Office of the Chief IT Architect (OCITA) M. Amodio/D. Bucci
NIHRFC0030 OCITA
Category: Standard June 2008

NIH Login to Support Single Sign-On Technologies for Web-Based Applications v1.4

Status of this Memo

This memo establishes a standard for the NIH architecture community. Distribution of this memo is unlimited.

Table of Contents

1 Introduction

The intent of this NIHRFC is two-fold. First, it is intended to set forth a standard for single sign-on and to prescribe the use of NIH Login for all new web-based applications requiring authentication. Second, the NIHRFC proposes the implementation of this recommendation through changes to the existing Identification and Authentication Brick (See http://enterprisearchitecture.nih.gov/ArchLib/AT/TA/IdentificationandAuthenticationBrick.htm).

By implementing this standard, NIH can evolve towards a more homogenous technical environment which will provide the following benefits:

- Reduced investment in duplicative technologies
- Ease of integration into existing technologies
- Reduced training investment
- Increased user convenience due to a consistency in authentication user experience across systems
- Improved security posture

2 NIH Login

2.1 Background

Single sign-on (SSO) is a method of access control that enables a user to authenticate once and gain access to the resources of multiple software systems in an enterprise. At the NIH, SSO is implemented through the NIH Login offering. NIH Login allows users to authenticate once and to be subsequently and automatically authenticated to other target systems when these are accessed — almost always without modification to the target systems. NIH Login also handles password change requests from target systems and may support post-sign-on automation for additional tasks. In addition, NIH Login provides Application Programming Interfaces (APIs) to allow for the creation of custom authentication screens which interface with the NIH Login software.

2.2 Business Objectives for NIH Login

Business objectives for implementing NIH Login included the following:

- **Promote Collaboration.** NIH Login allows an individual to access multiple applications or data sources securely and seamlessly.
- **Improved User Experience.** Users with fewer separate logins spend less time remembering, typing, and resetting passwords, resulting in higher productivity.
- **Remove Authentication Responsibilities from Application Providers.** Technologies like SSO ease the burden of authentication from the developers and allows them to concentrate on developing business solutions, including application-specific authorization schemes. This centralization of authentication also immediately provides feature enhancement (e.g. federated authentication) to participating applications.
- **Compliance and Audit.** NIH Login solutions allow you to centralize audit and hence monitor compliance.

- **Reduced Administration Burden.** Centralized administration of enterprise authentication frees individual application providers from supporting password resets and user provisioning.
- **More Robust Security.** Increased security, because users do not have "sticky notes" all over their desks displaying multiple passwords.
- **Ensure Person Deactivation.** It offers a more secure deactivation approach of people who leave NIH.
- **Reduced Application Development Time.** A shared authentication service eliminates the need for application developers to design, develop, test and implement the service in their applications.
- **Accommodation of non Microsoft Platforms**. Up to one third of NIH desktops are Apple Macs, and there are a significant number of Linux desktops.
- **Support Federated Authentication**. NIH Login solutions accommodate authentication with organizations outside of the NIH.

As with any technology, there are risks. The two major risks facing NIH Login are availability and interoperability between multiple systems and platforms. Below are the risks and mitigations:

- **Availability.** NIH Login is required to be available 7x24 with any uptime average of 99.95%. The implementation includes redundancy on site and extends redundancy and fail-over to a remote site.

- **Interoperability.** NIH Login strives to comply with open industry standards. The Login/Federation team actively interacts with members of various standard organizations such as Liberty Alliance, OASIS and the World wide web consortium (W3C.. By adhering to the standards and actively participating in those communities, we are able to quickly adapt to new and emerging standards.

3 NIH Login at the National Institutes of Health (NIH)

3.1 Introduction

NIH Login provides authentication to NIH web-based applications enterprise-wide, allowing one change (e.g., implement needed patches or new features) to immediately improve all systems.

3.2 History

Early in 2001, multiple programs at the NIH identified requirements for single sign-on and came to consensus that a single, enterprise implementation was the most cost-effective and architecturally correct approach. Initially, six Institutes and Centers (ICs), the NIH Business System (NBS), the Enterprise Human Resources and Payroll (eHRP) system, and the NIH Enterprise Directory (NED) programs participated.

