

Withdrawn Draft

Warning Notice

The attached draft document has been withdrawn and is provided solely for historical purposes. It has been followed by the document identified below.

Withdrawal Date August 21, 2024

Original Release Date December 16, 2022

The attached draft document is followed by:

Status Second Public Draft (2pd)

Series/Number NIST SP 800-63C-4 2pd

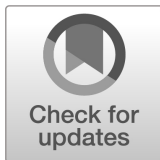
Title Digital Identity Guidelines: Federation and Assertions

Publication Date August 2024

DOI <https://doi.org/10.6028/NIST.SP.800-63c-4.2pd>

CSRC URL <https://csrc.nist.gov/pubs/sp/800/63/c/4/2pd>

Additional Information <https://www.nist.gov/identity-access-management/nist-special-publication-800-63-digital-identity-guidelines>



1 **NIST Special Publication**
2 **NIST SP 800-63C-4 ipd**

3 **Digital Identity Guidelines**
4 **Federation and Assertions**

5 Initial Public Draft

6 David Temoshok
7 Justin P. Richer
8 Yee-Yin Choong
9 James L. Fenton
10 Naomi Lefkovitz
11 Andrew Regenscheid

12 This publication is available free of charge from:
13 <https://doi.org/10.6028/NIST.SP.800-63c-4.ipd>

15 **NIST Special Publication**
16 **NIST SP 800-63C-4 ipd**
17 **Digital Identity Guidelines**
18 **Federation and Assertions**
19 **Initial Public Draft**

20 David Temoshok
21 Naomi Lefkovitz
22 *Applied Cybersecurity Division*
23 *Information Technology Laboratory*

24 Yee-Yin Choong
25 *Information Access Division*
26 *Information Technology Laboratory*


27 Andrew Regenscheid
28 *Computer Security Division*
29 *Information Technology Laboratory*

30 Justin P. Richer
31 *Bespoke Engineering*

32 James L. Fenton
33 *Altmode Networks*

34 This publication is available free of charge from:
35 <https://doi.org/10.6028/NIST.SP.800-63c-4.ipd>

36 December 2022



38 U.S. Department of Commerce
39 *Gina M. Raimondo, Secretary*

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

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

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

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

Authority

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

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

NIST Technical Series Policies

[Copyright, Fair Use, and Licensing Statements](#)
[NIST Technical Series Publication Identifier Syntax](#)

76 **Publication History**

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

79 **How to Cite this NIST Technical Series Publication**

80 Temoshok D, Richer JP, Choong YY, Fenton JL, Lefkovitz N, Regenscheid A (2022)
81 Digital Identity Guidelines: Federation and Assertions. (National Institute of Standards
82 and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-63C-4
83 ipd. <https://doi.org/10.6028/NIST.SP.800-63c-4.ipd>

84 **Author ORCID iDs**

85 David Temoshok: 0000-0001-6195-0331
86 Justin P. Richer: 0000-0003-2130-5180
87 Yee-Yin Choong: 0000-0002-3889-6047
88 James L. Fenton: 0000-0002-2344-4291
89 Naomi Lefkovitz: 0000-0003-3777-3106
90 Andrew Regenscheid: 0000-0002-3930-527X

91 **Public Comment Period**

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

93 **Submit Comments**

94 <mailto:dig-comments@nist.gov>

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

Reports on Computer Systems Technology

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

Abstract

These guidelines provide technical requirements for federal agencies implementing digital identity services and are not intended to constrain the development or use of standards outside of this purpose. This guideline focuses on the use of federated identity and the use of assertions to implement identity federations. Federation allows a given credential service provider to provide authentication attributes and (optionally) subscriber attributes to a number of separately-administered relying parties. Similarly, relying parties may use more than one credential service provider. This publication will supersede NIST Special Publication (SP) 800-63C.

Keywords

assertions; authentication; credential service provider; digital authentication; electronic authentication; electronic credentials; federations.

Note to Reviewers

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

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

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

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

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

Federation and Assertions

- What additional privacy considerations (e.g., revocation of consent, limitations of use) may be required to account for the use of identity and provisioning APIs that had not previously been discussed in the guidelines?

- Is the updated text and introduction of “bound authenticators” sufficiently clear to allow for practical implementations of federation assurance level (FAL) 3 transactions? What complications or challenges are anticipated based on the updated guidance?

General

- Is there an element of this guidance that you think is missing or could be expanded?
- Is any language in the guidance confusing or hard to understand? Should we add definitions or additional context to any language?
- Does the guidance sufficiently address privacy?
- Does the guidance sufficiently address equity?
 - What equity assessment methods, impact evaluation models, or metrics could we reference to better support organizations in preventing or detecting disparate impacts that could arise as a result of identity verification technologies or processes?
- What specific implementation guidance, reference architectures, metrics, or other supporting resources may enable more rapid adoption and implementation of this and future iterations of the Digital Identity Guidelines?
- What applied research and measurement efforts would provide the greatest impact on the identity market and advancement of these guidelines?

Reviewers are encouraged to comment and suggest changes to the text of all four draft volumes of the NIST SP 800-63-4 suite. NIST requests that all comments be submitted by 11:59pm Eastern Time on March 24, 2023. Please submit your comments to dig-comments@nist.gov. NIST will review all comments and make them available at the [NIST Identity and Access Management website](#). Commenters are encouraged to use the comment template provided on the [NIST Computer Security Resource Center website](#).

Call for Patent Claims

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

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

- a) assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or
- b) assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this ITL draft publication either:
 - i. under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or
 - ii. without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

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

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

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

Table of Contents

223	Table of Contents	
224	1. Purpose	2
225	2. Introduction	3
226	3. Definitions and Abbreviations	5
227	4. Federation Assurance Level (FAL)	6
228	4.1. Federation Assurance Level 1 (FAL1)	7
229	4.2. Federation Assurance Level 2 (FAL2)	8
230	4.3. Federation Assurance Level 3 (FAL3)	9
231	4.4. Requesting and Processing xALs	9
232	5. Federation	12
233	5.1. Trust Agreements	13
234	5.1.1. Bilateral Trust Agreements	15
235	5.1.2. Multilateral Trust Agreements	15
236	5.1.3. Proxied Federation	17
237	5.2. Registration	18
238	5.2.1. Manual Registration	19
239	5.2.2. Dynamic Registration	20
240	5.3. Authentication and Attribute Disclosure	21
241	5.3.1. IdP Allowlists of RPs	21
242	5.3.2. IdP Blocklists of RPs	22
243	5.3.3. IdP Runtime Decisions	22
244	5.3.4. RP Allowlists of IdPs	23
245	5.3.5. RP Blocklists of IdPs	23
246	5.3.6. RP Runtime Decisions	23
247	5.4. RP Subscriber Accounts	24
248	5.4.1. Provisioning Models	25
249	5.4.2. Attribute Synchronization	27
250	5.4.3. Provisioning APIs	28
251	5.4.4. Attribute Collection	29
252	5.4.5. Time-based Removal of RP Subscriber Accounts	29

253	5.5. Privacy Requirements	30
254	5.6. Reauthentication and Session Requirements in Federated Environments . .	31
255	5.7. Shared Signaling	32
256	6. Assertions	34
257	6.1. Assertion Binding	36
258	6.1.1. Bearer Assertions	36
259	6.1.2. Bound Authenticators	36
260	6.2. Assertion Protection	42
261	6.2.1. Assertion Identifier	43
262	6.2.2. Signed Assertion	43
263	6.2.3. Encrypted Assertion	43
264	6.2.4. Audience Restriction	44
265	6.2.5. Pairwise Pseudonymous Identifiers	44
266	6.3. Identity APIs	45
267	6.3.1. Attribute Providers	46
268	7. Assertion Presentation	48
269	7.1. Back-Channel Presentation	48
270	7.2. Front-Channel Presentation	51
271	7.3. Protecting Information	52
272	8. Security	53
273	8.1. Federation Threats	53
274	8.2. Federation Threat Mitigation Strategies	54
275	9. Privacy Considerations	56
276	9.1. Minimizing Tracking and Profiling	56
277	9.2. Notice and Consent	57
278	9.3. Data Minimization	58
279	9.4. Agency-Specific Privacy Compliance	58
280	9.5. Blinding in Proxied Federation	59
281	10. Usability Considerations	61
282	10.1. General Usability Considerations	62

283	10.2. Specific Usability Considerations	63
284	10.2.1. User Perspectives on Online Identity	63
285	10.2.2. User Perspectives of Trust and Benefits	64
286	10.2.3. User Mental Models and Beliefs	66
287	11. Equity Considerations	67
288	12. Examples	69
289	12.1. Security Assertion Markup Language (SAML)	69
290	12.2. Kerberos Tickets	70
291	12.3. OpenID Connect	71
292	References	72
293	General References	72
294	Standards	72
295	NIST Special Publications	73
296	Federal Information Processing Standards	73
297	Appendix A. Changelog	74

298 List of Tables

299	1. Federation Assurance Levels	7
300	2. Federation Threats	54
301	3. Mitigating Federation Threats	55
302	4. Proxy Characteristics	60

303 List of Figures

304	1. Federation Overview	12
305	2. Federation Authority	16
306	3. Federation Proxy	17
307	4. Manual Registration	19
308	5. Dynamic Registration	20
309	6. Just-In-Time Provisioning	25
310	7. Pre-Provisioning	26
311	8. Ephemeral Provisioning	27
312	9. IdP-Managed Bound Authenticators	38
313	10. RP-Managed Bound Authenticators	39
314	11. Binding Ceremony	41
315	12. Back-channel Presentation	49

316 13. Front-channel Presentation 51

317 **Acknowledgments**

318 The authors would like to thank their fellow collaborators on the current revision of this
319 special publication, Christine Abruzzi, Ryan Galluzzo, Sarbari Gupta, Connie LaSalle,
320 and Diana Proud-Madruga, as well as Kerrienne Buchanan for her contributions and
321 review. The authors would like to also acknowledge the past contributions of Donna F.
322 Dodson, Elaine M. Newton, Ray A. Perlener, W. Timothy Polk, Emad A. Nabbus, Paul A.
323 Grassi, Ellen M. Nadeau, Kristen K. Greene, Mary F. Theofanos, Sarah K. Squire, Jamie
324 M. Danker, Kaitlin Boeckl, Kat Megas, Ben Piccarreta, and Danna Gabel O'Rourke.

1. Purpose

This section is informative.

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

This document, SP 800-63C, provides requirements to identity providers (IdPs) and relying parties (RPs) of federated identity systems. Federation allows a given IdP to provide authentication attributes and (optionally) subscriber attributes to a number of separately-administered RPs through the use of federation protocols and assertions. Similarly, RPs can use more than one IdP as sources of identities.

2. Introduction

This section is informative.

Federation is a process that allows for the conveyance of authentication attributes and subscriber attributes across networked systems. In a federation scenario, the CSP provides a service known as an identity provider, or IdP. The IdP acts as a verifier for authenticators issued by the CSP. The IdP sends a message, called an assertion, about this authentication event to the RP. The RP receives the assertion provided by the IdP and uses it for authentication and authorization decisions, but the RP does not verify the authenticator directly.

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

In a federated identity scenario, the subscriber does not authenticate directly to the RP. Instead, the federation protocol defines a mechanism for an IdP to generate an assertion associated with a subscriber, generally in response to an explicit request from the RP. The IdP is responsible for authenticating the subscriber (though it may use session management as described in [SP800-63B], Sec. 7). The federation process allows the subscriber to obtain services from multiple RPs without the need to hold or maintain separate authenticators at each RP, a process sometimes known as *single sign-on*.

The subscriber is uniquely identified to the RP by a *federated identifier*, which is a logical combination of the *subject identifier* as asserted by the IdP as well as a unique identifier for the IdP itself. This multi-part identifier pattern is required because different IdPs manage their subject identifiers independently, and could therefore potentially collide in their choices of subject identifiers for different subjects. Therefore, it is imperative that an RP never process the subject identifier without taking into account which IdP issued that subject identifier.

An assertion includes a federated identifier for the subscriber, allowing association of the subscriber with their interactions with the RP over multiple authenticated sessions. Assertions may also include attribute values or derived attribute values that further characterize the subscriber and support authorization decisions at the RP. Additional attributes may also be available outside of the assertion as part of the larger federation protocol. These attribute values and derived attribute values are often used in determining access privileges for attribute-based access control (ABAC) or facilitating a transaction (e.g., providing a shipping address).

Federation requires relatively complex multiparty protocols that have subtle security and privacy requirements. When evaluating a particular federation structure, it may be

instructive to break it down into its component interactions: the subscriber to the IdP, the IdP to the RP, and the subscriber to the RP. Each party in a federation protocol bears specific responsibilities and expectations that must be fulfilled in order for the federated system to function as intended.

The IdP maintains a record for the subscriber that augments the *subscriber account* defined in [SP800-63A] with a set of federation-specific items, including but not limited to the following:

- One or more external subject identifiers, for use with a federation protocol
- A set of access rights, detailing which RPs can access which attributes of the subscriber account (such as runtime decisions by the subscriber)
- Federated account usage information
- Additional attributes collected or assigned by the IdP to the subscriber

The RP often maintains an *RP subscriber account* for the subscriber, which is derived from the augmented subscriber account information disclosed to the RP by the IdP. The RP subscriber account also contains information local to the RP itself, as described in Sec. 5.4.

The requirements in this document build on the requirements in the other volumes of these guidelines. Authentication between the subscriber and the IdP will be based on the authentication mechanisms presented in [SP800-63B], while the federation protocol will convey attributes to the RP established at the IdP using procedures in [SP800-63A] (along with other attributes).

The following table states which sections of the document are normative and which are informative:

- 1 Purpose *Informative*
- 2 Introduction *Informative*
- 3 Definitions and Abbreviations *Informative*
- 4 Federation Assurance Level (FAL) *Normative*
- 5 Federation *Normative*
- 6 Assertion *Normative*
- 7 Assertion Presentation *Normative*
- 8 Security *Informative*
- 9 Privacy Considerations *Informative*
- 10 Usability Considerations *Informative*
- 11 Equity Considerations *Informative*
- 12 Examples *Informative*
- References *Informative*

410 **3. Definitions and Abbreviations**

411 See [\[SP800-63\]](#), Appendix A for a complete set of definitions and abbreviations.

4. Federation Assurance Level (FAL)

This section is normative.

This section defines allowable *federation assurance levels* (FALs). The FAL describes requirements for securing federation transactions, including requirements on how relationships between IdPs and RPs are established and how assertions are presented and protected. These levels can be requested by an RP at runtime or required by the configuration of both the RP and the IdP for a given transaction. The FAL provides assurances for the RP receiving the assertion as well as assurances for the IdP creating the assertion to be used by an RP.

While many different federation implementation options are possible, the FAL is intended to provide clear guidance representing increasingly secure deployment options. See [\[SP800-63\]](#) for details on how to choose the most appropriate FAL.

Each FAL is characterized by a set of requirements that increase the security and complexity as the FAL increases. These requirements are listed here and expanded in other sections of this document:

Cryptographic Verifiability

The assertion presented in the federation protocol is traceable back to a specific IdP that issued it, and that connection can be verified with a cryptographic mechanism such as a digital signature or MAC. This also allows the RP to verify that the assertion was not modified or forged. This is required at all FALs.

Audience Restriction

The assertion presented in the federation protocol is targeted to a specific RP and the RP can verify that it is the intended audience of the assertion. This is required at all FALs.

Injection Protection

The RP is strongly protected from an attacker presenting an assertion in circumstances outside a current federation transaction request.

Trust Agreement

The IdP and RP have agreed to participate in a federation transaction with each other for the purposes of logging in the subscriber to the RP. This can be traced back to a static agreement between the parties or occur implicitly from the connection itself.

Registration

The IdP and RP have exchanged identifiers and key material to allow for the verification of assertions and other artifacts during future federation transactions.

Presentation

The assertion can be presented to the RP either on its own (as a bearer assertion) or in concert with a bound authenticator presented by the subscriber.

Table 1 provides a non-normative summary of aspects for each FAL. Each successive level subsumes and fulfills all requirements of lower levels (e.g., a federation process at FAL3 can be accepted at FAL2 or FAL1 since FAL3 satisfies all the requirements of these lower levels). Combinations not found in the Table 1 are possible but outside the scope of this volume.

Table 1. Federation Assurance Levels

FAL	Injection Protection	Trust Agreement	Registration	Presentation
1	Recommended	Dynamic or Static	Dynamic or Static	Bearer Assertion
2	Required	Static	Dynamic or Static	Bearer Assertion
3	Required	Static	Static	Assertion and Bound Authenticator

At all FALs, all assertions **SHALL** be used with a federation protocol as described in Sec. 5. All assertions **SHALL** comply with the detailed requirements in Sec. 6. All assertions **SHALL** be presented using one of the methods described in Sec. 7. Examples of assertions used in federated protocols include the ID Token in OpenID Connect [OIDC] and assertions written in the Security Assertion Markup Language [SAML].

At all FALs, the IdP **SHALL** employ appropriately tailored security controls (to include control enhancements) from the moderate or high baseline of security controls defined in [SP800-53] or equivalent federal (e.g., [FEDRAMP]) or industry standard.

4.1. Federation Assurance Level 1 (FAL1)

At FAL1, the assertion being generated by the IdP **SHALL** meet a core set of requirements defined in Sec. 6, including protection against modification or construction by an attacker by having the assertion contents signed by the IdP using approved cryptography. An RP **SHALL** verify the origin and integrity of the assertion upon receipt, as discussed in Sec. 6, ensuring that the assertion has originated from the expected source.

All assertions at FAL1 **SHALL** be audience-restricted to a specific RP or set of RPs, and the RP **SHALL** validate that it is one of the targeted RPs for the given assertion. The IdP **SHALL** ensure that any party holding the assertion, including the RP, is unable to impersonate the IdP at a non-targeted RP by protecting the assertion with a signature and key using approved cryptography. If the assertion is protected by a digital signature

using an asymmetric key, the IdP **MAY** use the same public and private key pair to sign assertions to multiple RPs. The IdP **MAY** publish its public key in a verifiable fashion, such as at an HTTPS-protected URL at a well-known location. If the assertion is protected by a keyed message authentication code (MAC) using a shared key, the IdP **SHALL** use a different shared key for each RP.

At FAL1, the trust agreement between the IdP and RP **MAY** be established entirely dynamically. For instance, the subscriber can identify their chosen IdP to the RP at runtime, allowing the RP to discover the IdP's parameters and register itself for use by the subscriber. The subscriber is prompted by the IdP to determine which attributes are released to the RP, and for what purposes. In this example, the trust between the IdP and RP is driven entirely by the desires and actions of the subscriber. Note that at FAL1, it is still possible for the trust agreement and registration to happen statically.

In existing federation protocols, FAL1 can be implemented with the OpenID Connect Implicit Client profile [OIDC-Implicit], the OpenID Connect Hybrid Client profile in [OIDC], or the SAML Web SSO [SAML-WebSSO] profile with no additional features. In each of these profiles, the assertion is signed by the IdP and the RP is identified in a portion of the assertion covered by the signature.

4.2. Federation Assurance Level 2 (FAL2)

All the requirements for FAL1 apply at FAL2 except where overridden by more specific or stringent requirements here.

At FAL2, the assertion **SHALL** also be strongly protected from being injected by an attacker. To accomplish this, the assertion **SHOULD** be presented using back channel presentation as discussed in Sec. 7.1, as in the OpenID Connect Basic Client profile [OIDC-Basic]. In this presentation method, the RP fetches the assertion directly from the IdP by using a single-use assertion reference, thereby preventing an attacker from injecting the assertions through an external access point. If front channel presentation is used as discussed in Sec. 7.2, additional injection protections **SHALL** be implemented by the RP.

Regardless of the presentation method used, injection attacks can be further mitigated by always requiring that the federation transaction start at the RP instead of being initiated by the IdP, thereby allowing the RP to associate an incoming assertion with a specific request that the subscriber initiated within a continuous session.

At FAL2, the trust agreement between the IdP and RP **SHALL** be established statically, including establishing limits of which attributes are made available to the RP and for what purpose. This trust agreement **MAY** be bilateral between the IdP and RP or **MAY** be managed through the use of a multilateral federation partnership. The registration **MAY** be dynamic, provided that the RP and IdP can prove their connection at runtime to the established trust agreement between them. Such methods for this proof vary by federation

protocol, but can include presentation of software attestations and proof of control over URLs at trusted domains.

