BUILDING BLOCK WHITE PAPER

PRIVACY-ENHANCED IDENTITY FEDERATION

Paul Grassi
Naomi Lefkovitz
Applied Cybersecurity Division

Kevin Mangold
Information Access Division

Information Technology Laboratory
National Institute of Standards and Technology

December 2016
petid-nccoe@nist.gov

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

NCCoE NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

The National Cybersecurity Center of Excellence (NCCoE) at the National Institute of Standards and Technology (NIST) works with industry, academic, and government experts to find practical solutions for businesses' most pressing cybersecurity needs. The NCCoE demonstrates how standards and best practices established by NIST and other organizations can be applied in technical reference architectures and serves as a collaboration hub where small businesses, market-leading companies, government agencies, and individuals from academia work together to address broad cybersecurity problems. To learn more about the NCCoE, visit https://nccoe.nist.gov. To learn more about NIST, visit http://www.nist.gov.

NCCoE building blocks address technology gaps that affect multiple industry sectors. They represent core capabilities that can and should be applied across industry cybersecurity and business use cases.

ABSTRACT

A *relying party (RP)* that accepts credentials from a *credential service provider (CSP)* (often referred to as an Identity Provider or IdP) to login to their website achieves a number of benefits for their users and for themselves. An RP does not need to directly manage credentials when utilizing a trusted third-party, allowing them to focus their efforts and assets (both financial and human) on their core business, while lowering costs associated with conducting identity proofing and authentication on their own. Users can utilize a credential of their choice at many sites, reducing the friction associated with unique logins for every website with which they interact. However, as an RP decides to accept credentials from a new CSP, a separate integration effort is required to establish the connection. As a result, while many possible architectures exist, the market has responded and a dominant entrant has emerged to facilitate the reuse of credentials between CSPs and RPs. Commonly referred to as an "identity broker," these entities resolve the repetitive cost an RP repeatedly endures when adding new credential choices to their customers.

An *identity broker* can provide business value to both RPs and CSPs since each RP and CSP only needs to integrate with the identity broker once. The value to the RP is quite simple—connect once (to the identity broker) and accept many types of credentials. Yet the identity broker, or any centralized architecture, may raise risks to individual privacy; such solutions, if deployed incorrectly, are in a significant position of power, as they create the potential to track or profile an individual's transactions. In addition, possible outcomes could include the identity broker gaining insight into user data it does not need in order to perform the operations desired by CSPs and RPs.

Privacy-enhancing technologies (PETs) are tools, applications, or automated mechanisms which—when built into software or hardware—reduce or eliminate adverse effects on individuals when their personal information is being collected and/or processed. PETs implemented by federated identity solutions can reduce the risk of superfluous exposure

of individuals' information to participant organizations that have no operational need for the information, as well as shrink the attack surface for unauthorized access.

This document describes the technical challenges unique to integrating PETs within identity federations. It suggests scenarios suited for exploring the tradeoffs of mitigating or accepting specific privacy risks. Ultimately, this project will result in a publicly available NIST Cybersecurity Practice Guide—a description of the practical steps needed to implement an example solution that addresses existing challenges in the current federated identity marketplace.

KEYWORDS

brokered identity management; digital identity; identity federation; identity management; privacy-enhancing technology

DISCLAIMER

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification does not represent an exhaustive list of commercially available technologies, is not intended to imply recommendation or endorsement by NIST, NSTIC, or NCCoE, nor is it intended to imply that the entities, materials or equipment are necessarily the best available option in the market.

COMMENTS ON NCCOE DOCUMENTS

Organizations are encouraged to review all draft publications during public comment periods and provide feedback. All publications from NIST's National Cybersecurity Center of Excellence are available at http://nccoe.nist.gov.

Comments on this publication may be submitted to: petid-nccoe@nist.gov

ACKNOWLEDGEMENTS

This work is made possible by the support of the NIST National Strategy for Trusted Identities in Cyberspace (NSTIC) National Program Office (NPO), the National Cybersecurity Center of Excellence (NCCoE), and the NIST Information Access Division (IAD).

CONTRIBUTORS

The authors gratefully acknowledge the contributions of:

- Ross J. Micheals (formerly from NIST IAD)
- William Fisher (NIST, National Cybersecurity Center of Excellence)
- Kristin Greene (NIST, Information Access Division)
- Sean Brooks (NIST, National Strategy for Trusted Identities in Cyberspace National Program Office)

IN MEMORIAM

This project and its related efforts are dedicated to, with the fondest memories of, Dr. Ross J. Micheals. This problem was unsolvable until it was brought to Ross's attention, who proved the problem could be solved using existing standards and technologies. Because of his efforts, we were able to expand on the effort and bring the work to the NCCoE to demonstrate a practical application. Ross had a passion for "digging in the weeds," developing software, learning the ins and outs of the latest technology, and never leaving any problem—no matter how obscure—unsolved. He loved sharing his findings with others. His expertise, guidance, and leadership made this project possible, and we hope our continuation of his work illustrates how impactful his work and thoughts have been.

TABLE OF CONTENTS

1. Executive Summary

A *relying party (RP)* that accepts credentials from a *credential service provider (CSP)* (often referred to as an Identity Provider or IdP) to login to their website achieves a number of benefits for their users and for themselves. An RP does not need to directly manage credentials when utilizing a trusted third-party, allowing them to focus their efforts and assets (both financial and human) on their core business, while lowering costs associated with conducting identity proofing and authentication on their own. Users can utilize a credential of their choice at many sites, reducing the friction associated with unique logins for every website with which they interact. However, as an RP decides to accept credentials from a new CSP, a separate integration effort is required to establish the connection.

Identity Federation in Action

Connect.Gov is a federal government solution that allows citizens to use the third party credential of their choice to interact with agency services. This approach simplifies agency and CSP integration and improves user privacy by eliminating the ability of CSPs to track user behavior. Any solution identified by this white paper could be applied to Connect.Gov.

As a result, while many possible architectures exist, the market has responded and a dominant entrant has emerged to facilitate the reuse of credentials between CSPs and RPs. Commonly referred to as an "identity broker," these entities resolve the repetitive cost an RP repeatedly endures when adding new credential choices to their customers.

An *identity broker* can provide business value to both RPs and CSPs since each RP and CSP only needs to integrate with the identity broker once. The value to the RP is quite simple—connect once (to the identity broker) and accept many types of credentials. Yet the identity broker, or any centralized architecture, may raise risks to individual privacy; such solutions, if deployed incorrectly, are in a significant position of power, as they create the potential to track or profile an individual's transactions. In addition, possible outcomes could include the identity broker gaining insight into user data it does not need in order to perform the operations desired by CSPs and RPs.