In mid-2002, the Netegrity (now Computer Associates (CA)) SiteMinder product was selected for implementation, with the NBS Program identified as the first consumer of NIH Login. This implementation required the cooperation of multiple CIT divisions in the summer of 2002 and

required all participating ICs and programs to migrate their current authentication stores to Microsoft Active Directory.

The first program, NIH Portal, was secured NIH Login in February 2003 (See http://cit.nih.gov/ProductsAndServices/WebServices/NihLogin.htm). NBS Travel, nVision and NED followed soon thereafter. Three applications were incorporated under NIH Login in 2003 and 40 applications were incorporated by 2005.

Over time, more and more applications utilized NIH Login for centralized authentication. Today, more than 200 separate applications utilize the single sign-on functionality that NIH Login provides.

NIH Login continues to evolve its functionality. For Example, NIH Login was recently extended to support federated authentication with organizations outside of the NIH (See http://enterprisearchitecture.nih.gov/About/Approach/FederatedAuthentication.htm). 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. Detailed information on federation technologies and architectures can be found in NIHRFC0028, "Federated Identity Bricks and Pattern."

New mandates also cause NIH Login to adapt and expand its capabilities and functionality. Examples of these mandates adopting federated identity include InCommon, Open ID, Active Directory Federation Services (ADFS), eAuthentication, and Homeland Security Presidential Directive 12 (HSPD-12).

3.3 Architecture

At a high level, NIH Login works in the following way:

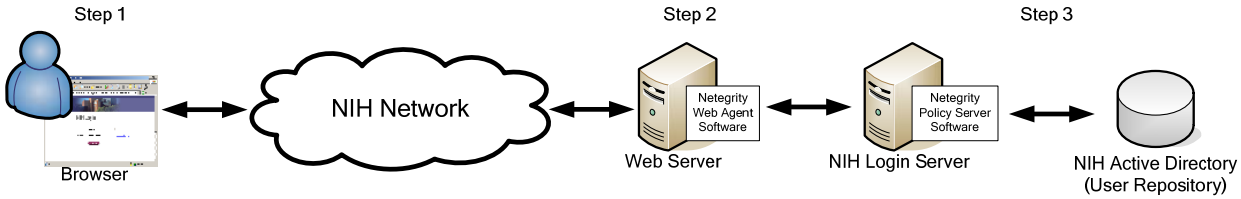

Step 1: An unauthenticated user enters the Uniform Resource Locator (URL) or Web Address for the web-based application he/she wishes to use (e.g. NBS).

Step 2: The Web Server hosting the web-based application (e.g. the NBS Web Server) receives and examines the user session and sees that the user has not yet logged-into the system. The Web Server redirects the user to the NIH Login browser page and the user enters his username and password information.

Step 3: The CA SiteMinder Web Agent Software intercepts the user's information and sends it to the NIH Login Server (Note: Applications which can support Security Access Markup Language (SAML) Assertions may not require the Web Agent. Many vendors already include SAML Assertions into their new products). For the sequence diagram of the SAML Assertions

process, please see Appendix A. There, the SiteMinder Policy Server Software authenticates the user's username and password against the NIH Active Directory user repository.

Once the NIH Login Server authenticates the user, the NIH Login Server sets the required parameters such as user's username and user information (e.g. full name, department) into the Web Server session and returns the session to the Web Server. The Web Server then utilizes the session information to look up the user's authorization roles in the web-based application's (e.g. NBS) unique authorization database, allowing the user to begin working with the web-based application. NIH Login can work with multiple levels of authentication, levels 1-5, including two-factor authentication.

For as long as the user maintains the current browser session, subsequent user logins to additional web-based applications work the same way, except that when the NIH Login Server receives the request, it checks to see if the user still has an active, authenticated session. If the user does have a session, NIH Login returns authentication credentials without requiring the user to re-login.

4 Identification and Authentication Brick

This standard establishes NIH Login as the required method of implementing authentication in new web-based applications at the NIH.

Authenticated identities are the basis for many other information security services. Therefore, NIH needs to:

- Verify user identity as the basis for access control to NIH resource
- Control individual user access to the resources and services provided by those systems
- Create an audit trail of individual user access or attempted access to those systems, resources and services.

Authentication services are crucial to access control and auditing services. If users' identities are not properly authenticated, NIH has no assurance that access to resources and services are properly controlled. In most situations, User ID and password combinations will provide an appropriate level of security if the User ID and password conform to NIH policy. However, NIH will implement stronger authentication for enterprise users with high system privileges (e.g. system, network and security administrators).

