

*A Great Plains Network Report
on*

Shibboleth

Prepared by

Rahul Deshmukh

Graduate Research Assistant

Great Plains Network

&

Kansas State University

Contents at a glance:

- 1. Introduction**
 - 1.1 Historical Context**
 - 1.2 Technical Context**
- 2. Shibboleth: Inner Workings**
 - 2.1 Origin**
 - 2.2 Target**
 - 2.3 WAYF**
- 3. Shibboleth via Screenshots**

References

Appendix I: Glossary of terms

Appendix II: Helpful URL's

GPN Reports are prepared as a service to assist GPN member universities to remain current on the latest technology, research trends and research opportunities. To learn more about the Great Plains Network, including membership and sponsorship opportunities, or to join the GPN distribution list please visit <http://www.greatplains.net> or contact Dr. Greg Monaco, greg@greatplains.net.

Special thanks to members of the Great Plains Network Middleware Initiative and, in particular, Professors Amy Apon (University of Arkansas), Gordon Springer (University of Missouri), and Byrav Ramamurthy (University of Nebraska), as well as Kathryn Huxtable (University of Kansas) and Denis Hancock (University of Missouri) for their patience and understanding in helping me to learn about Shibboleth. Thanks also to Professor Greg Monaco, GPN Director, for his careful reading of successive drafts of this manuscript.

1. Introduction: *Shibboleth* in Context

Shibboleth is a type of software, known as *middleware* [1], which is used in World Wide Web environments to help organizations protect online resources from access by *unauthorized* users. The word *shibboleth* has historical association with matters of identification and security apart from the information technology domain.

The purpose of this document is to provide a high level overview of Shibboleth and its associated components, and to provide the reader with some familiarity with the technical terms associated with Shibboleth and related technologies.

1.1 Historical Context

Shibboleth is a kind of linguistic password that identifies one as a member of a group. According to the history, two Semitic tribes, the Ephraimites and the Gileadites, have a great battle. The Gileadites defeat the Ephraimites, and set up a blockade to catch the fleeing Ephraimites. The Gileadite sentries ask each person to say the word *shibboleth*. The Ephraimites, who have no *sh* sound in their language, pronounced the word with an *s* and were thereby unmasked as the enemy and slaughtered. Thus, a person who violates a *shibboleth* can be rapidly identified as an outsider and immediately excluded from the group. In the English language the word Shibboleth means ***an arbitrary test or custom that distinguishes one group from another***, or a word or slogan identified with a particular group or party. [3]

1.2 Technical Context

Shibboleth is now the name of a standards-based, open source middleware [1] software that provides Web Single Sign On (SSO) [4] across or within organizational boundaries. It allows web sites to make informed authorization decisions for individual access of protected online resources in a privacy-preserving manner.

Most institutions have ***identity provider*** software, which controls access to protected web resources such as payroll records, school credentials, databases, on-line journals, web-enabled instruments, and so forth. A user goes to the home institution's web site, uses his/her name and password to self-identify or authenticate, and, then, gets access to the requested resource. Using Shibboleth, new web site resources, such as a parking permit application, can be easily added to the existing set of resources and take advantage of the existing authentication mechanism.

Even more interesting, resources may be shared across multiple institutions with the user signing on using his/her home institution name/password. This works because Shibboleth takes advantage of the concept of ***trust relationship***: If Institution A trusts Institution B, then Institution A will accept a user ***if that user*** is verified by Institution B.

For example...

For example, a student from the University of Kansas requests access to an online journal (a ***service***) located at the University of Missouri by accessing the University of Missouri's web page. The MU web page will redirect the student to a page containing a list of institutions. The student selects his home institution (KU) and provides his/her (KU) login name and password. After checking the credentials, KU

will notify the MU website, and pass information known as attributes (Attributes are things like **a student** at an institution, **enrolled in certain course, faculty member** or **staff member** in a particular department, or **working** on a certain grant or contract). Depending upon the attributes, MU may allow access to some or all of the resources.