Privacy-enhancing technologies (PETs) is a general term for a set of tools, applications or automated mechanisms which—when built into hardware or software—reduce or eliminate adverse effects on individuals when their personal information is being collected and/or processed. PETs implemented by identity brokers can reduce the risk of superfluous exposure of individuals' information to participant organizations that have no operational need for the information, as well as reduce vulnerabilities that could lead to unauthorized access.

This document describes the technical challenges unique to integrating PETs within identity federations. It suggests scenarios suited for exploring the tradeoffs of mitigating or accepting specific privacy risks. Ultimately, this project will result in a publicly available NIST Cybersecurity Practice Guide—a description of the practical steps needed to implement an example solution that addresses existing challenges in the current federated identity marketplace. NCCoE specifically seeks information technology and

cybersecurity product vendors, and open standards developers, as collaborators on the efforts to create a privacy-enhanced identity broker reference design and practice guide.

2. Business Value

As the National Strategy for Trusted Identities in Cyberspace (NSTIC), also referred to as Strategy stated,

> A secure cyberspace is critical to our prosperity. We use the Internet and other online environments to increase our productivity, as a platform for innovation, and as a venue in which to create new businesses 'Our digital infrastructure, therefore, is a strategic national asset, and protecting it—while safeguarding privacy and civil liberties—is a national security priority' and an economic necessity. By addressing threats in this environment, we will help individuals protect themselves in cyberspace and enable both the private sector and government to offer more services online. [1]

The NSTIC envisioned an identity ecosystem of federated identity solutions playing a key role in achieving a more secure cyberspace. Federated identity solutions, in which RPs accept third-party credentials from a CSP to login to their website, can provide a number of benefits. They minimize the number of digital credentials individuals need to access RP services, which can make it more convenient for individuals to use fewer, stronger credential options, such as multi-factor authentication. An RP that uses third-party credentials does not need to directly manage them, allowing them to focus on their core business and lower costs because CSPs will manage the identity proofing and authentication (and spread those costs across multiple RPs). CSPs can focus on offering secure and efficient identity proofing processes to strengthen trust in identities for higher assurance transactions across the internet.

However, each pairing of an RP with a CSP requires a separate integration effort. An identity broker, commonly used to solve these integration challenges, can provide business value to both RPs and CSPs since each RP and CSP only needs to integrate with the identity broker once. The identity broker also can provide mechanisms to apply technical and policy interoperability between RPs and CSPs.

Nevertheless, federated identity solutions raise new risks for the privacy of individuals and confidentiality of business information. The interoperability that provides the benefits described above can also create the potential for more tracking and profiling of individuals' transactions. The same interoperability can expose businesses, as the relationships between RPs and CSPs reveal who their customers are to each other; such exposure may be particularly problematic if the federation occurs within the same industry sector. In addition, the identity broker can become an appealing target to gain access to identity attributes being transmitted through the broker or to RP accounts.

[1] https://www.whitehouse.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf

Thus, participants in federated identity solutions—whether individuals or organizations—must be able to trust that the solutions are not going to reveal sensitive information, or they will not participate in identity federations.

PETs implemented in federated identity solutions can reduce the risk of superfluous exposure of individuals' information to participant organizations that have no operational need for the information, as well as shrink the attack surface for unauthorized access. Implementing such PETs will enable market differentiation for the adopters and increase trust in federation. Additionally, organizations may be subject to various privacy and security requirements under law or through trust frameworks. PETs can assist in demonstrating compliance, potentially with reduced costs over policy-based controls, with relevant privacy and security requirements.

Market demand within the private sector is not the only domain where business value can be attained. Governments also use federated identity services—and need to minimize the risk of privacy and civil liberties violations (or the international equivalent). A number of current solutions manage these risks via avoidance; they intentionally stay away from the transmission of attributes due to the privacy risks of unintentional disclosure. PETs can enable governments to derive the benefits of federated identity while minimizing potential violations of privacy and civil liberties that harm individuals and contribute to an overall breakdown in public trust.

3. Description

Purpose of the document

This document describes the specific privacy and cybersecurity goals unique to identity federations. To overcome these challenges, commercial software and open standards will be examined and utilized to establish a set of privacy enhancing technologies that can be applied to a variety of identity architectures. Not all of the goals of this paper may be desirable in certain communities, use cases, or by individuals. However, new privacy capabilities not available in today's federated solutions are critical to advancing the identity ecosystem, offering new tools to solution providers to build privacy into their technology, and allowing individuals to understand the privacy options available to them via technology, but not currently used in the market.

Audience

The intended audience of this document includes anyone with experience in identity management, privacy-enhancing technologies, cryptography, and their integration in solving real-world problems.

The NCCoE specifically seeks information technology and cybersecurity product vendors, and open standards developers, as collaborators on the efforts to create a privacy-enhanced identity broker reference design and practice guide.

The NCCoE will publish a Federal Register Notice (FRN) inviting vendors interested in collaborating on this effort.

Goals

The primary goal of this building block is to demonstrate how federated identity services, leveraging open, market dominant standards (or adopted profiles/extensions), can protect the attributes of a logged in user such that any entity that orchestrates or supports a federated transaction, honest or malicious, can never gain access to attribute information—while retaining an architecture in which RPs and CSPs do not know each other's organizational identities—i.e., *double-blind*. It is required that any approach utilized to achieve this goal can mitigate a broker-based man-in-the-middle attack. Specific goals are as follows:

Goal 1. **RP/CSP untraceability and unlinkability**. The federation prevents RPs and CSPs from learning each other's identities. Neither entity can track or link user activities beyond what is known from their direct relationship with the user.

Goal 2. **Participants in the federation, other than those the user approves, cannot access user attributes.** RPs obtain validated attributes (and sometimes self-asserted attributes) from authoritative CSPs. Users first consent to sharing the attribute from the CSP to the RP. Once the RP has the actual attribute value, they can use the information to fulfill their service requirements. A solution is required that allows for disclosure of an attribute value to the intended RP. In doing so, the double-blind must be retained; so utilizing an approach that "leaks" organizational identity (e.g., a public key), is not sufficient. In addition, any approach utilized must resist the threat of any valid intermediary accessing attribute values (e.g., man-in-the-middle attack).

Goal 3. **A compromised or malicious federation participant cannot impersonate a user.** A compromised entity (one that has been hacked or that becomes malicious of its own volition) might be able to satisfy the desired privacy enhancements, yet still be able to impersonate an end user. Controls must be established to reduce this threat.

