



UK Access Management Federation for
Education and Research

Technical Recommendations for Participants

Ian A. Young
15 October 2012

Version 1.3.1

Table of Contents

1 Introduction.....	3
1.1 Keeping Up To Date.....	3
1.2 Document Status.....	3
1.3 Changes in this Edition.....	3
1.4 Future Directions.....	4
2 Software.....	5
2.1 Choice of Software.....	5
2.2 Future Directions.....	6
3 SAML 1.1 Authentication Request and Response Profiles.....	7
3.1 Recommended Authentication Request Profile.....	7
3.2 Recommended Authentication Response Profiles.....	7
3.3 Future Directions.....	9
4 Metadata.....	10
4.1 Standard Metadata.....	10
4.2 Metadata Refresh.....	10
4.3 Metadata Signature Verification.....	11
4.4 Federation URI.....	12
4.5 Future Directions.....	12
5 Digital Certificates.....	14
5.1 Certificate Roles in the UK Federation.....	14
5.2 Recovery from Key Compromise.....	15
6 Discovery.....	16
6.1 Avoiding Discovery: Institutional Portals.....	16
6.2 Discovery by the Service Provider.....	16
6.3 Federation Central Discovery Service.....	17
6.4 Future Directions.....	18
7 Attribute Usage.....	19
7.1 Core Attributes.....	19
7.2 Attributes, Privacy and Data Protection.....	26
7.3 Subsidiary Attributes.....	26
7.4 Sources of Additional Attributes.....	26
7.5 Custom Attributes.....	27
7.6 Working Without Attributes.....	27
7.7 Future Directions.....	27
8 References.....	29

1 Introduction

This document provides technical recommendations for members of the UK Access Management Federation for Education and Research (the UK federation). Its primary audience is those technical staff involved in designing services or deploying software for use in the UK federation.

The information in this document is supplemented by that provided by the UK federation's web site.¹ In particular, the web site always contains the most up-to-date version of recommendations in areas of rapid change such as the suitability of certificate products or specific software versions.

A companion document, the *Federation Technical Specifications* ([UKFTS]), specifies the federation's technical architecture in more detail, including the rationale behind some of the technical choices made. Familiarity with the *Federation Technical Specifications* is not normally required for individual deployments; its primary audiences are developers of federation software and operators of partner federations.

The federation serves a broad constituency of member organisations with a wide range of uses for federated identity technologies. This document is therefore not prescriptive; rather, it aims to establish a common set of standards each of which, if followed, will increase an individual member's ability to inter-operate with other members of the federation.

1.1 Keeping Up To Date

Due to the rapidly changing nature of the software and standards associated with identity technologies, it will be necessary to update this document from time to time to reflect new developments. The latest version of this document can always be found on the federation web site (see [UKTRP]); federation members should review the latest version of this document periodically, and in any case whenever a new deployment is contemplated.

New editions of this and other federation technical documents, as well as other announcements thought to be relevant to federation members, are reported on the federation mailing list. The technical and administrative contacts listed for all entities registered with the federation are made members of the mailing list automatically; other addresses can be added to the list by request.

1.2 Document Status

This edition provides recommendations for the UK federation with effect from its date of publication as shown on the cover page.

1.3 Changes in this Edition

- Updated references to the Shibboleth web site, which has moved to a new domain.
- Remove some software-specific information.

¹ See <http://ukfederation.org.uk>

- Include validation fingerprints for the metadata signing certificate in use from November 2012.

1.4 Future Directions

Where appropriate, major sections of this document contain a sub-section called “Future Directions” describing likely future developments in the area under consideration. These notes are provided to allow members to incorporate this information into planning activities.

Several sections of this document have been identified as requiring frequent change in order to stay current. In future editions, it is intended that such information be provided on the UK federation's web site instead of in more static documentation.

2 Software

2.1 Choice of Software

The UK federation uses the Security Assertion Markup Language (SAML) standards² for the communication of authentication, entitlement and attribute information. The core of the federation is implemented using the Shibboleth software³ from the Shibboleth Consortium. It is recognised, however, that any particular software implementation may not be suitable for all participants, and federation members may deploy any software that meets their specific service goals.

It is likely that organisations which regularly update their implementations to use the latest version of the Shibboleth software will continue to benefit from the widest range of interoperability options with other federation members. Other software, however, may well be better suited to particular operating environments. It is the member organisation's responsibility to ensure that the software chosen for their deployment can interoperate with those other members of the federation that are important to their service aims.

Some remarks specific to particular software implementations are presented below; inclusion on this list does not constitute an endorsement and exclusion from this list does not constitute deprecation.

2.1.1 Shibboleth 1.3

At the end of June 2010, the Shibboleth 1.3 software became unsupported by the Shibboleth project. Because this means that no security updates will be available from that time, this should be regarded as the effective end of life date for this version of the software.

We strongly recommend against the use of Shibboleth 1.3 for new deployments, and it is no longer supported by the UK federation helpdesk.

In addition, any members still using Shibboleth 1.3 are recommended to upgrade to version 2.X as soon as possible. Not doing this exposes a site to a risk of compromise from security vulnerabilities discovered in Shibboleth 1.3.

2.1.2 Shibboleth 2.X

The latest release in the Shibboleth 2.X series is recommended for new Shibboleth deployments within the federation. This version is fully supported by the Shibboleth Consortium, including the provision of security updates in binary form for many target environments.

2.1.3 OpenAthens

The OpenAthens software, developed by Eduserv Athens, is known to be in production use by some members of the UK federation.

More information about the OpenAthens software can be found at:

<http://www.eduserv.org.uk/identity-access/products>

² See http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

³ See <http://shibboleth.net/>

2.1.4 Guanxi

The Guanxi software, developed by the UHI Millennium Institute in partnership with the University of Oxford and the University of Leeds, is known to be in production use by some members of the UK federation.

More information about the Guanxi software can be found at:

<http://www.guanxi.uhi.ac.uk/>

2.1.5 Stay Updated

Whichever software you choose, you should maintain it by, for example, applying security patches to it in a timely fashion. We strongly recommend upgrading your software as necessary, to stay current with software versions for which security patches are available from the vendor.

This applies not only to the identity and access management layer (Shibboleth or other software) but also to the underlying system software such as web servers, database and directory software, and operating systems.