Government-operated IdPs asserting authentication at FAL2 **SHALL** protect keys used for signing or encrypting those assertions with mechanisms validated at [FIPS140] Level 1 or higher.

4.3. Federation Assurance Level 3 (FAL3)

All the requirements at FAL1 and FAL2 apply at FAL3 except where overridden by more specific or stringent requirements here.

At FAL3, the subscriber **SHALL** authenticate to the RP by presenting an authenticator directly to the RP in addition to presenting an assertion. The authenticator presented is known as a *bound authenticator*, described in Sec. 6.1.2. For example, the subscriber goes through a federated login process at the IdP and RP, and the RP then prompts the subscriber for a bound authenticator that is associated with that RP subscriber account. The bound authenticator presented at FAL3 need not be the same authenticator used by the subscriber to authenticate to the IdP. The assertion is used to identify the subscriber to the RP while the bound authenticator gives very high assurance that the party attempting to log in is the subscriber identified in the assertion. FAL3 is not reached at the RP until the subscriber authenticates with the bound authenticator and the RP verifies that the authenticator presented is correctly bound to the RP subscriber account identified by the assertion.

At FAL3, the trust agreement and registration between the IdP and RP **SHALL** be established statically. All identifying key material and federation parameters for all parties (including the list of attributes sent to the RP) **SHALL** be fixed ahead of time, before the federated authentication process can take place. Runtime decisions **MAY** be used to further limit what is sent between parties in the federated authentication process (e.g., a runtime decision could opt to not disclose an email address even though this attribute was included in the parameters of the trust agreement).

IdPs asserting authentication at FAL3 **SHALL** protect keys used for signing or encrypting those assertions with mechanisms validated at [FIPS140] Level 1 or higher.

4.4. Requesting and Processing xALs

Since an IdP is capable of asserting the identities of many different subscribers with a variety of authenticators using a variety of federation parameters, the IAL, AAL, and FAL could vary across different federated logins, even to the same RP.

The RP **SHALL** be informed of the following information for each federated transaction:

- The IAL of the subscriber account being presented to the RP, or an indication that no IAL claim is being made

- The AAL of the currently active session of the subscriber at the IdP, or an indication that no AAL claim is being made
- The FAL of the federated transaction

The RP gets this xAL information from a combination of parameters in the trust agreement as described in [Sec. 5.1](#) and information included in the assertion as described in [Sec. 6](#). If the xAL is unchanging for all messages between the IdP and RP, the xAL information **SHALL** be included in the parameters of the trust agreement between the IdP and RP. If the xAL varies, the information **SHALL** be included as part of the assertion as discussed in [Sec. 6](#).

The IdP **MAY** indicate that no claim is made to the IAL or AAL for a given federation transaction. In such cases, no default value is assigned to the resulting xAL. That is to say, a federation transaction without an IAL declaration in either the trust agreement or the assertion is functionally considered to have “no IAL” and the RP cannot assume the account meets “IAL1”, the lowest numbered IAL described in this suite.

The RP **SHALL** determine the minimum IAL, AAL, and FAL it is willing to accept for access to any offered functionality. An RP **MAY** vary its functionality based on the IAL, AAL, and FAL of a specific federated authentication. For example, an RP can allow login at AAL2 for common functionality (e.g., viewing the status of a dam system) but require AAL3 be used for higher risk functionality (e.g., changing the flow rates of a dam system). Similarly, an RP could restrict management functionality to only certain subscriber accounts which have been identity proofed at IAL2, while allowing logins from all subscriber accounts regardless of IAL.

In a federation process, only the IdP has direct access to the details of the subscriber account, which determines the applicable IAL, and the authentication event at the IdP, which determines the applicable AAL. Consequently, the RP **SHALL** consider the IdP’s declaration of the IAL and AAL as the sole source of these levels for a given federated transaction.

The RP **SHALL** ensure that the federation transaction meets the requirements of the FAL declared in the assertion. For example, the RP needs to ensure the presentation method meets the injection protection requirements at FAL2 and above, and that the appropriate bound authenticator is presented at FAL3.

IdPs **SHALL** support a mechanism for RPs to specify a set of minimum acceptable xALs as part of the trust agreement and **SHOULD** support the RP specifying a more strict minimum set at runtime as part of the federation transaction. When an RP requests a particular xAL, the IdP **SHOULD** fulfill that request, if possible, and **SHALL** indicate the resulting xAL in the assertion. For example, if the subscriber has an active session that was authenticated at AAL1, but the RP has requested AAL2, the IdP needs to prompt the subscriber for AAL2 authentication to step up the security of the session at the IdP during

585 the subscriber's interaction at the IdP, if possible. The IdP sends the resulting AAL as part
586 of the returned assertion, whether it is AAL1 (the original session) or AAL2 (a stepped-up
587 authentication).

5. Federation

This section is normative.

In a federation protocol, a three-party relationship is formed between the subscriber, the IdP, and the RP, as shown in Figure 1.

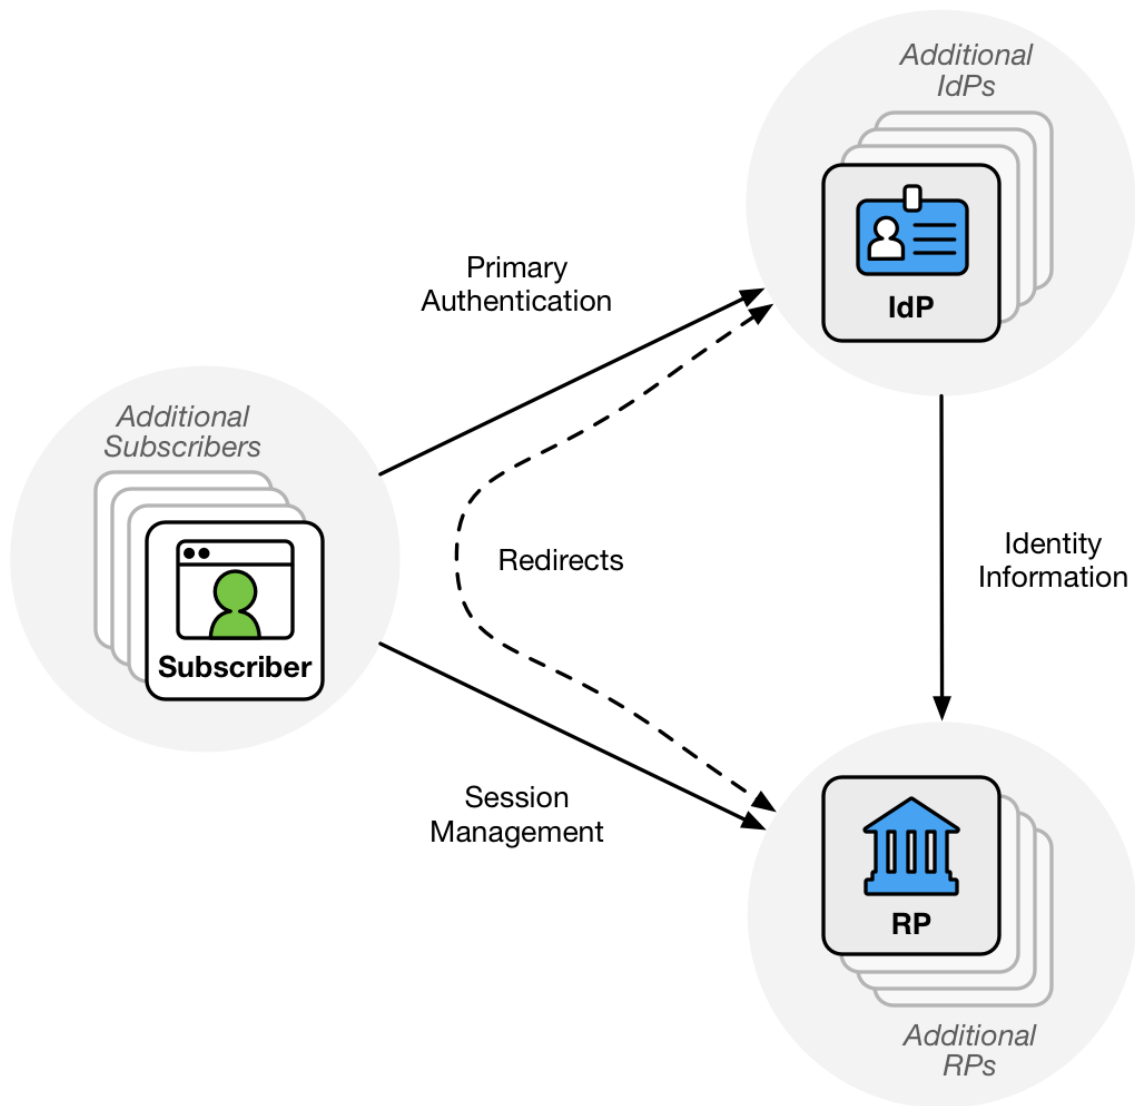


Figure 1. Federation Overview

A federation relationship between an IdP and RP is established in a multi-stage process:

1. First, the IdP and RP agree to enter into a trust agreement. This agreement can be bilateral between the parties, multilateral at the behest of an authority, or proxied through a trusted party. This step represents initial permission for the two systems in question to connect. Parameters of what can be requested and released are

597 established in this step, though the details of which attributes are released to a given
598 RP for a given subscriber can be deferred until a later stage.

599 2. Next, the IdP and RP perform registration to establish their trust at a protocol level,
600 allowing for information to be securely exchanged between the parties. While the
601 first step entails a policy decision representing a permission to connect, this step
602 entails establishment of credentials and identifiers representing the IdP and RP to
603 allow communication through the federation protocol. This stage can occur before
604 any subscriber tries to log in to the RP or as a response to a subscriber's attempt to
605 use an IdP at an RP.

606 3. Next, the IdP and RP determine that they want to engage in a federated
607 authentication transaction to authenticate the subscriber. As part of this, they
608 determine which attributes about the subscriber are to be passed from the IdP
609 to the RP during this transaction. The decision made in this step builds on the
610 trust agreement established in the first step and the identities of the RP and IdP
611 established in the second step.

612 4. Finally, the subscriber authenticates to the IdP and the result of that authentication
613 event is asserted to the RP across the network. The RP processes this assertion from
614 the IdP and establishes an authenticated session with the subscriber.

615 In this transaction, the IdP acts as the verifier of the subscriber's authenticators, as
616 described in [SP800-63B]. The authentication event information is carried from the IdP
617 to the RP through the use of an assertion, described in [Sec. 6](#). The IdP can also make
618 statements about identity attributes of the subscriber as part of this assertion or through a
619 secondary identity protocol protected by an authorized credential.

620 5.1. Trust Agreements

621 IdPs that provide authentication services and RPs that consume those services are known
622 as members of a federation. From an IdP's perspective, the federation consists of the RPs
623 that it serves. From an RP's perspective, the federation consists of the IdPs that it uses.
624 This section provides an overview of and requirements for common identity federation
625 models currently in use. In each model, relationships are established between members
626 of the federation. These relationships are established in either a bilateral or multilateral
627 fashion, as described in the following sections.

628 Trust agreements **SHALL** establish the following parameters:

- 629 • The set of attributes the IdP can make available to the RP
- 630 • The population of subscriber accounts the IdP can create assertions for
- 631 • The set of attributes the RP will request (a subset of the attributes made available)
- 632 • The purpose for each attribute requested by the RP

- The authorized party responsible for decisions regarding the release of subscriber attributes
- The means of informing subscribers about attribute release to the RP
- The xALs available from the IdP
- The xALs required by the RP

Trust agreements are able to be established either statically or dynamically. In a static establishment, there is often a legal or contractual agreement binding the parties to a set of expected behaviors, rights, and requirements. The parameters of static trust agreements **SHALL** be available to all parties in the agreement, including the operator of the IdP, the operator of the RP, and affected subscribers.

In dynamic trust establishment, in contrast, the trust agreement is implicitly defined when the RP and IdP first contact each other for the purposes of a subscriber's login. The expression of the parameters of a dynamic trust agreement is driven by the federation protocol in place, and are not usually tied to a contractual agreement between the federating parties. The parameters of a dynamic trust agreement **SHALL** be disclosed to the subscriber by the RP and the IdP during the federation transaction.

The *authorized party* in a trust agreement is the organization, person, or entity that is responsible for the specific release decisions covered by the trust agreement, including the release of subscriber attributes. For a static trust agreement, the authorized party **MAY** be the organization responsible for the IdP. In this case, consent to release attributes is decided for all subscribers and established by an allowlist as described in [Sec. 5.3.1](#), allowing for the disclosure of attribute information without direct decisions and involvement by the subscriber. A static trust agreement **MAY** stipulate that an individual, such as the subscriber, is to be prompted at runtime for consent to disclose attributes as discussed in [Sec. 5.3.3](#). Since dynamic trust agreements are established by subscriber actions, the authorized party in a dynamic trust agreement is always the subscriber. Disclosure of attributes in dynamic trust agreements **SHALL** be subject to a runtime decision from the subscriber and **SHALL NOT** be subject to an allowlist at the IdP.

For example, a static trust agreement is established for an organization (the IdP) connecting to an enterprise service (the RP) to be made available to all subscribers at the organization on an allowlist. The authorized party for this trust agreement is the organization. When a subscriber logs in to the enterprise service, they are not prompted with any runtime decisions regarding the service since the static trust agreement establishes this a priori. In a different scenario, another service is made available to all subscribers at the same organization, but the static trust agreement stipulates that the subscriber is the authorized party. When logging in to the service for the first time, each subscriber is prompted for their consent to release their attributes to the RP. In another scenario, a dynamic trust agreement is established implicitly when a subscriber goes to

access an RP that is otherwise unknown by their IdP. The RP informs the subscriber about the uses of all attributes being requested from the IdP, and the IdP prompts the subscriber for consent to release their attributes to the RP.

Establishment of a trust agreement is required for all federation transactions, even those in which the IdP and RP have a shared security domain or shared legal ownership. In such cases, the establishment of the trust agreement is an internal process that can be completed quickly.

During the course of a single federation transaction, it is important for the policies and expectations of the IdP and RP to be unambiguous for all parties involved. Therefore, there **SHOULD** be only one set of trust agreements in effect for a given transaction. This will usually be determined by the unique pair consisting of a single IdP and a single RP. However, these agreements could vary in other ways, such as an IdP and RP having different agreements for different populations of subscribers.

The existence of a trust agreement between two parties does not preclude the existence of other agreements for each party in the agreement to have with other parties. That is to say, an IdP can have (and generally does have) independent agreements with multiple RPs simultaneously, and an RP can likewise have independent agreements with multiple IdPs simultaneously.

5.1.1. Bilateral Trust Agreements

In a bilateral trust agreement, each potential pairing of an IdP and RP form a trust relationship with each other. In this model, the IdP and RP each act as their own authority and establish the other party as capable of performing its role within the federation.

The IdP **SHALL** disclose its supported IAL, AAL, and FAL levels to the RP. The IdP **MAY** disclose a subset of its capabilities to a given RP depending on the needs of the application, for example only disclosing to a low-risk RP that accounts are proofed at IAL1 or better.

The RP **SHALL** disclose its list of required attributes to the IdP, including its purpose for use of each requested attribute. The RP **SHALL** communicate its required IAL, AAL, and FAL to the IdP, including whether no claim is required for IAL or AAL.

The IdP **SHALL** transmit only those attributes that were explicitly requested by the RP. RPs **SHALL** include their requested attributes in their privacy risk assessment.

5.1.2. Multilateral Trust Agreements

In a multilateral trust agreement, the federated parties defer to a *federation authority* to assist in making federation trust decisions and to establish the working relationship between parties. In this model, the federation authority manages the membership of IdPs and RPs in the federation agreement. The federation authority conducts some level of

vetting on each party in the federation to verify compliance with predetermined standards that define the trust agreement. The level of vetting is unique to the use cases and models employed within the federation. This vetting is depicted in the left side of Figure 2.

Federation authorities approve IdPs to operate at certain IALs, AALs, and FALs. This information is used by relying parties, as shown in the right side of Figure 2, to determine which identity providers meet their requirements.

Federation authorities **SHALL** establish parameters regarding expected and acceptable IALs, AALs, and FALs in connection with the federated relationships they enable.

Federation authorities **SHALL** individually vet each participant in the federation to determine whether they adhere to their expected standards.

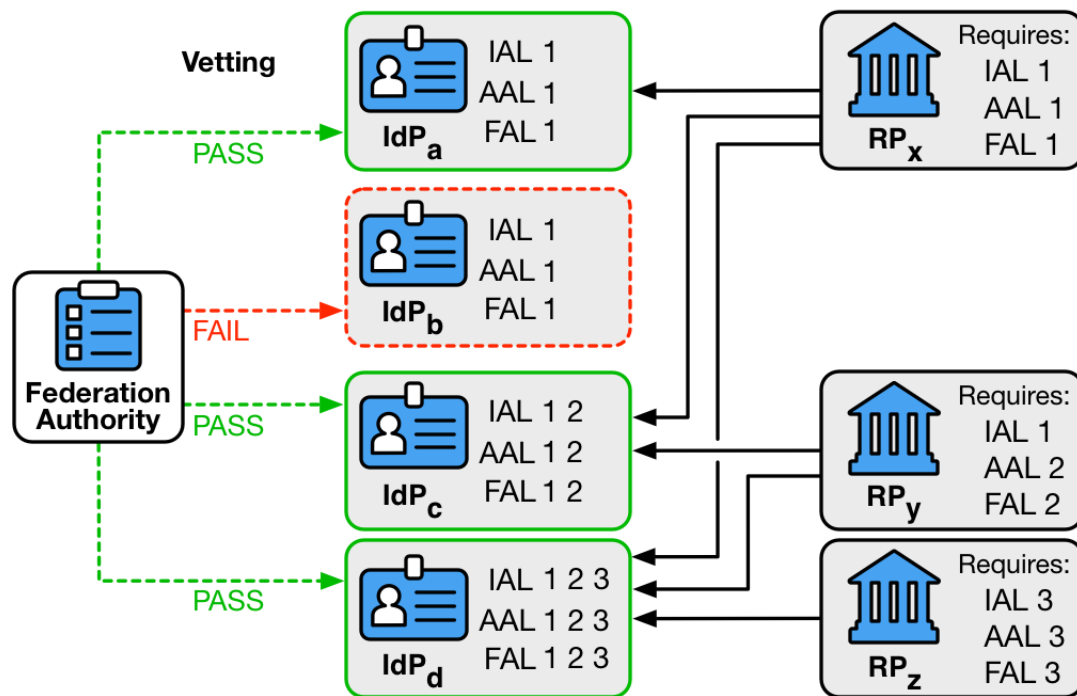


Figure 2. Federation Authority

Vetting of IdPs and RPs **SHALL** establish, as a minimum, that:

- Assertions generated by IdPs adhere to the requirements in [Sec. 6](#).
- RPs adhere to requirements for handling subscriber attribute data, such as retention, aggregation, and disclosure to third parties.
- RP and IdP systems use approved profiles of federation protocols.

Federation authorities **MAY** assist the technical connection and configuration process between members, such as by publishing configuration data for IdPs or by issuing software statements for RPs.

Most federations managed through authorities have a simple membership model: parties are either in the federation or they are not. More sophisticated federations **MAY** have multiple membership tiers that federated parties can use to tell whether other parties in the federation have been more thoroughly vetted. IdPs **MAY** decide that certain subscriber attributes are only releasable to RPs in higher tiers and RPs **MAY** decide to accept certain information only from IdPs in higher tiers.

5.1.3. Proxied Federation

In a proxied federation, all communication between the IdP and the RP is passed through an intermediary party in a way that prevents direct communication between the two parties. There are multiple methods to achieve this effect. Common configurations include:

- A third party that acts as a federation proxy (or *broker*)
- A network of nodes that distributes the communications and functions as a proxy between the endpoints

Where proxies are used, they function as an IdP on one side and an RP on the other. Therefore, all normative requirements that apply to IdPs and RPs **SHALL** apply to proxies in their respective roles.




Figure 3. Federation Proxy

A proxied federation model can provide several benefits. Federation proxies can simplify technical integration between the RP and IdP by providing a common interface for integration. Additionally, to the extent a proxy effectively blinds the RP and IdP from each other, it can provide some business confidentiality for organizations that want to guard their subscriber lists from each other. Proxies can also mitigate some of the privacy risks described in [Sec. 5.5](#) below.

