DCSS/DECA
NIHRFC0028
Category: Standard

V. Wampler, D. Bucci, and E. Harten
CIT
June 2008

NIH Federated Identity Bricks and Pattern v1.0

Status of this Memo

This document specifies a standard for the National Institutes of Health (NIH) and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Table of Contents

Approved June 25, 2008

1 Introduction

Federation is the name for the principles and technologies that make the negotiation "trust" that allows an individual person's identity and privileges to be portable across disparate domains. The purpose of this NIHRFC is to focus discussion on evolving Federated Identity service solutions that will affect the NIH Enterprise Architecture with respect to portable identities and secure access. This NIHRFC relies on NRFC0022, "Federal Authentication and Identity Management," for the definition of credentials relied upon for federation identities (also known as Assurance Level). The definitions of credentials will be used to define the assurance levels standards that create a trust among domain owners. Trusting a federation partner to authenticate its own users can only work if that partner has solid security and user-management practices. The proposed technology solutions in this document are intended as standards for the NIH architecture.

2 NIH Federation

The goal of NIH's Federated Identity service is to give a person the ability to use the same user name, password, or other personal identification to access multiple applications or data sources securely and seamlessly by relying on the identity provider's authentication process rather than NIH's. Federated Identity service is enabled through the use of open industry standards and/or openly published specifications. Use cases identified for business needs at the NIH include cross-domain, Web-based single sign-on; cross-domain user account provisioning; cross-domain entitlement management; and cross-domain user attribute exchange (point-to-point, uni- or bilateral, and many-to-many). Federated Identity service is not bound to any specific protocol, technology, implementation, or vendor. Federated identity's portability is achieved by adhering to widely accepted principles, which are described in the section below.

The federated identity market offers dozens of products and competing standards for cross-domain identity. The principles below define the nomenclature and standards that are to be followed and used within NIH's Federated Identity Service's. Many competing products and standards in the IT market provide barriers to the universality of Federated Identity services. NIH will be implementing a tiered technology approach to support interconnectivity. The goal is not to recognize a proprietary solution but rather to identify principles that shall be followed in delivering federation to NIH application owners.

2.1 Federation Principles

Assertions/Claims: Statements from an identity provider to a relying party that contain identity information about a person. Assertions may also contain verified attributes. These may be digitally signed objects or they may be obtained from a trusted source using a secure protocol. Assertions/claims are issued by a trusted authority to uniquely identify the person.

Authentication: The process of verifying a person's identity. Authentication precedes Authorization.

Authorization: The process of verifying that a known person has the authority to perform certain operations

Claims/Assertions Transformation: The process by which a token service can exchange the claims from an input token for a corresponding different set of claims in an output token.

Digital Identity: A set of characteristics by which a person is recognizable or distinguished in the digital realm. Identity is the set of information that correctly pertains to a physical person.

Discovery Service: Defines a process by which a centralized discovery service can provide a requesting service provider with the unique identifier of an identity provider that can authenticate a person. This is also known as “where are you from” or “WAYF.”

Federated Identity: The idea of having several independent identity providers, each able to assert that the claims that they know about a person are true and accurate. Federated identity does not involve the view of a central identity provider.

Federation Specifications: Applications or tools that communicate securely with services that implement the Web Services Federation Language (WS-Federation, WS-Security, WS-Trust) and OpenID specifications. The specifications may use Extensible Markup Language (XML), Simple Object Access Protocol (SOAP) and Web Services Description Language (WSDL) to provide mechanisms that enable authentication and authorization across different trust realms.

Federated Trust: The trust relationships between the parties involved in the secure exchange of persons’ identity information.

Identity Provider (IDP): Authoritative source for issuing and validating user identities and credentials for a set of people (trust broker, publisher)**.** Manages and asserts (to trusted Service Providers) a person’s attributes securely and accurately.

Service Provider (SP): Provides services for internal and external users via the World Wide Web. Acts as a relying party to validate credentials issued by a trusted identity provider. The relying party that offers access to on-line information, resources, or other services based on some aspect of the identity of users.

Relying Party: Processes identity credentials such as infocard, cardspace, PKI, etc. The requestor and eventual consumer of the token makes assertions/claims about an individual in order to uniquely identify the individual concerned. A relying party does not manage identity.

