



# **SERAPIS Project**

## ***Shibboleth Attribute Management***

**Arts and Humanities Data Service  
King's College London**

**Author: Sanjay Vivek**

**Contact: Mark Hedges**

**7 March 2007**

## Table Of Contents

<b>1.</b>	<b>Introduction .....</b>	<b>4</b>
<b>2.</b>	<b>Shibboleth IdP configuration .....</b>	<b>5</b>
2.1	IDP.xml .....	5
2.2	resolver.ldap.xml .....	5
2.3	arp.site.xml .....	6
<b>3.</b>	<b>Shibboleth SP configuration .....</b>	<b>7</b>
3.1	Shibboleth.xml .....	7
3.2	AAP.xml .....	7
<b>4.</b>	<b>Apache 2 configuration .....</b>	<b>9</b>
<b>5.</b>	<b>Using eduPersonTargetedID .....</b>	<b>10</b>
<b>6.</b>	<b>Using eduPersonEntitlement .....</b>	<b>11</b>
<b>7.</b>	<b>Federated Name Identifiers .....</b>	<b>13</b>
<b>8.</b>	<b>References.....</b>	<b>15</b>

## 1. Introduction

Shibboleth exercises access management and resource control via a set of policies, which define the control and access level to the resources. These configuration files alongside Apache 2's web server configuration file (`/opt/shibboleth-sp/etc/shibboleth/apache2.config`) provides the developer with rules and policies to protect web resources. Project SWISH [3] and SDSS [2] have written excellent guides for Shibboleth attribute management. These instructions were followed and adapted for the SERAPIS project to suit our specific requirements.

The steps required to secure and manage a protected resource are summarized below:

### Identity Provider:

- `/opt/shibboleth-idp/etc/idp.xml` only needs to be configured once.
- `/opt/shibboleth-idp/etc/resolver.ldap.xml` only needs to be configured once.
- `/opt/shibboleth-idp/etc/arp/arp.site.xml` needs adjustment for adding/deleting/updating affiliations and organizations (IdPs). This file releases attributes to a SP according to the configuration in place.

### Service Provider

- `/opt/shibboleth-sp/etc/shibboleth/shibboleth.xml` needs configuring whenever a new local directory within the local filesystem has to be protected
- `/opt/shibboleth-sp/etc/shibboleth/AAP.xml` needs adjustment for adding/deleting/updating affiliations and organizations (IdPs)
- `/opt/shibboleth-sp/etc/shibboleth/apache2.conf` needs adjustment for protected directories and adding/deleting/updating scoped affiliation

## 2. Shibboleth IdP configuration

### 2.1 IDP.xml

The main IdP configuration file is found at `/opt/shibboleth-idp/etc/idp.xml` and it defines 2 key elements as shown in red:

```
<IdPConfig
.....

AAUrl="https://idp.ahds.ac.uk:8443/shibboleth-idp/AA"
resolverConfig="file:/opt/shibboleth-idp/etc/resolver.ldap.xml"
defaultRelyingParty="urn:mace:ac.uk:sdss.ac.uk:federation:sdss"
providerId="https://idp.ahds.ac.uk/shibboleth">

.....

<ReleasePolicyEngine>
  <ArpRepository
implementation="edu.internet2.middleware.shibboleth.aa.arp.provider.FileSystemArpR
epository">
    <Path>file:/opt/shibboleth-idp/etc/arps/</Path>
  </ArpRepository>
</ReleasePolicyEngine>
```