See [Sec. 9.5](#) for further information on blinding techniques, their uses, and limitations.

Federations presented through a proxy **SHALL** be represented by the lowest FAL used during the proxied transaction. For example, if a proxy takes in an assertion from the IdP at FAL2 but presents a downstream assertion to the RP at FAL1, the entire transaction is

753 considered FAL1. Likewise if a federation takes in an assertion at FAL1 but presents a
754 downstream assertion to the RP at FAL3, the entire transaction is still considered FAL1.
755 The proxy **SHALL** communicate this aspect to the RP either at runtime or through pre-
756 configuration as part of the trust agreement.

757 **5.2. Registration**

758 Within federation protocols, protocol-specific information such as cryptographic
759 keys, system identifiers, service endpoint URLs, and required access rights need to be
760 established between the IdPs and RPs, allowing them to communicate securely with each
761 other. Furthermore, subscriber-facing information such as system display names and
762 home pages can be established to facilitate trust in and usability of the system. All of this
763 information is used to digitally and programmatically establish trust between the IdP and
764 RP within the scope of the federation protocol.

765 These exchanges of information happen in a pairwise fashion for each IdP and RP
766 communicating within a federation transaction, regardless of the trust agreement
767 underlying that transaction. The two phases of this process are commonly known as
768 *discovery* of the IdP by the RP and *registration* of the RP at the IdP. These processes
769 can happen in a manual, static fashion, where system administrators or developers enter
770 the information into the target systems, or in an automated, dynamic fashion, where the
771 systems themselves exchange information without direct human involvement.

5.2.1. Manual Registration

In the manual registration model, the operators of the IdP and RP manually provision configuration information about parties with which they expect to interoperate, prior to involvement of the subscriber.

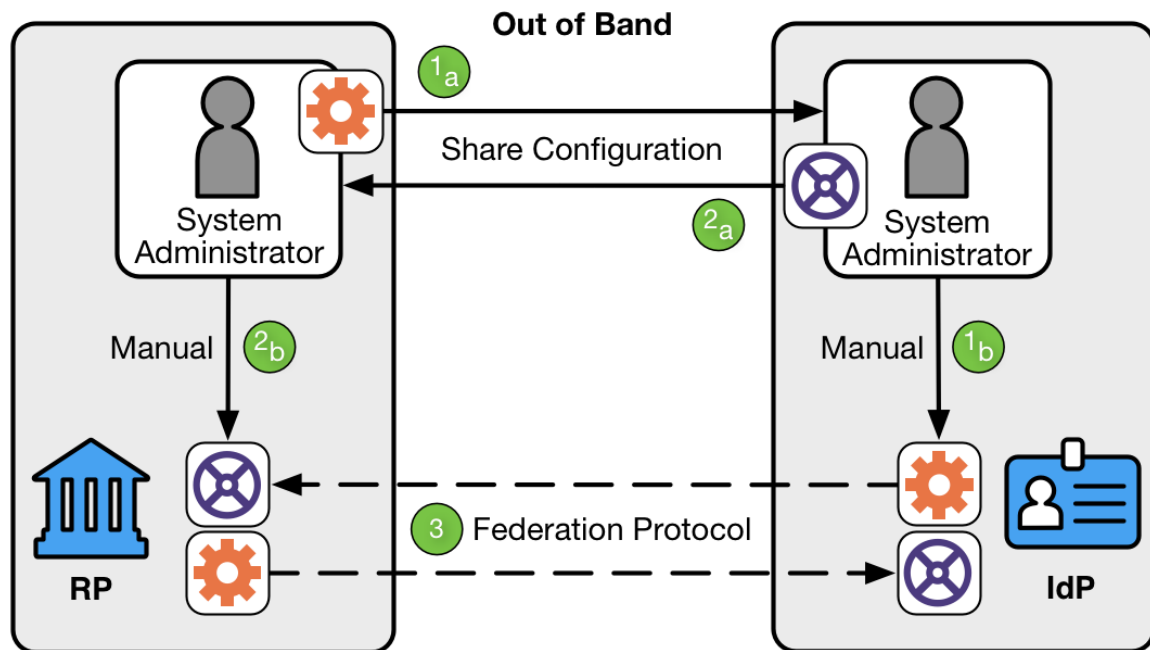


Figure 4. Manual Registration

As shown in Figure 4, manual registration involves three steps:

1. The RP's system administrator shares the RP's attributes with the IdP's system administrator, who associates those attributes with the RP.
2. The IdP's system administrator shares the IdP's attributes with the RP's system administrator, who associates those attributes with the IdP.
3. The IdP and RP then communicate using a standard federation protocol.

IdPs and RPs **MAY** act as their own authorities on who to federate with as in Sec. 5.1.1 or **MAY** externalize those authority decisions to an external party as in Sec. 5.1.2.

Protocols requiring the transfer of keying information **SHALL** use a secure method during the registration process to exchange keying information needed to operate the federated relationship, including any shared secrets or public keys. Any symmetric keys used in this relationship **SHALL** be unique to a pair of federation participants.

Federation relationships **SHALL** establish parameters regarding expected and acceptable IALs and AALs in connection with the federated relationship.

5.2.2. Dynamic Registration

In the dynamic registration model of federation, it is possible for relationships between members of the federation to be negotiated at the time of a transaction. This process allows IdPs and RPs to be connected together without manually establishing a connection between them using manual registration (See [Sec. 5.2.1](#)). IdPs that support dynamic registration **SHALL** make their configuration information (such as dynamic registration endpoints) available in such a way as to minimize system administrator involvement.

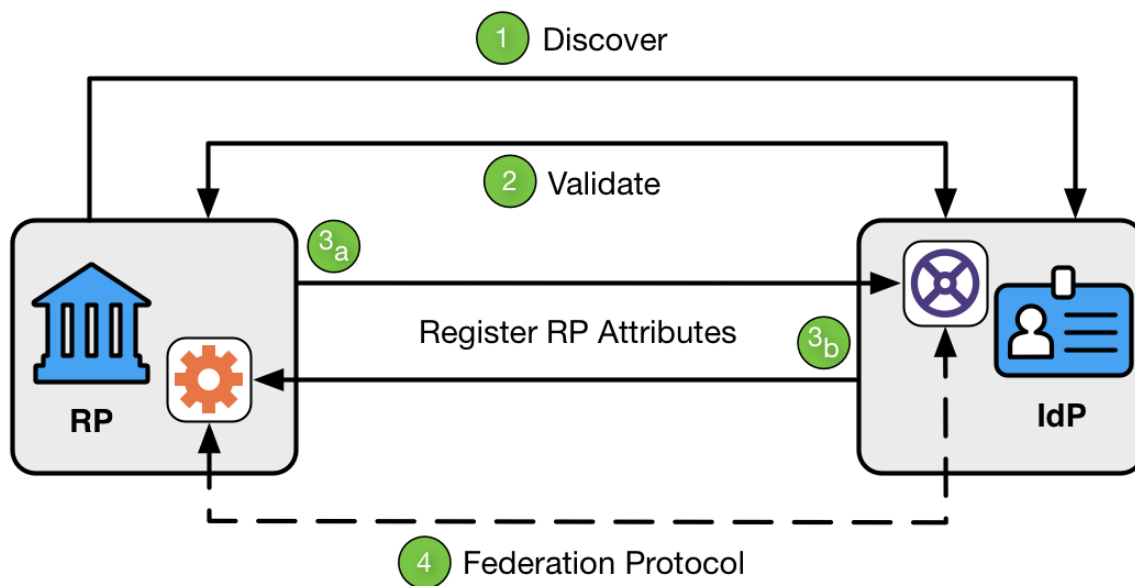


Figure 5. Dynamic Registration

As shown in [Figure 5](#), dynamic registration involves four steps:

1. Discover. The RP goes to a well-known location at the IdP to find the IdP's metadata.
2. Validate. The RP and IdP determine each other's validity. This can be accomplished through keying information, metadata, software statements, or other means.
3. Register RP attributes. The RP sends its attributes to the IdP, and the IdP associates those attributes with the RP.
4. Federation Protocol. The IdP and RP then communicate using a standard federation protocol.

Protocols requiring the transfer of keying information **SHALL** use a secure method during the registration process to establish such keying information needed to operate the federated relationship, including any shared secrets or public keys. Any symmetric keys used in this relationship **SHALL** be unique to a pair of federation participants.

IdPs **SHOULD** issue pairwise pseudonymous subject identifiers to dynamically registered RPs, as discussed in [Sec. 6.2.5](#).

Where possible, dynamic registration **SHOULD** be augmented by *software statements* anchored in their trust agreement. Software statements are lists of attributes describing the RP software, cryptographically signed by an authority (either the IdP itself, a federation authority as in [Sec. 5.1.2](#), or another trusted party). Software statements allow federated parties to cryptographically verify some attributes of an RP being dynamically registered without necessarily having all of the identifying information for that RP ahead of time. This cryptographically verifiable statement allows the connection to be established or elevated between the federating parties without relying solely on self-asserted attributes. (See [\[RFC7591\]](#) Sec. 2.3 for more information on one protocol's implementation of software statements.)

5.3. Authentication and Attribute Disclosure

Once the IdP and RP have entered into a trust agreement and have completed registration, the federation protocol can be used to pass subscriber attributes from the IdP to the RP. The decision of whether an authentication can occur or attributes may be passed **SHALL** be determined by the authorized party stipulated by the trust agreement, through use of an allowlist, a blocklist, or a runtime decision.

A subscriber's attributes **SHALL** be transmitted between IdP and RP only for identity federation transactions or support functions such as identification of compromised subscriber accounts as discussed in [Sec. 5.5](#). A subscriber's attributes are not to be transmitted for any other purposes, even when parties are allowlisted.

A subscriber's attributes **SHALL NOT** be used by the RP for purposes other than those stipulated in the trust agreement.

The subscriber **SHALL** be informed of the transmission of attributes to an RP. In the case where the authorized party is the organization, the organization **SHALL** make available to the subscriber the list of approved RPs and the associated sets of attributes sent to those RPs. In the case where the authorized party is the subscriber, the subscriber **SHALL** be prompted prior to release of attributes using a runtime decision at the IdP as described in [Sec. 5.3.3](#).

The IdP **SHALL** provide effective mechanisms for redress of subscriber complaints or problems (e.g., subscriber identifies an inaccurate attribute value). See [Sec. 10](#) on usability considerations for redress.

5.3.1. IdP Allowlists of RPs

In a static trust agreement, IdPs **MAY** establish allowlists of RPs authorized to receive authentication and attributes from the IdP without a runtime decision from the subscriber. When placing an RP on its allowlist, the IdP **SHALL** ensure that the RP abides by all

applicable provisions and requirements in the SP 800-63 guidelines. The IdP **SHALL** determine which identity attributes are passed to the allowlisted RP upon authentication. IdPs **SHALL** make allowlists available to subscribers as described in [Sec. 9.2](#).

IdP allowlists **SHALL** uniquely identify RPs through the means of domain names, cryptographic keys, or other identifiers applicable to the federation protocol in use. Any entities that share an identifier **SHALL** be considered equivalent for the purposes of the allowlist. For example, a wildcard domain identifier of “*.example.com” would match the domains “www.example.com”, “service.example.com”, and “unknown.example.com” equally. All three of these sites would be treated as the same RP for disclosure decisions using the allowlist. Allowlists **SHOULD** be as specific as possible to avoid unintentional impersonation of an RP.

5.3.2. IdP Blocklists of RPs

IdPs **MAY** establish blocklists of RPs not authorized to receive authentication assertions or attributes from the IdP, even if requested to do so by the subscriber. If an RP is on an IdP’s blocklist, the IdP **SHALL NOT** produce an assertion targeting the RP in question under any circumstances.

IdP blocklists **SHALL** uniquely identify RPs through the means of domain names, cryptographic keys, or other identifiers applicable to the federation protocol in use. Any entities that share an identifier **SHALL** be considered equivalent for the purposes of the blocklist. For example, a wildcard domain identifier of “*.example.com” would match the domains “www.example.com”, “service.example.com”, and “unknown.example.com” equally. All three of these sites would be treated as the same RP for decisions using the blocklist.

5.3.3. IdP Runtime Decisions

Every RP that is in a trust agreement with an IdP but not on an allowlist or a blocklist with that IdP **SHALL** be governed by a default policy in which runtime authorization decisions will be made by an authorized party identified by the trust agreement. In most circumstances, and for practical purposes, the authorized party is the subscriber; however, it is possible for an administrator or other party to be prompted on behalf of the subscriber. Note that in a dynamic trust agreement, only a runtime decision can be used to authorize the release of attributes.

In this mode of operation, the authorized party is prompted by the IdP during the federation transaction for their consent to provide an authentication assertion and release specific attributes to the RP on behalf of the subscriber. The IdP **SHALL** provide the authorized party with explicit notice and prompt them for positive confirmation before any attributes about the subscriber are transmitted to the RP. At a minimum, the notice **SHOULD** be provided by the party in the position to provide the most effective notice and obtain confirmation, consistent with [Sec. 9.2](#). The IdP **SHALL** disclose which attributes

will be released to the RP if the transaction is approved. If the federation protocol in use allows for optional attribute disclosure at runtime, the authorized party **SHALL** be given the option to decide whether to transmit specific attributes to the RP without terminating the federation transaction entirely.

To mitigate the risk of unauthorized exposure of sensitive information (e.g., shoulder surfing), the IdP **SHALL**, by default, mask sensitive information displayed to the authorized party. If the authorized party is the subscriber, the IdP **SHALL** provide mechanisms for the subscriber to temporarily unmask such information in order for the subscriber to view full values before transmission. For more details on masking, see [Sec. 10](#) on usability considerations.

An IdP **MAY** employ mechanisms to remember and re-transmit the exact attribute bundle to the same RP, remembering the authorized party's decision. This mechanism is associated with the subscriber account as managed by the IdP. If such a mechanism is provided, the IdP **SHALL** allow the authorized party to revoke such remembered access at a future time.

5.3.4. RP Allowlists of IdPs

RPs **MAY** establish allowlists of IdPs from which the RP will accept authentication and attributes without a runtime decision from the subscriber. When placing an IdP in its allowlist, the RP **SHALL** ensure that the IdP abides by the provisions and requirements in these guidelines.

RP allowlists **SHALL** uniquely identify IdPs through the means of domain names, cryptographic keys, or other identifiers applicable to the federation protocol in use.

5.3.5. RP Blocklists of IdPs

RPs **MAY** also establish blocklists of IdPs that the RP will not accept authentication or attributes from, even when requested by the subscriber. A blocklisted IdP can be otherwise in a valid trust agreement with the RP, for example if both are under the same federation authority.

RP blocklists **SHALL** uniquely identify IdPs through the means of domain names, cryptographic keys, or other identifiers applicable to the federation protocol in use.

5.3.6. RP Runtime Decisions

Every IdP that is in a trust agreement with an RP but not on an allowlist or a blocklist with that RP **SHALL** be governed by a default policy in which runtime authorization decisions will be made by the authorized party indicated in the trust agreement. In this mode, the authorized party is prompted by the RP to select or enter which IdP to contact for authentication on behalf of the subscriber. This process can be facilitated through use of a discovery mechanism allowing the subscriber to enter a human-facing identifier

such as an email address. This process allows the RP to programmatically select the appropriate IdP for that identifier.

The RP **MAY** employ mechanisms to remember the authorized party's decision to use a given IdP. Since this mechanism is employed prior to authentication at the RP, the manner in which the RP provides this mechanism (e.g., a browser cookie outside the authenticated session) is separate from the RP subscriber account described in [Sec. 5.4](#). If such a mechanism is provided, the RP **SHALL** allow the authorized party to revoke such remembered options at a future time.

5.4. RP Subscriber Accounts

It is common for an RP to keep a record representing a subscriber local to the RP itself, known as the *RP subscriber account*. The RP subscriber account can contain things like access rights at the RP as well as a cache of identity attributes for the subscriber. An active RP subscriber account is bound to one or more federated identifiers from the RP's trusted IdPs. Successful authentication of one of these federated identifiers through a federation protocol allows the subscriber to access the information and functionality protected by the RP subscriber account.

An RP subscriber account is *provisioned* when the RP has associated a set of attributes about the subscriber with a data record representing the subscriber account at the RP. The RP subscriber account **SHALL** be bound to at least one federated identifier, and a given federated identifier is bound to only one RP subscriber account at a given RP. The provisioning can happen prior to authentication or as a result of the federated authentication process, depending on the deployment patterns as discussed in [Sec. 5.4.1](#). Prior to being provisioned, the RP subscriber account does not exist and has no associated data record at the RP.

An RP subscriber account is *terminated* when the RP removes all access to the account at the RP. Termination **SHALL** include unbinding any federated identifiers and bound authenticators as well as removing attributes and information associated with the account except what is required for auditing and security purposes. An RP **MAY** terminate an RP subscriber account independently from the IdP for a variety of reasons, regardless of the current validity of the subscriber account from which it is derived.

An authenticated session **SHALL** be created by the RP only when the RP has processed and verified a valid assertion from the IdP that is the issuer of the federated identifier associated with the RP subscriber account. If the assertion also requires presentation of a bound authenticator at FAL3, the bound authenticator **SHALL** also be presented and processed before the RP subscriber account is associated with an authenticated session, as discussed in [Sec. 6.1.2](#). Before the federated assertion is processed and after termination of the authenticated session, the RP subscriber account is unauthenticated though it could still be provisioned.

5.4.1. Provisioning Models

The lifecycle of the provisioning process for an RP subscriber account varies depending on factors including the trust agreement discussed in [Sec. 5.1](#) and the deployment pattern of the IdP and RP. However, in all cases, the RP subscriber account **SHALL** be provisioned at the RP prior to the establishment of an authenticated session at the RP in one of the following ways:

Just-In-Time Provisioning

An RP subscriber account is created automatically the first time the RP receives an assertion with an unknown federated identifier from an IdP. Any identity attributes learned during the federation process, either within the assertion or through an identity API as discussed in [Sec. 6.3](#), **MAY** be associated with the RP subscriber account. Accounts provisioned in this way are bound to the federated identifier in the assertion used to provision them. This is the most common form of provisioning in federation systems, as it requires the least coordination between the RP and IdP. However, in such systems, the RP **SHALL** be responsible for managing any cached attributes it might have.

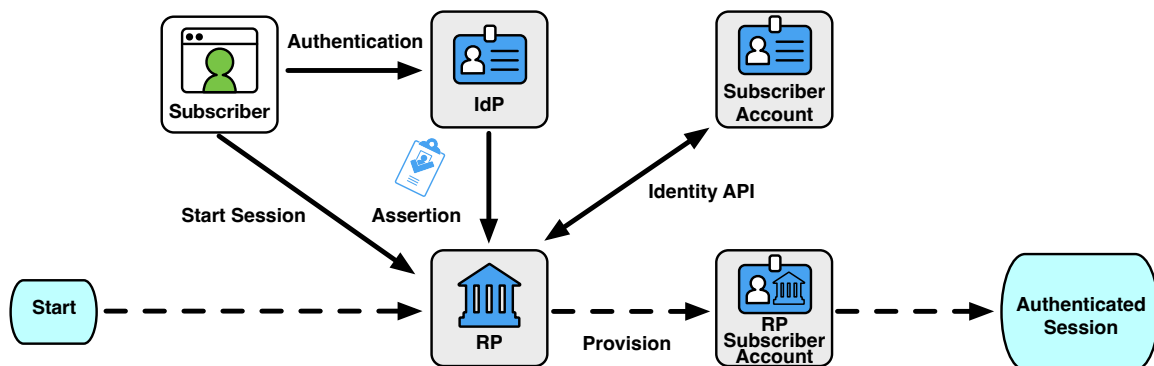


Figure 6. Just-In-Time Provisioning

Pre-provisioning

An RP subscriber account is created by the IdP pushing the attributes to the RP or the RP pulling attributes from the IdP. Pre-provisioning of accounts generally occurs in bulk through a provisioning API as discussed in [Sec. 5.4.3](#), as the provisioning occurs prior to the represented subscribers authenticating through a federated transaction. Pre-provisioned accounts **SHALL** be bound to a federated identifier at the time of provisioning. Any time a particular federated identifier is seen by the RP, the associated account can be logged in as a result. This form of provisioning requires infrastructure and planning on the part of the IdP and RP, but these processes can be facilitated by automated protocols. The RP also collects attributes about users who have not interacted with the RP system yet, which can cause privacy issues. Additionally, the IdP and RP must keep the set of provisioned accounts synchronized over time as discussed in [Sec. 5.4.2](#).