2.2 Future Directions

2.2.1 Shibboleth 3.0

The next major release of the Shibboleth software will be version 3.0. No release date is available at the time of writing.

Current plans for Shibboleth 3.0 involve only minor changes to system configuration, and in general Shibboleth 3.0 may be regarded as a continuation of the Shibboleth 2.X series. The change to the major version number is required by the Shibboleth team's version numbering policy as a result of changes to the API used by extensions to the software.

2.2.2 Specific Software Versions

As recommendations about specific software and versions need to change rapidly to keep up with developments, this information will be removed from future versions of this document. Current recommendations in this area will instead be available from the UK federation's web site.

2.2.3 Other Software

Member organisations should ensure that they are kept informed of the development roadmap associated with any software they use in connection with the federation.

3 SAML 1.1 Authentication Request and Response Profiles

The ability of federation members to interoperate with other federation members depends both on the software deployed and on the protocols and profiles which they use for communication. This section describes the profiles recommended for use within the federation.

3.1 Recommended Authentication Request Profile

The only authentication request profile currently recommended for use with SAML 1.1 within the UK federation is the Shibboleth authentication request profile⁴ as described in section 3.1.1 of [ShibProt].

All current federation SAML 1.1 identity providers implement this profile. The centralised “Where Are You From” (WAYF) discovery service provided by the UK federation also operates on the basis of this profile.

3.2 Recommended Authentication Response Profiles

3.2.1 SAML 1.1 Browser/POST With Attribute Pull

Older Shibboleth software (up to and including version 1.2) implemented only one authentication response profile, SAML 1.1 Browser/POST. This is defined in [SAMLBind] section 4.1.2; its use in Shibboleth is as described in [ShibProt] section 3.1.2.

We strongly recommend that all new members of the UK federation deploy software capable of making use of the SAML 1.1 Browser/POST profile, as this is the only authentication response profile known to be supported by all current federation entities supporting SAML 1.1.

Identity providers should always implement SAML 1.1 Browser/POST with *attribute pull*, which is to say in such a way that the authentication assertion is sent to the service provider without any accompanying attributes. This then causes the service provider to issue a separate attribute request over a protected and mutually authenticated channel, so that the transfer of attributes is both secure and known to be at the request of a particular identified party. This attribute exchange profile is described in [ShibProt] section 3.2.

During the attribute exchange operation, the service provider has the opportunity to indicate the attributes it is requesting through the use of the AttributeDesignator element. The Shibboleth service provider software can be configured to make use of this facility using corresponding AttributeDesignator elements in its configuration file. Note, however, that this is not the default configuration and many service providers therefore omit the AttributeDesignator element from their queries; such a query becomes a request for all attributes whose release is permitted by the identity provider’s attribute release policy for that service provider.

An identity provider is always responsible for protecting the privacy of its users through its choice of the attributes to be released to a particular service provider. Identity providers should never attempt to delegate that responsibility by relying on

⁴ urn:mace:shibboleth:1.0:profiles:AuthnRequest

appropriate AttributeDesignator elements being expressed by a service provider. Instead, identity providers should define appropriate attribute release policies for each service provider to which attributes containing personal data need to be released. The default attribute release policy should only allow the release of privacy preserving attributes.

We strongly recommend that SAML 1.1 Browser/POST is never implemented with *attribute push*, which causes the attributes for the subject to accompany the authentication assertion. This method of operation is insecure, as it transports the attributes in an unencrypted fashion through the user's browser. This makes possible attacks based on impersonating the service provider's identity, either through weaknesses in the browser's PKIX implementation, or through tricking the user in so-called "phishing" attacks. The attribute pull mechanism is not vulnerable to these particular attacks.

3.2.2 SAML 1.1 Browser/Artifact With Attribute Push

Shibboleth 1.3 introduced support for a second authentication response profile, Browser/Artifact, as defined in [SAMLBind] section 4.1.1; its use in Shibboleth is described in [ShibProt] section 3.1.3.

The advantages of the Browser/Artifact profile include faster authentication speed in certain circumstances, and a removal of the Browser/POST profile's need for ECMAScript support in the user's browser. Against this must be weighed the lack of widespread support for this profile by current federation members.

Browser/Artifact can be deployed with either attribute push or attribute pull without loss of security. However, Browser/Artifact with attribute pull causes two communications to be made back to the identity provider after the authentication assertion has been sent to the service provider, and is therefore much slower. We recommend the use of attribute push whenever the Browser/Artifact profile is employed.

An identity provider is always responsible for protecting the privacy of its users through its choice of the attributes to be released to a particular service provider. When using attribute push, an identity provider always releases all attributes included in the attribute release policy for the particular service provider. Therefore, identity providers should define appropriate attribute release policies for each service provider to which attributes containing personal data need to be released. The default attribute release policy should only allow the release of anonymous attributes.

We recommend deploying the Browser/Artifact profile if the software you are using supports it. Shibboleth 1.3 deployments, for example, provide this facility in their default configurations.

We do not recommend deploying entities capable of supporting only the Browser/Artifact profile.

3.2.3 Choice of Authentication Response Profile by Service Providers

The Shibboleth authentication request profile requires the service provider sending the authentication request to include the location of an assertion consumer service in the request. This location must match one of the `AssertionConsumerService` elements described in the federation metadata for

that service provider, and each such element also includes a Uniform Resource Identifier (URI) specifying the profile bound to that location.

If a service provider performs discovery locally, so that it knows the identity provider to which it is sending the authentication request, it may nominate any of its assertion consumer service locations that supports a profile known to be supported by that identity provider.

If the service provider does not perform local discovery, and instead makes use of centralised discovery services provided by the federation's WAYF, it must nominate an assertion consumer service location without advance knowledge of the identity provider which the request will be sent to. As not all federation identity providers currently support the Browser/Artifact profile, but all do support the Browser/POST profile, the service provider must always select the Browser/POST profile in this case.

3.3 Future Directions

3.3.1 Prevalence of SAML 1.1 Browser/Artifact

At the time of writing, the SAML 1.1 Browser/Artifact profile is supported by around 95% of all identity providers that support SAML 1.1. However, the recommendation that all SAML 1.1 entities remain capable of handling the SAML 1.1 Browser/POST profile is likely to stand for the foreseeable future, as it is unlikely that SAML 1.1 Browser/Artifact will ever be supported by every entity.