### 2.2 resolver.ldap.xml

The resolverConfig property in the idp.xml file points to an attribute resolver file (resolver.ldap.xml) that defines how attributes are found and matched using LDAP. The key elements are shown in red. This defines the entire chain from releasing attributes through to the Java plugin that obtains information from the LDAP directory.

```
<AttributeResolver.....>
  <SimpleAttributeDefinition id="urn:mace:dir:attribute-
def:eduPersonScopedAffiliation" smartScope="idp.ahds.ac.uk">
    <AttributeDependency requires="urn:mace:dir:attribute-
def:eduPersonAffiliation"/>
  </SimpleAttributeDefinition>

  <SimpleAttributeDefinition id="urn:mace:dir:attribute-
def:eduPersonPrincipalName" smartScope="idp.ahds.ac.uk">
    <DataConnectorDependency requires="directory"/>
  </SimpleAttributeDefinition>
```

```

.....
<JNDIDirectoryDataConnector id="directory">
  <Search filter="uid=%PRINCIPAL%">
    <Controls searchScope="SUBTREE_SCOPE"
      returningObjects="false" />
  </Search>
  <Property name="java.naming.factory.initial"
    value="com.sun.jndi.ldap.LdapCtxFactory" />
  <Property name="java.naming.provider.url"
    value="ldap://idp.ahds.ac.uk:389/ou=People,dc=ahds,dc=ac,dc=uk" />
  <Property name="java.naming.security.principal"
    value="cn=Manager,dc=ahds,dc=ac,dc=uk" />
  <Property name="java.naming.security.credentials" value="secret" />
</JNDIDirectoryDataConnector>
.....
</AttributeResolver>

```

### 2.3 arp.site.xml

On the IdP side, attributes are released to the Service Provider via the arp.xml file. Although it is possible to include a file for each eduPersonPrincipalName attribute in the federation, it is highly unfeasible. As such, the only attribute that we have any definite knowledge about is the eduPersonScopedAffiliation attribute because it is a mandatory attribute and with a setting of [staff/student/associate@ahds.ac.uk](mailto:staff/student/associate@ahds.ac.uk) (the scoping part is dependent on the IdP) as defined in Shibboleth. In this example, only [staff@ahds.ac.uk](mailto:staff@ahds.ac.uk) and [student@ahds.ac.uk](mailto:student@ahds.ac.uk), are released to the AHDS test Service Provider.

```

<AttributeReleasePolicy.....
  <Description>AHDS Test SP</Description>
  <Rule>
    <Target>
      <Requester>https://sp.xenophobe.ahds.ac.uk/shibboleth</Requester>
    </Target>
    <Attribute name=
      "urn:mace:dir:attribute-def:eduPersonScopedAffiliation">
      <Value release="permit">
        staff@.ahds.ac.uk</Value>
      <Value release="permit">
        student@.ahds.ac.uk</Value>
    </Attribute>
  </Rule>
</AttributeReleasePolicy>

```

### 3. Shibboleth SP configuration

#### 3.1 Shibboleth.xml

The main SP configuration file can be found at `/opt/shibboleth-sp/etc/shibboleth/shibboleth.xml` and contains the key elements highlighted in red. The example below shows three directories being protected by Shibboleth.

```
<SPConfig.....>
  <Local....>
    <RequestMapProvider .....
```

#### 3.2 AAP.xml

On the Shibboleth SP side, all attributes released from the IdP's AA are mapped onto HTTP Request headers and subsequently passed on to the application being protected by the Shibboleth SP. Before being mapped, the values are first filtered to see if they should be accepted by the application, and if so, which header each attribute should be mapped to. All this is handled by configuring the Attribute Acceptance Policy File ([AAP.xml](#)).

When a scope is presented in an attribute, the AAP is evaluated to verify if this particular scope matches the scope required in AAP. If the scope doesn't match, then the attribute is rejected.

Scoping is a way for both IdP and SP to get consensus about which organization domains are acceptable by the SP. A particular resource in SP is perhaps only accessible to users that belong to certain domains, e.g. only users from the physics domain (@physics.kcl.ac.uk) can access our physics report.

```
<AttributeAcceptancePolicy....>
  <AttributeRule
    Name="urn:mace:dir:attribute-def:eduPersonScopedAffiliation"
    Scoped="true" CaseSensitive="false" Header="Shib-EP-
    ScopedAffiliation"
    Alias="scopedaffiliation">

    <SiteRule Name="https://ahds.ac.uk/shibboleth">
      <Scope Accept="true">ahds.ac.uk</Scope>
      <Scope
        Type="regexp">^.+\.xenophobe\.ahds\.ac\.uk$</Scope>
        <Value>student</Value>
        <Value>staff</Value>
      </SiteRule>
    </AttributeRule>
  </AttributeAcceptancePolicy>
```