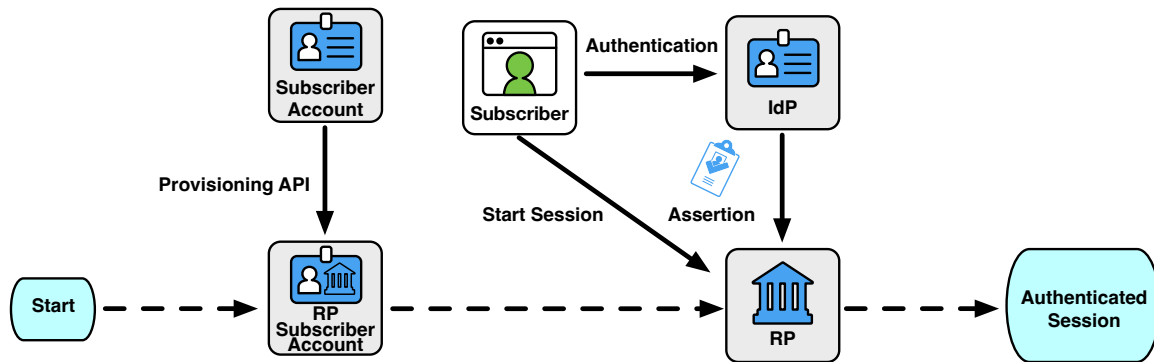


Figure 7. Pre-Provisioning

Ephemeral

An RP subscriber account is created when processing the assertion, but then the RP subscriber account is terminated when the authenticated session ends. This process is similar to a just-in-time provisioning, but the RP keeps no long-term record of the account when the session is complete, except what is required for audit and security purposes (such as access logs). This form of provisioning is useful for RPs that fully externalize access rights to the IdP, allowing the RP to be more simplified with less internal state. However, this pattern is not common because even the simplest RPs tend to have a need to track state within the application or at least keep a record of actions associated with the federated identifier.

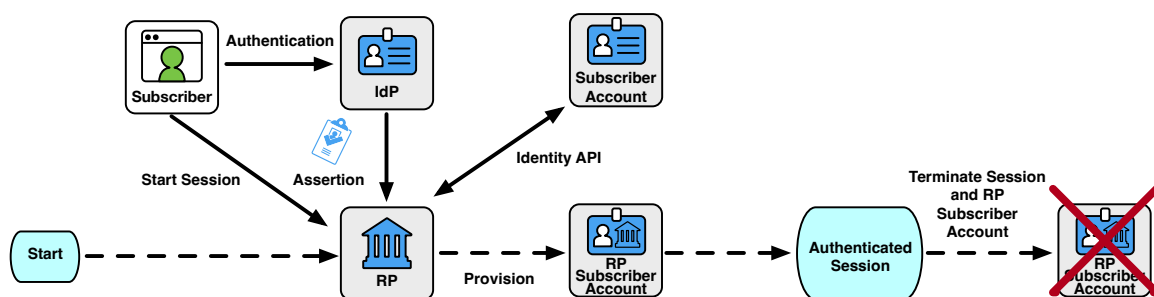


Figure 8. Ephemeral Provisioning

Other

Other RP subscriber account provisioning models are possible but the details of such models are outside the scope of these guidelines. The details of any alternative provisioning model **SHALL** be included in the privacy risk assessments of the IdP and RP.

All organizations **SHALL** document their provisioning model as part of their trust agreement.

5.4.2. Attribute Synchronization

In a federated process, the IdP and RP each have their own stores of identity attributes associated with the subscriber account. The IdP has a direct view of the subscriber account, but the RP subscriber account is derived from a subset of attributes from the subscriber account that are presented during the federation transaction. Therefore, it is possible for the IdP's and RP's attribute stores to diverge from with each other over time.

From the RP's perspective, the IdP is the authoritative source for any attributes that the IdP asserts as being associated with the subscriber account at the IdP. However, the RP **MAY** additionally collect, and optionally verify, other attributes to associate with the RP subscriber account. Sometimes, these attributes can even override what's asserted by the IdP. For example, if an IdP asserts a full display name for the subscriber, the RP can allow the subscriber to provide an alternative preferred name for use at the RP.

The IdP **SHOULD** signal downstream RPs when the attributes of a subscriber account available to the RP have been updated. This can be accomplished using shared signaling as described in [Sec. 5.7](#), through a provisioning API as described in [Sec. 5.4.3](#), or by providing a signal in the assertion (e.g., a timestamp indicating when relevant attributes were last updated, allowing the RP to determine that its cache is out of date).

The IdP **SHOULD** signal downstream RPs when a subscriber account is terminated, or when the subscriber account's access to an RP is revoked. This can be accomplished using shared signaling as described in [Sec. 5.7](#) or through a provisioning API as described in [Sec. 5.4.3](#). Upon receiving such a signal, the RP **SHALL** terminate the RP subscriber account and remove all personal information associated with the RP subscriber account, except what is required for audit and security purposes.

5.4.3. Provisioning APIs

As part of some proactive forms of provisioning, the RP can be given access to subscriber attributes through a general-purpose attribute API known as a *provisioning API*. This type of API allows an IdP to push attributes for a range of subscriber accounts, and sometimes allows an RP to query the attributes of these subscriber accounts directly. Since access to the API is granted outside the context of a federated transaction, access to the provisioning API for a given subscriber does not indicate to the RP that a given subscriber has been authenticated. See [Sec. 6, Assertions](#) for more information on how the federated authentication process is accomplished using assertions.

The attributes in the provisioning API available to a given RP **SHALL** be limited to only those necessary for the RP to perform its functions. As part of establishing the trust agreement, the IdP **SHALL** document when an RP is given access to a provisioning API including at least the following:

- the purpose for the access using the provisioning model;
- the set of attributes made available to the RP;
- whether the API functions as a push to the RP, a pull from the RP, or both; and
- the population of subscribers whose attributes are made available to the RP.

The IdP **SHALL** require authentication from the RP for any pull-based access to a provisioning API. The RP **SHALL** require authentication from the IdP for any push-based access to a provisioning API.

A provisioning API **SHALL NOT** be made available under a dynamic or implicit trust agreement. The IdP **SHALL NOT** make a provisioning API available to any RP outside of an established trust agreement. The IdP **SHALL** provide access to a provisioning API only as part of a federated identity relationship with an RP to facilitate federated transactions with that RP and related functions such as signaling revocation of the subscriber account. The IdP **SHALL** revoke an RP's access to the provisioning API

once access is no longer required by the RP for its functioning purposes or when the trust agreement is terminated.

Any provisioning API provided to the RP **SHALL** be under the control and jurisdiction of the IdP. External attribute providers **MAY** be used as information sources by the IdP to provide attributes through this provisioning API, but the IdP is responsible for the content and accuracy of the information provided by the referenced attribute providers.

When a provisioning API is in use, the IdP **SHALL** signal to the RP when a subscriber account has been terminated. When receiving such a signal, the RP **SHALL** terminate the associated RP subscriber account.

5.4.4. Attribute Collection

The RP **MAY** collect and maintain additional attributes from the subscriber beyond those provided by the IdP. These attributes are governed separately from any federation agreement since they are collected directly by the RP. All attributes associated with an RP subscriber account, regardless of their source, **SHALL** be removed when the RP subscriber account is terminated.

The RP **SHALL** disclose to the subscriber the purpose for collection of any additional attributes. These attributes **SHALL** be used solely for the stated purposes of the RP's functionality and **SHALL NOT** have any secondary use, including communication of said attributes to other parties.

An RP **SHALL** disclose any additional attributes collected, and their use, as part of its System of Records Notice (SORN). The RP **SHALL** provide an effective means of redress for the subscriber to update and remove these additionally-collected attributes from the RP subscriber account. See [Sec. 10](#) on usability considerations for redress.

5.4.5. Time-based Removal of RP Subscriber Accounts

Over time, an RP could accumulate RP subscriber accounts that are no longer accessible from the IdP. This poses a risk to the RP for holding personal information in the RP subscriber accounts, especially when a just-in-time provisioning model is in use and no shared signaling is available from the IdP to signal subscriber account termination as discussed in [Sec. 5.7](#). In such circumstances, the RP **SHOULD** employ a time-based mechanism to identify RP subscriber accounts for termination that have not been accessed after a period of time, for example, 120 days since last access.

When processing such an inactive account, the RP **SHALL** provide sufficient notice to the subscriber, if possible, about the pending termination of the account and provide the subscriber with an option to re-activate the account prior to its scheduled termination. Upon termination, the RP **SHALL** remove all personal information associated with the RP subscriber account, except what is required for audit and security purposes.

5.5. Privacy Requirements

The ultimate goal of a subscriber is to interact with and use the RP. Federation involves the transfer of personal attributes from a third party that is not otherwise involved in a transaction — the IdP. Federation also potentially gives the IdP broad visibility into subscriber activities and status. Accordingly, there are specific privacy requirements associated with federation.

Communication between the RP and the IdP could reveal to the IdP where the subscriber is conducting a transaction. Communication with multiple RPs allows the IdP to build a profile of subscriber transactions that would not have existed without federation. This aggregation could enable new opportunities for subscriber tracking and use of profile information that do not always align with subscribers' privacy interests.

If an IdP discloses information on subscriber activities at an RP to any party, or processes the subscriber's attributes for any purpose other than identity proofing, authentication, or attribute assertions (collectively "identity service"), related fraud mitigation, to comply with law or legal process, or, in the case of a specific user request, to transmit the information, the IdP **SHALL** implement measures to maintain predictability and manageability commensurate with the privacy risk arising from the additional processing. Measures **MAY** include providing clear notice, obtaining subscriber consent, or enabling selective use or disclosure of attributes. When an IdP uses consent measures, the IdP **SHALL NOT** make consent for the additional processing a condition of the identity service.

If the same subscriber account is asserted to multiple RPs, and those RPs communicate with each other, the colluding RPs could track a subscriber's activity across multiple applications and security domains. The IdP **SHOULD** employ technical measures, such as the use of pairwise pseudonymous identifiers described in [Sec. 6.2.5](#) or privacy-enhancing cryptographic protocols, to provide disassociability and discourage subscriber activity tracking and profiling between RPs.

An IdP **MAY** disclose information on subscriber activities to RPs for security purposes, such as communication of suspicious activity or a compromised subscriber account as described in [Sec. 5.7](#), if stated within the trust agreement. An RP **MAY** disclose information on subscriber activities to IdPs for security purposes, such as communication of suspicious activity or a compromised RP subscriber account, if stated within the trust agreement.

An IdP **SHOULD** signal subscriber account termination to RPs that have been provisioned with federated identifiers bound to that subscriber account using shared signaling as discussed in [Sec. 5.7](#). RPs that receive such a signal from the IdP **SHALL** terminate the RP subscriber account and remove all personal information associated with the RP subscriber account, except what is required for audit and security purposes.

The following requirements apply specifically to federal agencies:

- 1130 1. The agency **SHALL** consult with their Senior Agency Official for Privacy (SAOP)
1131 to conduct an analysis determining whether the requirements of the Privacy Act are
1132 triggered by the agency that is acting as an IdP, by the agency that is acting as an
1133 RP, or both (see [Sec. 9.4](#)).
- 1134 2. The agency **SHALL** publish or identify coverage by a System of Records Notice
1135 (SORN) as applicable.
- 1136 3. The agency **SHALL** consult with their SAOP to conduct an analysis determining
1137 whether the requirements of the E-Government Act are triggered by the agency that
1138 is acting as an IdP, the agency that is acting as an RP, or both.
- 1139 4. The agency **SHALL** publish or identify coverage by a Privacy Impact Assessment
1140 (PIA) as applicable.

1141 If the RP subscriber account lifecycle process gives the RP access to attributes through
1142 a provisioning API as discussed in [Sec. 5.4.3](#), additional privacy measures **SHALL** be
1143 implemented given the wide nature of information access. Specifically, it is possible
1144 for the attributes of a subscriber to be provided to an RP without the subscriber ever
1145 interacting with the RP in question. As a consequence, when a provisioning API is
1146 used, the IdP **SHALL** minimize the attributes made available to the RP. To prevent the
1147 transmission of attributes for users that will never use an RP, the IdP **SHALL** limit the
1148 population of subscriber accounts available via the provisioning API to the population of
1149 subscribers authorized to use the RP by the trust agreement.

1150 **5.6. Reauthentication and Session Requirements in Federated Environments**

1151 In a federated environment, the RP manages its sessions separately from any sessions
1152 at the IdP. The assertion is related to both sessions but its validity period is ultimately
1153 independent of them. In order for the IdP to create an assertion for the subscriber,
1154 the subscriber needs to establish an authenticated session with the IdP. To create an
1155 authenticated session at the RP, the RP needs to process a valid assertion from the IdP.

1156 Due to the distributed nature of a federated system, the subscriber's sessions with the
1157 IdP and with the RP terminate independently of each other. The RP **SHALL NOT** assume
1158 that the subscriber has an active session at the IdP past the issuance time of the assertion.
1159 The IdP **SHALL NOT** assume that termination of the subscriber's session at the IdP will
1160 propagate to any sessions that subscriber would have at downstream RPs. The RP and
1161 IdP **MAY** communicate session termination requests to other parties in the federation
1162 network, if supported by the federation protocol.

1163 At the time of a federated login request, the subscriber **MAY** have a pre-existing
1164 session at the IdP which **MAY** be used to generate an assertion to the RP. The IdP
1165 **SHALL** communicate any information it has regarding the time of the subscriber's
1166 latest authentication event at the IdP, and the RP **MAY** use this information in making
1167 authorization and access decisions. Depending on the capabilities of the federation

1168 protocol in use, the IdP **SHOULD** allow the RP to request that the subscriber repeat
1169 authentication at the IdP as part of a federation request.

1170 An RP requiring authentication through a federation protocol **SHALL** specify the
1171 maximum acceptable authentication age to the IdP, either through the federation protocol
1172 (if possible) or through the parameters of the trust agreement. The authentication age
1173 represents the time since the last authentication event in the subscriber's session at the
1174 IdP, and the IdP **SHALL** reauthenticate the subscriber if they have not been authenticated
1175 within that time period. The IdP **SHALL** communicate the authentication event time to
1176 the RP to allow the RP to decide if the assertion is sufficient for authentication at the RP
1177 and to determine the time for the next reauthentication event.

1178 If an RP is granted access to an identity API along with the assertion, the lifetime of
1179 the access to the identity API is independent from the lifetime of the assertion itself.
1180 Since access to the identity API is often combined with access to additional APIs, it is
1181 common for this access to be valid long after the assertion has expired and possibly after
1182 the session with the RP has ended, allowing the RP to access APIs on the subscriber's
1183 behalf while the subscriber is no longer present. As a consequence, the RP's ability to
1184 successfully fetch additional attributes through an identity API **SHALL NOT** be used to
1185 establish a session at the RP. Likewise, inability to access an identity API **SHOULD NOT**
1186 be used to end the session at the RP.

1187 See [SP800-63B], Sec. 7 for more information about session management requirements
1188 for both IdPs and RPs.

1189 **5.7. Shared Signaling**

1190 In some environments, it is useful for the IdP and RP to send information to each
1191 other outside of the federation transaction. These signals can communicate important
1192 changes in state between parties that would not be otherwise known. The use of any
1193 shared signaling **SHALL** be documented in the trust agreement between the IdP and
1194 RP. Signaling from the IdP to the RP **SHALL** require a static trust agreement. Signaling
1195 from the RP to the IdP **MAY** be used in a static or dynamic trust agreement.

1196 Any use of shared signaling **SHALL** be documented and made available to the authorized
1197 party stipulated by the trust agreement. This documentation **SHALL** include the events
1198 under which a signal is sent, the information included in such a signal (including any
1199 attribute information), and any additional parameters sent with the signal. The use of
1200 shared signaling **SHALL** be subject to privacy review under the trust agreement.

1201 The IdP **MAY** send a signal regarding the following changes to the subscriber account:

- 1202 • The account has been terminated.
- 1203 • The account is suspected of being compromised.

- 1204 • Attributes of the account, including identifiers other than the federated identifier
1205 (such as email address or certificate CN), have changed.
 - 1206 • The possible range of IAL, AAL, or FAL for the account has changed.
- 1207 The RP **MAY** send a signal regarding the following changes to the RP subscriber
1208 account:
- 1209 • The account has been terminated.
 - 1210 • The account is suspected of being compromised.
 - 1211 • An RP-managed bound authenticator is added.
 - 1212 • An RP-managed bound authenticator is removed.
- 1213 Additional signals from both the IdP and RP **MAY** be allowed subject to privacy and
1214 security review as part of the trust agreement.

6. Assertions

This section is normative.

An assertion used for authentication is a packaged set of attribute values or derived attribute values about or associated with an authenticated subscriber that is passed from the IdP to the RP in a federated identity system. Assertions contain a variety of information, including: assertion metadata, attribute values and derived attribute values about the subscriber, information about the subscriber's authentication at the IdP, and other information that the RP can leverage (e.g., restrictions and validity time window). While the assertion's primary function is to authenticate the user to an RP, the information conveyed in the assertion can be used by the RP for a number of use cases — for example, authorization or personalization of a website. These guidelines do not restrict RP use cases nor the type of protocol or data payload used to federate an identity, provided the chosen solution meets all mandatory requirements contained herein.

Assertions **SHALL** represent a discrete authentication event of the subscriber at the IdP and **SHALL** be processed as a discrete authentication event at the RP.

All assertions **SHALL** include the following attributes:

1. Subject identifier: An identifier for the party to which the assertion applies (i.e., the subscriber).
2. Issuer identifier: An identifier for the issuer of the assertion (i.e., the IdP).
3. Audience identifier: An identifier for the party intended to consume the assertion (i.e., the RP).
4. Issuance time: A timestamp indicating when the IdP issued the assertion.
5. Validity time window: A period of time outside of which the assertion **SHALL NOT** be accepted as valid by the RP for the purposes of authenticating the subscriber and starting an authenticated session at the RP. This is usually communicated by means of an expiration timestamp for the assertion in addition to the issuance timestamp.
6. Assertion identifier: A value uniquely identifying this assertion, used to prevent attackers from replaying prior assertions.
7. Signature: Digital signature or message authentication code (MAC), including key identifier or public key associated with the IdP, covering the entire assertion.
8. Authentication time: A timestamp indicating when the IdP last verified the presence of the subscriber at the IdP through a primary authentication event (if available).
9. IAL: Indicator of the IAL of the subscriber account being represented in the assertion, or an indication that no IAL is asserted.
10. AAL: Indicator of the AAL used when the subscriber authenticated to the IdP, or an indication that no AAL is asserted.

1251 11. FAL: An indicator of the IdP's intended FAL of the federation process represented
1252 by the assertion.

1253 If the assertion is used at FAL3 with a bound authenticator as described in [Sec. 6.1.2](#), the
1254 assertion **SHALL** include the following:

1255 1. Authenticator binding: The public key, key identifier, or other identifier of
1256 subscriber-held bound authenticator (for IdP-managed bound authenticators) or
1257 indicator that an RP-managed bound authenticator is required for verification of this
1258 assertion.

1259 Assertions **MAY** also include additional items, including the following information:

1260 1. Attribute values and derived attribute values: Information about the subscriber.
1261 2. Attribute metadata: Additional information about one or more subscriber attributes,
1262 such as those described in NIST Internal Report 8112 [[NISTIR8112](#)].

1263 Assertions **SHOULD** specify the AAL when an authentication event is being asserted and
1264 IAL when identity proofed attributes (or values derived from those attributes) are being
1265 asserted.

1266 All metadata within the assertion **SHALL** be validated by the RP upon receipt:

- 1267 • *Issuer verification*: ensuring the assertion was issued by the IdP the RP expects it to
1268 be from.
- 1269 • *Signature validation*: ensuring the signature of the assertion is valid and
1270 corresponds to a key belonging to the IdP sending the assertion.
- 1271 • *Time validation*: ensuring the expiration and issue times are within acceptable limits
1272 of the current timestamp.
- 1273 • *Audience restriction*: ensuring this RP is the intended recipient of the assertion.

