assertions across an open network, the verifier is responsible for ensuring that any sensitive subscriber information contained in the assertion can only be extracted by an RP that it trusts to maintain the information's confidentiality.

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 info claims may optionally be digitally signed.
- Kerberos tickets allow a ticket-granting authority to issue session keys to two authenticated parties using symmetric key based encapsulation schemes.

4.4.2 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 authenticated identity (pseudonymous or non-pseudonymous), the IAL, AAL, and FAL (FAL indicating the strength of the assertion protocol), and other factors to make authorization decisions. The verifier and the RP may be the same entity, or they may be separate entities. If they are separate entities, the RP normally receives an assertion from the verifier. The RP ensures that the assertion came from a verifier 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 and the corresponding risk assessment guidance supplement the *NIST Risk Management Framework* [NIST RMF] and its component special publications. This does not establish additional risk management processes for agencies. Rather, requirements contained herein provide specific guidance related to digital identity risk that agency RPs SHALL apply while executing all relevant RMF lifecycle phases.

5.1 Overview

In today's digital services, combining proofing, authenticator, and federation requirements into a single bundle sometimes has unintended consequences and can put unnecessary implementation burden on the implementing organization. It is quite possible that an agency can deliver the most effective set of identity services by assessing the risk and impacts of failures for each individual component of digital authentication, rather than as a single, all-encompassing LOA. To this end, these guidelines recognize that an authentication error is not a singleton that drives all requirements.

This volume details requirements to assist agencies in avoiding:

1. Identity proofing errors (i.e., a false applicant claiming an identity that is not rightfully theirs);
2. Authentication errors (i.e., a false claimant using a credential that is not rightfully theirs); and
3. Federation errors (i.e., an identity assertion is compromised).

From the perspective of an identity proofing failure, there are two dimensions of potential failure:

1. The impact of providing a service to the wrong subject (e.g., an attacker successfully proofs as someone else).
2. The impact of excessive identity proofing (i.e., collecting and securely storing more information about a person than is required to successfully provide the digital service).

As such, agencies SHALL assess the risk of proofing, authentication, and federation errors separately to determine the required assurance level for each transaction.

Section 5.3 provides impact categories specific to digital identity to assist in the overall application of the RMF.

Risk assessments 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. Once an agency has completed the overall risk assessment; selected

individual assurance levels for identity proofing, authentication, and federation (if applicable); and determined the processes and technologies they will employ to meet each assurance level, agencies SHALL develop a "Digital Identity Acceptance Statement", in accordance with SP 800-53 IA-1 a.1. See Section 5.5 for more detail on the necessary content of the Digital Identity Acceptance Statement.

5.2 Assurance Levels

An agency RP SHALL select, based on risk, the following individual 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 errors.
- 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 errors (i.e., a false claimant using a credential that is not rightfully theirs).
- FAL: The robustness of the assertion protocol the federation uses to communicate authentication and attribute information (if applicable) to an RP. FAL is optional as not all digital systems will leverage federated identity architectures. FAL is selected to mitigate potential federation errors (an identity assertion is compromised).

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

Table 5-1 Identity Assurance Levels

Identity Assurance Level
IAL1: At IAL1, attributes, if any, are self-asserted or should be treated as self-asserted.
IAL2: At IAL2, either remote or in-person identity proofing is required. IAL2 requires identifying attributes to have been verified in person or remotely using, at a minimum, the procedures given in SP 800-63A.
IAL3: At IAL3, in-person identity proofing is required. Identifying attributes must be verified by an authorized CSP representative through examination of physical documentation as described in SP 800-63A.

Table 5-2 Authenticator Assurance Levels

Authenticator 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(s) through a secure authentication protocol.
AAL2: AAL2 provides high confidence that the claimant controls authenticator(s) 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(s) registered to the subscriber. Authentication at AAL3 is based on proof of possession of a key through a cryptographic protocol. AAL3 is like AAL2 but also requires a "hard" cryptographic authenticator that provides verifier impersonation resistance.

Table 5-3 Federation Assurance Levels

Federation Assurance Level
FAL1: FAL1 permits the RP to receive a bearer assertion from an identity provider (IdP). The IdP must sign the assertion using approved cryptography.
FAL2: FAL2 adds the requirement that the assertion be encrypted using approved cryptography such that the RP is the only party that can decrypt it.
FAL3: FAL3 requires the subscriber to present proof of possession of a cryptographic key referenced in the assertion along with the assertion itself. The assertion must be signed using approved cryptography and encrypted to the RP using approved cryptography.

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

5.3 Risk and Impacts

This section provides details on the impact categories used to determine IAL, AAL, and FAL.

Potential Impact Categories: To determine the appropriate level of assurance of the user's asserted identity, agencies SHALL assess the potential risks and identify measures to minimize their impact.

Authentication, proofing, and federation errors with potentially worse consequences require higher levels of assurance. Business process, policy, and technology may help reduce risk.

Categories of harm and impact include:

1. Inconvenience, distress, or damage to standing or reputation;
2. Financial loss or agency liability;
3. Harm to agency programs or public interests;
4. Unauthorized release of sensitive information;
5. Personal safety; and
6. Civil or criminal violations.

Required assurance levels for digital transactions are determined by assessing the potential impact of each of the above categories using the potential impact values described in Federal Information Processing Standard (FIPS) 199 [FIPS 199].

The three potential impact values are:

1. Low impact,
2. Moderate impact, and
3. High impact.

5.3.1 Business Process vs. Online Transaction

The assurance level determination is only based on transactions that are part of a digital system. An online transaction may not be equivalent to a complete business process that requires offline processing, or online processing in a completely segmented system. In selecting the appropriate assurance levels, the agency should assess the risk associated with online transactions they are offering via the digital service, not the entire business process associated with the provided benefit or service. For example, in an online survey, personal information may be collected, but it is never made available online to the submitter after the information is saved. In this instance, it is important for the information to be carefully protected in backend systems, but there is no reason to identity proof or even authenticate the user providing the information for the purposes of their own access to the system or its associated benefits. The online transaction is solely a submission of the data. The entire business process may require a significant amount of data validation, without ever needing to know if the correct person submitted the information. In this scenario, there is no need for any identity proofing nor authentication.