Thus, users from one institution do not sign directly into the web site of a second, **service-provider** institution. Instead, **attributes** about the user are sent from the (home) identity-provider institution to the service provider. The service provider can then use these attributes to decide the type of access the user will have.

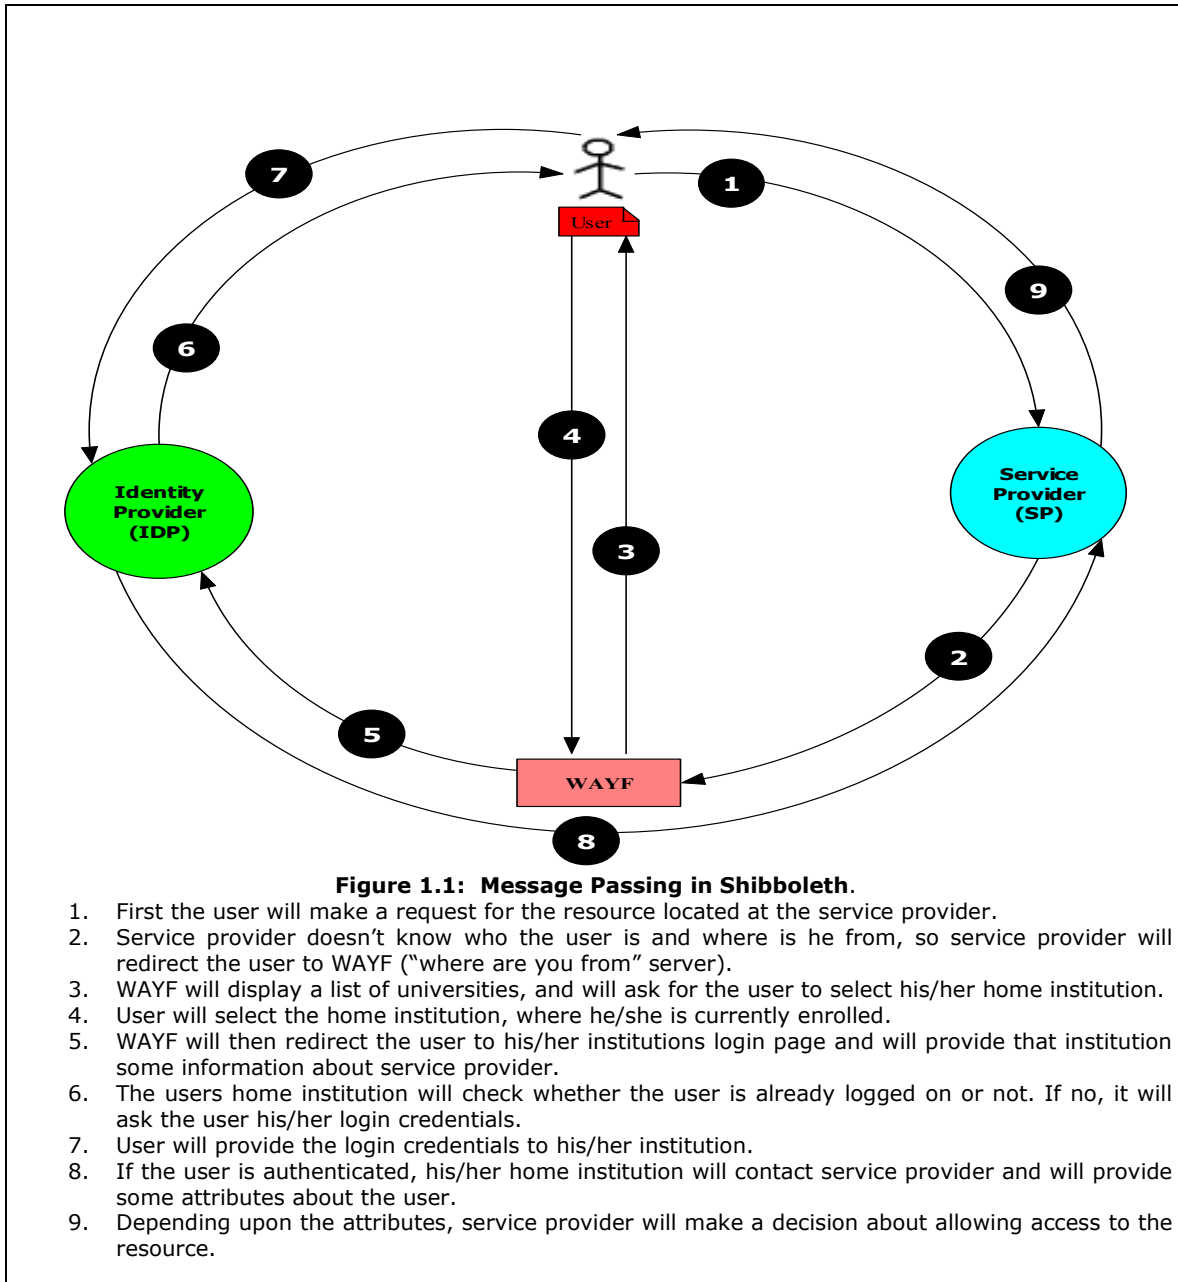
The benefits of Shibboleth are that

- Institutions can share resources with each other by passing attributes in a secure way,
- There is no need to maintain a user database at each site,
- Security is enhanced because only one account exists for all resources,
- User IDs do not have to be used to access a resource, thus privacy is maintained.

2. Shibboleth: Inner Workings

There are two major classes of service provided by Shibboleth: Origin (Identity Provider) and Target (Service Provider). The server on which web resources are located is known as the **target** or **service provider**, since it provides service to its users. The server that is responsible for user authentication is known as the **origin** or **identity provider**.

In addition to the origin and target, there is a central, trusted service known as a **WAYF** that operates to assist in the communication between user, origin and target. WAYF stands for "where are you from". It helps the user to identify his or her home institution from a list. Please refer to Figure 2.1 for a more complete depiction of the interaction between the user and each of these different elements.



2.1 Origin (Identity Provider)

The origin consists of four services: the Handle Service, the Attribute Authority, a directory service (user database) and a sign-on system (SSO). Please refer to Figure 1.2. Handle service and attribute authority are internal components of the origin and comes as a single installation package like MS Word, MS Excel in MS Office.