1274 An RP **SHALL** treat subject identifiers as not inherently globally unique. Instead, the
1275 value of the assertion's subject identifier is usually in a namespace under the assertion
1276 issuer's control. This allows an RP to talk to multiple IdPs without incorrectly conflating
1277 subjects from different IdPs.

1278 Assertions **MAY** include additional attributes about the subscriber. [Section 6.2.3](#) contains
1279 privacy requirements for presenting attributes in assertions. The RP **MAY** be given
1280 limited access to an identity API as discussed in [Sec. 6.3](#) along with the assertion, which
1281 the RP can use to fetch additional identity attributes for the subscriber.

1282 Although details vary based on the exact federation protocol in use, an assertion
1283 represents a discrete login event to the RP. The validity time window of an assertion
1284 is related to but separate from any session management at the IdP or RP. Specifically,
1285 an assertion is created during an authenticated session at the IdP, and processing an

assertion creates an authenticated session at the RP. After the IdP creates the assertion, the validity of the IdP's session is independent of the validity of the assertion. If a request comes to the IdP for a repeated authentication while the session is still valid at the IdP, this results in a new and separate assertion being created with its own validity time window. Similarly, after the RP consumes the assertion, the validity of the RP's session is independent of the validity of the assertion. Access granted to an identity API is likewise independent of the validity of the assertion or the lifetime of the authenticated session at the RP. See [Sec. 5.3](#) for more information on session management.

The assertion's validity time window is the time between its issuance and its expiration. This window needs to be large enough to allow the RP to process the assertion and create a local application session for the subscriber, but should not be longer than necessary for such establishment. Long-lived assertions have a greater risk of being stolen or replayed; a short assertion validity time window mitigates this risk. Assertion validity time windows **SHALL NOT** be used to limit the session at the RP. See [Sec. 5.3](#) for more information.

6.1. Assertion Binding

Assertion binding can be classified based on whether presentation by a claimant of an assertion is sufficient for binding to the party currently in session with the RP as the subscriber, or if the RP requires additional proof through the successful presentation of an authenticator bound to the subscriber.

6.1.1. Bearer Assertions

A bearer assertion can be presented by any party as proof of the bearer's identity. Similarly, a bearer assertion reference can be presented by any party to the RP and used by the RP to fetch an assertion; the assertion in this instance is also considered a bearer assertion. If an attacker can capture or manufacture a valid assertion or assertion reference representing a subscriber and can successfully present that assertion or reference to the RP, then the attacker could be able to impersonate the subscriber at that RP.

Note that mere possession of a bearer assertion or reference is not always enough to impersonate a subscriber. For example, if an assertion is presented in the back-channel federation model (described in [Sec. 7.1](#)), additional controls **MAY** be placed on the transaction (such as identification of the RP and assertion injection protections) that help further protect the RP from fraudulent activity.

6.1.2. Bound Authenticators

A bound authenticator is an authenticator presented to the RP by the subscriber alongside the assertion. In proving possession of the bound authenticator to the RP, the subscriber also proves with a certain degree of assurance that they are the rightful subject of the assertion. It is more difficult for an attacker to use a stolen assertion issued to a subscriber

since the attacker would need to steal the bound authenticator as well as the assertion and be able to present them together. Furthermore, use of a bound authenticator protects the RP against malicious or compromised IdPs through the use of independent authentication.

A bound authenticator **SHALL** be unique per subscriber at the RP such that two subscribers cannot present the same authenticator for their separate RP subscriber accounts. All bound authenticators **SHALL** be phishing resistant. Consequently, subscriber-chosen values such as a memorized secret cannot be used as bound authenticators. The RP **SHALL** accept authentication from a bound authenticator only in the context of processing an assertion. Consequently, the subscriber can not use a bound authenticator to log into the RP directly, bypassing the IdP in the process.

A bound authenticator can be managed by either the IdP or the RP under different circumstances, as detailed in the sections below. An FAL3 assertion contains an indication of whether the IdP expects the subscriber to present a specific IdP-managed bound authenticator or an RP-managed bound authenticator at the RP to reach FAL3.

6.1.2.1. IdP-Managed Bound Authenticators

When the bound authenticator is managed by the IdP as in [Fig. 9](#), a unique identifier for the authenticator (such as its public key) **SHALL** be included in the assertion presented to the RP. The RP **SHALL** prompt the subscriber to prove possession of the identified bound authenticator.

An IdP-managed bound authenticator **MAY** be distinct from the primary authenticator the subscriber uses to authenticate to the IdP. Bound authenticators managed at the IdP **SHALL** be phishing resistant and **SHALL** be independently dereferenceable by the RP based on a mutually-trusted security framework, such as a public-key infrastructure. When processing an IdP-managed bound authenticator for the first time, the RP **SHOULD** verify whether the authenticator being presented is appropriate to be associated with the subscriber account, such as through account resolution from the attributes in the authenticator's presented information.

For example, a subscriber could have a smart card loaded with a certificate, which is a multi-factor cryptographic device. Since the certificate can be presented to both the IdP and the RP, the IdP can include an identifier for the certificate in the FAL3 assertion to the RP. The RP would then prompt the subscriber to present the certificate from their smart card in order to reach FAL3.

“Holder of Key” (HoK) assertions are one example of IdP-managed bound authenticators, since the IdP knows the subscriber's key to be used at the RP and includes the key information in the assertion presented to the RP.

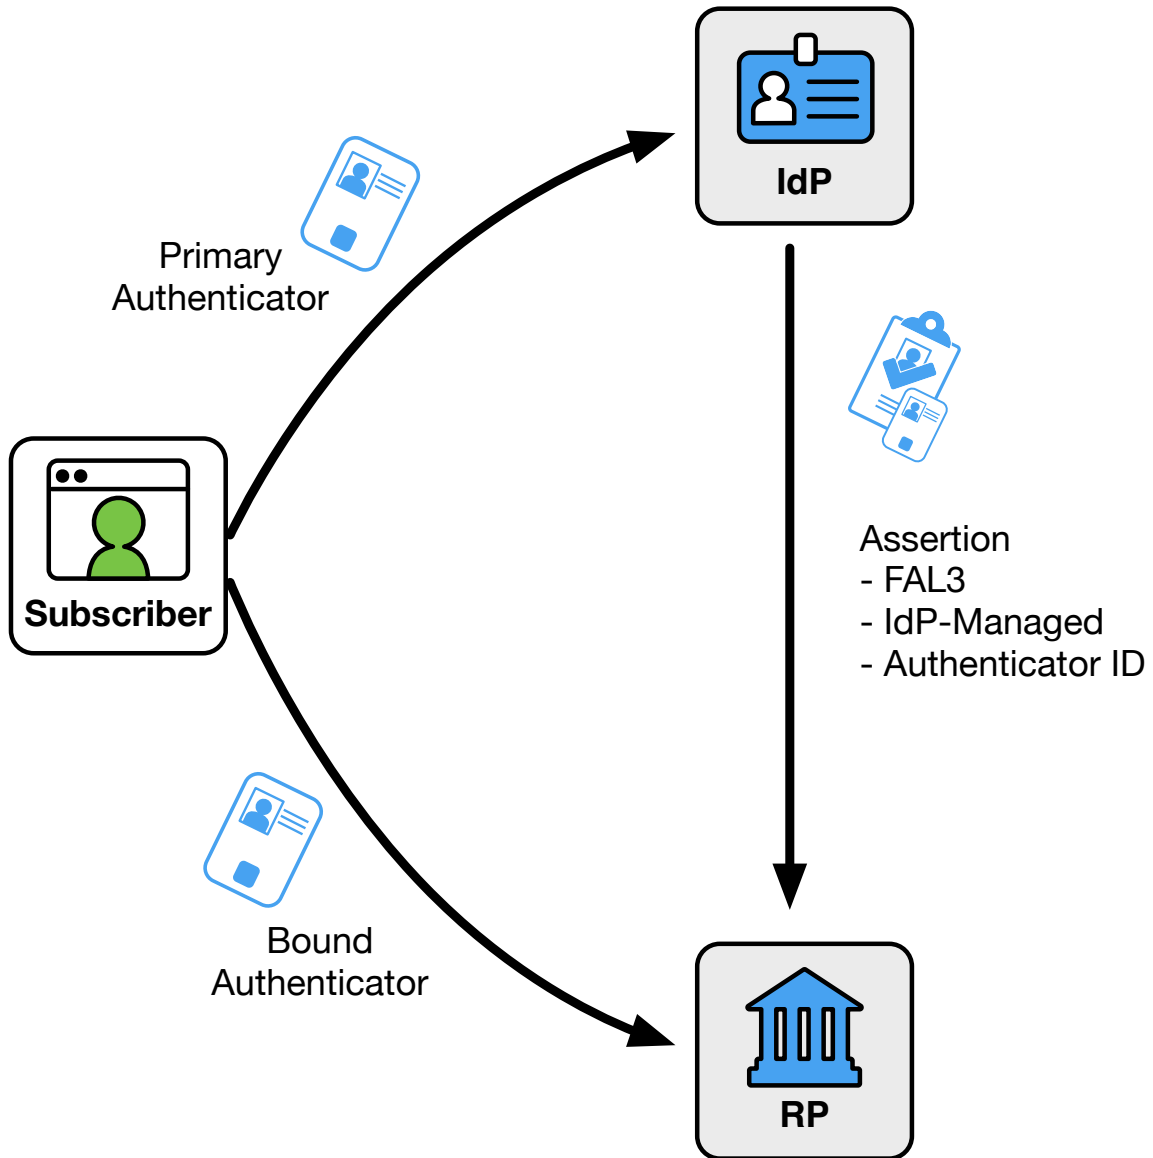


Figure 9. IdP-Managed Bound Authenticators

6.1.2.2. RP-Managed Bound Authenticators

When the bound authenticator is managed by the RP as in Fig. 10, the IdP **SHALL** include an indicator in the assertion that the assertion is to be used with a bound authenticator at FAL3. The unique identifier for the authenticator (such as its public key) **SHALL** be stored in the RP subscriber account.

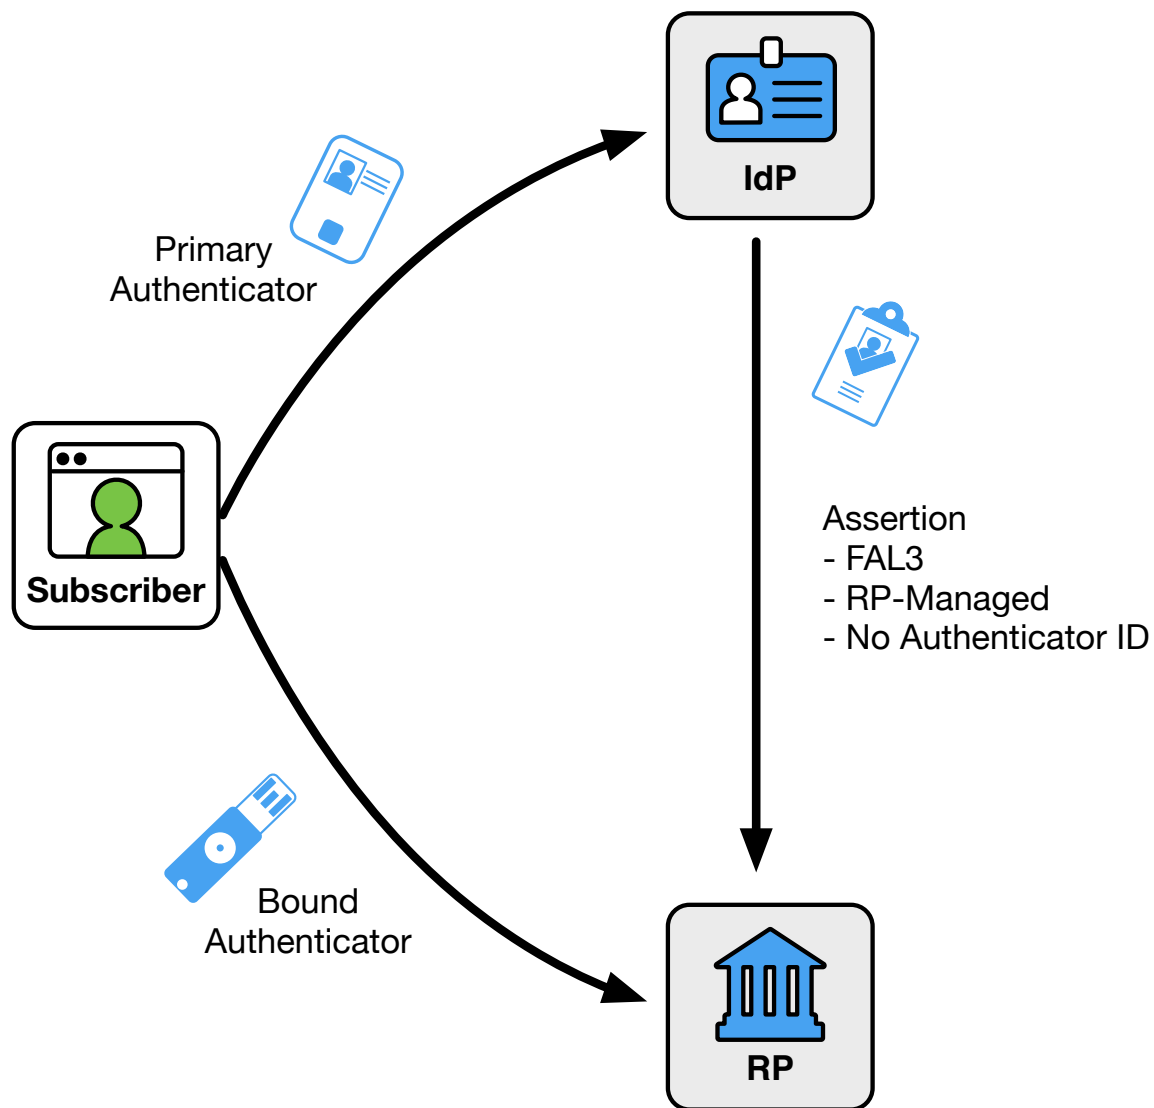


Figure 10. RP-Managed Bound Authenticators

Before an RP can successfully accept an FAL3 assertion, the RP subscriber account must include a bound authenticator. These authenticators can be provided by either the RP or the subscriber, with slightly different requirements applying to the initial binding of the authenticator to the RP subscriber account in each case.

For RP-provided authenticators, the administrator of the RP **SHALL** issue the

1368 authenticator to the subscriber directly for use with an FAL3 login. The administrator of
1369 the RP **SHALL** store a unique identifier for the bound authenticator in the RP subscriber
1370 account. The administrator of the RP **SHALL** determine through independent means that
1371 the party to which the authenticator is issued is the identified subject of the RP subscriber
1372 account.

1373 For subscriber-provided authenticators, if no bound authenticators are associated with
1374 the RP subscriber account, the RP **SHALL** perform a binding ceremony to establish the
1375 connection between the authenticator, the subscriber, and the RP subscriber account as
1376 shown in [Fig. 11](#). The RP **SHALL** first establish an authenticated session using federation
1377 with an assertion that meets all the other requirements of FAL3, including an indication
1378 that the assertion is intended for use at FAL3 with an RP-managed bound authenticator.
1379 The subscriber **SHALL** immediately be prompted to present and authenticate with the
1380 proposed authenticator. Upon successful presentation of the authenticator, the RP **SHALL**
1381 store a unique identifier for the authenticator (such as its public key) and associate this
1382 with the RP subscriber account associated with the federated identifier. If the subscriber
1383 fails to successfully present an appropriate authenticator, the binding ceremony fails. The
1384 binding ceremony session **SHALL** have a timeout of five minutes or less. The session
1385 used during the ceremony is not an authenticated session for the purposes of logging
1386 in. Upon successful completion of the binding ceremony, the RP **SHALL** immediately
1387 request a new assertion from the IdP at FAL3, including prompting the subscriber for the
1388 newly-bound authenticator.

1389 An RP **MAY** allow a subscriber to bind multiple subscriber-provided authenticators at
1390 FAL3. If this is the case, and the RP subscriber account has one or more existing bound
1391 authenticators, the binding ceremony makes use of the existing ability to reach FAL3.
1392 The subscriber **SHALL** first be prompted to present an existing bound authenticator to
1393 reach FAL3. Upon successful authentication, the RP **SHALL** immediately prompt the
1394 subscriber for the newly-bound authenticator.

1395 An RP **MAY** allow a subscriber to unbind a bound subscriber-provided authenticator
1396 from their RP subscriber account, thereby removing the ability to use that authenticator
1397 for FAL3. When a bound authenticator is unbound, the RP **SHALL** terminate all
1398 current FAL3 sessions for the subscriber and **SHALL** require reauthentication of the
1399 subscriber from the IdP. Note that in many cases, a subscriber will need to unbind a bound
1400 authenticator to account for a lost or compromised authenticator, and the subscriber will
1401 therefore not have access to the authenticator during the unbinding process.




Figure 11. Binding Ceremony
41

The RP **SHALL** notify the subscriber through an out-of-band mechanism, and **SHOULD** notify the IdP using a shared signaling system (see [Sec. 5.7](#)), if any of the following events occur:

- A new authenticator is bound to the RP subscriber account.
- An existing bound authenticator is unbound from the RP subscriber account.

For example, a subscriber could have a single factor cryptographic device as an authenticator. This authenticator uses name-based phishing resistance so the IdP and RP would see different keys when used in each location. The RP can use a binding ceremony as described here to allow the subscriber to use this device as a bound authenticator at FAL3. The RP will prompt the subscriber for this authenticator whenever it sees an assertion for this subscriber at FAL3 from the IdP.

6.1.2.3. Processing Bound Authenticators

When the RP receives an assertion associated with a bound authenticator, the subscriber proves possession of the bound authenticator directly to the RP. The primary authentication at the IdP and the federated authentication at the RP are processed separately. While the subscriber could use the same authenticator during the primary authentication at the IdP and as the bound authenticator at the RP, there is no assumption that these will be the same.

The following requirements apply to all assertions associated with a bound authenticator:

1. The subscriber **SHALL** prove possession of the bound authenticator to the RP, in addition to presentation of the assertion itself.
2. If the authenticator is managed at the IdP, reference to a given authenticator found within an assertion **SHALL** be trusted at the same level as all other information within the assertion.
3. If the authenticator is managed at the IdP, the assertion **SHALL NOT** include an unencrypted private or symmetric key to be used as an authenticator with the presentation.
4. The RP **SHALL** process and validate the assertion in addition to the bound authenticator.
5. Failure to authenticate with the bound authenticator **SHALL** result in an error at the RP.

6.2. Assertion Protection

Independent of the binding mechanism (discussed in [Sec. 6.1](#)) or the federation model used to obtain them (described in [Sec. 5.1](#)), assertions **SHALL** include a set of protections to prevent attackers from manufacturing valid assertions or reusing captured assertions at disparate RPs. The protections required are dependent on the details of the use case being considered, and specific protections are listed here.

6.2.1. Assertion Identifier

Assertions **SHALL** be sufficiently unique to permit unique identification by the target RP. Assertions **MAY** accomplish this by use of an embedded nonce, issuance timestamp, assertion identifier, or a combination of these or other techniques.

6.2.2. Signed Assertion

Assertions **SHALL** be cryptographically signed by the issuer (IdP). The RP **SHALL** validate the digital signature or MAC of each such assertion based on the issuer's key. This signature **SHALL** cover the entire assertion, including its identifier, issuer, audience, subject, and expiration.

The assertion signature **SHALL** either be a digital signature using asymmetric keys or a MAC using a symmetric key shared between the RP and issuer. Shared symmetric keys used for this purpose by the IdP **SHALL** be independent for each RP to which they send assertions, and are normally established during registration of the RP. Public keys for verifying digital signatures **SHALL** be transferred to the RP in a secure manner, and **MAY** be fetched by the RP in a secure fashion at runtime, such as through an HTTPS URL hosted by the IdP. Approved cryptography **SHALL** be used.

6.2.3. Encrypted Assertion

Encrypted assertions protect the contents of the assertion from being read by unintended parties, ensuring that only the targeted RP is able to read the assertion. Encrypting assertions provides two primary benefits: the assertion contents cannot be seen by any party other than the intended RP, and the assertion cannot be used by any RP other than the targeted one.