NIH Login shall be used by web-based applications for user authentication. If web servers would want to utilize their own authentications, they would need to install additional software such as WebCullis or CoreStreet to perform certificate validations when the user attempts to access a web resource. If an application utilizes NIH Login, that is eliminated completely.

Baseline Environment (Today)	Tactical Deployment (0-2 years)	Strategic Deployment (2-5 years)
■ NIH Login (username and password) ■ Application-specific user authentication based on databases including LDAP, RDBMSs ■ Application-specific user authentication including IP and MAC Addresses ■ Integrated Windows Authentication	■ NIH Login (username and password, PKI, SAML assertions)	■ NIH Login (username and password, PKI, SAML assertions)
Retirement Targets (Technology to eliminate)	**Containment (No new deployments)**	**Emerging (Technology to track)**
	■ Application-specific user authentication based on databases including LDAP, RDBMSs ■ Application-specific user authentication including IP and MAC Addresses ■ Integrated Windows Authentication	■ Biometrics which integrate with NIH Login ■ Smartcards which integrate with NIH Login ■ User-Centric Credentials Integrate with NIH Login
Comments		
■ The brick has been updated to specify just products instead of the combination of both products and standards as previously defined. ■ 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. ■ As the purpose of this NIHRFC is to standardize Identification and Authentication for NIH web-based applications through use of NIH Login, NIH Login is the only selection for Tactical and Strategic technologies and shall be used by new web-based applications requiring authentication functionality. ■ Integrated Windows Authentication was added to the Baseline Environment and has also been added to Containment. Microsoft applications do not scale, and they create trust issues. Integrated Security can only be used at NIH with Active Directory, and therefore, it will not be able to be utilized externally. ■ The NIH Login, itself, is the proposed standard and does not denote a specific supporting technology. ■ Currently, NIH Login utilizes CA SiteMinder.		

5 References

For additional information about the NIHRFC process and/or the NIH Enterprise Architecture, please visit http://enterprisearchitecture.nih.gov.

For additional information about the NIH Enterprise Architecture Identification and Authentication Brick, please visit http://enterprisearchitecture.nih.gov/ArchLib/AT/TA/IdentificationandAuthenticationBrick.htm.

For additional information about the NIH Single Sign-on via NIH Login, please visit http://cit.nih.gov/ProductsAndServices/WebServices/NihLogin.htm.

For additional information about the federated authentication capabilities available via NIH Login, please visit http://enterprisearchitecture.nih.gov/About/Approach/FederatedAuthentication.htm.

For additional information about the underlying technologies and architectures upon which NIH Federation is based, please read **NIHRFC0028**, "NIH Federation Identity Bricks and Pattern".

6 Contact

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

7 Security Considerations

Although this NIHRFC involves changes to security architecture procedures, the information contained in this document does not compromise security considerations at NIH.

8 Changes

Version	Date	Change	Authority	Author of Change
0.1	-	Original Template	N/A	Terrence Blair, NIH OCITA
0.2	12/21/2007	Edits	NIHRFC0001	Steve Thornton, NIHRFC Editor
0.3	12/28/2007	Edits		Terrence Blair, NIH OCITA
0.4	01/04/08	Edits		Terrence Blair, NIH OCITA
0.5	01/22/08	Applied NIHRFC number and draft stamp. Minor edits.	NIHRFC0001	Steve Thornton, NIHRFC Editor
0.6	01/24/08	Edits	NIHRFC0001	Terrence Blair, Matthew Amodio, NIH OCITA
0.7	2/28/08	-Addressed NIH community comments. -Referenced NIHRFC0028, "Federated Identity Bricks and Pattern"	NIHRFC0001	Terrence Blair, Matthew Amodio, NIH OCITA
0.8	3/19/2008	-Changed the title per the suggestion from the ITMC EA Subcommittee	ITMC EA Subcommittee	Matthew Amodio, NIH OCITA