3.3.2 Profiles for SAML 2.0

As well as the original SAML 1.1 standard, many SAML implementations now support the more modern SAML 2.0. Indeed, some more recent implementations have minimal or even no support for the older protocol.

Technical implementation and deployment profiles for SAML 2.0 operation within the UK federation can be found in [UKFTS]. Future editions of this document will include additional recommendations for the use of SAML 2.0.

For maximum compatibility, we recommend that all new deployments are made using software capable of both SAML 1.1 and SAML 2.0 operation. We believe that these protocols will exist in parallel for the foreseeable future.

4 Metadata

The federation publishes metadata describing participating entities. This metadata provides the information required for entities to know how to communicate with each other, and establishes a trust fabric permitting entities to verify each other's identities.

Note, however, that presence in the federation metadata alone should not be taken to imply particular behavioural guarantees. In particular:

- it is the responsibility of each identity provider to establish appropriate policies for attribute release based on their knowledge of individual service providers;
- it is the responsibility of each service provider to decide how much trust to place in the attributes presented by an identity provider based on their knowledge of the individual identity provider.

4.1 Standard Metadata

The federation's standard metadata is based on the metadata profile defined by the Shibboleth software. The Shibboleth profile is itself based on [SAML2Meta], [SAML1Meta-xsd] and [SAML1Meta], with additions defined in [ShibProt] section 3.4.

The additions and conventions defined for the federation metadata are described in detail in [UKFSTS].

The current federation standard metadata can be retrieved from the following location:

<http://metadata.ukfederation.org.uk/ukfederation-metadata.xml>

4.2 Metadata Refresh

The metadata published by the federation is regularly updated to include new entities, to describe changes to existing entities, and to remove old entities either because they have left the federation or because the entity has been reported as compromised.

Entities working with old copies of the federation metadata may therefore be unable to communicate with new federation members, be unable to communicate with members whose details have changed, and be vulnerable to attacks based on compromised entities. For these reasons, all federation members are strongly recommended to refresh the metadata used by their entities on a regular basis. A daily refresh operation should be regarded as normal.

Metadata refresh involves the following steps:

- retrieving the revised metadata from the publication location given above,
- verifying the authenticity of the revised metadata (see next section),
- replacing the metadata in use by the entity.

Most current federated identity software in use within the UK federation provides integrated functionality combining all three of these steps.

Users of other software may make use of the `xmlsectool` application provided by the OpenSAML project.⁵

4.3 Metadata Signature Verification

The security and reliable operation of each entity in the federation depends on using metadata which is both recent and authentic. The former requirement is met by regular metadata refresh, as described above; authenticity of the metadata is assured by verifying the digital signature on each downloaded metadata file before using the metadata.

The current signing certificate for the federation can be retrieved as a Base64-encoded X.509 certificate suitable for use with most current software from the following location:

`http://metadata.ukfederation.org.uk/ukfederation.pem`

Alternatively, the same certificate can be downloaded in the form of a Java Keystore file suitable for use with the `metadatatool` application from the following location:

`http://metadata.ukfederation.org.uk/ukfederation.jks`

Important note: the security of each federation member depends on the use of authentic metadata. In order to be sure that the metadata signature verification is being properly performed, it is first essential to verify that the correct signing certificate is being used in the verification operation. This may be achieved by checking the certificate's fingerprint "out-of-band", for example through a telephone call to the federation operator. It is *not* safe to assume that the certificate downloaded from the above location is itself authentic without performing this additional step.

We recommend that each federation member verifies the fingerprint of the federation signing certificate through a telephone call to the federation operator. The SHA-1 fingerprint of a Base64-encoded certificate can be obtained on many systems using the following command:

```
openssl x509 -noout -fingerprint -sha1 -in ukfederation.pem
```

The shorter but less secure MD5 fingerprint can be viewed by replacing "`-sha1`" in the above command with "`-md5`".

The SHA-1 and MD5 fingerprints of a federation signing certificate stored in a Java Keystore file may both be obtained using the following command:

```
keytool -list -v -keystore ukfederation.jks -alias ukfederation \
-storepass ukfederation
```

If you are unable to contact the federation operator directly, you can obtain a lower level of assurance as to the integrity of the downloaded certificate by comparing its fingerprints against the values given below. Note that the federation has used several certificates to represent the same key over time; any of these certificates are

⁵ See <https://wiki.shibboleth.net/confluence/display/SHIB2/XmlSecTool>

acceptable to the Shibboleth software, but the fingerprints do differ. Users of non-Shibboleth software may need to make use of the most recent version of the signing certificate.

The fingerprints for the version of the signing certificate in use from November 2012 are:

```
MD5: 3D:D8:EB:1B:89:6C:EA:D1:ED:39:FD:45:E1:5F:AD:74
SHA1: F9:7F:1A:1E:43:D3:D5:41:6D:C9:D5:0E:3B:6E:8F:5B:97:6C:4B:2E
```

The fingerprints for the version of the signing certificate in use from November 2010 to November 2012 are:

```
MD5: 91:76:33:AC:86:A3:21:D0:5E:8F:8A:E7:C1:2D:D7:D5
SHA1: 94:7F:5E:8C:4E:F5:E1:69:E7:DF:68:1E:48:AA:98:44:A5:41:56:EE
```

4.4 Federation URI

The following URI is used as the Name attribute of the outermost EntitiesDescriptor element in the federation metadata:

<http://ukfederation.org.uk>

This *federation URI* may be used to refer to the federation as a whole.

Use of the federation URI as a relying party identifier in attribute release policies (as is possible, for example, with the Shibboleth software) is NOT RECOMMENDED. Amongst other issues, use of the federation URI in this way assumes that the entity consumes an aggregate containing all of the federation metadata, and therefore prevents a transition to query-based metadata acquisition (see the Future Directions section for the Metadata Publication Service in [UKFTS]).

The federation URI should never be used in attribute release policies relating to personally identifying information (PII). For attributes which do not constitute PII, consider releasing to all authenticated relying parties instead.

4.5 Future Directions

4.5.1 Attribute Requirements for Service Providers

The SAML 2.0 metadata specification ([SAML2Meta]) enables the federation operator to publish details of the attributes used by service providers directly in the metadata. This will allow the service provider to announce the attributes it requires for basic operation and those it makes use of to provide additional services (e.g., user personalisation).

This facility is being investigated with a view to publishing the attribute requirements of service providers directly in the federation metadata.