Goal 4. **Minimization of user attributes.** Attributes are only provided when an RP requests them, not every time a user logs in to access RP services. In addition, RPs will request the minimum attributes to satisfy the transactions/services the user is accepting. The RP does not collect **all** attributes that may be needed based on service offering, just those required by the services a user is actually requesting. While this reduces the potential of exposing personal information, it alone does not alleviate the need to accomplish the first three privacy goals above.

Goal 5. **Users must explicitly consent** to disclosure of their attributes to an RP.

Goal 6. **Pseudonym unlinkability.** Entities that mediate identity transactions cannot track or link user pseudonyms across transactions, i.e., the mediation of multiple

transactions does not convey any information that could be used to infer transactions associated with a single user.

The above set of goals is intended to represent a comprehensive set of privacy goals that federated identity solutions may achieve based on requirements and demand. However, not all may be achievable due to the lack of industry standard, market-viable solutions. The solution developed under the auspices of this white paper will not be based on theoretical or academic solutions. Only those goals that can be achieved with market-available solutions will be considered.

Background

The economic and security benefits of strong authentication, increased demand and availability of reusable credentials, and the complexity of managing identities and accounts have resulted in an increase in online RPs that are willing to outsource authentication to trusted CSPs. The cost to manage credentials and comply with regulations associated with the collection and storage of identity data, the risk of users bailing out of the registration process, and the interoperability complexities associated with supporting multiple identity protocols are examples of business drivers to adopt identity federation.

Organizations that participate in a federation interoperate within a formal technical and policy trust framework. RPs realize savings and reduce complexity by shifting architectures, as illustrated in Figure 1. On the left, the RP establishes business, technological, and interoperability trust relationships with each CSP. On the right, the relationship is simplified with a single "broker."

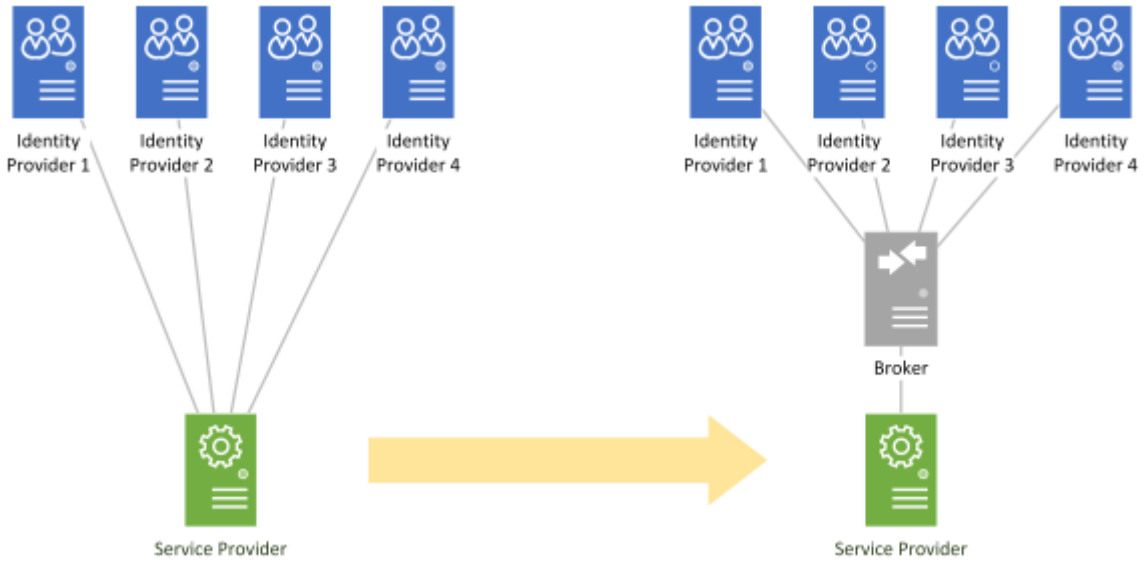

Figure 1. An RP migrates to a brokered identity management model. Instead of integrating with each CSP individually, it interfaces with a single broker.

In the context of this building block, a federated identity management solution serves the following essential functions:

1. alleviates the number, or complexity, of integrations required between RPs and CSPs
2. allows for protocol translation, reducing the number of protocols RPs and CSPs need to support
3. enables the privacy principles of untraceability and unlinkability by "blinding" the CSPs and RPs from each other

Unfortunately, despite the aforementioned benefits afforded by employing broker-like architectures, many protocols require explicit trust relationships. For example, Security Assertion Markup Language (SAML) metadata needs to be exchanged at design time, which typically includes public cryptographic keys to sign and encrypt messages (or portions of the message) as users authenticate to a CSP and access an RPs website. However, other protocols such as OpenID Connect allow for runtime discovery of CSPs and their associated public keys.

Consequently, it is necessary to employ additional security and privacy controls in collaboration with RPs and CSPs to ensure that as federated identity transactions are executed, the privacy principles expected by users are met. In complying with existing protocols, there is a risk that service providers will be in a position of power that erodes the security and privacy practices that are crucial to long-term market adoption.

Therefore, federated identity providers have unique privacy and cybersecurity challenges that must be overcome. In many identity management protocols, it is assumed that there is an explicit relationship, and direct connection, between the RP and the CSP. Many commonly used identity management protocols, such as SAML version 2.0 or OpenID Connect, were not specifically designed with unlinkability in mind. That is, as illustrated in Figure 2, a direct "trust" relationship is commonly established, a priori, to allow RPs and CSPs to directly communicate.

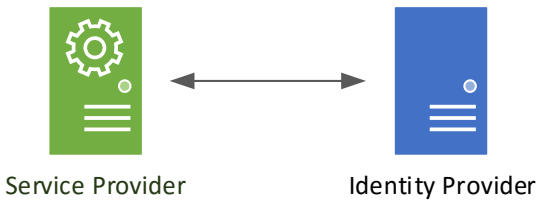

Figure 2. In many identity management protocols, there is a direct trust and communications relationship between an RP and a CSP.

"Weak unlinkability" through the use of pair-wise pseudonymous identifiers, can be achieved with both SAML and OpenID Connect, however any direct connection via TCP-IP allows either entity to discover the true organizational identity of its partner. With the

constraints of modern identity protocols, for a plurality of federations, the protection of user credentials and attributes must be maintained through:

- **Implicit trust relationships**[2]: The identity ecosystem should support explicit trust relationships wherever required (within a community) or requested (by a user). However, in the case of 3rd parties supporting identity transactions, intermediaries should be able to maintain chain of trust between RP, CSP, and user. In this case, the RP mutually trusts the broker and the broker mutually trusts the CSPs; CSPs and RPs can then indirectly trust one another through the transitive trust relationship maintained by the broker.
- **Transport layer and message security**: The RP and CSP use transport layer and message security to assure the integrity and confidentiality of credentials, user attributes, and/or security assertions (the specifics of what is communicated depends on the protocol employed).
- **Operational policies**: A federation provider would implement a host of technical controls (not policy or manual procedures) to help ensure the secure exchange of messages.