Resource Offering: Defines associations between a type of identity data and the service instance that provides information about obtaining access to the data (digital identity, claims, Security Assertion Markup Language (SAML) assertion, Security Token Service (STS)). Resource offerings can be stored as an attribute under a user’s profile using the Lightweight Directory Access Protocol (LDAP).

Identity Metasystem: The Identity Metasystem is an interoperable architecture for digital identity that assumes people will have several digital identities based on multiple underlying technologies, implementations, and providers. A process in which federations can initiate federated single-sign-on operations for a user at other federated domains (Lightweight Directory Access Protocol (LDAP), Identity Lifecycle management (ILM)). Using this approach, not only will individuals be put in control of their own identity, but organizations will be able to continue to use their existing identity infrastructure investments, choose the identity technology that works best for them, and more easily migrate from old technologies to new technologies without sacrificing interoperability with others.

Security Token Service (STS): Online identity management service that provides an identity selector that offloads authentication functions, irrespective of whether the agent signing in is a user logging-in to Web sites and services or a site or service owner who needs to authenticate users.

Token: A system object representing the subject of an access control operation. It is a corroborated set of claims, cryptographically signed (to ensure that the contents of the token have not been tampered with) by a trusted third party (such as a bank, credit card company, insurance company, etc.) who is able to assert that the information contained within the token is accurate. It can also be something that the claimant possesses and controls, typically a private key or password that is used to verify the claimant's use of his or her credential to claim an identity.

Two-factor Authentication: Requires two forms of identification in order to access a system, for example, a personal identification number (PIN) and a credit card. There are three forms of identification "factors" generally in use today:

1. Something a person knows: password, PIN, etc.
2. Something a person has: credit card, smartcard, hardware token.
3. Something a person is: biometric information such as fingerprint, retina scan, etc.

In most systems, at least two of the above "factors" are needed to identify an individual.

User Centric Identity: The user is in the middle of the data transaction and has a consistent user experience. The user decides which identity and data gets utilized and with which relying party.

2.2 Federation Technologies

For the implementation of Federated Identity Services at NIH, existing solutions, in particular NIH Login and Active Directory (AD), will be utilized. No new software need be purchased. NIH Login and AD will be configured to support the IT industry Federation Open Standards. The current recognized federation standards include SAML, WS-Federation, Liberty Alliance, WS-Trust and Shibboleth which software vendors utilize to communicate via the principles listed above. Any technology or software that follows both the principles identified above, as well as other NIH or Federal requirements, has the opportunity to become a part of future implementations or upgrades of the Federated Identity solution. The following sections identify currently known or anticipated Federation technologies that follow the federation principles detailed above:

- Identity Provider
- Authentication/Authorization
- Protocols

2.2.1 Identity Provider

Baseline Environment (Today)	**Tactical Deployment (0-2 years)**	**Strategic Deployment (2-5 years)**
■ Active Directory (NIH, NIH External) ■ Oracle Internet Directory (Commons) ■ LDAP ■ Apple Open Directory ■ eDirectory	■ Active Directory Federated Services ■ Identity Provider (NED, InCommon, Active Directory Application Mode) ■ Identity Lifecycle Manager (ILM) ■ LDAP ■ Directory Services (edDirectory, Apple Open Directory, ActiveDirectory)	■ Identity Provider (NED, ibroker (identity service provider for identity services))
Retirement Targets (Technology to eliminate)	**Containment (No new deployments)**	**Emerging (Technology to track)**
	■ Oracle Internet Directory (Commons, other NIH Oracle directories) ■ NIH local directories duplicating additional identities	■ OpenID ■ OASIS ■ Liberty Alliance
Comments		

2.2.2 Authentication/Authorization

Baseline Environment (Today)	Tactical Deployment (0-2 years)	Strategic Deployment (2-5 years)
■ CA Etrust Site Minder ■ Secure ID/RSA ■ HHS PKI ■ Active Directory services ■ Local database authentication	■ Active Directory Federated Services 1.0 ■ CA Etrust Site Minder Resource (SAML, claims/tokens, STS) ■ HHS PKI ■ Information Card (infocard, CardSpace) ■ Shibboleth ■ Active Directory	■ Active Directory Federation Services 2.0 ■ CA Etrust SOA ■ Information Card (infocard, CardSpace) ■ Public Key Infrastructure (PKI) ■ Higgins
Retirement Targets (Technology to eliminate)	**Containment (No new deployments)**	**Emerging (Technology to track)**
	■ Local database authentication	■ Two-factor authentication ■ Three-factor authentication
Comments		