4.5.2 Removal of Support for Signing Certificate in JKS Format

The UK federation has in the past published each version of the federation signing certificate both as a Base64-encoded PEM file and as a Java Keystore (JKS) file.

The JKS file is only required for *new* installations of Shibboleth 1.3 identity provider software; existing installations can continue to use the JKS file that was current at the time of deployment indefinitely.

Shibboleth 1.3 reached its end of life in June 2010, and no new installations of this software are anticipated (they are certainly NOT RECOMMENDED). Therefore, starting from the federation signing key recertification in late 2012, the JKS file variant of the signing certificate will no longer be provided.

5 Digital Certificates

5.1 Certificate Roles in the UK Federation

The protocols and profiles used within the UK federation make extensive use of X.509 certificates to carry the public keys used for various purposes. These certificates can be broken down into two classes according to their function.

5.1.1 Browser-facing Certificates

Browser-facing certificates are visible only to a user's browser. The browser-facing certificates are:

- The identity provider's SSL server certificate seen by browsers (for example, on a user login page),
- The service provider's SSL server certificate seen by browsers (for example, on actual site pages being protected by Shibboleth and at assertion consumer service endpoints),
- Any SSL server certificates seen by browsers during the discovery process (for example, on local discovery services or at institutional portals, see section 6 below).

Browser-facing certificates are not part of the federation's trust fabric; this means that they do not appear in federation metadata and that the federation operator is not normally aware of the certificates that you use in this role.

You can use certificates from any source as browser-facing certificates; the main constraint is that the issuer of the certificate (often a commercial Certification Authority, or CA) is accepted as trusted by the user's browser.

5.1.2 Trust Fabric Certificates

Trust fabric certificates are visible only to the identity provider and service provider software; they are never seen by the user's browser. Only certain certificate products are acceptable for this purpose; see below. The trust fabric certificates are:

- The certificate for the identity provider's XML signing key pair for SAML services,
- The certificate for the identity provider's SSL server key pair for SAML services,
- The certificate for the service provider's SSL client key pair for SAML services,
- The certificate for the service provider's XML signing key pair for SAML requests.

Trust fabric certificates are, as the name implies, part of the federation's trust fabric. This means that they appear in federation metadata and that you need to include information about your entity's trust fabric certificates in your entity registration.

Most modern SAML software generates a long-lived, self-signed trust fabric certificate during the installation process. We recommend that you use this automatically generated trust fabric certificate for your entity unless you have a specific reason not to do so.

The federation also accepts metadata registrations making use of trust fabric certificates issued by commercial or private certification authorities. Current information about the options available for trust fabric certificates is available from the UK federation web site.⁶

5.2 Recovery from Key Compromise

If any of the private keys associated with a federation entity leave the control of the owners of that entity, they should be regarded as permanently compromised. Should this happen, the following steps must be taken immediately:

- Any Certification Authorities that have acted as issuers for certificates associated with the compromised key should be notified. This will allow the CAs to revoke all affected certificates.
- The federation must be notified of the compromised key, and of all affected certificates. The federation will make an immediate announcement to the federation mailing list, and rebuild the federation metadata to temporarily exclude all affected entities.

Recovering from any system compromise is a complex process, often involving rebuilding the affected systems. After the system has been re-secured, it is then necessary for the compromised entity to at least generate a new key pair and have new certificates signed by the appropriate CA.

Because of the way certificates are handled in the Shibboleth software, recovery may also involve changing the DNS name of the affected entity before generating new certificates. Determination of the exact steps required will be made on a case-by-case basis by the federation operator for each compromise as it occurs.

⁶ See <http://www.ukfederation.org.uk/content/Documents/GetCertificate>

6 Discovery

The SAML protocols are most often used in a *service provider first* mode of operation, in which the user visits a site providing protected content before providing that site with an authentication assertion.

In this situation, the service provider must send the user to their identity provider bearing an authentication request message. The problem of correctly determining the identity provider to which the user should be sent is referred to as the *discovery problem*.

6.1 Avoiding Discovery: Institutional Portals

It is possible to avoid the discovery problem entirely by setting up an institutional portal which makes authentication requests to the organisation's identity provider on behalf of each selected service provider. Such a portal can greatly improve the user experience for members of an organisation with interest in a common set of resources; for example, the students enrolled in a particular class.

URLs suitable for use in such institutional portals can be easily captured by performing a service provider first access to a resource, following the process through to the identity provider and extracting the resulting URL from the browser's address bar. It is normally necessary to remove the `time` parameter from such a URL.

One disadvantage of this technique is that it does not adapt to changes in the service provider's configuration. It is recommended that identity providers wishing to make use of this technique make arrangements with the service providers concerned to be informed in advance of any changes that would affect them. It is also possible for service providers running Shibboleth 1.3 or later to set up Session Initiator locations to give identity providers building this kind of institutional portal a more stable, and significantly simplified, interface.

6.2 Discovery by the Service Provider

The discovery process can be completed by the service provider using a number of different approaches. For example, the Shibboleth 1.3 software sets a *discovery cookie* on each successful authentication; this can be used on subsequent visits by the client to suggest the most likely identity provider for that particular user. Other heuristics include comparing the user's IP address against a table of known IP address ranges for different institutions, and making the same resource available at different URLs for different client institutions.

A service provider may also make use of the federation metadata to display a list from which the user may select their identity provider; many service providers will be able to restrict this list to the identity providers for those institutions that are known to be clients of the service. This approach is sometimes referred to as a "local WAYF".

Performing the discovery process at the service provider is likely to provide a better user experience than the alternative of delegating the process to another entity with less knowledge of the specific service. Service providers are therefore recommended to consider implementing at least partial discovery whenever possible.

6.3 Federation Central Discovery Service

When discovery cannot be avoided through techniques such as institutional portals, and cannot be performed for whatever reason by an individual service provider, recourse may be made to a centralised discovery service provided by the federation. Such a discovery service is often referred to as a WAYF, because the question that it asks is simply: “Where Are You From?”

The federation provides a reliable discovery service hosted on multiple servers at geographically distributed co-location sites. Full details of the service can be found in [UKFTS].

Service providers can make use of the central discovery service through two separate protocols: the newer *Identity Provider Discovery Service Protocol and Profile* (the “DS protocol”) described in [IdPDisco], and the older “WAYF protocol” described in [ShibProt].