Another example where the assessed risk could differ if the agency evaluated the entire business process rather than the online transaction requirements is a digital service that accepts résumés to apply for open job postings. In this use case, the digital service allows an individual to submit – or at least does not restrict an individual from submitting – a résumé on behalf of anyone else, and in subsequent visits to the site, access the résumé for various purposes. Since the résumé information is available to the user in later sessions, and is likely to contain personal information, the agency must select an AAL that requires MFA, even though the user self-asserted the personal information. In this case, the requirements of [EO 13681] apply and the application

must provide at least AAL2. However, the identity proofing requirements remain unclear. The entire business process of examining a résumé and ultimately hiring and onboarding a person requires a significant amount of identity proofing. The agency needs a high level of confidence that the job applicant is in fact the subject of the résumé submitted online if a decision to hire is made. Yet this level of proofing is not required to submit the résumé online. Identity proofing is not required to complete the digital portion of the transaction successfully. Identity proofing the submitter would create more risk than required in the online system as excess personal information would be collected when no such information is needed for the portion of the hiring process served by the digital job application portal and may reduce usability. Therefore, the most appropriate IAL selection would be 1. There is no need to identity proof the user to successfully complete the online transaction. This decision for the online portal itself is independent of a seemingly obvious identity proofing requirement for the entire business process, lest a job be offered to a fraudulent applicant.

5.3.2 Impacts per Category

This section defines the potential impacts for each category of harm. Each assurance level, IAL, AAL, and FAL (if accepting or asserting a federated identity) SHALL be evaluated separately.

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

Potential impact of inconvenience, distress, or damage to standing or 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.

Potential impact of financial loss:

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

Potential impact of harm to agency programs or public interests:

- Low: at worst, a limited adverse effect on organizational operations or assets, or public interests. Examples of limited adverse effects are: (i) mission capability degradation to the extent and duration that the organization is able to perform its primary functions with noticeably reduced effectiveness, or (ii) minor damage to organizational assets or public interests.

- Moderate: at worst, a serious adverse effect on organizational operations or assets, or public interests. Examples of serious adverse effects are: (i) significant mission capability degradation to the extent and duration that the organization is able to perform its primary functions with significantly reduced effectiveness; or (ii) significant damage to organizational assets or public interests.
- High: a severe or catastrophic adverse effect on organizational operations or assets, or public interests. Examples of severe or catastrophic effects are: (i) severe mission capability degradation or loss of to the extent and duration that the organization is unable to perform one or more of its primary functions; or (ii) major damage to organizational assets or public interests.

Potential impact of unauthorized release of sensitive information:

- 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 FIPS 199.
- 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 FIPS 199.
- 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 FIPS 199.

Potential impact to personal safety:

- Low: at worst, minor injury not requiring medical treatment.
- Moderate: at worst, moderate risk of minor injury or limited risk of injury requiring medical treatment.
- High: a risk of serious injury or death.

The potential impact of civil or criminal violations is:

- Low: at worst, a risk of civil or criminal violations of a nature that would not ordinarily be subject to enforcement efforts.
- Moderate: at worst, a risk of civil or criminal violations that may be subject to enforcement efforts.
- High: a risk of civil or criminal violations that are of special importance to enforcement programs.

5.4 Risk Acceptance and Compensating Controls

The SP 800-63 suite specifies baseline requirements for digital identity services based on assurance level. Agencies SHOULD implement identity services per the requirements in these guidelines and SHOULD consider additional techniques and technologies to further secure and privacy-enhance their services.

Agencies MAY determine alternatives to the NIST-recommended guidance, for the assessed *x*ALs, based on their mission, risk tolerance, existing business processes, special considerations for certain populations, availability of data that provides similar mitigations to those described in this suite, or due to other capabilities that are unique to the agency.

Agencies SHALL demonstrate comparability of any chosen alternative, to include any compensating controls, when the complete set of applicable SP 800-63 requirements is not implemented. For example, an agency may choose a National Information Assurance Partnership (NIAP) protection profile over FIPS, where the profile is equivalent to or stronger than the FIPS requirements. That said, agencies SHALL NOT alter the assessed *x*AL based on agency capabilities. Rather, the agency MAY adjust their implementation of solutions based on the agency's ability to mitigate risk via means not explicitly addressed by SP 800-63 requirements. The agency SHALL implement procedures to document both the justification for any departure from normative requirements and detail the compensating control(s) employed.

This guidance addresses only those risks associated with authentication and identity proofing errors. NIST Special Publication 800-30, Risk Management Guide for Information Technology Systems [SP 800-30] recommends a general methodology for managing risk in federal systems.

5.5 Digital Identity Acceptance Statement

Agencies SHOULD include this information in existing artifacts required to achieve a SA&A.

The statement SHALL include, at a minimum:

1. Assessed *x*AL,
2. Implemented *x*AL,
3. Rationale, if implemented *x*AL differs from assessed *x*AL,
4. Comparability demonstration of compensating controls when the complete set of applicable 800-63 requirements are not implemented, and
5. If not accepting federated identities, rationale.

5.6 Migrating Identities

As these guidelines are revised, CSPs SHALL consider how changes in requirements affect their user population. In some instances, the user population will be unaffected, yet in others, the CSP will require users undergo a transitional activity. For example, CSPs may request users — upon initial logon since last revision — to supply additional proofing evidence to adhere to new IAL requirements. This SHALL be a risk-based decision, made in context of the CSP, any RPs that use the CSP, mission, and the population served. The following considerations serve only as a guide to agencies when considering the impacts of requirements changes:

1. If the RP is experiencing identity-related fraud, a migration may prove beneficial. If not, migration may not be an added value.