Despite these protections, since identity management protocols are not explicitly "blinding," federation providers may have access to unencrypted security assertions and user attributes, and have the ability to link user transactions across RPs and CSPs.

As illustrated in Figure 3, if an identity protocol does not explicitly recognize the role (or entity) of a federation intermediary, then entities acting as brokers must act as a CSP to the actual RP, and an RP to the actual CSP. Any privacy enhancing technologies must be implemented in such a manner that they are compatible with this model.

[2] Trust can be managed in many different ways, such as accreditation and trust frameworks. This paper does not assume any specific model for asserting or conveying trust, and therefore still expects strong privacy controls to be built inherently into all federated identity transaction.

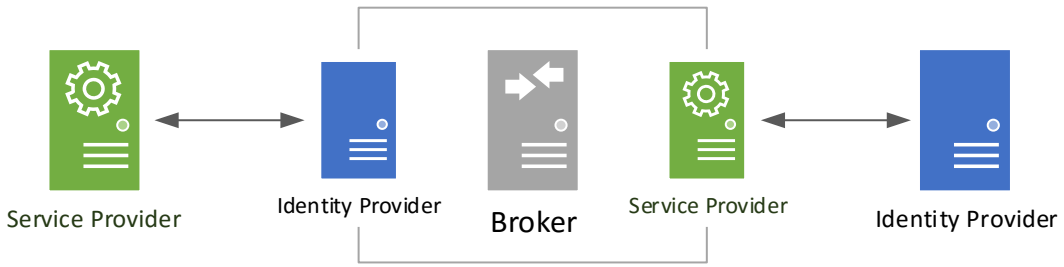

Figure 3. Identity Broker-Based Relationship Model.

Scope

This building block will demonstrate how federated identity architectures can use profiles and/or extensions of market dominant protocols, such as SAML and OpenID Connect, to implement the privacy enhancements discussed in the Goals section above. Identification of the challenges to implementing these privacy enhancements is an inherent part of the building block's scope; those enumerated in this document are only a starting point for a larger collaboration effort with the private sector. This effort will include the deployment of the infrastructure required to simulate a federated architecture, the use of multiple authenticators, and the inclusion of appropriate, publicly available and proven cryptographic algorithms.

With respect to cybersecurity, this particular building block focuses only on the challenges unique to the entities that facilitate federated identity architectures. How the attributes are protected at rest, and used by RPs and CSPs, is out of scope. Authorization, and any use of fine-grained access control, to include attribute-based access control (ABAC), is also not in scope at this time.

Assumptions

The following foundational assumptions have been made to achieve the goals stated in this white paper:

1. The technologies, algorithms, standards, and processes available in the market may need to be profiled or extended to satisfy the goals of this building block.
2. Components identified in this building block are relatively high-level. For simplicity, the white paper treats each RP, CSP, or identity broker as a standalone, single entity. In reality, however, each actor in a production system may itself be a system of systems—comprising other components. For example, behind the abstraction of a CSP could be security token services, identity stores, and/or multifactor authentication technologies. Entities are scoped so that the building block can concentrate specifically on those challenges unique to enhancing privacy.

3. Authentication can be separated from attribute delivery, effectively creating an environment where a valid 3rd party intermediary learns nothing about an identity, with the exception of a set of transactions assigned to a pseudonymous identifier. This does not obviate the need to identify methods to protect user attributes and the organizational identities of participants in a transaction. Therefore, this paper does not consider separating identity from attributes as a solution to existing challenges.
4. The goal of this building block is to consider how to augment existing, market dominant protocols; it is *not* to develop or research new protocols. However, we recognize that changes to existing protocols and profiles may be necessary to fulfill the building block's privacy enhancement requirements. In addition, the solution must use existing commercial-off-the-shelf technology.

4. Scenarios

Federated Logon Overview and Example

In a federated logon, an RP trusts the identity assertions issued by a CSP to allow users to access their system. Federated sign-on is not a new concept; in fact, many popular websites allow users to access their services using third party credentials, such as e-mail or social networking accounts.

Consider the following example of a real-world implementation of federated logon:

1. Alice wishes to access the National Institutes of Health publication database, *PubMed.* Alice browses to the PubMed website and is presented with the screen shown in Figure 4.

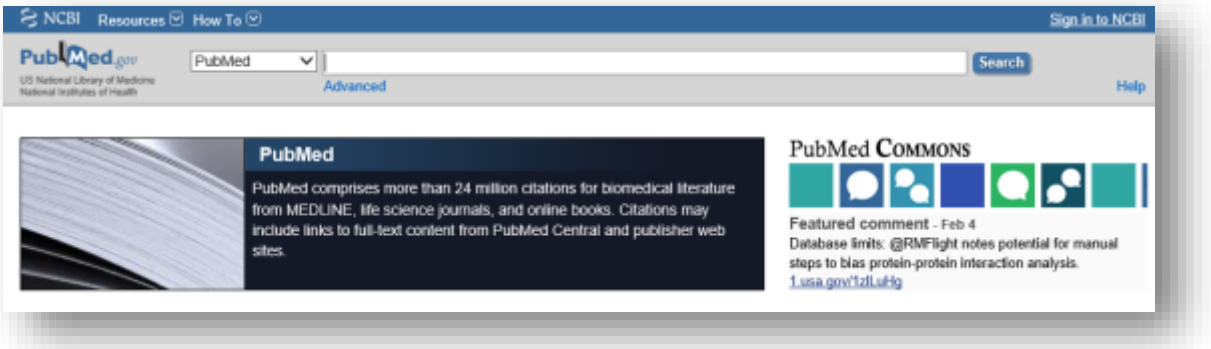

Figure 4. PubMed landing page. Note the "Sign in to NCBI" link in the upper right corner.