Use of the DS protocol is recommended if the service provider is capable of implementing it. In this case, the service provider should be configured to use the DS protocol with the following discovery location:

<https://wayf.ukfederation.org.uk/DS>

The legacy WAYF protocol is only recommended for service providers which can not implement the DS protocol. It is much less flexible, and in particular restricts the entities to use of the SAML 1.1 protocol, even when both of them are capable of using SAML 2.0.

To configure a service provider to use the legacy WAYF protocol, use the following URL:

<https://wayf.ukfederation.org.uk/WAYF>

Any authentication request sent to this location must:

- use the Shibboleth authentication request profile
urn:mace:shibboleth:1.0:profiles:AuthnRequest as defined in [ShibProt] section 3.1.1,
- contain a `shire` parameter referring to an assertion consumer service endpoint bound to the SAML 1.1 Browser/POST profile.

6.3.1 Deprecated Endpoints

The following endpoint location was originally implemented to allow service providers to specify that the user should be able to choose from a list containing all identity providers present in the federation metadata, instead of just those intended for production use:

<https://wayf.ukfederation.org.uk/all.wayf>

This functionality has now been incorporated into the central discovery service's user interface (in the form of a “Search over All Sites” link at the bottom of the page) so that it is now possible to access any identity provider from any service provider.

The behaviour of this endpoint is therefore now identical to that of the “**WAYF**” endpoint described above and its use is **NOT RECOMMENDED**.

6.3.2 Use of Undocumented Endpoints

Service providers **MUST NOT** use any endpoints at the central discovery service which are not listed above or in the corresponding section of [UKFTS]. In particular, endpoints derived from the transient locations shown in a browser's address bar **MUST NOT** be used with the CDS, as they are not guaranteed to remain operational.

6.4 Future Directions

6.4.1 Federation Central Discovery Service

Members of the federation can expect to see the central federation discovery service experience improve incrementally as experience is gained with the live federation deployment.

7 Attribute Usage

7.1 Core Attributes

A core set of attributes has been identified that identity providers are recommended to support, and that service providers should consider when setting attribute requirements. There are two reasons for making these recommendations:

- to advise identity providers of the attributes commonly required by service providers as a condition for authorising access – a failure to supply these attributes is likely to result in a refusal of service from some service providers;
- to advise service providers of the attributes which identity providers are likely to be willing to supply – some institutions may be unable to supply attributes other than those in the recommended set.

Attributes in the core set have been chosen to be versatile, and should be sufficient for the great majority of applications.

The following are defined as core attributes; their individual use is described in the subsections following:

- *eduPersonScopedAffiliation*. This attribute indicates the user's relationship (e.g., staff, student, etc.) with the organisation. For many applications, examination of this attribute is sufficient to determine whether the user has sufficient privilege to access the resource.
- *eduPersonTargetedID*. If a service provider is presented only with the affiliation of an anonymous subject, as provided by *eduPersonScopedAffiliation*, it cannot provide service personalisation or usage monitoring across sessions. These capabilities are enabled by the *eduPersonTargetedID* attribute, which provides a persistent user pseudonym, distinct for each service provider.
- *eduPersonPrincipalName*. This attribute is used where a persistent user identifier, consistent across different services, is required. It often corresponds to the user's single sign-on (SSO) name, and may be useful for securing both internal institutional services and external services where access control lists are used.
- *eduPersonEntitlement*. This attribute enables an organisation to assert that a user satisfies an additional set of specific conditions that apply for access to a particular resource. A user may possess different values of the *eduPersonEntitlement* attribute relevant to different resources.

The core attributes are defined in the *eduPerson* specification ([*eduPerson03*], [*eduPerson06*], [*eduPerson08*]) and are exchanged using the MACE-Dir Attribute Profile for SAML 1.x, as described in [*MACEAttr*], section 2. Further information on the use of each of these attributes is given below.

7.1.1 Security Domains (Scopes)

The first three of the core attributes are structured as *scoped* attributes, and share a common syntax: *local-part@security-domain*, where *local-part* is attribute-specific

and *security-domain* is a dotted string. The security domain contains a DNS name that the federation operator has verified is registered to the identity provider's owner (or, in the case of an outsourced identity provider, the identity provider's institutional client).

While the security domain has the appearance of a DNS name, it is not constrained to the semantics of a DNS name. In particular, for historical reasons the UK has issued pairs of DNS names to many institutions, and in DNS terms these are equivalent. For example, the University of Edinburgh has been issued both `edinburgh.ac.uk` and `ed.ac.uk`. As Shibboleth security domains, however, these names are distinct and cannot be used interchangeably. This is a potential source of configuration problems, which can be readily avoided if the organisation selects just one of its DNS names at all times (including when registering with other federations). This has no implications for the user interface, as security domains are used only in machine-to-machine exchanges. In our example, the University of Edinburgh has chosen to use the `ed.ac.uk` scope exclusively.

Institutions making use of outsourced identity providers are strongly recommended to use scopes based on domain names owned by themselves rather than names allocated by the identity provider of which they are a client. This allows for future flexibility in identity provision for the organisation: migration from one outsourced identity provider to another, or from an outsourced identity provider to in-house provision, is much more difficult when an organisation does not have control over its own scope.⁷

Institutions in the HE/FE sector are recommended to use their principal institutional domain name as their scope.

All schools in the UK have a `.sch.uk` domain name⁸ suitable for use as a scope. Note that it is not necessary for the school to be using this domain name on the web or elsewhere in order for it to be used as a scope: the only requirement is that the federation operator can be satisfied that the domain name is registered to the school in question.

Although the `.sch.uk` name is recommended for use in most cases, it may also be appropriate for some schools to be given names under the domain name of a Local Authority or Regional Broadband Consortium in order to leverage existing methods used by LAs and RBCs to uniquely identify schools.

The federation operator is responsible for verifying that a federation member is authorised to make assertions for each security domain it registers before adding this information to the federation metadata. Shibboleth service provider software ensures that the only values of security domain which can be asserted by an identity provider are those present in the federation metadata for that identity provider.

7.1.2 eduPersonScopedAffiliation

This attribute enables an organisation to assert its relationship with the user. This addresses the common case where a resource is provided on a site licence basis,

⁷ Migrating from one identity provider is not simple even when the scope can remain unchanged: in particular, values of `eduPersonTargetedID` are relative to the issuing entity, and would become invalid after any such migration without significant co-ordination between identity providers and service providers. SAML 2.0 introduces new functionality that may help to address this issue in the future.

