

# Introduction to Shibboleth

Amy Apon and Kurt Landrus

Computer Science and Computer Engineering Department  
University of Arkansas, Fayetteville, AR 72701

## Abstract

*A key issue in distributed computing is the authentication and authorization of users to remote resources. This paper is an introduction to Shibboleth, a project and software package of Internet2/MACE. Shibboleth is designed to authorize a user to access a remote web-based resource through the use of the login and attribute information that is maintained at the home institution of the user. It allows user privacy to be maintained and can be used to enforce levels of access based on user characteristics. For example, if members of a class or research group need access to a remote data or computing resource, Shibboleth can be used to grant this access based on group membership. Shibboleth provides an architecture for trust relationships between users and resources.*

## 1. Introduction

With the widespread growth of computational resources, data archives, research tools, courseware, and information sources, many researchers, students, teachers, and information technology workers have the need to share and access resources that are located and owned by a different institution than the individual's home institution. By building federations, institutions that own resources can share them, and there is opportunity for technology transfer, greater learning, and improved ability and access to research tools and products.

Shibboleth is an Internet2 architecture, policy framework, and practical technology to support sharing of web resources and services [1]. It is a product of the Middleware Architecture Committee for Education (MACE) committee of Internet2, and is an open-source standards-based inter-institutional resource sharing mechanism.

The primary use of Shibboleth is to provide access to a remote resource in an authenticated manner and subject to access controls at the remote resource. While most other remote authentication mechanisms require that a user reveal user identification and a password to the remote institution, Shibboleth has the capability to control the access to a resource without this information. The current version of Shibboleth, version 1.2.1a, works with web-based resources.

Shibboleth is named for the first known password. The word comes from the Hebrew *shbal*, to grow, increase, or

flow. It was used by Jephthah as a test-word to identify the Ephraimites, who were unable to pronounce the "sh" (Judges 12:1-6) at the fords of the Jordan.

Shibboleth protects resources in the same way that a user ID and password can protect resources. However, Shibboleth protection is based on group membership, or attributes, rather than on the identification of a particular user. As an analogy, consider that only officers can enter the officer's club at the base. Any officer can enter, but only officers, and it is not necessary to show an individual ID. If a resource requires it, then Shibboleth has the flexibility to require that individual identification information be provided.

An example can illustrate the versatility and a usage scenario of Shibboleth. For example, suppose that the University of Arkansas in Fayetteville (UA) would like to share access to its supercomputing cluster and terascale storage facility to authorized users at the University of Arkansas for Medical Sciences in Little Rock (UAMS). Suppose that the potential users at UAMS are members of the biomedical informatics research group and the resource is to be made available for a particular research project. With Shibboleth, users at UAMS would login using their institutional login ID and password. The user attribute, "member of the biomedical informatics research group" would be passed to the resource at UA. Through the Shibboleth infrastructure, the UA computing resource would acquire the attribute (e.g., group membership) that it needs from UAMS. It would test the value of the attribute (e.g., biomedical informatics research group) to verify that the user has been allowed to access the resource, and then the user would have access. In this process the user's identity does not need to be revealed, but only the attributes, such as group membership and institution affiliation, that are required to access the resource. Additionally, the UAMS user has the option not to allow those attributes to be sent to UA. The UA resource may grant or deny access based on the access control policy that it has established.

In this scenario, the user does not have to reveal his or her personal identity to the research resource. The user could desire this, for example, if the researcher is performing access on a large data resource for AIDS and the researcher does not want to reveal his or her personal identity as an AIDS researcher. Only the group

membership information is sent, and the user's privacy is protected.

Shibboleth allows users to access many types of resources but to maintain fewer accounts and passwords per user and between institutions. Access may be based on roles instead of hard-coding of user names and passwords.

## 2. Key Shibboleth Concepts

With Shibboleth, a user need not have an account ID or reveal personally identifying information to gain access to a remote resource. Rather, Shibboleth requires only a single user ID and password at a home institution for access to resources on potentially thousands of remote resources at numerous institutions. Shibboleth works among institutions that choose to band together in federations of trust and using commonly trusted Certificate Authorities.