2.2.3 Protocols

Baseline Environment (Today)	Tactical Deployment (0-2 years)	Strategic Deployment (2-5 years)
■ Kerberos ■ NT token/cookie	■ Security Assertion Markup Language (SAML) ■ Secure Token Service ■ Kerberos ■ SOAP ■ WS* Token	■ Security Assertion Markup Language (SAML) ■ Secure Token Service ■ Kerberos ■ WS* Token ■ SOAP
Retirement Targets (Technology to eliminate)	**Containment (No new deployments)**	**Emerging (Technology to track)**
	■ NT Token/cookie	■ WS* Token

Comments
■ Tactical and Strategic products were selected to leverage NIH's investment in products that are a proven fit for NIH's known future needs. Leveraging baseline products in the future will minimize the operations, maintenance, support and training costs for new products. ■ Evolving open source products are Emerging because open source developers have done a better job of modularizing their software, making it more feasible to combine components in order to produce a desired solution. ■ Containment items listed were generated from discussions with two Domain teams. Identity Providers should be geared towards the investment in SSO and Federated Identity and not to continue NIH-developed, duplicate identity silos. ■ Authentication and Authorization are tracked together in Federated Identity Services. Authentication is verifying an identity and authorization is taking that authenticated identity and verifying access. In Federated Identity, the data to accommodate both actions is encapsulated within the same task.

3 References

National Institutes of Health, NRFC0022, Federal Identity and Authentication Management: http://enterprisearchitecture.nih.gov/NR/rdonlyres/F62905D1-E25E-4952-A4B1-FAFA67FFC12C/0/NRFC0022.pdf (Assurance Level)

Security Assertion Markup Language v2.0 technical overview: http://www.oasis-open.org/committees/download.php/20645/sstc-saml-tech-overview-2%200-draft-10.pdf

Federation. Microsoft TechNet: http://msdn2.microsoft.com/en-us/library/ms730908.aspx

eAuthentication: www.cio.gov/eauthentication

Liberty Alliance Project: www.projectliberty.org/liberty/content/download/423/2832/file/tutorialv2.pdf

OASIS Security SAML: http://www.oasisopen.org/committees/tc_home.php?wg_abbrev=security

4 Contact

To contact the NIHRFC Editor, send an email message to EnterpriseArchitecture@mail.nih.gov.

5 Security Considerations

This NIHRFC raises no security issues.

6 Changes

Version	Date	Change	Authority	Author of Change
0.1	8/10/07	Original Template		Valerie Wampler Debbie Bucci Elgin Harten
0.2	10/31/07	Edited wording and moved text		Valerie Wampler Debbie Bucci Jane Small
0.3	11/7/07	Added text under Bricks to indicate that exiting solutions will be employed and no new software is needed to implement Federated identity		Valerie Wampler
0.4	12/28/07	Edited text before brick clarifying intent	NRFC Comments	Valerie Wampler
0.5	12/28/07	Edited brick spell out abbreviations not defined, corrected grammar, added PKI	NRFC Comments	Valerie Wampler
0.6	12/28/07	Federation Principles: arranged alphabetically, spelled out abbreviations, clarified sections, punctuation	NRFC Comments	Valerie Wampler
0.7	1/4/08	NIH Federation: updated intro for clarity and definition	NRFC Comments	Valerie Wampler