1.0	4/14/2008	-ARB approved on 4/2/2008. -Changed author's address.	ARB	Steve Thornton, NIHRFC Editor
1.1	5/29/2008	-Added Debbie Bucci as Technical Writer -Addressed comments made after approval date -Added background information -Added Integrated Windows Authentication to Baseline and Containment in Brick	NIHRFC0001	Matthew Amodio, NIH OCITA
1.2	6/10/2008	-Added benefits into Section 2.2	NIHRFC0028	Matthew Amodio, NIH OCITA
1.3	6/10/2008	-Added Appendix A SAML Assertions Process sequence diagram -Added reference link in Section 3.3 to Appendix A	NIHRFC0001	Matthew Amodio, NIH OCITA
1.4	6/18/2008	-Added Risks and Mitigations for NIH Login	NIHRFC0001	Matthew Amodio and Debbie Bucci, NIH OCITA

9 Authors' Addresses

Matthew Amodio
National Institutes of Health
10401 Fernwood Road
MSC 4806
Bethesda, Maryland 20817
Phone: 301-402-1088
Email: EnterpriseArchitecture@mail.nih.gov

Debbie Bucci
National Institutes of Health
2 Democracy Plaza, 322
MSC 5476
Bethesda, Maryland 20817
Phone: 301-435-7546
Email: bucci@exchange.nih.gov

10 Appendix A – Security Access Markup Language (SAML) Assertions Process

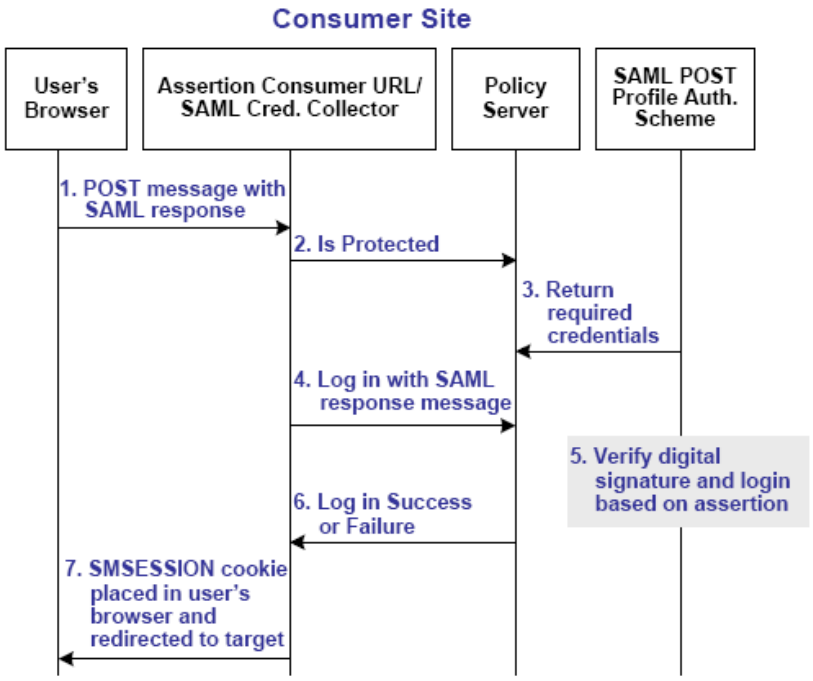

Unless otherwise stated, the following process takes place at the consumer site:

To use the SAML POST profile for passing assertions, the assertion generator at the producer site needs to sign the SAML response that contains the assertion. The assertion consumer at the consumer site needs to verify that signature.

To accomplish these tasks, you must set up a key database for each Policy Server that is responsible for signing, verification or both. The key database is a flat-file key and certificate database that lets you manage and retrieve keys and certificates required to sign and validate SAML responses used with SAML POST profile authentication

1. A user's browser POSTs an HTML form to the Assertion Consumer URL (which is the URL for the SAML credential collector). This form contains a SAML response message and target URL originally generated at the producer.

2. The SAML credential collector makes a call to the Policy Server to determine if the target resource is protected.

3. The Policy Server replies that the target URL is protected by the SAML POST profile authentication scheme. This indicates to Federation Web Services application that a signed response from the POSTed form is the expected credential for the login call.

4. The SAML credential collector makes a login call to the Policy Server, passing the digitally signed SAML response as credentials.

5. The SAML POST profile authentication scheme verifies the signature and other fields of the response and the assertion.

6. If the checks succeed and the user is found in the directory, then authentication succeeds. If any of the checks fail, authentication fails.

7. Assuming login succeeds, the SAML credential collector sets an SMSESSION cookie, which it puts in the user's browser, and then redirects the user to the target resource. If the login fails, the credential collector redirects the user to the configured No Access URL.