When encrypting assertions, the IdP **SHALL** encrypt the contents of the assertion using either the RP's public key or a shared symmetric key. Shared symmetric keys used for this purpose by the IdP **SHALL** be independent for each RP to which they send assertions, and are normally established during registration of the RP. Public keys for encryption **SHALL** be securely transferred to the IdP and **MAY** be fetched by the IdP in a secure fashion at runtime, such as through an HTTPS URL hosted by the RP.

All encryption of assertions **SHALL** use approved cryptography.

When personally-identifiable information is included in the assertion and the assertion is handled by intermediaries such as a browser, the federation protocol **SHALL** encrypt assertions to protect the sensitive information in the assertion from leaking to unintended parties. For example, a SAML assertion can be encrypted using XML-Encryption, or an OpenID Connect ID Token can be encrypted using JSON Web Encryption (JWE).

6.2.4. Audience Restriction

Assertions **SHALL** use audience restriction techniques to allow an RP to recognize whether or not it is the intended target of an issued assertion. All RPs **SHALL** check that the audience of an assertion contains an identifier for their RP to prevent the injection and replay of an assertion generated for one RP at another RP.

6.2.5. Pairwise Pseudonymous Identifiers

In some circumstances, it is desirable to prevent the subscriber account from being easily linked at multiple RPs through use of a common identifier. A pairwise pseudonymous identifier (PPI) allows an IdP to provide multiple distinct federated identifiers to different RPs for a single subscriber account. This prevents different RPs from colluding together to track the subscriber using the federated identifier.

6.2.5.1. General Requirements

When using pairwise pseudonymous identifiers within the assertions generated by the IdP for the RP, the IdP **SHALL** generate a different federated identifier for each RP as described in [Sec. 6.2.5.2](#) below.

When PPIs are used with RPs alongside attributes, it may still be possible for multiple colluding RPs to re-identify a subscriber by correlation across systems using these identity attributes. For example, if two independent RPs each see the same subscriber identified with different pairwise pseudonymous identifiers, they could still determine that the subscriber is the same person by comparing the name, email address, physical address, or other identifying attributes carried alongside the pairwise pseudonymous identifier in the respective assertions. Privacy policies **SHOULD** prohibit such correlation, and pairwise pseudonymous identifiers can increase effectiveness of these policies by increasing the administrative effort in managing the attribute correlation.

Note that in a proxied federation model, the initial IdP may be unable to generate a pairwise pseudonymous identifier for the ultimate RP, since the proxy could blind the IdP from knowing which RP is being accessed by the subscriber. In such situations, the pairwise pseudonymous identifier is generally established between the IdP and the federation proxy itself. The proxy, acting as an IdP, can itself provide pairwise pseudonymous identifiers to downstream RPs. Depending on the protocol, the federation proxy may need to map the pairwise pseudonymous identifiers back to the associated identifiers from upstream IdPs in order to allow the identity protocol to function. In such cases, the proxy will be able to track and determine which pairwise pseudonymous identifiers represent the same subscriber at different RPs. The proxy **SHALL NOT** disclose the mapping between the pairwise pseudonymous identifier and any other identifiers to a third party or use the information for any purpose other than federated authentication, related fraud mitigation, to comply with law or legal process, or in the case of a specific user request for the information.

6.2.5.2. Pairwise Pseudonymous Identifier Generation

Pairwise pseudonymous identifiers **SHALL** contain no identifying information about the subscriber. They **SHALL** also be unguessable by a party having access to some information identifying the subscriber. Pairwise pseudonymous identifiers **MAY** be generated randomly and assigned to subscribers by the IdP or **MAY** be derived from other subscriber information if the derivation is done in an irreversible, unguessable manner (e.g., using a keyed hash function with a secret key).

Normally, the identifiers **SHALL** only be known by and used by one pair of endpoints (e.g., IdP-RP). An IdP **MAY** generate the same identifier for a subscriber at multiple RPs at the request of those RPs, provided:

- The trust agreement stipulates a shared pseudonymous identifier for a specific family of RPs;
- The authorized party consents to and is notified of the use of a shared pseudonymous identifier;
- Those RPs have a demonstrable relationship that justifies an operational need for the correlation, such as a shared security domain or shared legal ownership; and
- All RPs sharing an identifier consent to being correlated in such a manner (i.e., one RP cannot request to have another RP's PPI without that other RP's knowledge and consent).

The RPs **SHALL** conduct a privacy risk assessment to consider the privacy risks associated with requesting a common identifier. See [Sec. 9.2](#) for further privacy considerations.

The IdP **SHALL** ensure that only intended RPs are correlated; otherwise, a rogue RP could learn of the pseudonymous identifier for a set of correlated RPs by fraudulently posing as part of that set.

6.3. Identity APIs

Attributes about the subscriber, including profile information, **MAY** be provided to the RP through a protected *attribute API* known as the *identity API*. The RP is granted limited access to the identity API during the federation transaction, in concert with the assertion. For example, in OpenID Connect, the UserInfo Endpoint provides a standardized identity API for fetching attributes about the subscriber. This API is protected by an OAuth 2.0 Access Token, which is issued to the RP along with OpenID Connect's assertion, the ID Token. The use of identity APIs along with federation assertions has several advantages for the overall security, privacy, and efficiency of the federation system.

By making attributes available at an identity API, the IdP no longer has to use the assertion to convey as much information to the RP. This not only means that sensitive attributes do not have to be carried in the assertion itself, it also makes the assertion

smaller and easier to process by the RP. The contents of the assertion can then be limited to essential fields (e.g., unique subject identifiers) and information about the immediate authentication event being asserted.

The RP often caches attributes provided by the IdP in an RP subscriber account, discussed in [Sec. 5.4](#). Attributes provided in the assertion are passed on every login, and since the RP does not know the identity of the subscriber before the attribute is requested, the IdP is incentivized to include as much information as possible in the assertion itself. However, most of a subscriber's attributes will not change in between subsequent logins, making this information redundant. As a consequence, most of these more-stable attributes can instead be made available through an identity API that is called by the RP only when necessary. The IdP can indicate in the assertion when the last time the subscriber's attributes have been updated in the subscriber account, allowing the RP to decide if it needs to fetch the attributes anew or if those in the RP subscriber account are sufficient.

Access to the identity API **SHALL** be time limited. The time limitation is separate from the validity time window of the assertion and the lifetime of the authenticated session at the RP. Access to an identity API by the RP without an associated valid assertion **SHALL NOT** be sufficient for the establishment of an authenticated session at the RP.

A given identity API deployment is expected to be capable of providing attributes for all subscribers for whom the IdP can create assertions. However, when access to the identity API is granted within the context of a federation transaction, the attributes provided by an identity API **SHALL** be associated with only the single subscriber identified in the associated assertion. If the identity API is hosted by the IdP, the returned attributes **SHALL** include the subject identifier for the subscriber. This allows the RP to positively correlate the assertion's subject to the returned attributes. Note that when access to an attribute API is provided as part of pre-provisioning of RP subscriber accounts as discussed in [Sec. 5.4.1](#), the RP is usually granted blanket access to the identity API outside the context of the federated transaction and these requirements do not apply.

6.3.1. Attribute Providers

While most attribute APIs used in federation are hosted as part of the IdP, it is also possible for the IdP to grant access to identity APIs hosted by external attribute providers. These services provide attributes about the subscriber in addition to those made available directly from the IdP.

When the IdP grants access to an attribute provider, the IdP is making an explicit statement that the information returned from the attribute provider is associated with the subscriber identified in the associated assertion. For the purposes of the trust agreement, the IdP is the responsible party for the accuracy and content of the attribute API.

The attributes returned by the attribute provider are assumed to be independent of those returned directly from the IdP, and as such **MAY** use different identifiers, formats, or

1586 schemas. The RP **SHALL** verify that the identified attribute provider is capable of
1587 providing the kinds of attributes that are present, under the auspices of the applicable
1588 trust agreement.

1589 For example, an IdP could provide access to a subscriber's medical license information as
1590 part of the federation process. Instead of the IdP asserting the license status directly, the
1591 IdP provides the RP access to a record for the subscriber at a medical licensure agency.
1592 The RP can make a strong association between the current subscriber and the license
1593 record, even though the license record will not likely use the same subject identifier that
1594 the IdP does in this case.

7. Assertion Presentation

This section is normative.

Depending on the specifics of the protocol, the RP and the IdP communicate with each other in two ways, which leads to two different ways in which an assertion can be passed from the IdP to the RP:

- The *front channel*, through redirects involving the subscriber and the subscriber's browser; or
- The *back channel*, through a direct connection between the RP and IdP, not involving the subscriber directly.

There are tradeoffs with each model, but each requires the proper validation of the assertion. Assertions **MAY** also be proxied to facilitate federation between IdPs and RPs using different presentation methods, as discussed in detail in [Sec. 5.1.3](#).

7.1. Back-Channel Presentation

In the *back-channel* presentation model, the subscriber is given an assertion reference to present to the RP, generally through the front channel. The assertion reference itself contains no information about the subscriber and **SHALL** be resistant to tampering and fabrication by an attacker. The RP presents the assertion reference to the IdP, usually along with authentication of the RP itself, to fetch the assertion.

As shown in [Figure 12](#), the back-channel presentation model consists of three steps:

1. The IdP sends an assertion reference to the subscriber through the front channel.
2. The subscriber sends the assertion reference to the RP through the front channel.
3. The RP presents the assertion reference and its RP credentials to the IdP through the back channel. The IdP validates the credentials and returns the assertion.

The assertion reference:

1. **SHALL** be limited to use by a single RP.
2. **SHALL** be single-use.
3. **SHALL** be time limited, and **SHOULD** have a lifetime of no more than a small number of minutes in length.
4. **SHALL** be presented along with authentication of the RP to the IdP.
5. **SHALL** contain at least 128 bits of entropy.

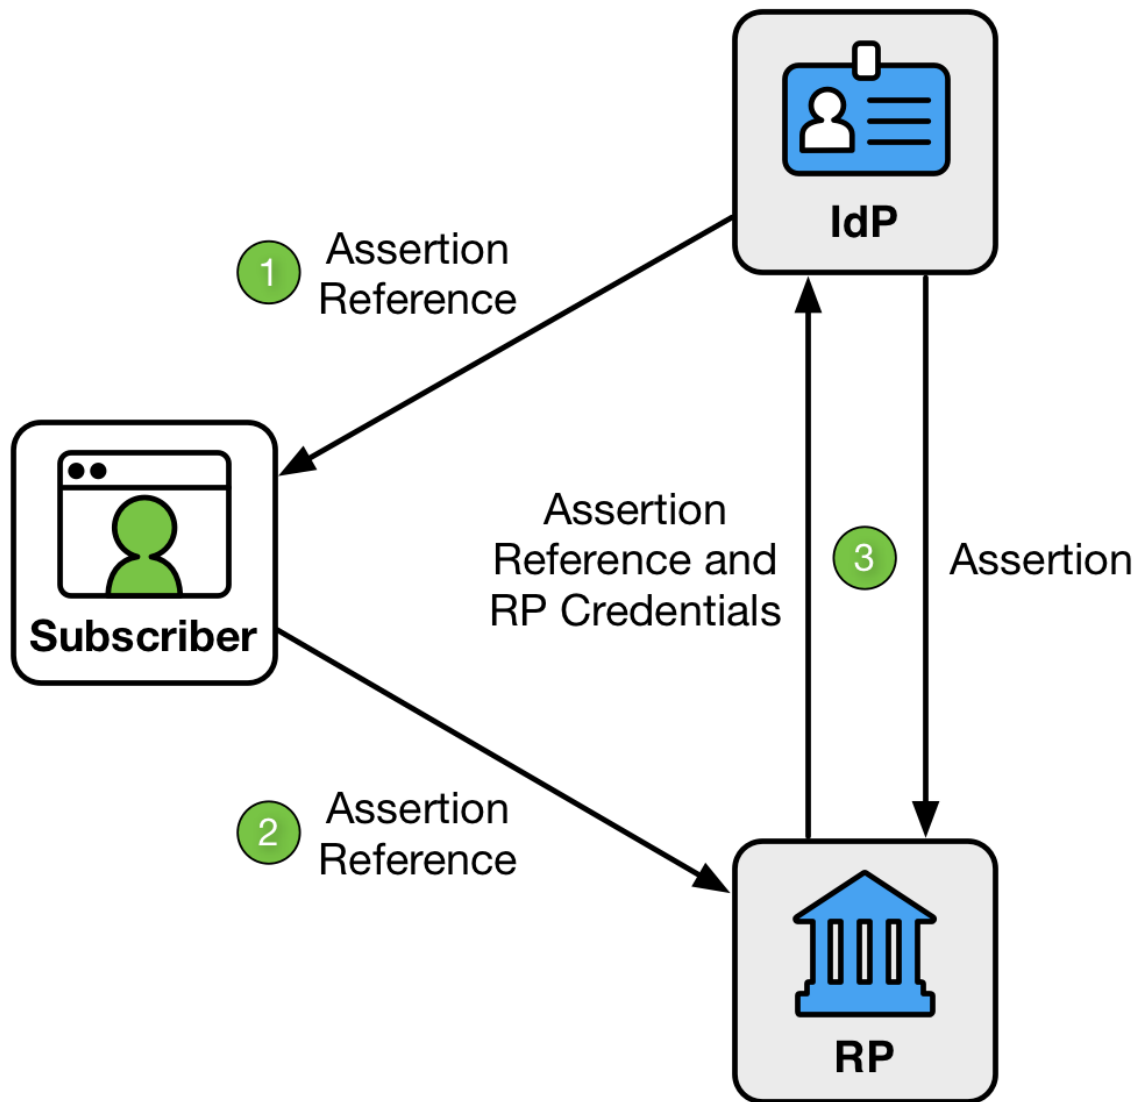


Figure 12. Back-channel Presentation

1625 In this model, the RP directly requests the assertion from the IdP, minimizing chances of
1626 interception and manipulation by a third party (including the subscriber themselves).

1627 This method also facilitates the RP querying the IdP for additional attributes about the
1628 subscriber not included in the assertion itself, since back-channel communication can
1629 continue to occur after the initial authentication transaction has been completed without
1630 sending the user back to the IdP. This query occurs using an identity API, as described in
1631 [Sec. 6.3](#).

1632 More network transactions are required in the back-channel method, but the information
1633 is limited to only those parties that need it. Since an RP is expecting to get an assertion
1634 only from the IdP directly, the attack surface is reduced. Consequently, it is more difficult
1635 to inject assertions directly into the RP and this presentation method is recommended for
1636 FAL2 and above.

1637 The RP **SHALL** protect itself against injection of manufactured or captured assertion
1638 references by use of cross-site scripting protection or other accepted techniques.

1639 Conveyance of the assertion reference from the IdP to the subscriber, as well as from
1640 the subscriber to the RP, **SHALL** be made over an authenticated protected channel.

1641 Conveyance of the assertion reference from the RP to the IdP, as well as the assertion
1642 from the IdP to the RP, **SHALL** be made over an authenticated protected channel.

1643 When assertion references are presented, the IdP **SHALL** verify that the party presenting
1644 the assertion reference is the same party that requested the authentication. The IdP can
1645 do this by requiring the RP to authenticate itself when presenting the assertion reference
1646 to the IdP or through other similar means (see [\[RFC7636\]](#) for one protocol's method of
1647 dynamic RP verification).

1648 Note that in a federation proxy described in [Sec. 5.1.3](#), the IdP audience restricts the
1649 assertion reference and assertion to the proxy, and the proxy restricts any newly-created
1650 assertion references or assertions to the downstream RP.

7.2. Front-Channel Presentation

In the *front-channel* presentation model, the IdP creates an assertion and sends it to the subscriber after successful authentication. The assertion is presented by the subscriber to authenticate to the RP, usually through mechanisms within the subscriber's browser such as redirects.

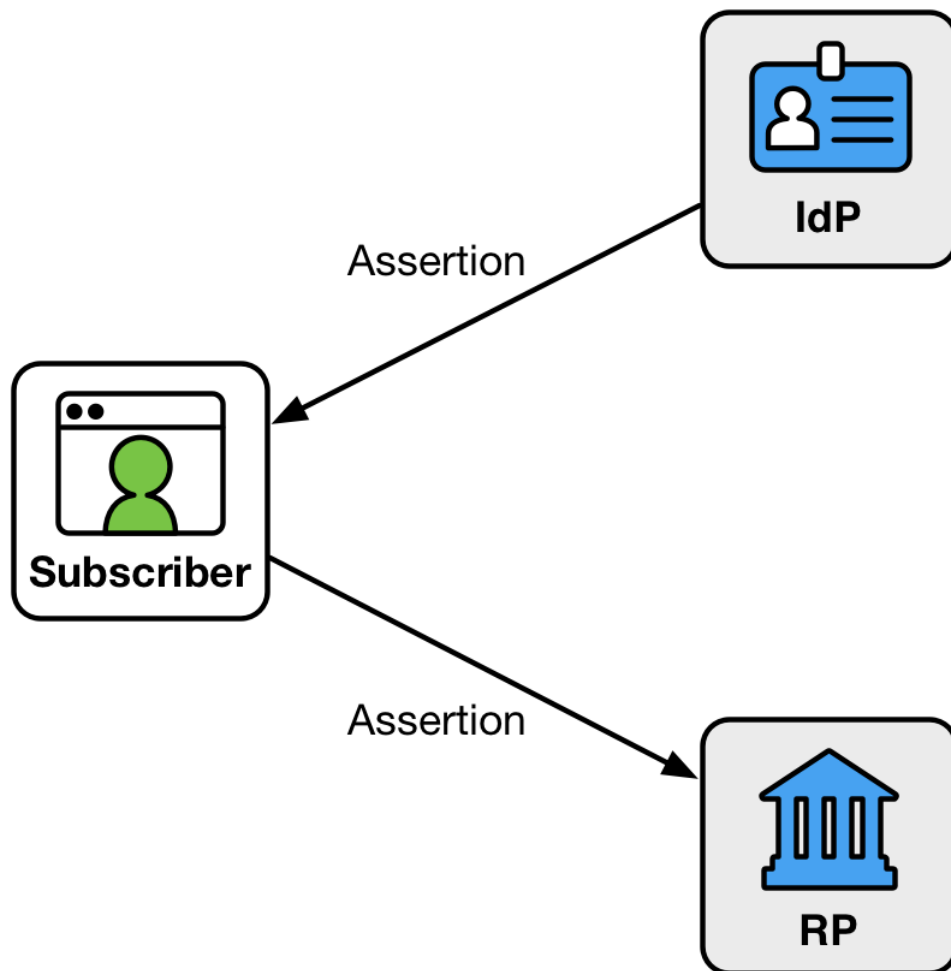


Figure 13. Front-channel Presentation

An assertion is visible to the subscriber in the front-channel method, which could potentially cause leakage of system information included in the assertion. Further, it is possible but more awkward in this model for the RP to query the IdP for additional attributes after the presentation of the assertion using an identity API, as described in [Sec. 6.3](#).

Since the assertion is under the subscriber's control, the front-channel presentation method also allows the subscriber to submit a single assertion to unintended parties,

perhaps by a browser replaying an assertion at multiple RPs. Even if the assertion is audience-restricted and rejected by unintended RPs, its presentation at unintended RPs could lead to leaking information about the subscriber and their online activities. Though it is possible to intentionally create an assertion designed to be presented to multiple RPs, this method can lead to lax audience restriction of the assertion itself, which in turn could lead to privacy and security breaches for the subscriber across these RPs. Such multi-RP use is not recommended. Instead, RPs are encouraged to fetch their own individual assertions.

The RP **SHALL** protect itself against injection of manufactured or captured assertions by use of cross-site scripting protection and other accepted techniques.

Conveyance of the assertion from the IdP to the subscriber, as well as from the subscriber to the RP, **SHALL** be made over an authenticated protected channel.

Note that in a federation proxy described in [Sec. 5.1.3](#), the IdP audience restricts the assertion to the proxy, and the proxy restricts any newly-created assertions to the downstream RP.

7.3. Protecting Information

Communications between the IdP and the RP **SHALL** be protected in transit using an authenticated protected channel. Communications between the subscriber and either the IdP or the RP (usually through a browser) **SHALL** be made using an authenticated protected channel.

Note that the IdP may have access to information that may be useful to the RP in enforcing security policies, such as device identity, location, system health checks, and configuration management. If so, it may be a good idea to pass this information along to the RP within the bounds of the subscriber's privacy preferences described in [Sec. 9.2](#).