⁸ See <http://www.nominet.org.uk/registrants/schools/>

and the only access requirement is that the user is a *bona fide* member of the organisation, or a specific school or faculty within it.

The attribute is multi-valued (that is, a user can have more than one value for the attribute), and is structured as a scoped attribute, with the form *affiliation@security-domain*, where *affiliation* is one of a number of prescribed categories of user. The concept of *security-domain* is as described above (often taken as institutional DNS name).

7.1.2.1 eduPersonScopedAffiliation in the HE/FE Sector

The following table identifies the permitted values of `eduPersonScopedAffiliation` and provides the recommended interpretation for them in UK higher and further education. In particular, it indicates which category of user is typically regarded as authorised to access licensed materials according to the relevant JISC Model Licence⁹.

Defined value	Authorised User	Notes
<code>student</code>	yes	Undergraduate or postgraduate
<code>staff</code>	yes	UK term for all staff
<code>faculty</code>	yes	US term to distinguish teaching staff
<code>employee</code>	yes	Other than <code>staff/faculty</code> (e.g., contractor)
<code>member</code>	yes	Comprises all the categories named above
<code>affiliate</code>	no	Relationship short of full <code>member</code>
<code>alum</code>	no	Alumnus (graduate)
<code>library-walk-in</code>	yes	General library privileges

In general, other categories of user such as Honorary Staff or Visiting Scholar, who are treated as members with normal institutional privileges, would be assigned the value `member`. The value `affiliate` is defined as applying to those with whom the organisation has some dealings, but to whom no set of general membership privileges are extended. This could be applied to those with a short-term association with the organisation which is less close than `member`. Whether an affiliate is considered an authorised user for a specific service may vary from case to case.

Where a computer identity is assigned to a walk-in user, the identity provider must ensure that the user is physically present on approved premises before providing any authentication assertions for that user. This may be accomplished by IP address checking or by any other means.

7.1.2.2 eduPersonScopedAffiliation in the Schools Sector

The following table identifies the permitted values of `eduPersonScopedAffiliation` and provides the recommended interpretation for them in the UK schools sector.

⁹ See <http://www.jisc-collections.ac.uk/Help-and-information/How-Model-Licences-work/NESLi2-Model-Licence-/>

Defined value	Notes
student	Pupil, student, learner
staff	All staff
faculty	Teaching staff
employee	Non-teaching staff
member	Comprises all the categories named above
affiliate	Relationship short of full member
alum	Alumnus (ex pupil)
library-walk-in	General library privileges

7.1.2.3 Generating and Interpreting eduPersonScopedAffiliation

Several values of eduPersonScopedAffiliation are regarded as being “contained” within other values: for example, the `student` value is contained within `member`.

It is recommended that identity providers have the ability either to maintain these multiple values for a given individual, or otherwise provide the ability to release either value as appropriate for a particular service provider. For example, although some service providers might require the release of the more specific `student` value, a different service provider that only requires the less specific `member` value should only be sent the less specific value. Releasing `student` in this case gives the service provider more information about the user than is required, raising privacy and data protection concerns.

Despite the recommendation above that identity providers should be conservative in what they send, service providers are recommended to be liberal in what they accept. For example, a service provider requiring `member` affiliation should also accept `student`, `staff`, etc. as alternatives.

Deployers of service provider entities should note that the precise meaning of the affiliation values may vary slightly from identity provider to identity provider. If this is not acceptable for the particular application being deployed and a precisely delimited category of users must be identified, a better solution may be the definition by the service provider of a custom value for the `eduPersonEntitlement` attribute (see below).

Members working with international partners should note that cultural differences mean that the `staff` and `employee` values may have different meanings to members of some other federations. The best solution to this issue is for service providers to avoid the use of these values when possible, and to explicitly confirm that they have a common understanding of the meaning of these values if they are required.

7.1.3 eduPersonTargetedID

Important note: the definition of the eduPersonTargetedID attribute has changed between [eduPerson03] and [eduPerson06], as have best practices surrounding its use. The recommendations here take the [eduPerson03] definition and the value encoding described under “Legacy name and Syntax” in [MACEAttr] section 2.3.2.1.2. However, recommendations are provided that should allow a smooth transition to the newer definition as appropriate.

A service provider may use eduPersonTargetedID to support aspects of its service that depend on recognising the same user from session to session. The most common use is to enable service personalisation, to record user preferences such as stored search expressions across user sessions. A secondary use is to enable tracking of user activity, to make it easier to detect systematic downloading of content or other suspected breaches of licence conditions.

The attribute enables an organisation to provide a persistent, opaque, user identifier to a service provider. For each user, the identity provider presents a different value of eduPersonTargetedID to each service provider to which the attribute is released. The attribute is defined as multi-valued (with one value for each service provider to which eduPersonTargetedID is released), though only a single value is ever released at a time. It is structured as a scoped attribute, with the form *pseudonym@security-domain*. The pseudonym is guaranteed to be unique within the context of the *security-domain*.

7.1.3.1 Generating eduPersonTargetedID

The eduPerson specification requires that a value of eduPersonTargetedID once assigned to a user for a given service provider shall never be reassigned to another user. Users and service providers should note, however, that not all identity providers may be able to guarantee that a user will always present the same value of eduPersonTargetedID; indeed, identity providers may offer their users the ability to generate new values of eduPersonTargetedID if they feel their privacy has been compromised.

There are two ways in which an identity provider may implement eduPersonTargetedID:

1. *Algorithmic*. This generates the pseudonym part of the eduPersonTargetedID value algorithmically from other attributes. This avoids the need for the identity provider to store the attribute value, as it can simply be regenerated dynamically as required.

This has the disadvantage (for the end user and the service provider) that the value will change if any of the source attributes or the algorithm employed changes. Consequently, any user personalisation data such as stored search expressions would be lost. The user would also be unable to alter or delete any previously registered service alert requests.

2. *Storage*. An alternative solution is to store all values of eduPersonTargetedID ever issued. When a new value is required, this database is checked to prevent reassignment. Current values of eduPersonTargetedID are stored with the corresponding user entry. This is the most reliable way to ensure that the constraint on reassignment of values of eduPersonTargetedID is satisfied.

