UK Access Management Federation for Education and Research

Federation Technical Specifications

Ian A. Young
3 February 2012

Version 1.3

Table of Contents

1 Introduction

This document specifies the detailed technical architecture of the UK Access Management Federation for Education and Research (the UK federation).

Familiarity with this document is not normally required for individual deployments; its primary audiences are developers of federation software and operators of partner federations. A companion document, the *Technical Recommendations for Participants* ([UKTRP]), provides specific technical recommendations for members of the federation based on these specifications.

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 [UKFTS]); 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 UK 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 describes the UK federation with effect from its date of publication as shown on the cover page.

1.3 Notation

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in [RFC 2119].

Conventional XML namespace prefixes are used throughout this document to stand for their respective namespaces as follows:

Prefix	XML Namespace	Defined in
`ds:`	`http://www.w3.org/2000/09/xmldsig#`	[XMLSig]
`idpdisc:`	`urn:oasis:names:tc:SAML:profiles:SSO:idp-discovery-protocol`	[IdPDisco]
`md:`	`urn:oasis:names:tc:SAML:2.0:metadata`	[SAML2Meta]
`saml2:`	`urn:oasis:names:tc:SAML:2.0:assertion`	[SAML2Core]
`saml2p:`	`urn:oasis:names:tc:SAML:2.0:protocol`	[SAML2Core]
`ukfedlabel:`	`http://ukfederation.org.uk/2006/11/label`	This document.
`wayf:`	`http://sdss.ac.uk/2006/06/WAYF`	This document.

This document uses the following typographical conventions in text:

- `<prefix:XMLElement>` to signify an XML element. If the prefix is omitted, "`md:`" can be assumed.
- `XMLAttribute` to signify an XML attribute. Attributes accompanied by values are written as `XMLAttribute="value"`.

1.4 Changes in this Edition

- Added a "Notation" section to centralise the definition of conventional XML namespace prefixes, explain typographical conventions and introduce the RFC2119 requirements vocabulary.
- Change to using `<Element>` instead of `Element` to represent XML elements in text. This convention is also used by the SAML standards, and helps to distinguish elements from attributes.
- Adjust references to pages in the Shibboleth wiki, which has moved to a new domain.
- Removed documentation of `ukfedlabel` elements which are no longer present in published metadata (`<ukfedlabel:SDSSPolicy>`, `<ukfedlabel:DeletedEntity>`).
- Removed discussion of *Organization convention A*. Remove specific references by name to Organization convention B, as it is now the only convention in use.
- Added a new section defining the nature of aggregates published through the MPS.
- Document the use of the MIME media type for SAML metadata defined in [SAML2Meta].
- Signing certificate in JKS form will not be provided after the next signing key re-certification.
- Added a new section defining the Central Discovery Service.
- Added a new section defining a SAML V2.0 Browser SSO Implementation Profile.
- Added a new section defining a SAML V2.0 Browser SSO Deployment Profile.
- Updated [MDQuery] reference to the latest revision.
- Added new references required for the other new content.

1.5 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.

2 Trust Fabric

One of the roles of the metadata published by the UK federation is to allow the federation to act as a broker of technical trust between members. This is enabled by including `<KeyDescriptor>` elements for each entity, with each `<KeyDescriptor>` representing a credential (in the form of an RSA keypair) held by the entity.

A federation member registering an entity can include `<KeyDescriptor>` elements compatible with either or both of two independent trust mechanisms:

- The trust mechanism originally adopted by the UK federation refers to keys by name rather than by value. This mechanism depends on the use of X.509 certificates issued by a limited number of qualified certification authorities, along with PKIX path validation performed at run time.

- More recently, the UK federation has also supported the direct embedding of key values (in the form of X.509 certificates with any origin, containing the public key part of the credential) in entity metadata.

The PKIX-based trust mechanism, although still supported with a limited collection of qualified certification authorities, has not aged well:

- Embedded key material is required for some important SAML 2.0 features, such as XML encryption of SAML messages.

- The short lifetime of certificates issued by commercial certification authorities presents an additional maintenance workload for members and the federation helpdesk when those certificates must be embedded in federation metadata.

- PKIX validation in an inter-federation environment requires federations to accept partner federations' trust roots, resulting in large trust root collections. Experience with the very large collections of trust roots embedded in common browser software does not augur well for this approach.

- Commercial certification authorities have much less stability in terms of their certificate hierarchies than was previously believed, resulting in frequent dequalification of certificate products from the list supported by the UK federation.

These and other reasons have led to the PKIX-based trust mechanism falling out of favour internationally, and being gradually replaced in most environments by the direct embedding approach as defined in the [SAML2MIOP] specification.

2.1 Verifying Entity Credentials

There are a number of circumstances in which entities present credentials which must be verified by a relying party:

- Authentication responses issued by an IdP to an SP using the Browser/POST profile are signed using a credential which must then be verified by the SP. In this case, the SP locates the information required for the verification in the IdP entity's `<IDPSSODescriptor>`.