Additional attributes about the user **MAY** be included outside of the assertion itself by use of authorized access to an identity API as discussed in [Sec. 6.3](#). Splitting user information in this manner can aid in protecting user privacy and allow for limited disclosure of identifying attributes on top of the essential information in the authentication assertion itself.

The RP **SHALL**, where feasible, request derived attribute values rather than full attribute values as described in [Sec. 9.3](#). The IdP **SHALL** support derived attribute values to the extent possible.

1695 **8. Security**

1696 *This section is informative.*

1697 Since the federated authentication process involves coordination between multiple
1698 components, including the CSP which now acts as an IdP, there are additional
1699 opportunities for attackers to compromise federated identity transactions. This section
1700 summarizes many of the attacks and mitigations applicable to federation.

1701 **8.1. Federation Threats**

1702 As in non-federated authentication, attackers' motivations are typically to gain access
1703 (or a greater level of access) to a resource or service provided by an RP. Attackers may
1704 also attempt to impersonate a subscriber. Rogue or compromised IdPs, RPs, user agents
1705 (e.g., browsers), and parties outside of a typical federation transaction are potential
1706 attackers. To accomplish their attack, they might intercept or modify assertions and
1707 assertion references. Further, two or more entities may attempt to subvert federation
1708 protocols by directly compromising the integrity or confidentiality of the assertion data.
1709 For the purpose of these types of threats, any authorized parties who attempt to exceed
1710 their privileges are considered attackers.

Table 2. Federation Threats

Federation Threats/Attacks	Description	Examples
Assertion Manufacture or Modification	The attacker generates a false assertion	Compromised IdP asserts identity of a claimant who has not properly authenticated
	The attacker modifies an existing assertion	Compromised proxy that changes AAL of an authentication assertion
Assertion Disclosure	Assertion visible to third party	Network monitoring reveals subscriber address of record to an outside party
Assertion Repudiation by the IdP	IdP later claims not to have signed transaction	User engages in fraudulent credit card transaction at RP, IdP claims not to have logged them in
Assertion Repudiation by the Subscriber	Subscriber claims not to have performed transaction	User agreement (e.g., contract) cannot be enforced
Assertion Redirect	Assertion can be used in unintended context	Compromised user agent passes assertion to attacker who uses it elsewhere
Assertion Reuse	Assertion can be used more than once with same RP	Intercepted assertion used by attacker to authenticate their own session
Assertion Substitution	Attacker uses an assertion intended for a different subscriber	Session hijacking attack between IdP and RP

1711 8.2. Federation Threat Mitigation Strategies

1712 Mechanisms that assist in mitigating the above threats are identified in [Table 3](#).

Table 3. Mitigating Federation Threats

Federation Threat/Attack	Threat Mitigation Mechanisms	Normative Reference(s)
Assertion Manufacture or Modification	Cryptographically sign the assertion at IdP and verify at RP	4.1, 6
	Send assertion over an authenticated protected channel authenticating the IdP	7.1, 7.2
	Include a non-guessable random identifier in the assertion	6.2.1
Assertion Disclosure	Send assertion over an authenticated protected channel authenticating the RP	7.1, 7.2
	Encrypt assertion for a specific RP (may be accomplished by use of a mutually authenticated protected channel)	6.2.3
Assertion Repudiation by the IdP	Cryptographically sign the assertion at the IdP with a key that supports non-repudiation; verify signature at RP	6.2.2
Assertion Repudiation by the Subscriber	Issue assertions with bound authenticators; proof of possession of bound authenticator verifies subscriber's participation to the RP	6.1.2
Assertion Redirect	Include identity of the RP ("audience") for which the assertion is issued in its signed content; RP verifies that they are intended recipient	6, 7.1, 7.2
Assertion Reuse	Include an issuance timestamp with short validity period in the signed content of the assertion; RP verifies validity	6, 7.1, 7.2
	RP keeps track of assertions consumed within a configurable time window to ensure that a given assertion is not used more than once.	6.2.1
Assertion Substitution	Ensure that assertions contain a reference to the assertion request or some other nonce that was cryptographically bound to the request by the RP	6
	Send assertions in the same authenticated protected channel as the request, such as in the back-channel model	7.1

9. Privacy Considerations

This section is informative.

9.1. Minimizing Tracking and Profiling

Federation offers numerous benefits to RPs and subscribers, but requires subscribers to have trust in the federation participants. [Sec. 5](#) and [Sec. 6.2.5](#) cover a number of technical requirements, the objective of which is to minimize privacy risks arising from increased capabilities to track and profile subscribers. For example, a subscriber using the same IdP to authenticate to multiple RPs allows the IdP to build a profile of subscriber transactions that would not have existed absent federation. The availability of such data makes it vulnerable to uses that may not be anticipated or desired by the subscriber and may inhibit subscriber adoption of federated services.

[Sec. 5.5](#) requires IdPs to use measures to maintain the objectives of predictability (enabling reliable assumptions by individuals, owners, and operators about PII and its processing by an information system) and manageability (providing the capability for granular administration of PII, including alteration, deletion, and selective disclosure) commensurate with privacy risks that can arise from the processing of attributes for purposes other than identity proofing, authentication, authorization, or attribute assertions, related fraud mitigation, or to comply with law or legal process [[NISTIR8062](#)].

IdPs may have various business purposes for processing attributes, including providing non-identity services to subscribers. However, processing attributes for different purposes from the original collection purpose can create privacy risks when individuals are not expecting or comfortable with the additional processing. IdPs can determine appropriate measures commensurate with the privacy risk arising from the additional processing. For example, absent applicable law, regulation or policy, it may not be necessary to get consent when processing attributes to provide non-identity services requested by subscribers, although notices may help subscribers maintain reliable assumptions about the processing (predictability). Other processing of attributes may carry different privacy risks that call for obtaining consent or allowing subscribers more control over the use or disclosure of specific attributes (manageability). Subscriber consent needs to be meaningful; therefore, when IdPs do use consent measures, they cannot make acceptance by the subscriber of additional uses a condition of providing the identity service.

Consult the SAOP if there are questions about whether the proposed processing falls outside the scope of the permitted processing or the appropriate privacy risk mitigation measures.

[Sec. 5.5](#) also encourages the use of technical measures to provide disassociability (enabling the processing of PII or events without association to individuals or devices beyond the operational requirements of the system) and prevent subscriber activity

tracking and profiling [NISTIR8062]. Technical measures, such as those outlined in Sec. 5.1.3 for proxied federation and Sec. 6.2.5 for pairwise pseudonymous identifiers, can increase the effectiveness of policies by making it more difficult to track or profile subscribers beyond operational requirements.

9.2. Notice and Consent

To build subscriber trust in federation, subscribers need to be able to develop reliable assumptions about how their information is being processed. For instance, it can be helpful for subscribers to understand what information will be transmitted, which attributes for the transaction are required versus optional, and to have the ability to decide whether to transmit optional attributes to the RP. Accordingly, Sec. 5.1 requires that positive confirmation be obtained from the authorized party before any attributes about the subscriber are transmitted to any RP. In determining when a set of RPs should share a common pairwise pseudonymous identifier as in Sec. 6.2.5.2, the IdP considers the subscriber's understanding of such a grouping of RPs and the role of notice in assisting such understanding. An effective notice will take into account user experience design standards and research, as well as an assessment of privacy risks that may arise from the information processing. There are various factors to be considered, including the reliability of the assumptions subscribers may have about the processing and the role of different entities involved in federation. However, a link to a complex, legalistic privacy policy or general terms and conditions that a substantial number of subscribers do not read or understand is never an effective notice.

Sec. 5.1 does not specify which party should provide the notice. In some cases, a party in a federation may not have a direct connection to the subscriber in order to provide notice and obtain consent. Although multiple parties may elect to provide notice, it is permissible for parties to determine in advance, either contractually or through trust framework policies, which party will provide the notice and obtain confirmation, as long as the determination is being based upon factors that center on enabling the subscriber to pay attention to the notice and make an informed choice.

If an IdP is using an allowlist of RPs as described in Sec. 5.3, any RPs on that list are not presented to the subscriber during an authentication transaction. Since the IdP does not provide notice to the subscriber at runtime, the IdP makes its list of allowlisted RPs available to the subscriber so that the subscriber can see which RPs on the allowlist have access to which of the subscriber's attributes in an authentication transaction. Since IdPs can not share a subscriber's authentication information or attributes with an allowlisted RP outside of an authentication transaction involving the subscriber (see Sec. 5.5), the existence of an RP on a list of IdPs does not indicate that the subscriber's information will be shared. However, if the subscriber logs into any of the allowlisted RPs using the IdP, the attributes indicated will be shared as part of the authentication transaction.

If a subscriber's runtime decisions at the IdP were stored in the subscriber account by the IdP to facilitate future transactions, the IdP also needs to allow the subscriber to view

and revoke any RPs that were previously approved during a runtime decision. This list includes information on which attributes were approved. Similarly, if a subscriber's runtime decisions at the RP are stored in some fashion, the RP also needs to allow the subscriber to view and revoke any IdPs that were approved during a runtime decision.

9.3. Data Minimization

Federation enables the data exposed to an RP to be minimized, which can yield privacy protections for subscribers. Although an IdP may collect additional attributes beyond what the RP requires for its use case, only those attributes that were explicitly requested by the RP are to be transmitted by the IdP. In some instances, an RP does not require a full value of an attribute. For example, an RP may need to know whether the subscriber is over 13 years old, but has no need for the full date of birth. To minimize collection of potentially sensitive PII, the RP may request a derived attribute value (e.g., Question: Is the subscriber over 13 years old? Response: Y/N or Pass/Fail). This minimizes the RP's collection of potentially sensitive and unnecessary PII. Accordingly, [Sec. 7.3](#) requires the RP to, where feasible, request derived attribute values rather than full attribute values. To support this RP requirement IdPs are, in turn, required to support a derived attribute value.

9.4. Agency-Specific Privacy Compliance

[Section 5.5](#) identifies agency requirements to consult their SAOP to determine privacy compliance requirements. It is critical to involve the agency's SAOP in the earliest stages of digital authentication system development to assess and mitigate privacy risks and advise the agency on compliance obligations such as whether the federation triggers the Privacy Act of 1974 or the E-Government Act of 2002 requirement to conduct a PIA. For example, if the agency is serving as an IdP in a federation, it is likely that the Privacy Act requirements will be triggered and require coverage by either a new or existing Privacy Act system of records since credentials would be maintained at the IdP on behalf of any RP it federates with. If, however, the agency is an RP and using a third-party IdP, digital authentication may not trigger the requirements of the Privacy Act, depending on what data passed from the RP is maintained by the agency at the RP (in such instances the agency may have a broader programmatic SORN that covers such data).

The SAOP can similarly assist the agency in determining whether a PIA is required. These considerations should not be read as a requirement to develop a Privacy Act SORN or PIA for use of a federated credential alone. In many cases it will make the most sense to draft a PIA and SORN that encompasses the entire digital authentication process or includes the digital authentication process as part of a larger programmatic PIA that discusses the program or benefit the agency is establishing online access.

Due to the many components of digital authentication, it is important for the SAOP to have an awareness and understanding of each individual component. For example, other privacy artifacts may be applicable to an agency offering or using federated IdP or RP

services, such as Data Use Agreements, Computer Matching Agreements, etc. The SAOP can assist the agency in determining what additional requirements apply. Moreover, a thorough understanding of the individual components of digital authentication will enable the SAOP to thoroughly assess and mitigate privacy risks either through compliance processes or by other means.

9.5. Blinding in Proxied Federation

While some proxy structures — typically those that exist primarily to simplify integration — may not offer additional subscriber privacy protection, others offer varying levels of privacy to the subscriber through a range of blinding technologies. Privacy policies may dictate appropriate use of the subscriber attributes and authentication transaction data (e.g., identities of the ultimate IdP and RP) by the IdP, RP, and the federation proxy.

Technical means such as blinding can increase effectiveness of these policies by making the data more difficult to obtain. A proxy-based system has three parties, and the proxy can be used to hide information from one or more of the parties, including itself. In a double-blind proxy, the IdP and RP do not know each other's identities, and their relationship is only with the proxy. In a triple-blind proxy, the proxy additionally does not have insight into the data being passed through it. As the level of blinding increases, the technical and operational implementation complexity may increase. Since proxies need to map transactions to the appropriate parties on either side as well as manage the keys for all parties in the transaction, fully triple-blind proxies are very difficult to implement in practice.

Even with the use of blinding technologies, a blinded party may still infer protected subscriber information through released attribute data or metadata, such as by analysis of timestamps, attribute bundle sizes, or attribute signer information. The IdP could consider additional privacy-enhancing approaches to reduce the risk of revealing identifying information of the entities participating in the federation.

The following table illustrates a spectrum of blinding implementations used in proxied federation. This table is intended to be illustrative, and is neither comprehensive nor technology-specific.

Table 4. Proxy Characteristics

Proxy Type	RP knows IdP	IdP knows RP	Proxy can track subscriptions between RP and IdP	Proxy can see attributes of Subscriber
Non-Blinding Proxy with Attributes	Yes	Yes	Yes	Yes
Non-Blinding Proxy	Yes	Yes	Yes	N/A
Double Blind Proxy with Attributes	No	No	Yes	Yes
Double Blind Proxy	No	No	Yes	N/A
Triple Blind Proxy with or without Attributes	No	No	No	No

10. Usability Considerations

This section is informative.

Ergonomic of Human-System Interaction — Part 11: Usability: Definitions and Concepts [ISO/IEC9241-11] defines usability as the “extent to which a system, product or service can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use.” This definition focuses on users, goals, and context of use as key elements necessary for achieving effectiveness, efficiency and satisfaction. A holistic approach considering these key elements is necessary to achieve usability.

From the usability perspective, one of the major potential benefits of federated identity systems is to address the problem of user fatigue associated with managing multiple authenticators. While this has historically been a problem with usernames and passwords, the increasing need for users to manage many authenticators — whether physical or digital — presents a usability challenge.

While many other approaches to authentication have been researched extensively and have well-established usability guidelines, federated identity is more nascent and, therefore, lacks the depth and conclusiveness of research findings. As ongoing usability research matures, usability guidelines for federated identity systems will have stronger supporting data. For example, additional data is needed to support guidance on the translation of technical attribute names and values into user-friendly language.

As stated in the usability sections in 800-63A and 800-63B, overall user experience is critical to the success of any authentication method. This is especially true for federated identity systems as federation is a less familiar user interaction paradigm for many users. Users’ prior authentication experiences may influence their expectations.

The overall user experience with federated identity systems should be as smooth and easy as possible. This can be accomplished by following usability standards (such as the ISO 25060 series of standards) and established best practices for user interaction design.

Note: In this section, the term “users” means “claimants” or “subscribers.”
The terms “entity” and “entities” refer to the parties of federated systems.

Guidelines and considerations are described from the users’ perspective.

Accessibility differs from usability and is out of scope for this volume. [Section508] was enacted to eliminate barriers in information technology and requires federal agencies to make their electronic and information technology public content accessible to people with disabilities. Refer to Section 508 law and standards for accessibility guidance.

10.1. General Usability Considerations

Federated identity systems should:

- Minimize user burden (e.g., frustration, learning curve)
 - Minimize the number of user actions required.
 - Allow users to quickly and easily select among multiple subscriber accounts with a single IdP. For example, approaches such as [Account Chooser](#) allow users to select from a list of subscriber accounts they have accessed in the recent past, rather than start the federation process by selecting their IdP from a list of potential IdPs.
 - Balance minimizing user burden with the need to provide sufficient information to enable users to make informed decisions.
- Minimize the use of unfamiliar technical jargon and details (e.g., users do not need to know the terms IdP and RP if the basic concepts are clearly explained).
- Strive for a consistent and integrated user experience across the IdP and RP.
- Help users establish an understanding of identity by providing resources to users such as graphics, illustrations, FAQs, tutorials and examples. Resources should explain how users' information is treated and how transacting parties (e.g., RPs, IdPs, and brokers) relate to each other.
- Provide clear, honest, and meaningful communications to users (i.e., communications should be explicit and easy to understand).
- Provide users online services independent of location and device.
- Make trust relationships explicit to users to facilitate informed trust decisions. Trust relationships are often dynamic and context dependent. Users may be more likely to trust some IdPs and RPs with certain attributes or transactions more than others. For example, users may be more hesitant to use federated identity systems on websites that contain valuable personal information (such as financial or health). Depending on the perceived sensitivity of users' personal data, users may be less comfortable with social network providers as IdPs since people are often concerned with the broadcasting nature of social networking implementations.
- Follow the usability considerations specified in [\[SP800-63A\]](#), Sec. 9 for any user-facing information.
- Clearly communicate how and where to acquire technical assistance. For example, provide users with information such as a link to an online self-service feature, chat sessions or a phone number for help desk support. Avoid redirecting users back and forth among transacting parties (e.g., RPs, IdPs, and brokers) to receive technical assistance.
- Perform integrative and continuous usability evaluations with representative users and realistic tasks in an appropriate context to ensure success of federated identity systems from the users' perspectives.

10.2. Specific Usability Considerations

This section addresses the specific usability considerations that have been identified with federated identity systems. This section does not attempt to present exhaustive coverage of all usability factors related to federated identity systems. Rather it is focused on the larger, more pervasive themes in the usability literature, primarily users' perspectives on identity, user adoption, trust, and perceptions of federated identity space. In some cases, implementation examples are provided. However, specific solutions are not prescribed. The implementations mentioned are examples to encourage innovative technological approaches to address specific usability needs. See standards for system design and coding, specifications, APIs, and current best practices (such as OpenID and OAuth) for additional examples. Implementations are sensitive to many factors that prevent a one-size-fits-all solution.

10.2.1. User Perspectives on Online Identity

Even when users are familiar with federated identity systems, there are different approaches to federated identity (especially in terms of privacy and the sharing of information) that make it necessary to establish reliable expectations for how users' data are treated. Users and implementers have different concepts of identity. Users think of identity as logging in and gaining access to their own private space. Implementers think of identity in terms of authenticators and assertions, assurance levels, and the necessary set of identity attributes to provide a service. Given this disconnect between users' and implementers' concepts of identity, it is essential to help users form an accurate concept of identity as it applies to federated identity systems. A good model of identity provides users a foundation for understanding the benefits and risks of federated systems and encourage user adoption and trust of these systems.

Many properties of identity have implications for how users manage identities, both within and among federations. Just as users manage multiple identities based on context outside of cyberspace, users must learn to manage their identity in a federated environment. Therefore, it must be clear to users how identity and context are used. The following factors should be considered:

- Provide users the requisite context and scope in order to distinguish among different user roles. For example, whether the user is acting on their own behalf or on behalf of another, such as their employer.
- Provide users unique, meaningful, and descriptive identifiers to distinguish among entities such as IdPs, RPs, and accounts. Any such user-facing identifiers are likely to be in addition to identifiers used by the underlying protocols, which are not normally exposed to the user.
- Provide users with information on data ownership and those authorized to make changes. Identities, and the data associated with them, can sometimes be updated and changed by multiple actors. For example, some healthcare data is updated and

- 1970 owned by the patient, while some data is only updated by a hospital or doctor's
1971 practice.
- 1972 • Provide users with the ability to easily verify, view, and update attributes. Identities
1973 and user roles are dynamic and not static; they change over time (e.g., age, health,
1974 and financial data). The ability to update attributes or make attribute release
1975 decisions may or may not be offered at the same time. Ensure the process for how
1976 users can change attributes is well known, documented, and easy to perform.
 - 1977 • Provide users means for updating data, even if the associated entity no longer exists.
 - 1978 • Provide users means to delete their identities completely, removing all information
1979 about themselves, including transaction history. Consider applicable audit, legal, or
1980 policy constraints that may preclude such action. In certain cases, full deactivation
1981 is more appropriate than deletion.
 - 1982 • Provide users with clear, easy-to-find, site/application data retention policy
1983 information.
 - 1984 • Provide users with appropriate anonymity and pseudonymity options, and the
1985 ability to switch among such identity options as desired, in accordance with an
1986 organization's data access policies.
 - 1987 • Provide means for users to manage each IdP to RP connection, including complete
1988 separation as well as the removal of RP access to one or more attributes.