Handle Service (HS)

The Handle Service is a form of web service provided by shibboleth using JAVA. HS is responsible for providing a pointer (handle) that is submitted over target. The WAYF

will notify the user request to origin's handle service along with the information about service provider. HS will then determine whether the user has already logged on to the origin or not by consulting its SSO. If they have not they will be authenticated. The Handle will be produced and passed to the targets SHIRE.

Attribute Authority (AA)

The AA is responsible for responding the requests from the targets SHAR (shibboleth attribute requester) service, the target will send the handle sent by the HS, this is then used to look for policies used to determine which attributes can be released, based on these policies the AA will then query the directory service for the attributes of the user, these are then sent to the targets SHAR.

Directory Service

The directory service is the organizations user database, normally in the form of an LDAP directory

Sign-On System (SSO)

The sign-on system provides authentication for the user, this is separate from Shibboleth, but is triggered by the HS component of the origin when a user has not yet logged on. This trigger is made using standard web server authentication methods. [4]

2.2 Target (Service Provider)

The target consists of three software components that are installed on the target server: the SHIRE (Shibboleth Indexical Reference Establisher), the SHAR (Shibboleth Attribute Requester), and the RM (Resource Manager), please see Figure 1.2.

SHIRE

The SHIRE is responsible for redirecting users request for resource towards WAYF so that user can select his home institution and get authenticated. Shire is also responsible for receiving handle sent by origins HS. After receiving the handle SHIRE will pass the handle to shibboleth attribute requester (SHAR).