7.1.3.2 Interpreting eduPersonTargetedID

Although eduPersonTargetedID as described here is structured as a scoped attribute, this approach presents future compatibility problems due to the change in definition of this attribute between [eduPerson03] and [eduPerson06], along with the different value encodings described in [MACEAttr] section 2.3.2.1 and its subsections.

In order to be forwards-compatible with the new definition of eduPersonTargetedID, service providers should always treat an eduPersonTargetedID value as a triple composed of the following components:

- the entity name of the identity provider that created the value (this is not contained in the scoped value, but can be determined in the context of the attribute assertion as a whole),
- the entity name of the service provider or group for which the value was created (again, not contained in the scoped value, but a property of the service provider itself),
- the opaque string value that forms the local-part of the scoped value.

Note that in this revised view, the security-domain part of the scoped value is *not* part of the eduPersonTargetedID value, although it will be closely related to the first component in many circumstances.

On receipt of a scoped eduPersonTargetedID value, a service provider may either use it in conjunction with the two implicit entity name components described above, or decompose it to retrieve the local-part, and then combine it with the other components to form a new value. In the latter case, it is recommended that the new value is formed by concatenating the following elements:

- the entity name of the identity provider
- a single ‘!’ character
- the entity name of the service provider
- a single ‘!’ character
- the opaque string value that forms the local-part of the scoped value

The resulting string will contain the same value as the Shibboleth 1.3 service provider software would present to an application on receipt of the SAML 2.0-based persistent identifier encoding of eduPersonTargetedID now recommended by [MACEAttr].

7.1.4 eduPersonPrincipalName

This attribute is used where a persistent user identifier, consistent across all services, is required and typically corresponds to the identifier which a user presents when authenticating to local institutional services (i.e., the user’s single-signon name or “netID”). The attribute is single-valued and structured as a scoped attribute, with the form *local-name@security-domain*. The *security-domain* component has the same semantics as the corresponding component in eduPersonScopedAffiliation. The *local-name* is guaranteed to be unique within the context of the *security-domain*.

It is recommended that a value of `eduPersonPrincipalName` previously associated with one individual should never be reassigned to another individual. Non-reuse may be assured by deriving `eduPersonPrincipalName` from a (non-repeating) staff number or student matriculation number, though care should be taken to ensure that any implicit information is not inadvertently leaked; for example, age may be encoded as part of the matriculation number. As in the case of `eduPersonTargetedID`, users and service providers should be aware that identity providers may not always be able to guarantee to present the same value of `eduPersonPrincipalName`.

7.1.5 `eduPersonEntitlement`

Values of `eduPersonEntitlement` take the form of a URI, most frequently using the “http” or “urn” schemes. For example:

`http://publisher.example.com/contract/GL123`

`urn:mace:ac.uk:sdss.ac.uk:entitlement:emol.sdss.ac.uk:restricted`

`http://ukfederation.org.uk/entitlements/example`

The meaning of a given value of `eduPersonEntitlement` is normally defined by a service provider. In the case of a value using the “http” scheme, it is recommended that the value resolve to a document giving the definition of the value. Having defined the meaning of the attribute value, the service provider then invites some or all identity providers to express that value for those users who satisfy the definition. In this way the service provider can delegate to the identity provider some or all of the responsibility for authorisation of access to a particular resource. Typically, this is used to assert entitlements over and above those enjoyed by other members of the organisation; for example, “Entitled to access the restricted material present in the Med123 resource”. In this case, the service provider trusts the organisation to verify that the user satisfies the (arbitrarily complex) authorisation conditions associated with the entitlement. This often involves an additional licence clause, where the organisation undertakes to assign the `eduPersonEntitlement` values according to agreed criteria.

Institutions are encouraged to consider the use of locally-defined values of `eduPersonEntitlement` to control access to local services. Such values are for internal use only, to model useful aspects of internal administrative operation, such as roles (e.g., “Member of the parking committee”) or specific authorisations (e.g., “Authorised to raise orders up to £1,000 in value”). Although the values are not released to external partners, a side-effect of using them should be to increase the trust an external service provider is likely to place in the identity and attribute assertions made by an organisation which relies on these same mechanisms for its internal administration.

7.1.5.1 Storing and Releasing `eduPersonEntitlement`

Because a particular value of `eduPersonEntitlement` often represents an entitlement to access a specific resource, identity providers should be capable of associating any number of entitlements with an individual user.

However, such entitlements may represent personal or even sensitive personal data about the individual. It is therefore important to control the release of individual values of `eduPersonEntitlement` closely, so that only service providers with a legitimate need for any given value of `eduPersonEntitlement` will have that value

released to them. For example, values defined by a particular service provider should normally only be released back to that same service provider.

7.2 Attributes, Privacy and Data Protection

UK data protection law and the normal institutional obligation to preserve user privacy both require that information identifying individuals only be exchanged when strictly necessary. For most applications the attributes eduPersonScopedAffiliation or eduPersonTargetedID should be sufficient. Since these do not permit identification of an individual they should not raise privacy or data protection concerns. Identity providers should therefore expect to provide one or both of these attributes in most circumstances; service providers should normally request only these and other privacy preserving attributes. Any exchange of eduPersonPrincipalName will require both parties to comply with the data protection principles set out in the Act.¹⁰

7.3 Subsidiary Attributes

The core attributes described here should be sufficient for most circumstances, and service providers are recommended to require only these attributes whenever possible in order to gain compatibility with the maximum number of identity providers.

However, it is recognised that it may become necessary for the federation to list small numbers of additional attributes that, while not likely to be implemented universally enough to be recommended as core attributes, are nevertheless of use to sufficient federation members for a standard definition to be useful. Such *subsidiary attributes* will be defined here; none are currently defined.

7.4 Sources of Additional Attributes

Where the core and subsidiary attribute groups defined by the federation do not meet the particular needs of regional and subject-based groups, it is possible for such groups to define and use their own attribute groups. In these cases, it is strongly recommended that service providers and identity providers make use of existing attribute definitions from the following sources before defining custom attributes:

- the eduPerson object class ([eduPerson06, eduPerson03]),
- the person and organizationalPerson object classes (X.521),
- the inetOrgPerson object class (RFC2798).

All attributes should be encoded according to the recommendations of [MACEAttr] sections 2.2 and 2.3.

Note that inclusion in the above list does not imply endorsement by the federation of the use of any specific attributes from the listed object classes. Federation members should carefully consider the privacy and data protection implications of any attribute definition before making use of it.