2. She clicks **Sign in to NCBI** and sees the web page shown in Figure 5.

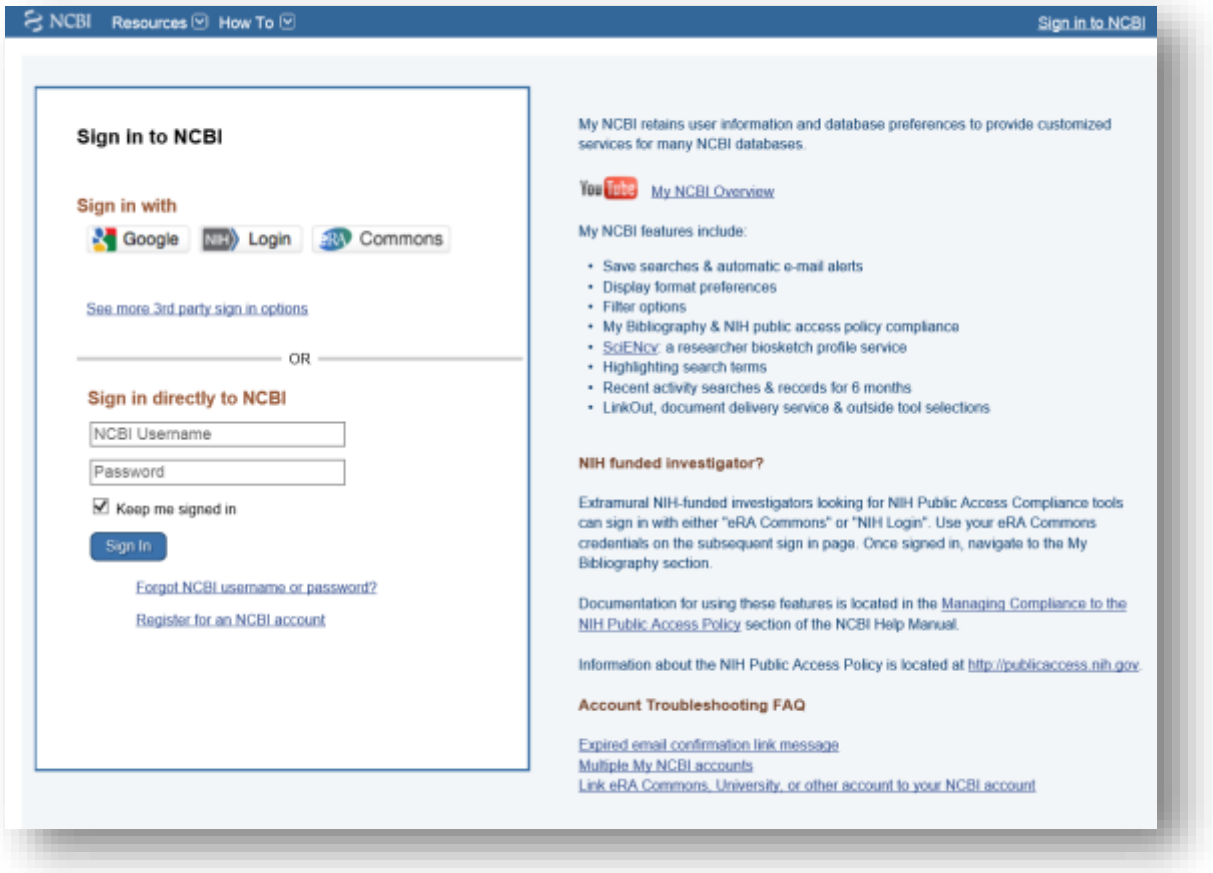

Figure 5. PubMed sign-on page. Users can logon with a direct username and password or use a "third-party option."

3. Alice has the ability to choose a PubMed username and password to logon. She has the option to sign in with a PubMed account **and** a variety of third-party credentials. At the time of writing this document, PubMed allowed for logon with over 90 third-party CSPs.

The following scenarios establish incremental capabilities to achieve the goals of this white paper. Also starting with this baseline scenario, it is important that the system design maintains a flexibility for 'a variety of transactions', as asked by NSTIC, including 'anonymous with validated attributes' and 'pseudonymous without attributes' transactions. Thus, to exemplify the former, an implemented system should enable, without need for structural changes, supporting transactions where attributes are validated (e.g., an age range), but neither the broker nor the RP receive from the CSP (nor are in a position to infer from the elements of the transaction) any (persistent) user pseudonym that would remain persistent beyond said transaction.

Also, the system should support transactions where the RP and the broker are able to signal that no attribute or attribute validation is required for a transaction.

Scenario 1. Baseline: Authentication and Attribute Delivery Given an Identity Broker

In the first scenario, the building block will demonstrate user authentication and attribute delivery, as illustrated in PubMed walkthrough, to achieve the previously specified Goal 1 (untraceability and unlinkability).

In the example, the RP, PubMed, was responsible for implementing and maintaining the technology and policy relationships with their third-party CSPs (the left side of Figure 1). In the baseline scenario, we replace these relationships with a single integration with the broker (the right-hand side of Figure 1). This baseline scenario is intended to capture the essence of the migration from dedicated, multiple CSP connections, to a concept of operations based on an "outsourced," brokered CSP integration concept of operations.

The baseline scenario is a required step to simulate an identity broker along with a set of RPs and CSPs, satisfying the initial double-blinding. The goal of this scenario would be to mimic, as much as possible, a system that closely matches the technical control typically in place today—that is, no additional attribute or credential protection other than what is afforded by the native protocols and policies.

In summary, the first scenario is establishing what currently exists in the market—*RP acceptance of a CSPs credentials via an identity broker*.

Scenario 2. Authentication and Confidential Attribute Delivery Given an Honest-But-Curious Broker

In Scenario 2, Goal 1, Goal 4, and Goal 5 are achieved. Any entity that is supporting federation between the CSP and RP to achieve these privacy goals is assumed to be an *honest but curious (HBC)* adversary. The "honest but curious" adversary model means that the target protocol is implemented correctly (the entity is honest), but might look at the information passing through it in an attempt to learn information (it is curious). This is analogous to a situation in which an attacker has gained access to a system and can read information passing through it, but cannot change that information.

To achieve these characteristics, building block participants will need to identify threats unique to this scenario, as well as design specialized mitigations to eliminate or reduce the potential risk of these threats. Threat identification, mitigation, and technological cost/benefit analyses will be among the core building block collaboration activities.

Scenario 3. Authentication and Confidential Attribute Delivery Given a Malicious Identity Broker

In Scenario 3, additional controls are applied to Scenario 2 to achieve Goal 3. In this scenario, however, we assume that federation participants other than the CSP and RP might be compromised. A malicious broker is one example that could actively seek to exploit architectural or security vulnerabilities in order to disrupt the overall system's

ability to maintain confidentiality, information integrity, or system availability. This is analogous to a situation in which an attacker has gained access to the broker and can covertly inject their own behaviors. Protection in the face of a malicious broker, particularly one that exfiltrates sensitive information silently, is a significant cybersecurity challenge.