SHAR

The SHAR will use the handle given by the SHIRE along with the address of the AA to request attributes about the user, once these attributes have been passed to the SHAR, AAP's (Attribute Acceptance Policy) will be used to provide validation and analysis before passing the attributes to the RM.

RM

The Resource Manager is responsible for making access decisions based on the attributes sent to it by the SHAR.

2.3 WAYF (Where Are You From?)

In a **federation** where many organizations share a **trust relationship**, the federation may manage a WAYF server. The WAYF [5] provides a mechanism for allowing users to be forwarded to the correct home organization. This is normally presented in the form of a web page with a drop down list. Once the user selects an

organization, the user is forwarded to that institution's origin server/identity provider to complete the authentication process. Two organizations which manage WAYF servers are InCommon and InQueue. InCommon [6] is a production level WAYF and there is an annual fee, whereas InQueue [7] is a test WAYF, maintained for free by Internet2.




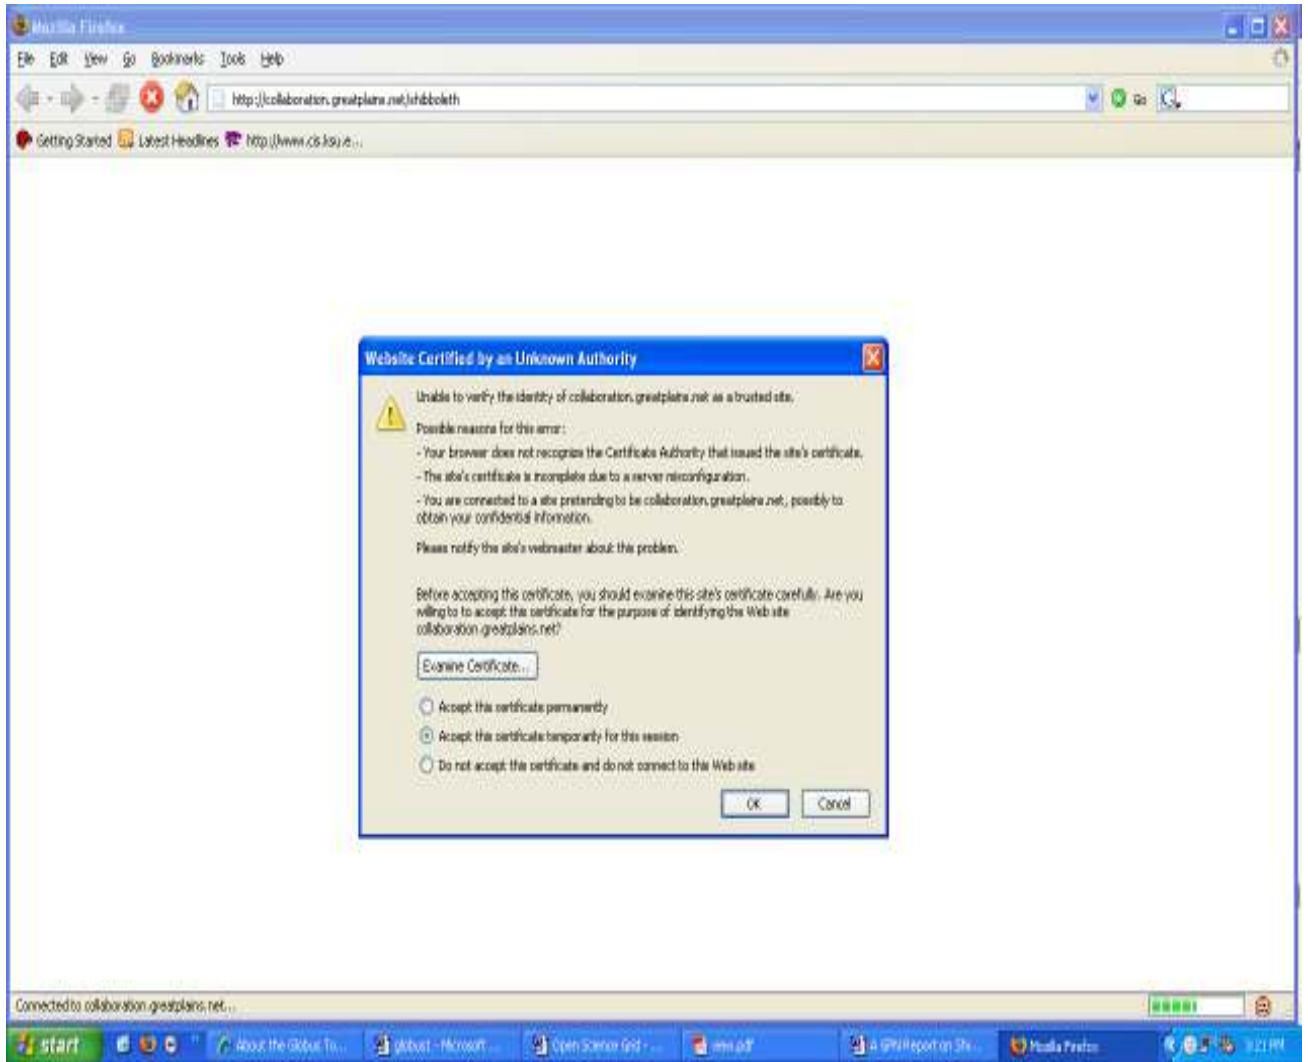
Figure 1.2: Shibboleth Internal Components*