2. New, stronger, or user-friendly authentication options are added to individual AALs the CSP could issue new authenticators or allow users to register authenticators they already have.
3. Federation requirements may or may not have a user impact. For example, consent requirements or infrastructure requirements could necessitate an infrastructure or protocol upgrade.
4. Addition or removal of *x*ALs may not require a migration, but would trigger a new risk assessment to determine if a change is necessary for the RP.

The guidance does not prescribe that any migration needs to occur, only that it be considered as revisions are released. It is up to the CSP and RP, based on their risk tolerance and mission, to determine the best approach.

6 Selecting Assurance Levels

This section is normative.

The risk assessment results are the primary factor in selecting the most appropriate levels. This section details how to apply the results of the risk assessment with additional factors unrelated to risk to determine the most advantageous *x*AL selection.

First, compare the risk assessment impact profile to the impact profiles associated with each assurance level, as shown in Table 6-1 below. To determine the required assurance level, find the lowest level whose impact profile meets or exceeds the potential impact for every category analyzed in the risk assessment.

Table 6-1 Maximum Potential Impacts for Each Assurance Level

	Assurance Level		
Impact Categories	**1**	**2**	**3**
Inconvenience, distress or damage to standing or reputation	Low	Mod	High
Financial loss or agency liability	Low	Mod	High
Harm to agency programs or public interests	N/A	Low/Mod	High
Unauthorized release of sensitive information	N/A	Low/Mod	High
Personal Safety	N/A	Low	Mod/High
Civil or criminal violations	N/A	Low/Mod	High

In analyzing risks, the agency SHALL consider all of the expected direct and indirect results of an authentication failure, including the possibility that there will be more than one failure, or harms to more than one person or organization. The definitions of potential impacts contain some relative terms, like "serious" or "minor," whose meaning will depend on context. The agency SHOULD consider the context and the nature of the persons or entities affected to decide the relative significance of these harms. Over time, the meaning of these terms will become more definite as agencies gain practical experience with these issues. The analysis of harms to agency programs or other public interests depends strongly on the context; the agency SHOULD consider these issues with care.

It is possible that the assurance levels may differ across IAL, AAL, and FAL. For example, suppose an agency establishes a "health tracker" application in which users submit personal

information in the form of personal health information (PHI). In line with the terms of EO 13681 requiring "that all agencies making personal data accessible to citizens through digital applications require the use of multiple factors of authentication," the agency is required to implement MFA at AAL2 or AAL3.

EO 13681 also requires agencies employ "an effective identity proofing process, as appropriate" when personal information is released. This does not mean that proofing at IAL2 or IAL3 (to match the required AAL) is necessary. In the above example, there may be no need for the agency system to know the actual identity of the user. In this case, an "effective proofing process" would be to not proof at all, therefore the agency would select IAL1. This allows the user of the health tracker system to be pseudonymous.

Despite the user being pseudonymous, the agency should still select AAL2 or AAL3 for authentication because a malicious actor could gain access to the user's PHI by compromising the account.

> Note: An agency can accept a higher assurance level than those required in the table above. For example, in a federated transaction, an agency can accept an IAL3 identity if their application is assessed at IAL2. The same holds true for authenticators: stronger authenticators can be used at RPs that have lower authenticator requirements. However, RPs will have to ensure that this only occurs in federated scenarios with appropriate privacy protections by the CSP such that only attributes that have been requested by the RP and authorized by the subscriber are provided to the RP and that excessive personal information does not leak from the credential or an assertion. See the privacy considerations in SP 800-63C for more details.

> Note: The upshot of potentially having a different IAL, AAL, and FAL within a single application stems from the fact that this document no longer supports the notion of an overall LOA — the "low watermark" approach to determining LOA no longer applies. An application with IAL1 and AAL2 should not be considered any less secure or privacy-enhancing than an application with IAL2 and AAL2. The only difference between these applications is the amount of proofing required, which may not impact the security and privacy of each application. That said, if an agency incorrectly determines the *x*AL, security and privacy could very well be impacted.

6.1 Selecting IAL

The IAL decision tree in Figure 6-1 combines the results from the risk assessment with additional considerations related to identity proofing services to allow agencies to select the most appropriate identity proofing requirements for their digital service offering.

The IAL selection does not mean the digital service provider will need to perform the proofing themselves. More information on whether an agency can federate is provided in Section 7.