Within a federation, different institutions may have radically different local authentication infrastructures. Typical campus-level infrastructures for identity management and authentication may include LDAP or Windows Active Directory Services, for example. Federated identity services allow organizations using disparate authentication and authorization methods to interoperate. The Shibboleth protocol operates over these existing services rather than replacing them. Institutions agree to grant access based on user attributes, such as student at an institution, enrolled in a certain course, faculty member, staff member in a particular department, or working on a certain grant or contract. Shibboleth allows users to specify which attributes may be released to which institutions or to access which remote resources.

Shibboleth is a profile of the SAML protocol defined by OASIS [2]. The Internet2 Shibboleth implementation is one implementation of the Shibboleth Architecture. Shibboleth is based on common security technologies, including public key cryptography, public key certificates and certificate authorities.

Key to the operation of Shibboleth is the separation of the step of authentication and the step of authorization to use a resource. A user is *authenticated* when you are sure that the user is who he/she claims to be (e.g., that user logs in to an account with a password). With Shibboleth the home institution is responsible for the authentication of a user. A user is *authorized* to use a resource if he/she is allowed to have access to it. With Shibboleth, authorization is granted by a remote resource based on the attribute values that have been sent by the home institution. Authorization always implies authentication.

A stated goal of Shibboleth is to move to *active privacy*. The traditional approach to privacy is *passive privacy*. With a passive privacy, a user passes identity to a service, and then has no control over the use of the information and how it may or may not be revealed by the

server. To comply with privacy, services often have significant regulatory requirements. With the active approach to privacy, a user (through the home security domain) can pass attributes to services that are not necessarily sufficient to identify the user. If the attributes are sufficient to reveal an individual's identity, then user can decide whether or not to release them.

The Shibboleth architecture uses two kinds of trust. When two institutions agree to exchange attribute information and to grant access based on the values of certain attributes then *collaborative trust* is being used. Collaborative trust can be between two institutions, or can be used within virtual organizations that include many real organizations that have agreed to cooperate with a common need for resource access. Collaborative trust must be arranged through relationship and negotiation. Various policies must be established, including policies about what attributes will be required for sharing, what values these attributes may take, what level of security is expected for the enterprise infrastructure, and other policies and procedures. Typical attributes for higher education institutions have been identified and collected into standard schema [3].

In addition to this collaborative trust Shibboleth also utilizes *hierarchical trust networks* and Certificate Authorities (CAs). When messages are exchanged in the Shibboleth protocol, the executing entities ensure that the other is really who it thinks it is through the use of mutual certificate-based authentication.

With Shibboleth, institutions agree to *broker* authentication and authorization between the user and the service provider. A local campus trusts the authentication mechanism of another campus, and also trusts that the attributes sent to it regarding a user have been maintained securely and have not been forged or altered. An institution trusts the recipient of attributes to not misuse or alter them as well.

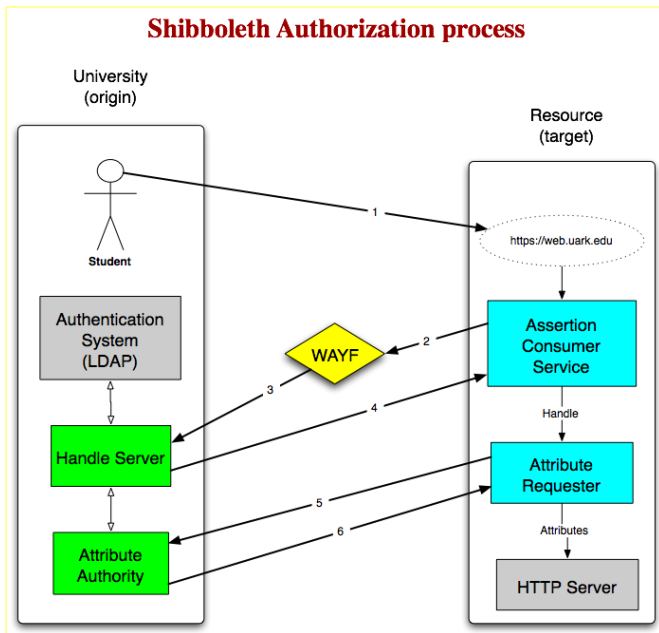
## 3. Shibboleth Architecture and Protocol

Shibboleth consists of: 1) a resource or service provider (target) that is protected via access control mechanisms, 2) an identity provider (origin) such as a university authentication infrastructure, and 3) a Where Are You From (WAYF) server, that may be operated by the service provider but may also be operated by a federation.