## 4. Apache 2 configuration

The Alias previously defined in the AttributeAcceptancePolicy can be used in conjunction with the Shibboleth configuration file, [/opt/shibboleth-sp/etc/shibboleth/apache2.conf](#) to provide authorization control over the secure directories. This can be achieved by using the Location element as shown:

```
<Location /staffonly>
  AuthType shibboleth
  ShibRequireSession On
  require scopedaffiliation staff@ahds.ac.uk
</Location>

<Location /studentonly>
  AuthType shibboleth
  ShibRequireSession On
  require scopedaffiliation student@ahds.ac.uk
</Location>

<Location /secure>
  AuthType shibboleth
  ShibRequireSession On
  require scopedaffiliation staff@ahds.ac.uk student@ahds.ac.uk
</Location>
```

The configuration above simply states that /staffonly is only accessible by users who have supplied their eduPersonScopedAffiliation attribute AND this attribute matches [staff@ahds.ac.uk](#). The same holds true for /studentonly but when the attribute matches [student@ahds.ac.uk](#) while /secure is accessible by both “[staff@ahds.ac.uk](#)” and “[student@ahds.ac.uk](#)” users.

## 5. Using eduPersonTargetedID

As described in the Shibboleth Glossary:

**Persistent Identifier (eduPersonTargetedID):** This special identifier type allows an IdP and SP to preserve a single identifier for one principal across all current and future transactions involving that principal without revealing or using that principal's identity. These identifiers are opaque to SP's and vary per SP, protecting against collaborative Denial of Privacy attacks, while preserving a 1:1 guaranteed mapping for liability and preference management purposes.

It is therefore a persistent opaque identifier, which enables service personalisation without the SP knowing who the user is, when released by the IdP.

A Shibboleth identity provider can generate the opaque eduPersonTargetedID attribute automatically from some other stored attribute that holds the user id in the clear, such as eduPersonPrincipalName, by editing [resolver.ldap.xml](#) and uncommenting or adding:

```
<PersistentIDAttributeDefinition id="urn:mace:dir:attribute-def:eduPersonTargetedID"
scope="ahds.ac.uk" sourceName="eduPersonPrincipalName">
  <DataConnectorDependency requires="directory"/>
  <Salt>XXXXXXXXXXXXXXXXXXXXXXXXXXXX</Salt>
</PersistentIDAttributeDefinition>
```

The <Salt> is a constant, arbitrary value that you should choose once and keep secret. The value must be at least 16 characters long, otherwise the software will silently ignore it and expect the value to be supplied from a Java keystore. The Salt value is used to generate the persistent opaque identifier from the scope and some other attribute, normally the user id (eduPersonPrincipalName). Its purpose is to prevent attempts to work back from the opaque identifier to the user's identity by combining knowledge of the scope and the hash function used with an exhaustive search of the possible user ids.

The default Shibboleth attribute release policy does not release eduPersonTargetedID. You must therefore manually edit the [ARP.xml](#) file to enable this feature:

```
<Attribute name="urn:mace:dir:attribute-def:eduPersonTargetedID">
  <AnyValue release="permit"/>
</Attribute>
```

## 6. Using eduPersonEntitlement

Identity providers may assert a particular eduPersonEntitlement value to indicate that an individual user should have access to a resource. When a user's IdP such entitlements, the business rules that evaluate a user's attributes to determine eligibility are evaluated there. The Service Provider does not know of the user's attributes beyond their entitlement. The IdP and SP has to come to a common agreement about the values of the entitlement and this is usually done out of band.

To enable /secure to be only accessed by users who have an entitlement of *urn:ahds.ac.uk:staff*, the following have to be done:

### On IdP side

1. Update the LDAP store to include **eduPersonEntitlement: urn:ahds.ac.uk:depositAllAccess** for the user.
2. Configure the **apache.config** file:

```
<Location /secure>
  AuthType shibboleth
  ShibRequireSession On
  require entitlement urn:ahds.ac.uk:depositAllAccess
</Location>
```

3. Configure the **ARP.xml** file to release the eduPersonEntitlement attribute because this is not released optionally:

```
<Attribute name="urn:mace:dir:attribute-def:eduPersonEntitlement">
  <AnyValue release="permit"/>
</Attribute>
```