1989 **10.2.2. User Perspectives of Trust and Benefits**

1990 Many factors can influence user adoption of federated identity systems. As with any
1991 technology, users may value some factors more than others. Users often weigh perceived
1992 benefits versus risks before making technology adoption decisions. It is critical that
1993 IdPs and RPs provide users with sufficient information to enable them to make informed
1994 decisions. The concepts of trust and tiers of trust — fundamental principles in federated
1995 identity systems — can drive user adoption. Finally, a positive user experience may also
1996 result in increased user demand for federation, triggering increased adoption by RPs.

1997 This sub-section is focused primarily on user trust and user perceptions of benefits versus
1998 risks.

1999 To encourage user adoption, IdPs and RPs need to establish and build trust with users and
2000 provide them with an understanding of the benefits and risks of adoption. The following
2001 factors should be considered:

- 2002 • Allow users to control their information disclosure and provide explicit consent
2003 through the appropriate use of notifications (see [Sec. 9.2](#)). Balancing the content,
2004 size, and frequency of notifications is necessary to avoid thoughtless user click-
2005 through.

- 2006 • For attribute sharing, consider the following:
 - 2007 – Provide a means for users to verify those attributes and attribute values that
 - 2008 will be shared. Follow good security practices (see [Sec. 7.3](#) and [Sec. 8](#)).
 - 2009 – Enable users to consent to a partial list of attributes, rather than an all-or-
 - 2010 nothing approach. Allow users some degree of online access, even if the user
 - 2011 does not consent to share all information.
 - 2012 – Allow users to update their consent to their list of shared attributes.
 - 2013 – Minimize unnecessary information presented to users. For example, do
 - 2014 not display system generated attributes (such as pairwise pseudonymous
 - 2015 identifiers) even if they are shared with the RP as part of the authentication
 - 2016 response.
 - 2017 – Minimize user steps and navigation. For example, build attribute consent into
 - 2018 the protocols so they're not a feature external to the federated transaction.
 - 2019 Examples can be found in standards such as OAuth or OpenID Connect.
 - 2020 – Provide effective and efficient redress methods such that a user can recover
 - 2021 from invalid attribute information claimed by the IdP or collected by the RP.
 - 2022 – Minimize the number of times a user is required to consent to attribute sharing.
 - 2023 Limiting the frequency of consent requests avoids user frustration from
 - 2024 multiple requests to share the same attribute.
- 2025 • Collect information for constrained usage only, and minimize information
- 2026 disclosure (see [Sec. 9.3](#)). User trust is eroded by unnecessary and superfluous
- 2027 information collection and disclosure or user tracking without explicit user consent.
- 2028 For example, only request attributes from the user that are relevant to the current
- 2029 transaction, not for all possible transactions a user may or may not access at the RP.
- 2030 • Clearly and honestly communicate potential benefits and risks of using federated
- 2031 identity to users. Benefits that users value include time savings, ease of use,
- 2032 reduced number of passwords to manage, and increased convenience.

2033 User concern over risk can negatively influence willingness to adopt federated identity
2034 systems. Users may have trust concerns, privacy concerns, security concerns, and single-
2035 point-of-failure concerns. For example, users may be fearful of losing access to multiple
2036 RPs if a single IdP is unavailable, either temporarily or permanently. Additionally, users
2037 may be concerned or confused about learning a new authentication process. In order to
2038 foster the adoption of federated identity systems, the perceived benefits must outweigh
2039 the perceived risks.

10.2.3. User Mental Models and Beliefs

Users' beliefs and perceptions predispose them to expect certain results and to behave in certain ways. Such beliefs, perceptions, and predispositions are referred to in the social sciences as mental models. For example, people have a mental model of dining out which guides their behavior and expectations at each establishment, such as fast food restaurants, cafeterias, and more formal restaurants. Thus, it is not necessary to be familiar with every establishment to understand how to interact appropriately at each one.

Assisting users in establishing good and complete mental models of federation allows users to generalize beyond a single specific implementation. If federated identity systems are not designed from users' perspectives, users may form incorrect or incomplete mental models that impact their willingness to adopt these systems. The following factors should be considered:

- Clearly explain the working relationship and information flow among the transacting parties (e.g., RPs, IdPs, and brokers) to avoid user misconceptions. Use the actual names of the entities in the explanation rather than using the generic terms IdPs and RPs.
 - Provide prominent visual cues and information so that users understand why seemingly unrelated entities have a working relationship. For example, users may be concerned with mixing online personal activities with government services due to a lack of understanding of the information flow in federated identity systems.
 - Provide prominent visual cues and information to users about redirection when an RP needs to redirect control from their site to an IdP. For example, display RP branding within the IdP user interface to inform users when they are logging in with their IdP for access to the destination RP.
- Provide users with clear and usable ways (e.g., visual assurance) to determine the authenticity of the transacting parties (e.g., RPs, IdPs, and brokers). This will also help to alleviate user concern over leaving one domain for another, especially if the root domain changes (e.g., .gov to .com). For example, display the URL of the IdP so that the user can verify that they are not being phished by a malicious site.
- Provide users with clear information, including visual cues, regarding implicit logins and explicit logouts. Depending on the implementation, logging into an RP with a federated account may authenticate users to both the IdP and RP. Users may not realize that ending their session with the RP will not necessarily end their session with the IdP; users will need to explicitly "log out" of the IdP. Users require clear information to remind them if explicit logouts are required to end their IdP sessions.

11. Equity Considerations

This section is informative.

Equitable access to the functions of IdPs and RPs is an essential element of a federated identity system. The ability for all subscribers to authenticate reliably is required to provide equitable access to government services, even when using federation technology, as specified in Executive Order 13985, *Advancing Racial Equity and Support for Underserved Communities Through the Federal Government* [EO13985]. In assessing equity risks, IdPs and RPs should consider the overall user population served by their federated identity service. Additionally, IdPs and RPs further identify groups of users within the population whose shared characteristics can cause them to be subject to inequitable access, treatment, or outcomes when using that service. The Usability Considerations provided in [Sec. 10](#) should also be considered to help ensure the overall usability and equity for all persons using federated identity services.

In its role as the verifier, the IdP needs to be aware of equity considerations related to identity proofing, attribute validation, and enrollment as enumerated in [\[SP800-63A\]](#) Sec. 11 and equity considerations concerning authenticators as enumerated in [\[SP800-63B\]](#) Sec. 11. An RP offering FAL3 will also need to be aware of these same authenticator considerations when processing bound authenticators, whether the authenticators are managed at the IdP or RP.

Since the federation process takes place over a network protocol between multiple active parties, the experience of authenticating using the federation system may present equity problems, such as the following examples:

- Completing the entire federated transaction without timing out may be difficult for subscribers without a reliable network connection, such as those in rural areas.
- It may be difficult to provide informed consent for a runtime decision regarding the release of attributes for subscribers with intellectual, developmental, learning, or neurocognitive difficulties.
- Systems with sufficient processing power, network access, and other features required to interact with both the IdP and the RP simultaneously may be difficult to afford for some subscribers.
- Subscribers that share devices may find allowlist-based systems difficult to manage securely, as other users of the device could silently gain unintended access to an RP through a session still active at the IdP.
- It could be prohibitively difficult to re-establish an account at the RP for subscribers who lose access to their IdP for any of a variety of reasons.

Normative requirements have been established requiring IdPs and RPs to mitigate the problems in this area that are expected to be most common. However, normative

2114 requirements are unlikely to have anticipated all potential equity problems. Potential
2115 equity problems also will vary for different applications. Accordingly, IdPs and RPs need
2116 to provide mechanisms for subscribers to report inequitable authentication requirements
2117 and to advise them on potential alternative authentication strategies.

2118 This guideline allows the binding of additional federated identifiers to an RP subscriber
2119 account to minimize the risk of IdP access loss (see [Sec. 5.4](#)). However, a subscriber
2120 might find it difficult to have multiple IdP accounts that are acceptable to the RP at the
2121 same time. This inequity can be addressed by having the RP having its own account
2122 recovery process that allows for the secure binding and unbinding of multiple federated
2123 identifiers from the RP subscriber account.

2124 RPs need to be aware that not all subscribers will necessarily have access to the same
2125 IdPs. The RPs can institute locally authenticated accounts for such subscribers, and later
2126 allow binding of those accounts to federated identifiers.

12. Examples

This section is informative.

Three types of assertion technologies are discussed below: SAML assertions, Kerberos tickets, and OpenID Connect tokens. This list is not inclusive of all possible assertion technologies, but does represent those commonly used in federated identity systems.

12.1. Security Assertion Markup Language (SAML)

SAML is an XML-based framework for creating and exchanging authentication and attribute information between trusted entities over the internet. As of this writing, the latest specification for SAML is SAML v2.0, issued 15 March 2005.

The building blocks of SAML include:

- The Assertions XML schema, which defines the structure of the assertion.
- The SAML Protocols, which are used to request assertions and artifacts (the assertion references used in the indirect model described in [Sec. 7.1](#)).
- The Bindings, which define the underlying communication protocols (such as HTTP or SOAP), and can be used to transport the SAML assertions.

The three components above define a SAML profile that corresponds to a particular use case such as “Web Browser SSO”.

SAML Assertions are encoded in an XML schema and can carry up to three types of statements:

- *Authentication statements* include information about the assertion issuer, the authenticated subscriber, validity period, and other authentication information. For example, an Authentication Assertion would state the subscriber “John” was authenticated using a password at 10:32pm on 06-06-2004.
- *Attribute statements* contain specific additional characteristics related to the subscriber. For example, subject “John” is associated with attribute “Role” with value “Manager”.
- *Authorization statements* identify the resources the subscriber has permission to access. These resources may include specific devices, files, and information on specific web servers. For example, subject “John” for action “Read” on “Webserver1002” given evidence “Role”.

Authorization statements are beyond the scope of this document and will not be discussed.

12.2. Kerberos Tickets

The Kerberos Network Authentication Service [RFC4120] was designed to provide strong authentication for client/server applications using symmetric-key cryptography on a local, shared network. Extensions to Kerberos can support the use of public key cryptography for selected steps of the protocol. Kerberos also supports confidentiality and integrity protection of session data between the subscriber and the RP. Even though Kerberos uses assertions, it was designed for use on shared networks and, therefore, is not truly a federation protocol.

Kerberos supports authentication of a subscriber over a network using one or more IdPs. The subscriber implicitly authenticates to the IdP by demonstrating the ability to decrypt a random session key encrypted for the subscriber by the IdP. (Some Kerberos variants also require the subscriber to explicitly authenticate to the IdP, but this is not universal.) In addition to the encrypted session key, the IdP also generates another encrypted object called a Kerberos ticket. The ticket contains the same session key, the identity of the subscriber to whom the session key was issued, and an expiration time after which the session key is no longer valid. The ticket is confidentiality and integrity protected by a pre-established key that is shared between the IdP and the RP during an explicit setup phase.

To authenticate using the session key, the subscriber sends the ticket to the RP along with encrypted data that proves that the subscriber possesses the session key embedded within the Kerberos ticket. Session keys are either used to generate new tickets or to encrypt and authenticate communications between the subscriber and the RP.

To begin the process, the subscriber sends an authentication request to the Authentication Server (AS). The AS encrypts a session key for the subscriber using the subscriber's long-term credential. The long-term credential may either be a secret key shared between the AS and the subscriber, or in the PKINIT variant of Kerberos, a public key certificate. Most variants of Kerberos based on a shared secret key between the subscriber and IdP derive this key from a user-generated password. As such, they are vulnerable to offline dictionary attacks by passive eavesdroppers, unless Flexible Authentication Secure Tunneling (FAST) [RFC6113] or some other tunneling and armoring mechanism is used.

In addition to delivering the session key to the subscriber, the AS also issues a ticket using a key it shares with the Ticket Granting Server (TGS). This ticket is referred to as a Ticket Granting Ticket (TGT), since the verifier uses the session key in the TGT to issue tickets rather than to explicitly authenticate the verifier. The TGS uses the session key in the TGT to encrypt a new session key for the subscriber and uses a key it shares with the RP to generate a ticket corresponding to the new session key. The subscriber decrypts the session key and uses the ticket and the new session key together to authenticate to the RP.

When Kerberos authentication is based on passwords, the protocol is known to be vulnerable to offline dictionary attacks by eavesdroppers who capture the initial user-to-KDC exchange. Longer password length and complexity provide some mitigation

to this vulnerability, although sufficiently long passwords tend to be cumbersome for users. However, when Kerberos password-based authentication is used in a FAST (or similar) tunnel, a successful attacker-in-the-middle attack is additionally required in order to perform the dictionary attack.

12.3. OpenID Connect

OpenID Connect [OIDC] is an internet-scale federated identity and authentication protocol built on top of the OAuth 2.0 authorization framework and the JSON Object Signing and Encryption (JOSE) cryptographic system.

OpenID Connect builds on top of the OAuth 2.0 authorization protocol to enable the subscriber to authorize the RP to access the subscriber's identity and authentication information. The RP in both OpenID Connect and OAuth 2.0 is known as the client.

In a successful OpenID Connect transaction, the IdP issues an ID Token, which is a signed assertion in JSON Web Token (JWT) format. The client parses the ID Token to learn about the subscriber and primary authentication event at the IdP. This token contains at minimum the following information about the subscriber and authentication event:

- `iss` - An HTTPS URL identifying the IdP that issued the assertion.
- `sub` - An IdP-specific subject identifier representing the subscriber.
- `aud` - An IdP-specific audience identifier, equal to the OAuth 2.0 client identifier of the client at the IdP.
- `exp` - The timestamp at which the ID Token expires and after which **SHALL NOT** be accepted the client.
- `iat` - The timestamp at which the ID Token was issued and before which **SHALL NOT** be accepted by the client.

In addition to the ID Token, the IdP also issues the client an OAuth 2.0 access token which can be used to access the UserInfo Endpoint at the IdP. This endpoint returns a JSON object representing a set of attributes about the subscriber, including but not limited to their name, email address, physical address, phone number, and other profile information. While the information inside the ID Token is reflective of the authentication event, the information in the UserInfo Endpoint is generally more stable and could be more general purpose. Access to different attributes from the UserInfo Endpoint is governed by the use of a specially-defined set of OAuth scopes, `openid`, `profile`, `email`, `phone`, and `address`. An additional scope, `offline_access`, is used to govern the issuance of refresh tokens, which allow the RP to access the UserInfo Endpoint when the subscriber is not present. Access to the UserInfo Endpoint is structured as an API and may be available when the subscriber is not present. Therefore, access to the UserInfo Endpoint is not sufficient for proving a subscriber's presence and establishing an authenticated session at the RP.

References

This section is informative.

General References

[EO13985] Executive Order 13985, *Advancing Racial Equity and Support for Underserved Communities Through the Federal Government*, January 25, 2021, available at: <https://www.federalregister.gov/d/2021-01753>.

[FEDRAMP] General Services Administration, *Federal Risk and Authorization Management Program*, available at: <https://www.fedramp.gov/>.

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

[NISTIR8112] NIST Internal Report 8112 (Draft), *Attribute Metadata*, available at: <https://pages.nist.gov/NISTIR-8112/NISTIR-8112.html>.

[Section508] Section 508 Law and Related Laws and Policies (January 30, 2017), available at: <https://www.section508.gov/manage/laws-and-policies/>.

Standards

[ISO/IEC9241-11] International Standards Organization, ISO/IEC 9241-11 Ergonomic of Human-System Interaction — Part 11: Usability: Definitions and Concepts. ISO, Geneva, Switzerland, 2018, available at: <https://www.iso.org/standard/63500.html>.

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

[OIDC-Basic] N. Sakimura, J. Bradley, M. Jones, B. de Medeiros, C. Mortimore, *OpenID Connect Basic Client Implementer's Guide 1.0*, April 18, 2022. Available at: https://openid.net/specs/openid-connect-basic-1_0.html

[OIDC-Implicit] N. Sakimura, J. Bradley, M. Jones, B. de Medeiros, and C. Mortimore, *OpenID Connect Implicit Client Implementer's Guide 1.0*, September 14, 2022. Available at: https://openid.net/specs/openid-connect-implicit-1_0.html

[RFC4120] IETF, *The Kerberos Network Authentication Service (V5)*, RFC 4120, DOI 10.17487/RFC4120, July 2005, <https://doi.org/10.17487/RFC4120>.

[RFC6113] IETF, *A Generalized Framework for Kerberos Pre-Authentication*, RFC 6113, DOI 10.17487/RFC6113, April 2011, <https://doi.org/10.17487/RFC6113>.

[RFC7591] IETF, *OAuth 2.0 Dynamic Client Registration Protocol*, RFC 7591, DOI

2269 10.17487/RFC7591, July 2015, <https://doi.org/10.17487/RFC7591>.

2270 **[RFC7636]** IETF, *Proof Key For Code Exchange*, RFC 7636, DOI 10.17487/RFC7636,
2271 September 2015, <https://doi.org/10.17487/RFC7636>.

2272 **[SAML]** OASIS, *Security Assertion Markup Language (SAML) V2.0 Technical Overview*,
2273 March 2008, available at: [https://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-](https://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-tech-overview-2.0.html)
2274 [tech-overview-2.0.html](https://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-tech-overview-2.0.html).

2275 **[SAML-WebSSO]** OASIS, *Profiles for the OASIS Security Assertion Markup Language*
2276 *(SAML) V2.0*, 15 March 2005. Available at: [https://docs.oasis-open.org/security/saml/v2.](https://docs.oasis-open.org/security/saml/v2.0/saml-profiles-2.0-os.pdf)
2277 [0/saml-profiles-2.0-os.pdf](https://docs.oasis-open.org/security/saml/v2.0/saml-profiles-2.0-os.pdf)

2278 **NIST Special Publications**

2279 NIST 800 Series Special Publications are available at: [https://csrc.nist.gov/publications/](https://csrc.nist.gov/publications/sp800)
2280 [sp800](https://csrc.nist.gov/publications/sp800). The following publications may be of particular interest to those implementing
2281 systems of applications requiring digital authentication.

2282 **[SP800-53]** NIST Special Publication 800-53 Revision 4, *Recommended Security*
2283 *and Privacy Controls for Federal Information Systems and Organizations*, April 2013
2284 (updated January 22, 2015), <https://dx.doi.org/10.6028/NIST.SP.800-53r4>.

2285 **[SP800-63]** NIST Special Publication 800-63-4, *Digital Identity Guidelines*, November
2286 2022, <https://doi.org/10.6028/NIST.SP.800-63-4.ipd>.

2287 **[SP800-63A]** NIST Special Publication 800-63B-4, *Digital Identity Guidelines:*
2288 *Enrollment and Identity Proofing*, November 2022, [https://doi.org/10.6028/NIST.SP.800-](https://doi.org/10.6028/NIST.SP.800-63a-4.ipd)
2289 [63a-4.ipd](https://doi.org/10.6028/NIST.SP.800-63a-4.ipd).

2290 **[SP800-63B]** NIST Special Publication 800-63B-4, *Digital Identity Guidelines:*
2291 *Authentication and Lifecycle Management*, November 2022, [https://doi.org/10.6028/](https://doi.org/10.6028/NIST.SP.800-63b-4.ipd)
2292 [NIST.SP.800-63b-4.ipd](https://doi.org/10.6028/NIST.SP.800-63b-4.ipd).

2293 **Federal Information Processing Standards**

2294 **[FIPS140]** Federal Information Processing Standard Publication 140-3, *Security*
2295 *Requirements for Cryptographic Modules*, March 22, 2019, [https://doi.org/10.6028/NIST.](https://doi.org/10.6028/NIST.FIPS.140-3)
2296 [FIPS.140-3](https://doi.org/10.6028/NIST.FIPS.140-3).

2297 **Appendix A. Changelog**

2298 *This appendix is informative.* It provides an overview of the changes to SP 800-63C since
2299 its initial release.

- 2300 • Added discussion of equity considerations and requirements.
- 2301 • Established trust agreements and registration as discrete steps in the federation
2302 process.
- 2303 • All FALs have requirements around establishment of trust agreements and
2304 registration.
- 2305 • FAL definitions no longer have encryption requirements; encryption is triggered by
2306 passing PII in an assertion through an untrusted party regardless of FAL.
- 2307 • FAL2 requires injection protection.
- 2308 • FAL3 allows more general bound authenticators including RP-managed
2309 authenticators, in addition to classical holder-of-key.
- 2310 • Communication of IAL/AAL/FAL required.
- 2311 • Updated language to be more inclusive.
- 2312 • Added definition and discussion of RP subscriber accounts.
- 2313 • Added attribute provisioning models and discussion.