The service provider site is required to provide an enterprise infrastructure that includes a resource that is (currently) accessible via a web server (e.g., Apache or IIS). Shibboleth components of the service provider site include an assertion consumer service, an attribute requestor, and a resource manager in addition to the resource.

The identity provider site is required to provide enterprise infrastructure that includes a user authentication

mechanism and an attribute repository. Shibboleth components of the identity provider site include a handle server and the attribute authority.



**Figure 1: Shibboleth Components**

The steps of the Shibboleth protocol can be described at a high-level as shown in Figure 1 and as follows:

1. A user requests service at a remote Resource.
2. The Shibboleth Assertion Consumer Service intercepts the request.
3. The request is forwarded to a Where Are You From (WAYF) server that asks for the user's home organization and redirects the request to it. The Authentication System at the user's home institution asks for a user ID and password. Then the Shibboleth Handle Server generates a temporary name, or "handle," for the user.
4. The handle is sent back to the Resource site.
5. The Resource site doesn't know who this handle represents, only that a trusted institution issued it. The Shibboleth Attribute Requestor asks the home institution for the unknown person's attributes.
6. The Attribute Authority returns the attributes to the Resource site. If the attributes meet the requirements for access then access is granted.

The whole process of using the computer resource at a remote institution happens without revealing the user's true identity, without any personal information leaving the home institution, and without an administrator at the remote site assigning a user ID and password.

InQueue is a federation operated by Internet2 that can be used by institutions that are becoming familiar with Shibboleth [4]. InQueue operates a WAYF and provides test resources that an institution may access to test the functionality of its Shibboleth identity provider. Another federation, InCommon, is becoming established as a Shibboleth federation for production use [5].

#### 4. Summary and the Future of Shibboleth

Shibboleth is an architecture that allows users to access resources at remote institutions using the authentication information from their home institution. Shibboleth allows institutions to federate and share resources using collaborative trust relationships. Shibboleth relieves password overhead and enhances computer privacy.

The GridShib Project [6] may encourage the adoption of Shibboleth by uniting protocols that are being used in the Grid computing community with the scalability and privacy benefits of Shibboleth.

Shibboleth is being actively adopted by institutions within the U.S. and around the world. More information about activities and workshops related to Shibboleth is available at the NMI-EDIT website [7].

#### Acknowledgments

This work is supported in part by funding from Educause to the Great Plains Network Consortium through the GPN Extending the Reach (ETR) project [8]. Funding is on behalf of the NMI-EDIT consortium of Internet2, Educause, and SURA, and with support from several statewide university systems and regional networks.

#### References

- [1] The Shibboleth Project, <http://shibboleth.internet2.edu/>
- [2] Security Assertion Markup Language (SAML) 1.1 Specification, OASIS, November 2003.
- [3] InCommon Federation: Common Identity Attributes, <http://www.incommonfederation.org/docs/policies/federate-dattributes.html>
- [4] InQueue, <http://inqueue.internet2.edu/>
- [5] InCommon, <http://www.incommonfederation.org/>
- [6] GridShib – A Policy Controlled Attribute Framework, <http://grid.ncsa.uius.edu/GridShib/>
- [7] NMI-EDIT, Identity and Access Management for Higher Education and Research, <http://www.nmi-edit.org/>
- [8] Great Plains Network Consortium Extending the Reach Project, <http://archie.csce.uark.edu/gpn/>
- [9] R. L. "Bob" Morgan, Scott Cantor, Steven Carmody, Walter Hoehn, and Ken Klingenstein, "Federated Security: The Shibboleth Approach," *Educause Quarterly*, Vol. 27, No. 4, November 4, 2004, pp. 12-17.