1. The user will make a request for a resource located on service provider (SP). The SHIRE is the component on the SP that handles the request.
2. The SHIRE does not recognize the user and redirects the user to WAYF.
3. WAYF will present a list to the user containing his/her home institutions and will ask to select his/her institution.
4. The user will select his institution from the list.
5. User will be redirected to his/her home institution to get authenticated and authorized. WAYF does this by contacting HS of identity provider (IDP) at the home institution.
6. HS will check whether the user is logged on or not. If not, the HS will trigger SSO system to get user authenticated. User will be asked for login name and password.
7. User will enter his/her login name and password.
8. User will be verified for login credentials and will be authenticated.
9. HS will then send handle directly to SP's SHIRE without contacting WAYF--HS knows on which resource user made a request. Handle is nothing but a symbol of user authentication and trust between the institutions. By using a handle, the user's identity is protected.
10. SHIRE will pass the handle to SHAR.
11. SHAR will contact the attribute authority of IDP and will pass the handle, asking attributes of the user.
12. AA will look for user attributes in the database and will send the required attributes to SHAR. SHAR will check the attributes according to the attribute acceptance policy.
13. SHAR will pass the attributes to resource manager. Depending upon the attributes RM will decide whether to grant access or not.
14. User will be notified and allowed to access the resource after authentication.