- During SOAP callbacks from the SP to the IdP (whether for attribute query or for artifact resolution) both the IdP and SP present credentials (normally through the TLS handshake) which must then be verified by the other party:
    - The SP locates the information required to verify the IdP's credential within the role descriptor element associated with the endpoint to which the callback is being made:
        - For attribute query callbacks, in the IdP entity's `<AttributeAuthority>`.
        - For artifact resolution callbacks, in the IdP entity's `<IDPSSODescriptor>`.
    - The IdP locates the information required to verify the SP's credential in the SP entity's `<SPSSODescriptor>`.

When a credential is to be verified, the first step is to collect the appropriate verification information, in the form of a set of `<KeyDescriptor>` elements, from the appropriate role descriptor. Note that in the case of an IdP, the `<IDPSSODescriptor>` and `<AttributeAuthority>` will usually contain the same set of `<KeyDescriptor>` elements, but that this should never be assumed. Only the `<KeyDescriptor>` elements from the role descriptor associated with the particular endpoint in use should be considered.

For verification purposes, all `<KeyDescriptor` elements with an explicit `use="encryption"` attribute should now be discarded. If no `<KeyDescriptor` elements remain, the verification has failed. UK federation metadata will normally contain, within each role descriptor, at least one `<KeyDescriptor` element whose use includes `signing` either explicitly or implicitly through an absent `use` attribute.

For compatibility reasons, `<KeyDescriptor>` elements in IdP role descriptors will always include explicit `use` attributes in UK federation metadata. However, this should never be assumed by software and the case of an omitted `use` attribute should always be handled correctly by regarding the credential within the `<KeyDescriptor>` as valid for both signing and encryption purposes.

`<KeyDescriptor>` elements in SP role descriptors may or may not include explicit `use` attributes; again, no assumption about the presence of an explicit `use` attribute should be made by software relying on UK federation metadata.

Verification against the set of `<KeyDescriptor>` elements associated with an entity acting in a particular role can succeed if verification against any of the `<KeyDescriptor>` elements succeeds: a failure to verify requires that verification against every appropriate `<KeyDescriptor>` elements fails independently. One implication of this is that software is at liberty to perform tests against the set of `<KeyDescriptor>` elements in any order; one performance optimisation would be to cache information about which `<KeyDescriptor>` was successfully verified during a previous operation.

[SAML2Meta] defines the `<KeyDescriptor>` element as always containing a single `<ds:KeyInfo>` element, but goes into no more detail. UK federation metadata supports two alternative models of credential verification:

- If the entity's credential can be verified using direct key trust verification, the `<ds:KeyInfo` will contain one or more `<ds:X509Data>` elements, each of which will contain exactly one `<ds:X509Certificate>` element.

- If the entity's credential can be verified using PKIX trust verification, the `<ds:KeyInfo>` will contain one or more `<ds:KeyName>` elements.

Each `<KeyDescriptor>` in UK federation metadata may support one of the verification models, or it may support both (when the certificate embedded in metadata could also be verified against the federation's PKIX trust roots). As with the set of `<KeyDescriptor>`s, verification against a single `<KeyDescriptor>` succeeds when verification can be performed against either of the available models; failure to verify a credential under one model has no significance if it can be verified under the other model. Similarly, when more than one alternative is available under a given model within a particular `<KeyDescriptor>`, all alternatives must be exhausted before verification against that particular `<KeyDescriptor>` should be regarded as having failed.

As with multiple `<KeyDescriptor>` elements, one implication of this is that the information within an individual `<KeyDescriptor>` may be considered in any order without affecting the outcome. We recommend, however, verifying a `<KeyDescriptor>` (or all available `<KeyDescriptor>`s, when appropriate) using the direct key scheme first before falling back to the PKIX scheme, which has a much higher computational burden due to the requirement to verify potentially long chains of certificates.

2.1.1 Verification using the Direct Key scheme

See:

- Shibboleth 2 implementation:
  https://wiki.shibboleth.net/confluence/display/SHIB2/ExplicitKeyTrustEngine

- Shibboleth 1 implementation:
  https://wiki.shibboleth.net/confluence/display/SHIB/BasicTrustEngine

The direct key verification scheme corresponds to the [SAML2MIOP] *SAML V2.0 Metadata Interoperability Profile*. This means that an X.509 certificate embedded in metadata is treated as a convenient wrapper for a cryptographic public key, with none of the additional semantics associated with X.509 certificates. In particular, such a certificate is not subject to PKIX path validation or to checks against its expiry.

The [SAML2MIOP] profile requires that all runtime decisions are made solely on the basis of key comparisons. One way to perform such checks is to extract the public key from the metadata certificate and compare it against the key extracted from the certificate presented by the claimant (after, of course, verifying that the claimant has cryptographically demonstrated its possession of the corresponding private key). However, in some circumstances a performance optimisation is available by comparing the certificate presented by the claimant directly against the certificate included in metadata, as these will frequently be identical. However, failure of such a comparison has no significance but to signal that key extraction and direct key comparison will be necessary.