Version	Date	Change	Authority	Author of Change
0.8	1/8/08	Edit for grammar and format	Valerie Wampler	Katherine Matthews
0.9	1/17/08	Edited Identity Provider table	Helen Schmitz	Valerie Wampler
0.10	1/29/08	Various edits	Helen Schmitz/Jane Small	Valerie Wampler
0.11	2/1/2008	-Reapplied NIHRFC labels; updated version numbers; corrected appendix numbering; updated status of this memo section.	NIHRFC0001	Steve Thornton, NIHRFC Editor
0.12	4/16/2008	-Added pattern description and brought diagram into alignment with pattern template.	NIHRFC0001 and ITMC-EA Subcommittee	Steve Thornton, NIHRFC Editor
0.13	05/12/2008	-Updated brick to correct version control problem. Specifically several technologies were categorized as retirement that should have been in containment.	NIHRFC0001 and ITMC-EA Subcommittee	Steve Thornton, NIHRFC Editor
0.14	06/04/08	-updated brick comments to include description of containment and authentication/authorization with Federated Identities; edited containment to be less application specific; added current baseline in authentication; removed transport from table as a duplication of protocol	Late comments from NIHRFC discussion	Wampler, Valerie Author
1.0	06/26/2008	ARB Approved	ARB	Steve Thornton, NIHRFC Editor

7 Authors' Addresses

Valerie Wampler
National Institutes of Health
12 South Drive
MSC 5649
Bethesda, MD 20892–5649
Phone: 301.402.7169
Email: wamplerv@mail.nih.gov

Debbie Bucci
National Institutes of Health
6707 Democracy Blvd
MSC 5676
Bethesda, MD 20892–5476
Phone: 301.435.7546
Email: bucci@mail.nih.gov

Elgin Harten
National Institutes of Health
12 South Drive
MSC 5649
Bethesda, MD 20892–5649
Phone: 301.594.6667
Email: hartenel@mail.nih.gov

Appendix A: Federation Pattern

This pattern presents a logical workflow for Federated Identity when a user attempts to access a protected Information Technology resource at NIH.

When the access is detected, the protected resource checks to see if the user has been authenticated. If authenticated, the protected resource checks to see if the user is authorized based on information it has available. If authorization is successful, then the application is invoked; if not, the process ends.

If the user is not already authenticated, the access is routed to either the user-centric store using the info card or PIV certificate method or routed to a directory service to select the Identity Provider. The Identity Provider (IdP) attempts to authenticate the user. If this is not successful, the process ends.

If the authentication is successful, the IdP forwards a pre-defined set of attributes for use by the relying application/relying party/service provider. A check is made to see if the account is known at NIH. If not, the service updates the local store with the identifying information forwarded by the IdP.

The attribute information is then forwarded to the application/relying party/service provider to make an authorization decision based on the information provided. If the user is authorized to access the application, then the application is invoked; if not, the process ends.