Scenario 3 will focus on preventing a malicious broker that:

1. initiates its own authorization or attribute query request without permission from a user or RP
2. "phishes" an end user's credentials by pretending to be a CSP
3. impersonates the end user by replaying identity assertions
4. attempts to perform a man-in-the-middle attack to obtain the encryption keys that would enable decryption of user attributes intended for the RP, and/or obtain user pseudonyms intended solely for the RP or the CSP

Like Scenario 2, a core building block activity will be to identify additional threats, mitigations, and their technological cost/benefit.

Scenario 4. Authentication Pseudonym Unlinkability

This scenario seeks innovative ways, if market-possible, to prevent the broker from observing persistent user pseudonyms in transit between the CSP and the RP, including preventing federation participants from inferring any user pseudonym linkable to other transactions, while also ensuring that user pseudonyms are transformed appropriately for each RP and retaining an architecture in which RPs and CSPs do not know each other's organizational identities.

Summary

Table 1 provides a summary of the scenarios. A checkmark indicates that the scenario includes the corresponding requirement.

Requirement	Scenario			
	1	2	3	4
Federated authentication and attribute delivery via an identity broker	✓	✓	✓	
Scenario implements the desired security characteristics		✓	✓	
Identity Broker is an "honest but curious" adversary		✓		
Identity Broker is a "malicious" adversary			✓	
Identify unique threats, mitigations, and cost/benefit tradeoffs		✓	✓	
Prevent pseudonym linkability				✓

Table 1. Summary of Scenarios. A checkmark indicates that the scenario fulfills the corresponding requirement.

5. Current Building Block Challenges

RPs wish to accept third-party credentials so that (a) they themselves do not have to manage user credentials, and (b) they reduce the abandonment rate due to requiring users to create another account they may not want (unfortunately, often a username and password). An identity broker can provide business value to an RP (and CSPs alike) by specializing in integration, policy harmonization, service, and CSP "matchmaking."

The NSTIC envisions an Identity Ecosystem that, "will provide multi-faceted privacy protections," that are built into the technologies that provide authentication and federation services. The strategy specifically advocates the use of privacy-enhancing technical standards that, "minimize the transmission of unnecessary information and eliminate the superfluous 'leakage' of information that can be invisibly collected by third parties. Such standards will also minimize the ability to link credential use among multiple RPs, thereby preventing them from developing a complete picture of an individual's activities online."

Typical identity federations that leverage broker architectures to achieve some privacy objectives have conflicting requirements under this viewpoint. On one hand, the broker needs information about all of the entities involved in a particular transaction so that it can help guarantee the integrity and confidentiality of the transaction, as well as the information that is contained within the transaction. Yet, the Strategy also advocates unlinkability—individual behavior should not be observable among the participants of a trust framework or federation. In addition, RP's may want to know the actual organization asserting attributes in order to make appropriate risk-based decisions. For example, if an RP needs sensitive data such as medical or financial information, but the ecosystem of CSPs expand beyond these specific communities, how can the RP trust that this information came from legitimate CSPs capable of asserting medical or financial information? A trust framework can mitigate this, but RPs may want additional runtime information about a CSP's authority to assert specialized attribute data. Can such requirements be managed in a privacy-preserving manner?

Another challenge that is apparent in multi-party identity federated architectures, but compounded with the inclusion of built-in privacy controls, is in the realm of auditing. There needs to be a mechanism to "put the pieces back together again" in the case that something goes wrong. Reconstructing a single or set of transactions should be possible in order to provide forensic information about the context of a transaction, but should not be possible by a single entity. In addition, as the participants in a federation reconstruct transaction information, it must be done in the way that does not expose the data of persons or organizations that are not subject to forensic review. The architecture should allow for the reconstruction of specific events without causing all events and the information that composed the transaction to be revealed.

6. Desired Solution Objectives

Below is a list of target characteristics based on the Goals section. The omission of any security or privacy engineering objective from the complete set is not an indication that the federated identity architecture may not have characteristics of the omitted objective. Any information system needs to maintain all of the objectives to some degree, but this building block is designed to demonstrate capabilities for the specific objectives listed below. In addition, any solution objective must be met with market-available solutions.

Functional Objectives

Table 2 - Functional Objectives

Functional Objective	Example Capabilities
Identity federation	• users can choose from a number of CSPs • CSPs can be dynamically discovered, while maintaining privacy goals
Protocol translation	• federation providers can transform an input protocol to a different output protocol, and vice versa • encrypted and signed data in one protocol can be migrated, transformed, or converted to another protocol without access to plaintext and without breaking the chain of trust of originator of message
Reduce risk of user profiling	• CSP does not have knowledge of RP identity • RP does not have knowledge of CSP identity • any other participant in the federation does not know identity of user conducting transaction
Reduce risk of transactional profiling	Across any two transactions mediated by a 3rd party: • a change of RP is hidden from the CSP (i.e., for two transactions involving the same CSP, the CSP does not learn whether the RP is the same or has changed) • a change of CSP is hidden from the RP (i.e., for two transactions involving the same RP, the RP does not learn whether the CSP is the same or has changed) • no entity can link users across transactions (i.e., a broker cannot tell whether users in any two transactions are the same or different, regardless of the involved RPs and CSPs)

Security Objectives

Table 3 - Security Objectives

Security Objective	Example Capabilities
Confidentiality	• 3rd parties do not have plaintext access to user credentials or attributes either at rest, or in motion • 3rd parties will never have access to decryption keys • a malicious man-in-the-middle attack will not result in a breach of personal data of the authenticated user • unauthorized access to transactional data, even encrypted, is not possible
Integrity	• RP is assured that the data has not been modified by the hub or a malicious 3rd party • RP is assured that the data is provided by a valid CSP • RP is assured that a malicious 3rd party cannot impersonate a valid user and/or reuse prior, valid assertions

Privacy Engineering Objectives

NIST has developed three draft privacy engineering objectives for the purpose of facilitating the development and operation of privacy-preserving information systems: predictability, manageability, and disassociability. These objectives are designed to enable system designers and engineers to build information systems that are capable of achieving their functional purpose while implementing an organization's privacy goals and supporting the management of privacy risk. As with the above security objectives, these privacy objectives are core characteristics of information systems.

- ***Predictability*** is the enabling of reliable assumptions by individuals, owners, and operators about personal information and its processing by an information system.
- ***Manageability*** is providing the capability for granular administration of personal information including alteration, deletion, and selective disclosure.
- ***Disassociability*** is enabling the processing of personal information or events without association to individuals or devices beyond the operational requirements of the system.