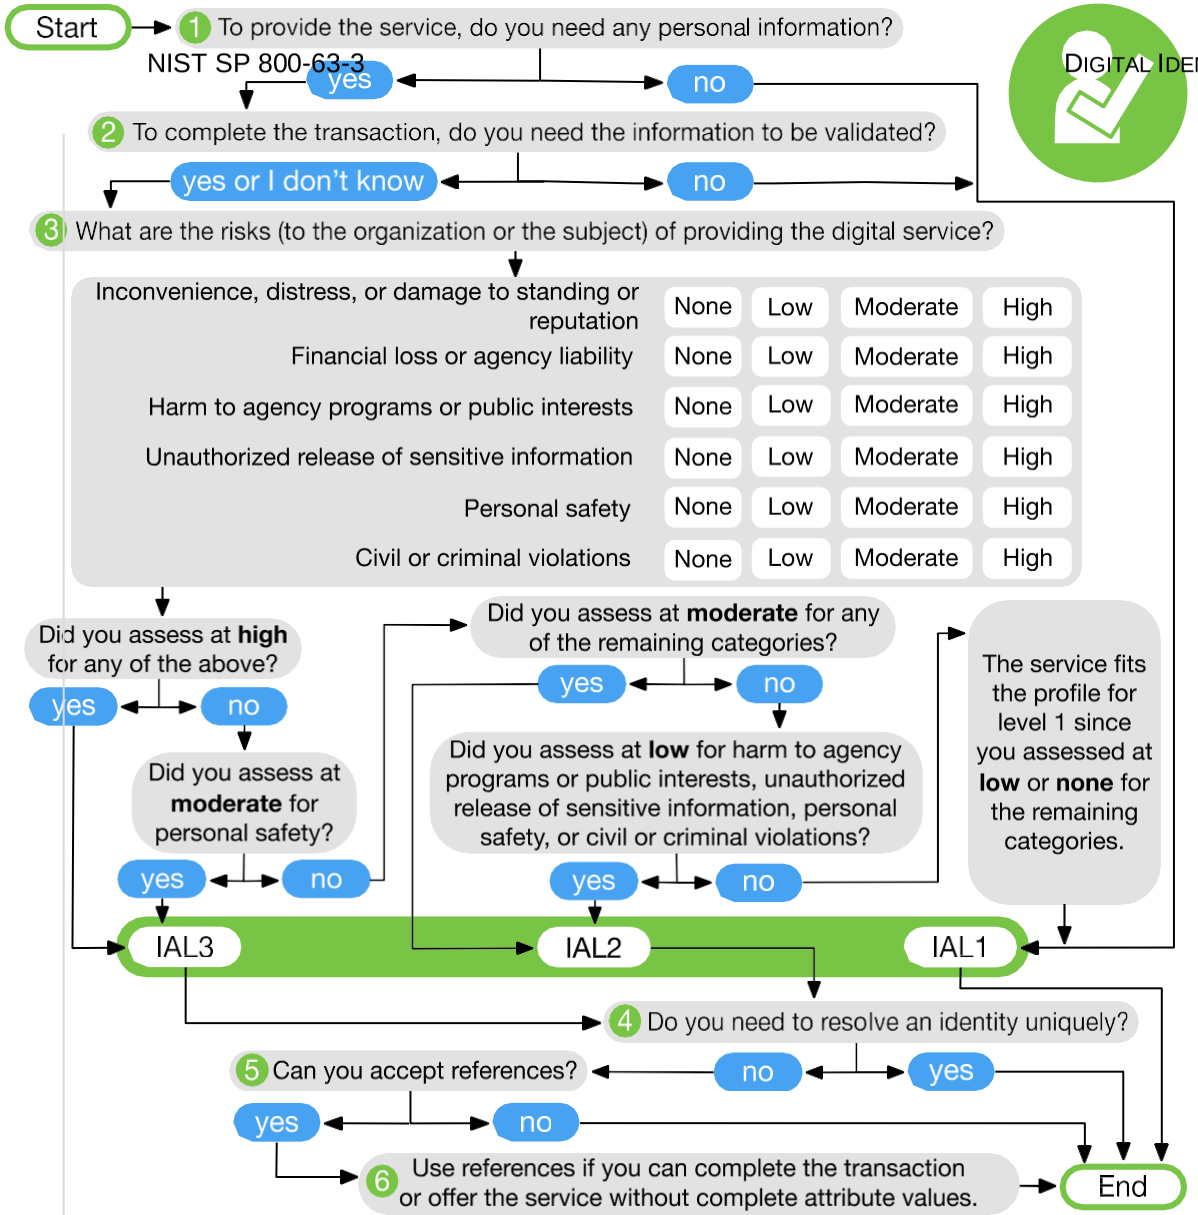

Figure 6-1 Selecting IAL

1363

1 To provide the service, do you need any personal information?
The risk assessment and IAL selection can be short circuited by answering this question first. If the service does not require any personal information to execute any digital transactions, the system can operate at IAL1.
2 To complete the transaction, do you need the information to be validated?
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 even a single validated and verified attribute is needed, then the provider will need to accept attributes that have been IAL2 or IAL3 proofed. Again, the selection of IAL can be short circuited to IAL1 if the agency can deliver the digital service with self-asserted attributes only.
3 What are the risks (to the organization or the subject) of providing the digital service?
At this point, the agency understands that some level of proofing is required. Step 3 is intended to look at the potential impacts of an identity proofing failure to determine if IAL2 or IAL3 is the most appropriate selection. The primary identity proofing failure an agency may encounter is accepting a falsified identity as true, therefore providing a service or benefit to the wrong or ineligible person. In addition, proofing, when not required, or collecting more information than needed is a risk in and of itself. Hence, obtaining verified attribute information when not needed is also considered an identity proofing failure. This step should identify if the agency answered Step 1 and 2 incorrectly, realizing they do not need personal information to deliver the service. Risk should be considered from the perspective of the organization and to the user, since one may not be negatively impacted while the other could be significantly harmed. Agency risk management processes should commence with this step.
4 Do you need to resolve an identity uniquely?
Step 4 is intended to determine if the personal information required by the agency will ultimately resolve to a unique identity. In other words, the agency needs to know the full identity of the subject accessing the digital service, and pseudonymous access, even with a few validated and verified attributes, is not possible. If the agency needs to uniquely identify the subject, the process can end. However, the agency should consider if Step 5 is of value to them, as the acceptance of claims will reduce exposure to the risk of over collecting and storing more personal information than is necessary.

1364

5 Can you accept references? Step 5 focuses on whether the digital service can be provided without having access to full attribute values. This does not mean all attributes must be delivered as claims, but this step does ask the agency to look at each personal attribute they have deemed necessary, and identify which can suffice as claims and which need to be complete values. A federated environment is best suited for receiving claims, as the digital service provider is not in control of the attribute information to start with. If the application also performs all required identity proofing, claims may not make sense since full values are already under the digital service provider's control.
6 Use references if you can complete the transaction or offer the service without complete attribute values. If the agency has reached Step 6, claims should be used. This step identifies the digital service as an excellent candidate for accepting federated attribute references from a CSP (or multiple CSPs), since it has been determined that complete attribute values are not needed to deliver the digital service.