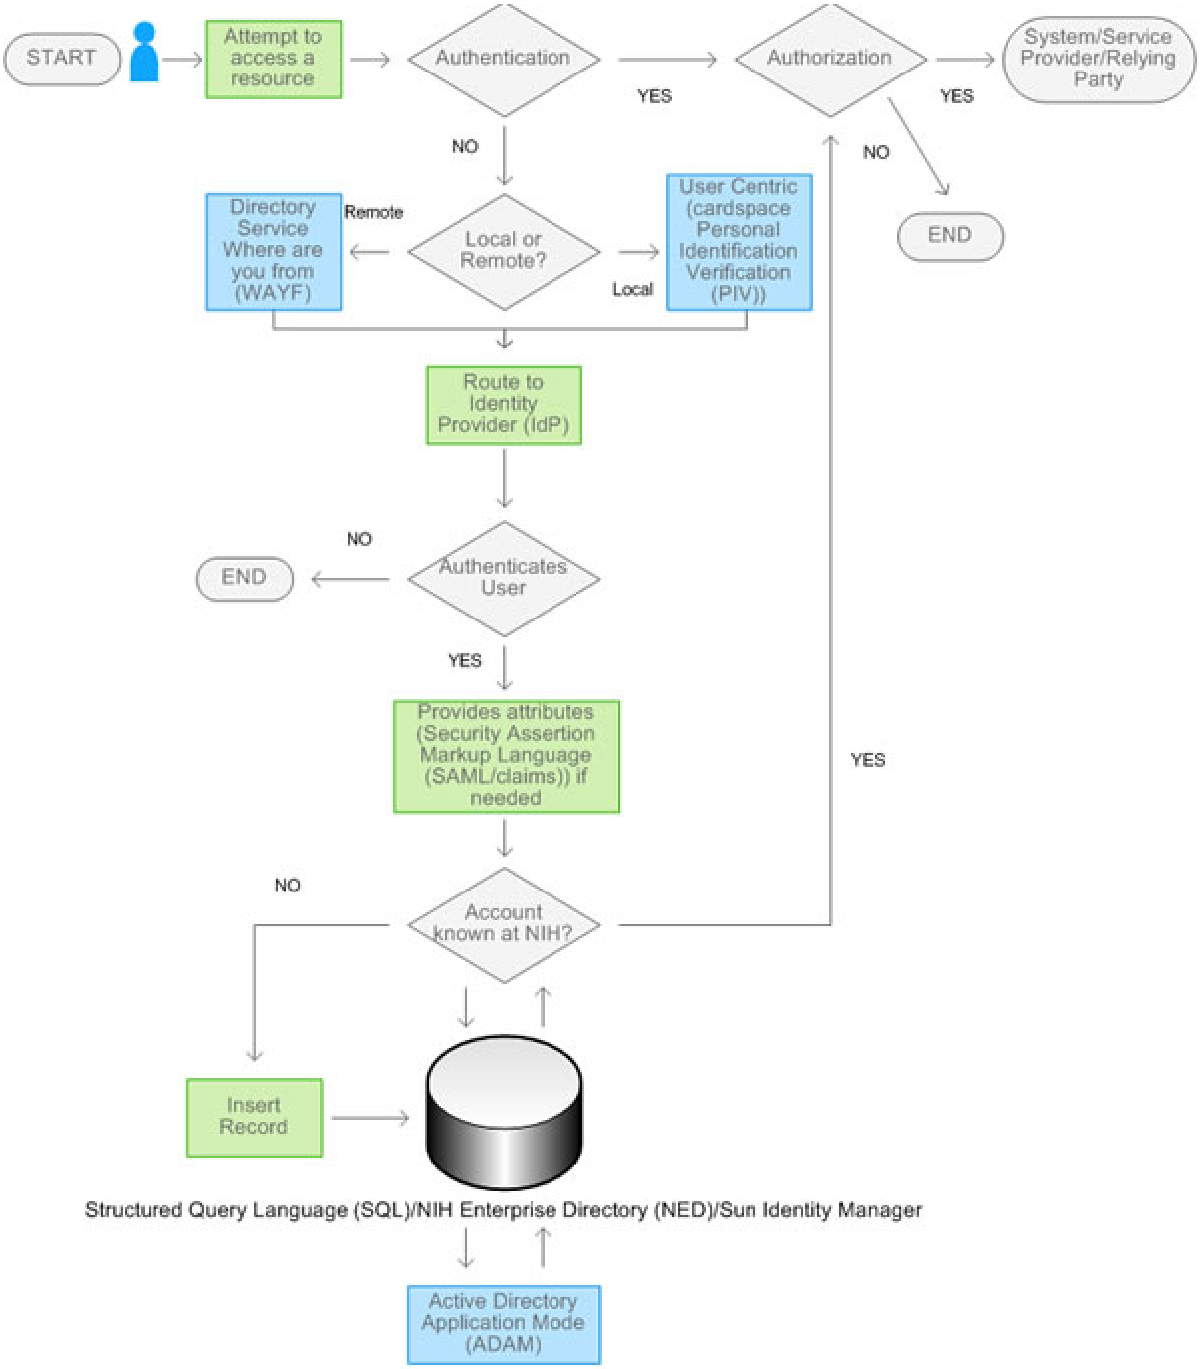
START
Attempt to access a resource
Authentication
YES
Authorization
YES
System/Service Provider/Relying Party
NO
NO
END
Directory Service Where are you from (WAYF)
Remote
Local or Remote?
Local
User Centric (cardspace Personal Identification Verification (PIV))
Route to Identity Provider (IdP)
NO
END
Authenticates User
YES
Provides attributes (Security Assertion Markup Language (SAML/claims)) if needed
YES
NO
Account known at NIH?
Insert Record
Structured Query Language (SQL)/NIH Enterprise Directory (NED)/Sun Identity Manager
Active Directory Application Mode (ADAM)