*Based on a figure by Kathryn Huxtable, University of Kansas

3 Shibboleth via Screenshots

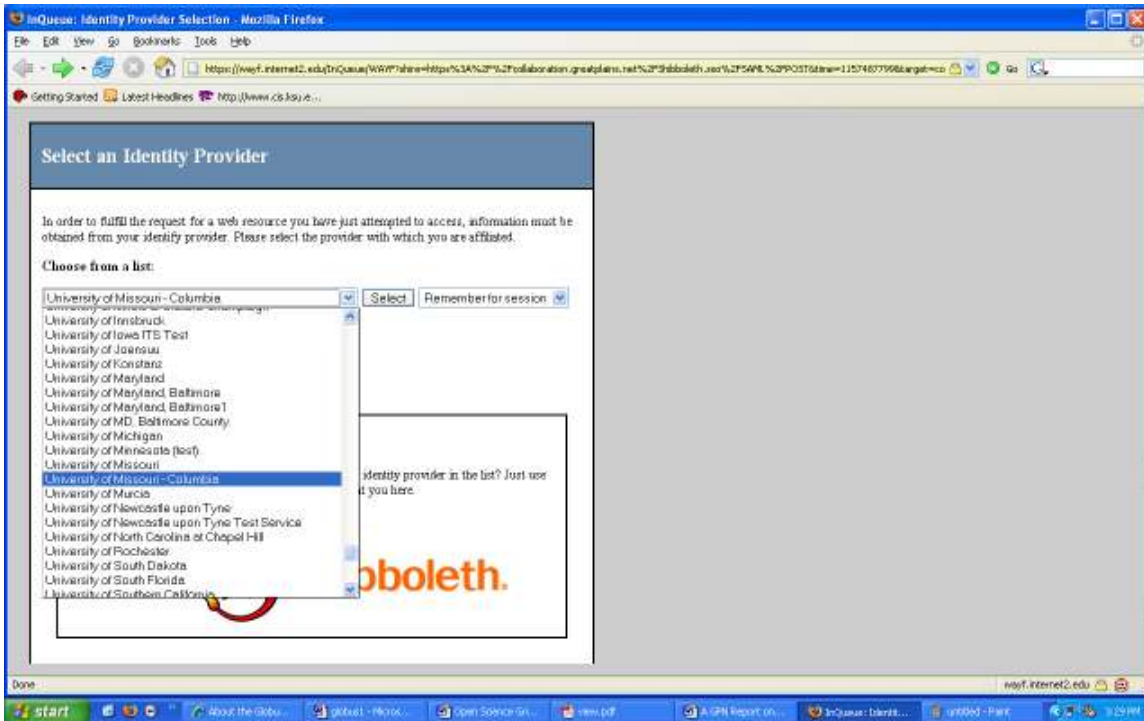
1. User makes a request to a resource in the form of web page located on <https://collaboration.greatplains.net/shibboleth>