1365
1366
1367 > Note: Agencies should also consider their constituents' demographics when selecting the
1368 > most appropriate proofing process. While not a function of IAL selection, certain
1369 > proofing processes may be more appropriate for some demographics than others.
1370 > Agencies will benefit as this type of analysis ensures the greatest opportunity for their
1371 > constituents to be proofed successfully.
1372

1373 ### 6.2 Selecting AAL

1374
1375 The AAL decision tree in Figure 6-2 combines the results from the risk assessment with
1376 additional considerations related to authentication to allow agencies to select the most
1377 appropriate authentication requirements for their digital service offering.

1378 The AAL selection does not mean the digital service provider will need to issue authenticators
1379 themselves. More information on whether the agency can federate is provided in Section 7.

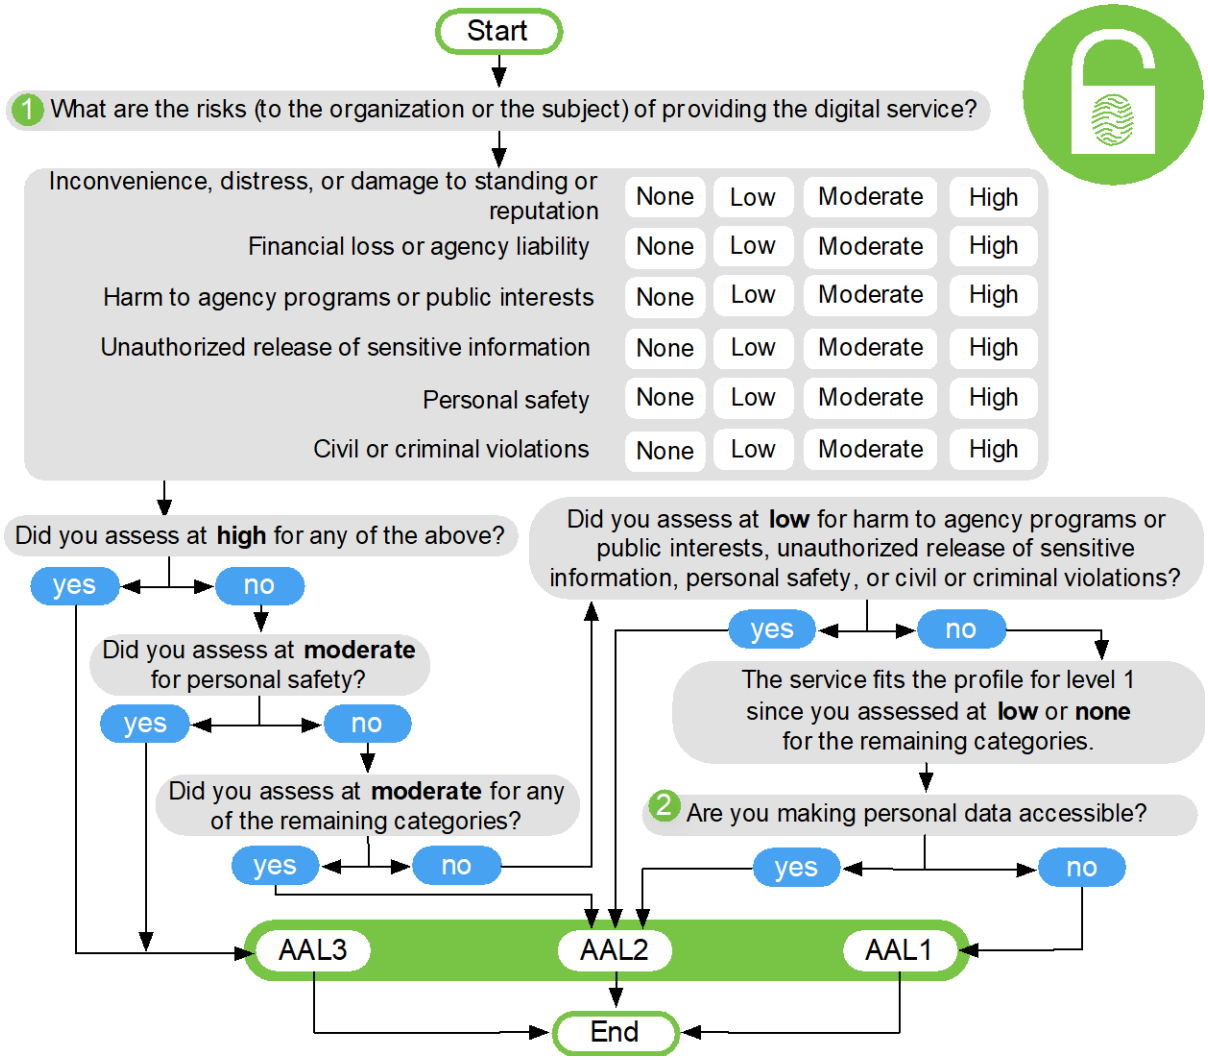

Figure 6-2 Selecting AAL

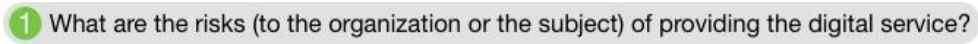

Step 1 asks agencies to look at the potential impacts of an authentication failure. In other words, what would occur if an unauthorized user accessed one or more valid user accounts? Risk should be considered from the perspective of the organization and to a valid user, since one may not be negatively impacted while the other could be significantly harmed. Agency risk management processes should commence with this step.