¹⁰ The Data Protection Act 1998, see: <http://www.legislation.gov.uk/ukpga/1998/29/contents>

7.5 Custom Attributes

The expectation for any newly invented attribute must be that it will not be widely implemented by members of the federation. It is therefore recommended that federation members only define new attributes as a last resort when no suitable definition exists elsewhere.

It is strongly recommended that any new attribute definitions follow the SAML attribute naming conventions of [MACEAttr] section 2.2, and the value encoding conventions of [MACEAttr] section 2.3.

Federation members should carefully consider the privacy and data protection implications of any newly invented attribute.

7.6 Working Without Attributes

Most Shibboleth service providers make authorisation decisions on the basis of a collection of attributes issued by the identity provider in respect of the authenticated user. It is, however, possible to authorise access for any user authenticated by a particular identity provider: any authentication statements from that identity provider are therefore given equal weight.

This authorisation model is recommended for use only when the service provider has specific assurance that the identity provider in question only issues authentication assertions for individuals acceptable to the service provider.

Authorisation without attributes is not recommended for general use within the federation, where:

- Institutional identity providers often provide identities for individuals who are only indirectly connected with the organisation, such as contractors.
- Many institutions may share the same outsourced identity provider.

Instead, scoped attributes such as eduPersonScopedAffiliation should be used to establish the individual's relationship with the organisation, and to distinguish between organisations making use of the same shared identity provider.

7.7 Future Directions

7.7.1 Unique Learner Number

The concept of a Learner Registration Service, with an accompanying Unique Learner Number (ULN), has been developed in the UK and is now undergoing operational trials.¹¹

The possibility of defining the Unique Learner Number as a subsidiary (but not core) attribute for the federation is under consideration.

7.7.2 New Definition for eduPersonTargetedID

The definition of eduPersonTargetedID has always been problematic due to the dependency of the value used on the identity of the service provider; values of eduPersonTargetedID are not expected to be stored along with other values in a

¹¹ See <http://www.miap.gov.uk/uniquelearnernumbers.htm>

conventional attribute store. To address this, the formal definition changed significantly between [eduPerson03] and [eduPerson06]; the new usage is clarified by [MACEAttr] section 2.3.2.1.

The recommendations presented in this document rely on the [eduPerson03] definition of eduPersonTargetedID, but if followed in full allow a smooth transition to the newer definition using either of the specified name and syntax combinations given in [MACEAttr].

This approach has been taken with regard to the current composition of the federation. It is likely that as federation members upgrade and use of the newer encoding of eduPersonTargetedID becomes more widely practical, that in turn:

- this document will change to emphasise the new terminology over the old;
- both forms of eduPersonTargetedID will be recommended for acceptance by service providers;
- both forms of eduPersonTargetedID will be recommended for generation by identity providers;
- finally, the [MACEAttr]-recommended form of eduPersonTargetedID may become the form recommended for use within the federation.

8 References

- [eduPerson03] Internet2 Middleware Architecture Committee for Education, Directory Working Group (MACE-Dir). *EduPerson Object Class Specification (200312)*. Document ID Internet2-mace-dir-eduPerson-200312. See <http://middleware.internet2.edu/eduperson/docs/internet2-mace-dir-eduperson-200312.html>
- [eduPerson06] Internet2 Middleware Architecture Committee for Education, Directory Working Group (MACE-Dir). *EduPerson Object Class Specification (200604)*. Document ID internet2-mace-dir-eduPerson-200604. See <http://middleware.internet2.edu/eduperson/docs/internet2-mace-dir-eduperson-200604.html>
- [eduPerson08] Internet2 Middleware Architecture Committee for Education, Directory Working Group (MACE-Dir). *EduPerson Object Class Specification (200806)*. Document ID internet2-mace-dir-eduPerson-200806. See <http://middleware.internet2.edu/eduperson/docs/internet2-mace-dir-eduperson-200806.html>
- [IdPDisco] OASIS Committee Specification, *Identity Provider Discovery Service Protocol and Profile*, March 2008. See <http://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-idp-discovery.pdf>
- [MACEAttr] S. Cantor et al. *MACE-Dir SAML Attribute Profiles*. Internet2 Middleware Architecture Committee for Education, Directory Working Group (MACE-Dir), April 2006. Document ID internet2-mace-dir-saml-attributes-200604. See <http://middleware.internet2.edu/dir/docs/internet2-mace-dir-saml-attributes-200604.pdf>
- [SAML1Meta] G. Whitehead and S. Cantor, *SAML 1.x Metadata Profile*. OASIS SSTC, March 2005. Document ID sstc-saml1x-metadata-cd-01. See <http://www.oasis-open.org/committees/security/>
- [SAML1Meta-xsd] S. Cantor et al., *SAML 1.x Metadata Profile Schema*. OASIS SSTC, March 2005. Document ID sstc-saml1x-metadata. See <http://www.oasis-open.org/committees/security/>
- [SAML2Meta] S. Cantor et al., *Metadata for the OASIS Security Assertion Markup Language (SAML) V2.0*. OASIS SSTC, March 2005. Document ID sstc-saml-metadata-2.0. See <http://www.oasis-open.org/committees/security/>
- [SAMLBind] E. Maler et al. *Bindings and Profiles for the OASIS Security Assertion Markup Language (SAML)*. OASIS, 2 September 2003. Document ID oasis-sstc-saml-bindings-1.1. See <http://www.oasis-open.org/committees/security/>
- [ShibProt] S. Cantor et al. *Shibboleth Architecture: Protocols and Profiles*. Internet2-MACE, September 2005. Document ID internet2-mace-shibboleth-arch-protocols-200509. See <https://wiki.shibboleth.net/confluence/download/attachments/2162702/internet2-mace-shibboleth-arch-protocols-200509.pdf>

- [UKFTS] *UK Access Management Federation for Education and Research:
Federation Technical Specifications.*
See <http://www.ukfederation.org.uk/>
- [UKPROC] *UK Access Management Federation for Education and Research:
Federation Operator Procedures.* Document ID ST/AAI/UKF/DOC/005.
See <http://www.ukfederation.org.uk/>
- [UKTRP] *UK Access Management Federation for Education and Research:
Technical Recommendations for Participants.* This document.
See <http://www.ukfederation.org.uk/>