### On the SP side

1. Configure **AAP.xml** to state which IdPs the SP is prepared to accept the entitlement from. For example, to only accept entitlements from *idp.ahds.ac.uk*:

```
<AttributeRule Name="urn:mace:dir:attribute-def:eduPersonEntitlement"
Header="Shib-EP-Entitlement" Alias="entitlement">

    <SiteRule Name="https://idp.ahds.ac.uk/shibboleth">
        <Value Type="regexp">^urn:ahds.ac.uk:.$</Value>
    </SiteRule>

</AttributeRule>
```

2. To accept attributes entitlements from ANY IdP:

```
<AttributeRule Name="urn:mace:dir:attribute-def:eduPersonEntitlement"
Header="Shib-EP-Entitlement" Alias="entitlement">

    <AnySite>
        <Value Type="regexp">^urn:ahds.ac.uk:.$</Value>
    </AnySite>

</AttributeRule>
```

## 7. Federated Name Identifiers

The AHDS Catalogue and Delivery system is a web-based application for searching and downloading collections in the AHDS repository. These collections are in many cases freely available; however, some of them are subject to restrictions, and currently users are obliged to use a local registration and logon system to access these collections. During this project, Shibboleth was used to replace this local system, so that the catalogue is protected by the Shibboleth SP, and any IdP in the UK federation can access the restricted collection. Users can choose to authenticate themselves either by Shibboleth or the local login system, in which case the user has to complete a registration form before accessing the restricted resource.

Due to license agreements between the AHDS and the collection owners, certain collections are restricted to registered users, and every user that downloads the collection has to be recorded. The obvious choice for identifying the user in a Shibboleth context is the eduPersonPrincipalName attribute. This poses a problem, however, since it is likely that IdPs will only release the ePSA attribute. As mentioned in the UK Federation documentation [4], a SP should only acquire and an IdP should only release the EPPN attribute for good and sufficient reasons. This is largely because, of the four core attributes identified by the Federation, EPPN is the only one that identifies users and as such generates personal data subject to the Data Protection Act. Under normal UK Federation conventions, the user's EPPN attribute is in the form of [user@domain](#), which is in format similar to an email address, and it is likely that many IdP administrators will equate the EPPN attribute to the user's email address.

An alternative would be to require an IdP to release the eduPersonTargetedID (EPTID) attribute instead of the EPPN, and to use the EPTID to identify users. The EPTID is an opaque identifier, and is unique for any given combination of EPPN, IdP and SP. It is not necessarily persistent however as it is calculated at run-time from both the EPPN attribute and a salt value, as described in Section 5. As mentioned above, EPPNs are often username-based and are subject to change and re-assignment, and salt values could also be changed, either of which would result in the EPTID changing.

Another solution would be to define a new EPPN-like attribute, which is both opaque and persistent, in a similar fashion to the swissEduPersonUniqueID attribute used by SWITCH [1]. However, this would have to be agreed by all the institutions in the federation, and would place an additional and unacceptable burden on the IdP administrators, who would have to deal with an extra attribute and an ever-growing data store.

In the case where an IdP does not release any attributes that can be used to identify a user, it would still be possible for the user to register and logon using our local registration system. However, this rather defeats the object of using Shibboleth; the intention behind leaving the local registration system in place was that some of our users are private researchers not affiliated with any institution or IdP. For the moment,



the prototype Shibbolised catalogue requires that an IdP releases either the EPPN or EPTID attribute, if the user is to access restricted resources using Shibboleth-based authentication. If this proves to be an impractical assumption once the catalogue becomes live, the approach will be revisited.

## 8. References

1. Authorization Attribute Specification for SWITCH, [http://tim-3.ethz.ch/docu/AAI\\_Attr\\_Specs.pdf](http://tim-3.ethz.ch/docu/AAI_Attr_Specs.pdf)
2. SDSS Attribute Usage, <http://www.sdss.ac.uk/content/Documents/AttributeUsage>
3. SWISH Attribute Management Guide, [http://gilead.ex.ac.uk/swish/index.php?option=com\\_content&task=view&id=40&Itemid=9](http://gilead.ex.ac.uk/swish/index.php?option=com_content&task=view&id=40&Itemid=9)
4. Technical Recommendations for Participants, UK Federation. <http://www.ukfederation.org.uk/library/uploads/Documents/technical-recommendations-for-participants.pdf>