MFA is required when any personal information is made available online. Since the other paths in this decision tree already drive the agency to an AAL that requires MFA, the question of personal information is only raised at this point. That said, personal information release at all AALs should be considered when performing the risk assessment. An important point at this step is that the collection of personal information, if not made available online, does not need to be validated or verified to require an AAL of 2 or higher. Release of even self-asserted personal information requires account protection via MFA. Even though self-asserted information can be falsified, most users will provide accurate information to benefit from the digital service. As such, self-asserted data must be protected appropriately.

6.3 Selecting FAL

The FAL decision tree in Figure 6-3 combines the results from the risk assessment with additional considerations related to federation to allow agencies to select the most appropriate requirements for their digital service offering.

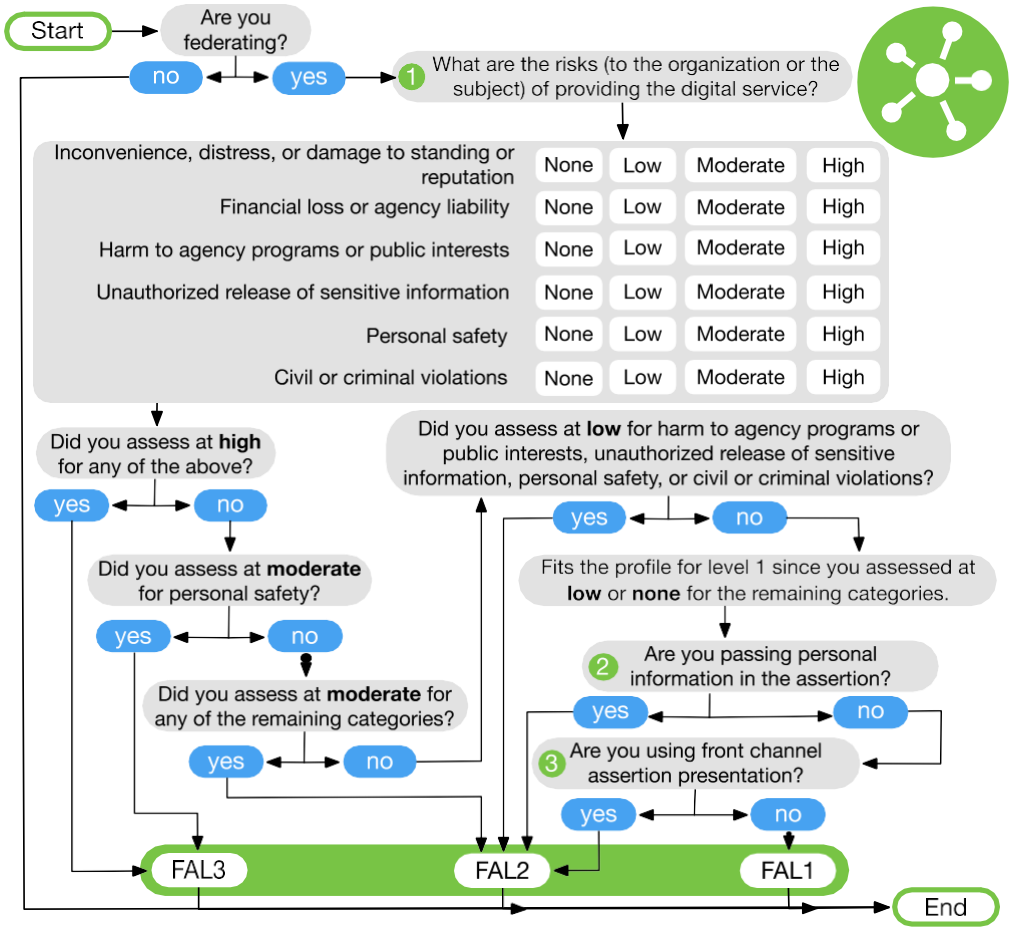

Figure 6-3 Selecting FAL

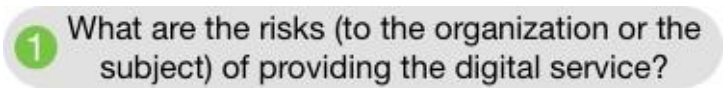

Step 1 asks agencies to look at the potential impacts of a federation failure. In other words, what would occur if an unauthorized user could compromise an assertion? Examples of compromise include use of assertion replay to impersonate a valid user or leakage of assertion information through the browser. Risk should be considered from the perspective of the organization and to the subscriber, since one may not be negatively impacted while the other could be significantly harmed. Agency risk management processes should commence with this step.

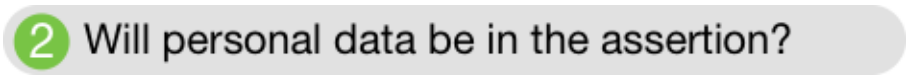

FAL2 is required when any personal information is passed in an assertion. Personal information release at all FALs should be considered when performing the risk assessment. FAL2 or higher is required when any personal information is contained in an assertion, as the audience and encryption requirements at FAL1 are not sufficient to protect personal information from being released. Release of even self-asserted personal information requires assertion protection via FAL2. Even though self-asserted information can be falsified, most users will provide accurate information to benefit from the digital service. However, when personal information is available to the RP via an authorized API call, such information need not be included in the assertion itself. Since the assertion no longer includes personal information, it need not be encrypted and this FAL requirement does not apply.

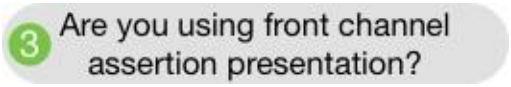

RPs should use a back-channel presentation mechanism as described in SP 800-63C, Section 7.1 where possible as such mechanisms allow for greater privacy and security. Since the subscriber handles only an assertion reference and not the assertion itself, there is less chance of leakage of attributes or other sensitive information found in the assertion to the subscriber's browser or other programs. As the RP directly presents the assertion reference to the IdP, the IdP can often take steps to identify and authenticate the RP during this step. Furthermore, as the RP fetches the assertion directly from the IdP over an authenticated protected channel, there are fewer opportunities for an attacker to inject an assertion into an RP.