2. User is redirected to WAYF



3. User selects home institution (University of Missouri, Columbia) from the list



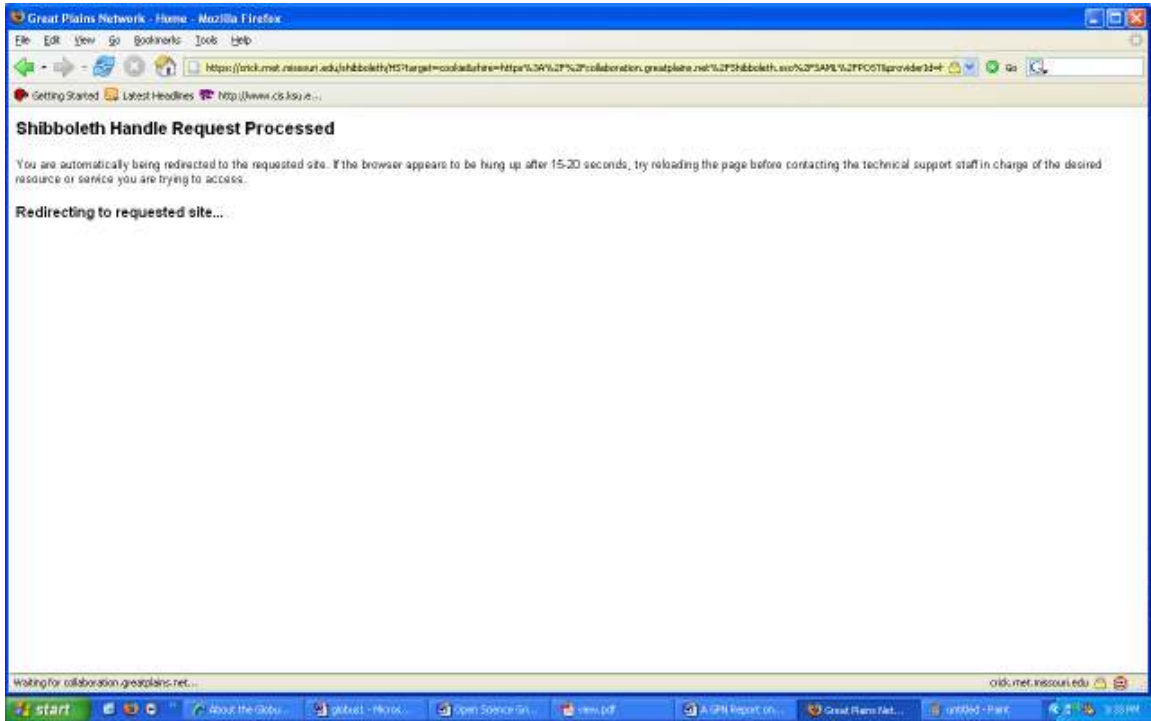
4. WAYF is contacted



5. User is redirected to home institution (University of Missouri, Columbia IDP) to get authenticated



6. Acknowledgement that user is authenticated and attributes are passed to the initial web site



6. Access to the initial resource is now permitted <http://collaboration.greatplains.net/shibboleth>



REFERENCES

[1] Middleware

www.internet2.edu/pubs/200404-IS-MW.pdf

[2] Shibboleth Description

<http://kidderminster.ac.uk/kc-rolo/documents/KCROLOShibbolethGuidev1.doc>

[3] Historical context of shibboleth

<http://www.answers.com/topic/shibboleth>

[4] Web Single sign on

<http://www.avenuea-razorfish.com/articles/SSOApproachPaper.pdf>

[5] WAYF Technical Discussion

<http://www.incommonfederation.org/technical.html>

[6] InCommon

<http://www.incommonfederation.org>

[7] InQueue

<http://inqueue.internet2.edu>

Appendix I Glossary of Terms

Attributes

Attributes in shibboleth context are nothing but the features that identifies a person for example as a student, or faculty member, the department in which he/she studies/works, the courses enrolled/teaching.

Authentication

Authentication is a process of identifying valid users by checking their login credentials against stored database. The authentication is used to deny access for invalid users. In private and public computer networks (including the Internet), authentication is commonly done through the use of logon passwords.

Authorization

Authorization is a next step to the authentication. Authorization determines a level of access to a valid user. The resource can be shared by one person and cannot by other depending upon the access level. For example the instructor of a course can modify the contents of the course but the student can only read the contents though both are having access to the content.

Federation

A federation in Shibboleth is the collective term for a group of Identity Providers and Service Providers sharing a trust relationship, the federation defines the policies that each member must adhere to, for example an IDP may have to use a certain certificate for it to be allowed into the federation for security reasons, attributes required to be released are also defined, if these rules are not met then the institution is not trusted so cannot be part of the federation, the federation therefore requires administration staff to write these policies and take care of the registration duties.

Identity Provider

Please see Origin

Middleware

Middleware is a layer of software between the network and the applications. [1] This software provides services such as identification, authentication, authorization, directories, and security. The Internet2 Middleware Initiative (I2-MI) promotes standardization and interoperability and is working toward the deployment of core middleware services at internet2.

Open Source

Open source is an idea to make users software available for free of cost. One can view the source code and can change the code and can redistribute the software without charging money.

Origin

A server that has identity information about users, usually, the user's home institution.

Service Provider

Please see Target

SSO (Single Sign On)

Single sign-on (SSO) is a specialized form of software authentication that enables a user to authenticate once and gain access to the resources of multiple software systems.

Target

A web server that has some web resources located on it and ready to share with others.

WAYF

A web server that is used by Shibboleth to determine where a user is from (home institution).

Appendix II Helpful URLs

1. Shibboleth Home

<http://shibboleth.internet2.edu/>

2. Shibboleth Wiki

<https://authdev.it.ohio-state.edu/twiki/bin/view/Shibboleth/WebHome>

3. Single Sign On

<http://www.opengroup.org/security/sso/>