Table 4 - Privacy Objectives

Privacy Engineering Objective	Example Capabilities

Predictability	• enables user, RP, CSP, and identity broker to assume that the identity broker does not have access to user identity attributes • enables user, RP, CSP, and identity broker to assume that the CSP cannot process information about the user's relationship with the RP • enables user, RP, CSP, and identity broker to assume that the RP cannot process information about the user's relationship with the CSP
Manageability	• only the user can choose to disclose their attribute information to an RP • a user can see their attribute values at a CSP prior to release to an RP, and have a mechanism to dispute inaccuracies prior to release • a user can selectively choose which relationships can be directly associated, e.g., a user can state that the CSP^{1} and RP^{A} may communicate directly with each other and forgo any intermediary
disassociability	• the identity broker can transmit identity attributes from a CSP to an RP without being able to access them • the RP can accept an authentication assertion and identity attributes without associating a user to a CSP • the CSP can transmit an authentication assertion and identity attributes without associating a user to an RP

This is not an exhaustive list; it highlights those features that are particularly salient to the unique challenges to this domain. In addition, these characteristics will need to be balanced with the risk level. For example, it might be acceptable (e.g., for specific security or operational reasons) to allow an RP to know the identity of the CSP while still blocking broker access to plaintext user attributes. As stated previously, a goal of this building block is to understand the nature of these tradeoffs among the configuration space of various protections.

7. Relevant Standards, Specifications, and Guidance

- NIST Special Publication 800-63 Revision 2: Electronic Authentication Guideline
- Organization for the Advancement of Structured Information Standards (OASIS) Security Assertion Markup Language (SAML) v2.0 Standard
- OpenID Connect Core
- Draft NISTIR 8062 - Privacy Risk Management for Federal Information Systems

- OAuth 2.0 Specification
- Federal Information Processing Standards 140-2, Special Requirements for Cryptographic Modules
- JavaScript Object Signing and Encryption (JOSE)
- XML Encryption
- XML Signature

8. Security Control Mapping

This table maps the necessary objectives of the commercial products that the NCCoE will apply to this cybersecurity challenge to the applicable standards and best practices described in the Framework for Improving Critical Infrastructure Cybersecurity (CSF), and other NIST activities. This exercise is meant to demonstrate the real-world applicability of standards and best practices, but does not imply that products that meet these objectives will achieve a given industry's requirements for regulatory approval or accreditation.

Objectives Objective	Cybersecurity Standards and Best Practices						
	CSF Function	**CSF Category**	**CSF Subcategory**	**NIST 800-53-4**	**IEC/ISO27001**	**SANS/CSC**	**CSF CCMv3.0.1**
Identity federation	Protect	Access	PR.AC-1 PR.AC-5	IA-4 SC-23	A.9.4.2 A.13.1.1 A.13.2.3	16-2 16-15 17-7	IAM-09 AIS-01 AIS-02 EKM-03 STA-0
		Data Security	PR.DS-2				
		Protective Technologies	PR.PT-4				
Protocol translation	Protect	Access	PR.AC-5	AC-4 SC-8 SC-23 SI-10	A.13.1.1 A.13.2.3	6-2	AIS-01 AIS-02 AIS-03 AIS-04 DSI-01 DSI-03 EKM-03 EKM-04 STA-03
		Data Security	PR.DS-2				
		Protective Technologies	PR.PT-4				

Objectives Objective	Cybersecurity Standards and Best Practices						
	CSF Function	**CSF Category**	**CSF Subcategory**	**NIST 800-53-4**	**IEC/ISO27001**	**SANS/CSC**	**CSF CCMv3.0.1**
Confidentiality	Protect	Access	PR.AC-1 PR.AC-4	AC-3 AC-5 AC-6 SC-8 SC-13	A.9.2 A.9.4.1 A.10 A.13.1.2 A.13.2.3 A.14.1.2 A.14.1.3	12-1 15-1 15-4 17-2 17-3 17-7 17-9 17-10 17-12 17-13 17-15	AIS-01 DSI-03 EKM-02 EKM-03 EKM-04 IAM-05 IAM-09 IAM-12 IAM-13
		Data Security	PR.DS-2 PR.DS-5				
		Protective Technologies	PR.PT-4				

Objectives Objective	Cybersecurity Standards and Best Practices						
	CSF Function	**CSF Category**	**CSF Subcategory**	**NIST 800-53-4**	**IEC/ISO27001**	**SANS/CSC**	**CSF CCMv3.0.1**
Disassociability Triple Blinding	Protect	Data Security	PR.DS-2 PR.DS-5 PR.DS-6	AC-4 AC-8 AC-14 AC-23 CM-5 IA-4 SC-4 SC-8 SC-12 SC-13 SC-17 SC-26 SC-30 SI-16	A.10 A.12.2 A.12.6.1 A.13.1.2 A.13.2.3 A.14.1.2 A.14.1.3	5-6 15-1 15-4 17-2 17-3 17-7 17-9 17-10 17-12 17-13 17-15	AIS-01 AIS-04 DSI-01 DSI-02 DSI-03 EKM-02 EKM-03 EKM-04 IAM-06 IAM-09

Objectives Objective	Cybersecurity Standards and Best Practices						
	CSF Function	**CSF Category**	**CSF Subcategory**	**NIST 800-53-4**	**IEC/ISO27001**	**SANS/CSC**	**CSF CCMv3.0.1**
Predictability Integrity	Protect	Data Security	PR.DS-2	AC-8 AC-14 AC-23 IA-4 SA-13 SA-18 SC-7 SC-11 SC-13 SC-17 SI-4 SI-7 SI-12	A.10 A.13.1.2 A.13.2.3 A.14.1.2 A.14.1.3	17-2 17-3 17-7 17-9 17-10 17-12 17-13 17-15	AIS-01 AIS-03 DSI-02 DSI-03 DSI-04 IAM-05 IAM-09 EKM-02 EKM-03 EKM-04 IVS-01 IVS-06 IVS-09 IVS-12 TVM-01
		Information Protection Processes and Procedures	PR.IP-6				

9. High-Level Architecture

The following is a high-level diagram of the candidate building block architecture. This architecture captures the various actors at a *system of systems* level; each RP and CSP could comprise a variety of additional components.