All FALs require assertions to have a baseline of protections, including signatures, expirations, audience restrictions, and others enumerated in SP 800-63C. When taken together, these measures make it so that assertions cannot be created or modified by an unauthorized party, and that an RP will not accept an assertion created for a different system.

6.4 Combining *x*ALs

This guideline introduces a model where individual *x*ALs can be selected without requiring parity to each other. While options exist to select varying *x*ALs for a system, in many instances the same level will be chosen for all *x*ALs.

The ability to combine varying *x*ALs offers significant flexibility to agencies, but not all combinations are possible due to the nature of the data collected from an individual and the

authenticators to protect that data. Table 6-2 details valid combinations of IAL and AAL to ensure personal information remains protected by MFA.

Table 6-2 Acceptable Combinations of IAL and AAL

	AAL1	AAL2	AAL3
IAL1: Without personal data	Allowed	Allowed	Allowed
IAL1: With personal data	**NO**	Allowed	Allowed
IAL2	**NO**	Allowed	Allowed
IAL3	**NO**	Allowed	Allowed

Note: Per Executive Order 13681 [EO 13681], the release of personal data requires protection with MFA, even if the personal data is self-asserted and not validated. When the transaction does not make personal data accessible, authentication may occur at AAL1, although providing an option for the user to choose stronger authentication is recommended. In addition, it may be possible at IAL1 to self-assert information that is not personal, in which case AAL1 is acceptable.

7 Federation Considerations

This section is informative.

This guideline and its companion volumes are agnostic to the authentication and identity proofing architecture an agency selects. However, there are scenarios an agency may encounter that make identity federation potentially more efficient and effective than establishing identity services local to the agency or individual applications. The following list details scenarios where, if any apply, the agency may consider federation a viable option. This list does not take into consideration any economic benefits or weaknesses of federation vs. localized identity architectures.

Federate authenticators when:

1. Potential users already have an authenticator at or above required AAL.
2. Multiple credential form factors are required to cover all possible user communities.
3. Agency does not have infrastructure to support authentication management (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.

Federate attributes when:

1. Pseudonymity is required, necessary, feasible, or important to stakeholders accessing the service.
2. Access to the service only requires a partial attribute list.
3. Access to the service only requires at least one attribute reference.
4. The agency 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), such that agency does not need to locally persist the data.

8 References

This section is informative.

8.1 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.

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

[ESIG] Federal CIO Council, *Use of Electronic Signatures in Federal Organization Transactions*, January 25, 2013, available at: https://cio.gov/wp-content/uploads/downloads/2014/03/Use_of_ESignatures_in_Federal_Agency_Transactions_v1-0_20130125.pdf.

[FISMA] *Federal Information Security Modernization Act of 2014*, available at: https://www.congress.gov/bill/113th-congress/senate-bill/2521.

[HSPD-12] Department of Homeland Security, *Homeland Security Presidential Directive 12: Policy for a Common Identification Standard for Federal Employees and Contractors*, August 27, 2004, available at: https://www.dhs.gov/homeland-security-presidential-directive-12.

[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.

[M-04-04] OMB Memorandum M-04-04, *E-Authentication Guidance for Federal Agencies*, December 16, 2003, available at: https://georgewbush-whitehouse.archives.gov/omb/memoranda/fy04/m04-04.pdf.

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

[NIST RMF] Risk Management Framework Overview, available at http://csrc.nist.gov/groups/SMA/fisma/framework.html.

[RFC 5280] IETF, *Internet X.509 Public Key Infrastructure Certificate and CRL Profile*, RFC 5280, DOI 10.17487/RFC5280, May 2008, https://doi.org/10.17487/RFC5280.

[Steiner] Steiner, Peter. “On the Internet, nobody knows you’re a dog”, *The New Yorker*, July 5, 1993.

8.2 Standards

[BCP 195] 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, https://doi.org/10.17487/RFC7525.

[Canada] Government of Canada, *Standard on Identity and Credential Assurance*, February 1, 2013, available at: https://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=26776.

[ISO 9241-11] 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, B., Jones, M., de Medeiros, B., and C. Mortimore, *OpenID Connect Core 1.0 incorporating errata set 1*, November, 2014, available at: https://openid.net/specs/openid-connect-core-1_0.html.

8.3 NIST Special Publications

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

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

[SP 800-37] NIST Special Publication 800-37 Revision 1, *Guide for Applying the Risk Management Framework to Federal Information Systems, A Security Life Cycle Approach*, February 2010 (updated June 5, 2014), https://doi.org/10.6028/NIST.SP.800-37r1.

[SP 800-52] NIST Special Publication 800-52 Revision 1, *Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations*, April 2014 https://doi.org/10.6028/NIST.SP.800-52r1.

[SP 800-53A] NIST Special Publication 800-53A Revision 4, *Assessing Security and Privacy Controls in Federal Information Systems and Organizations, Building Effective Assessment Plans*, December 2014 (updated December 18, 2014), https://doi.org/10.6028/NIST.SP.800-53Ar4.

8.4 Federal Information Processing Standards

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

[FIPS 201] Federal Information Processing Standard Publication 201-2, *Personal Identity Verification (PIV) of Federal Employees and Contractors*, August 2013, https://doi.org/10.6028/NIST.FIPS.201-2.

Appendix A—Definitions and Abbreviations

This section is normative.

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.

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 man-in-the-middle (MitM), impersonation, and session hijacking.

Address of Record

The validated and verified location (physical or digital) where an individual can receive communications using approved mechanisms.

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, which 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, which 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.

Attribute

A quality or characteristic ascribed to someone or something.

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. Attribute bundles are synonymous with OpenID Connect scopes [OpenID Connect Core 1.0].

Attribute Reference

A statement asserting a property of a subscriber without necessarily containing identity information, independent of format. For example, for the attribute "birthday," a reference could be "older than 18" or "born in December."

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."

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) [BCP 195] 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

Verifying the identity of a user, process, or device, often as a prerequisite to allowing access to a system's resources.

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 re-authenticate 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 Protocol Run

An exchange of messages between a claimant and a verifier that results in authentication (or authentication failure) between the two parties.

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 previous editions of SP 800-63, this was referred to as a *token*.

Authenticator 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.

Authorization Component

A set of data items issued to an RP by an IdP during an identity federation transaction that grants the RP authorized access to a set of APIs (e.g., an OAuth access token). This credential can be separate from the assertion provided by the federation protocol (e.g., an OpenID Connect ID Token).

Authorize

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

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.

Biometrics

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

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.

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.

While common usage often assumes that the subscriber maintains the credential, these guidelines also use the term to refer to electronic records maintained by the CSP that establish binding between the subscriber's authenticator(s) and identity.

Credential Service Provider (CSP)

A trusted entity that issues or registers subscriber authenticators and issues electronic credentials to subscribers. A CSP may be an independent third party or issue credentials for its own use.

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 where the secret is a cryptographic key.

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 SP 800-57 Part 1.

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 Credential

A credential issued based on proof of possession and control of an authenticator associated with a previously issued credential, so as not to duplicate the identity proofing process.

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: Enabling the processing of PII or events without association to individuals or devices beyond the operational requirements of the system.

Diversionary

In regards to KBV, a multiple-choice question for which all answers provided are incorrect, requiring the applicant to select an option similar to “none of the above.”

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.

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: http://www.nist.gov/itl/fips.cfm

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".

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:

One-way - It is computationally infeasible to find any input that maps to any pre-specified output; and

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 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)

The party that manages the subscriber's primary authentication credentials and issues assertions derived from those credentials. This is commonly the CSP as discussed within this document suite.

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 SP 800-63C Section 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.

Man-in-the-Middle Attack (MitM)

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.

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, man-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 SHALL 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 protocol run or by penetrating a system and stealing security files) that he/she is able to analyze in a system of his/her own choosing.

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, Personally Identifiable Information 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.

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 Credentials

Credentials that cannot be disclosed by the CSP because the contents can be used to compromise the authenticator.

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*, involve measurement and analysis of anatomical characteristics or involuntary or voluntary reactions, in order to determine if a biometric sample is being captured from a living subject present at the point of capture.

Protected Session

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

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

Pseudonym

A name other than a legal name.

Pseudonymity

The use of a pseudonym to identify a subject.

Pseudonymous Identifier

A meaningless but unique number that does not allow the RP to infer anything regarding the subscriber but which does permit the RP to associate multiple interactions with the subscriber's claimed identity.

Public Credentials

Credentials that describe the binding in a way that does not compromise the authenticator.

Public Key

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

Public Key Certificate

A digital document issued and digitally signed by the private key of a certificate authority that binds an identifier to a subscriber to a public key. The certificate indicates that the subscriber identified in the certificate has sole control and access to the private key. See also RFC 5280.

Public Key Infrastructure (PKI)

A set of policies, processes, server platforms, software, and workstations used for the purpose of administering certificates and public-private key pairs, including the ability to issue, maintain, and revoke public key certificates.

Re-authentication

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 the subscriber's authenticator(s) and credentials or a verifier's assertion of a claimant'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.

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 or CSP in lieu of the subscriber's authentication credentials.

Session Hijack Attack

An attack in which the attacker is able to insert himself or herself 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.

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

A party who has received a credential or authenticator from a CSP.

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.

Token

See Authenticator.

Token Authenticator

See Authenticator Output.

Token Secret

See Authenticator Secret.

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 RFC 5246. 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 [SP 800-52], 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

Per ISO/IEC 9241-11: 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.

Verifier

An entity that verifies the claimant's identity by verifying the claimant's possession and control of one or two authenticators using an authentication protocol. To do this, the verifier may also need to validate credentials that link the authenticator(s) to the subscriber's identifier and check their status.

Verifier Impersonation

A scenario where the attacker impersonates the verifier in an authentication protocol, usually to capture information that can be used to masquerade as a subscriber to the real verifier. In previous editions of SP 800-63, authentication protocols that are resistant to verifier impersonation have been described as "strongly MitM resistant".

Supervised Remote 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.

Weakly Bound Credentials

Credentials that are bound to a subscriber in a manner than can be modified without invalidating the credential.

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.

Table A.2 Abbreviations

Abbreviation	Term
ABAC	Attribute Based Access Control
AAL	Authenticator Assurance Level
AS	Authentication Server
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

2286

Abbreviation	Term
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
IAL	Identity Assurance Level
IM	Identity Manager
IdP	Identity Provider
IoT	Internet of Things
ISO/IEC	International Organization for Standardization/International Electrotechnical Commission

2287

Abbreviation	Term
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
MitM	Man-in-the-Middle
MitMA	Man-in-the-Middle Attack
MFA	Multi-Factor Authentication

2288

Abbreviation	Term
N/A	Not Applicable
NARA	National Archives and Records Administration
OMB	Office of Management and Budget
OTP	One-Time Password
PAD	Presentation Attack Detection
PHI	Personal Health Information
PIA	Privacy Impact Assessment
PII	Personally Identifiable Information
PIN	Personal Identification Number
PKI	Public Key Infrastructure
PL	Public Law

2289

Abbreviation	Term
PSTN	Public Switched Telephone Network
RA	Registration Authority
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

Abbreviation	Term
TEE	Trusted Execution Environment
TGS	Ticket Granting Server
TGT	Ticket Granting Ticket
TLS	Transport Layer Security
TPM	Trusted Platform Module
VOIP	Voice-Over-IP