Figure 6. High-Level Architecture

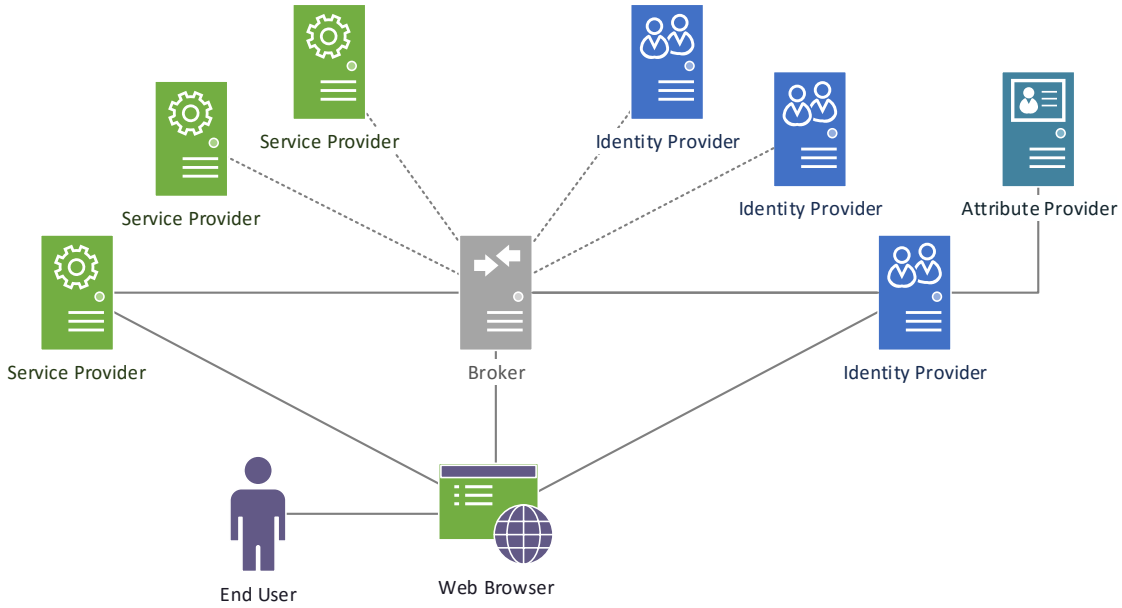

It is important to note that a single solution may not exist, and that innovation and collaboration within the private sector may identify solutions that require additional components and/or standards than those already identified.

10. Component List

The following list is an example of the components that might comprise a final building block solution. This list is only a starting point; specific components will be identified through future vendor collaborations.

- RP hosts (physical or virtual) and instances
- CSP hosts (physical or virtual) and instances
- Identity Federation Manager host(s) (physical or virtual) and instance(s)
- Attribute provider hosts (physical or virtual) and instance(s) (optional)
- User agent/host with web browser
- Multi-factor credentials
- Cryptographic Module(s), to include any necessary key management system(s).
- Network, computer, and storage infrastructure to support the above

APPENDIX A – ACRONYMS AND ABBREVIATIONS

The following are acronyms commonly used in the context of identity management and may be helpful for readers of this and related National Cybersecurity Center of Excellence materials.

ABAC	Attribute-Based Access Control
BB	Building Block
FICAM	Federal Identity, Credential, and Access Management
FR	Federal Register
HBC	Honest But Curious
Id *or* ID	Identity
CSP	Credential service provider
IETF	Internet Engineering Task Force
IT	Information Technology
LOA	Level of Assurance
MFA	Multi-factor Authentication
NCCoE	National Cybersecurity Center of Excellence
NIST	National Institute of Standards and Technology
NSTIC	National Strategy for Trusted Identities in Cyberspace
OASIS	Organization for the Advancement of Structured Information Standards
OMB	Office of Management and Budget
PET	Privacy-Enhancing Technologies
PKI	Public Key Infrastructure
RFC	Request for Comment
RP	Relying Party
SAML	Security Assertion Markup Language

APPENDIX B – GLOSSARY

This building block, where possible, leverages external authoritative sources of terms for identity, credential and access management. The table below outlines terms as they are used within the context of this building block.

Term	Definition	Source
access control	a process by which use of system resources is regulated according to a security policy and is permitted only by authorized entities (users, programs, processes, or other systems) according to that policy	Internet Engineering Task Force (IETF) Request for Comments (RFC) 4949
assertion	a statement from a verifier to a relying party that contains identity information about a subscriber. Assertions may also contain verified attributes. Assertions may be digitally signed objects or they may be obtained from a trusted source by a secure protocol	NIST Special Publication 800-63-2
assurance level	a measure of trust or confidence in an authentication mechanism in terms of four levels: Level 1 - little or no confidence; Level 2 - some confidence; Level 3 - high confidence; Level 4 - very high confidence	Office of Management and Budget (OMB) Memorandum M-04-04
attribute	a claim of a named quality or characteristic inherent in or ascribed to someone or something	NIST Special Publication 800-63-2
attribute based access control (ABAC)	a policy-based access control solution that uses attributes assigned to subjects, resources or the environment to enable access to resources and controlled information sharing	Authorization and Attribute Services Committee Glossary
authentication	the process of establishing confidence in the identity of users or information systems	NIST Special Publication 800-63-2

credential	an object or data structure that authoritatively binds an identity (and optionally, additional attributes) to a token possessed and controlled by a subscriber	NIST Special Publication 800-63-2
federation	a trust relationship between discrete digital credential service providers (CSPs) that enables a relying party to accept credentials from an external credential service provider in order to make access control decisions; provides path discovery and secure access to the credentials needed for authentication; federated services typically perform security operations at run-time using valid NPE credentials	Federal Identity, Credential, and Access Management (FICAM)
identity	a set of attributes that uniquely describe an entity within a given context	Modified from NIST Special Publication 800-63-2
Multi-factor authentication	Combining two or more authentication factors to logon to an authentication system. Allowable factors include "something you know," "something you have," and "something you are."	
credential service provider (CSP)	a trusted entity that issues or registers subscriber tokens and generates subscriber credentials	Modified from NIST Special Publication 800-63-2
password	a secret that a claimant memorizes and uses to authenticate his or her identity	NIST Special Publication 800-63-2
privacy-enhancing technologies	a set of tools, applications, or mechanisms which—when integrated in information systems—enables the mitigation of risks of adverse effects on individuals from the processing of their personal information within the information systems.	NIST

public key infrastructure	a set of policies, processes, server platforms, software and workstations used for the purpose of administering certificates and public-private key pairs, including the ability to issue, maintain, and revoke public key certificates	NIST Special Publication 800-63-2
Relying Party (RP)	an entity that relies upon the subscriber's token and credentials or a verifier's assertion of a claimant's identity, typically to process a transaction or grant access to information	NIST Special Publication 800-63-2
Unlinkable	assures that two or more authentication or attribute assertion transactions cannot be determined to be related to the same individual by the CSP and RP.	
Untraceable	assures that an attacker is unable to identify or infer the existence of a transaction and the identities of the entities that initiate or participate in the transaction.	