[SAML2MIOP] allows keys to be represented using either `<ds:X509Certificate>` or `<ds:KeyValue>` elements. At present, UK federation metadata does not make use of `<ds:KeyValue>`. It is however possible that `<ds:KeyValue>` elements may be introduced at a later date and developers are recommended to implement support for this as part of support for [SAML2MIOP].

UK federation metadata currently contains only RSA public keys, and support of other public key cryptosystems (such as elliptic curve cryptosystems, or DSA keys) is not envisaged in the near future.

2.1.2 Verification using the PKIX scheme

See:

- Shibboleth 2 implementation:
  https://wiki.shibboleth.net/confluence/display/SHIB2/PKIXTrustEngine

- Shibboleth 1 implementation:
  https://wiki.shibboleth.net/confluence/display/SHIB/ShibbolethTrustEngine

The PKIX verification scheme is a profile developed for the Shibboleth software which relies on PKIX path validation from an end entity certificate presented by the claimant to a "key authority" declared in the metadata. This scheme has never been formally standardised, but is intended to be similar in broad outline to X.509 certificate handling as performed in other contexts.

One result of the lack of a formal specification for this validation scheme is that although the documentation referred to above may be of assistance, the final test of compatibility with the PKIX scheme is to demonstrate interoperability against a selection of deployments of the Shibboleth software.

Validation succeeds if all of the following are true:

- the claimant demonstrates possession of the private key corresponding to the public key contained in the presented certificate

- PKIX path validation can be performed from the end entity certificate to one of the federation's key authorities

- one of the `<ds:KeyName>` elements associated with the entity acting in the appropriate role matches the presented certificate

`<ds:KeyName>` values may match in a number of different ways. The most common is a direct match against the CN component of the presented certificate's DN, but others are also possible (see the references above to the Shibboleth trust engine implementations).

3 Metadata Usage and Extensions

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.

The federation's standard metadata format 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. These standards leave the meaning of some constructs undefined to allow flexibility, and allow extensions to the metadata to be defined to meet unforeseen requirements. This document therefore specifies the UK federation's particular uses of the standardised constructs, and documents the extensions to the standards which are used in the federation's published metadata.

3.1 UK Federation Label Namespace

The following XML namespace is defined for use in UK federation metadata:

http://ukfederation.org.uk/2006/11/label

The conventional prefix used for this namespace is "`ukfedlabel`".

All elements defined in this namespace will take the form of simple labels which are either present or absent in a particular context. Labels may be either XML elements (with or without attributes) or simple attributes.

An XML Schema document describing the label namespace is available through the federation helpdesk. Only those elements of this namespace which appear in metadata published by the UK federation are described here.

Note that although the identifier for the label namespace contains its date of definition, additional elements may be added to this namespace at any time.

3.1.1 UK Federation Member Label

If an entity is owned by a member in good standing of the UK federation, the following element will be added to the `<Extensions>` element of the entity's `<EntityDescriptor>` element:

```
<ukfedlabel:UKFederationMember/>
```

The presence of this element indicates that the owner of the entity has agreed to be bound by the UK federation's Rules of Membership [UKROM].

3.1.2 Accountable Users Label

The federation's Rules of Membership allow for a member to assert to the federation operator that a given identity provider entity provides for user accountability (see [UKROM] section 6.1). A member making such an assertion must comply with all the requirements of section 6 of the Rules.

If such an assertion has been made to the federation operator in respect of an entity, the following element will be added to the `<Extensions>` element of that entity's `<EntityDescriptor>` element:

```
<ukfedlabel:AccountableUsers/>
```

Note that the assertion of user accountability is made by the federation member alone; it is not verified by the federation operator.

3.2 SDSS Federation WAYF Namespace

UK federation metadata currently makes use of an XML namespace originally defined by the SDSS federation:

http://sdss.ac.uk/2006/06/WAYF

The conventional prefix used for this namespace is "`wayf`".

This namespace is used solely to label identity provider entities in order to hide them from the normal (filtered) federation "Where Are You From" (WAYF) service. This is done by adding the following element to the `<EntityDescriptor>`'s `<Extensions>` element:

```
<wayf:HideFromWAYF/>
```

The different central federation WAYF services are described in section 6.3 of [UKTRP].

3.3 `<EntityDescriptor>` Element

3.3.1 `ID` Attribute

Each `<EntityDescriptor>` element is given a unique `ID` attribute, formed by concatenating the two letters "`uk`" and six decimal digits, such as "`uk000123`". This attribute value is used as a name for the individual `<EntityDescriptor>` by the federation operator as part of the operational procedures of the federation.

During the transition from the SDSS federation to the UK federation, it was always the case that:

- Entities which appeared in both the SDSS federation metadata and the UK federation metadata had `ID` attribute values of `uk000199` or lower.
- Entities which only appeared in the UK federation metadata had `ID` attribute values of `uk000200` or higher.

This numerical convention will not necessarily be observed in the future, although present practice is to give all new entities `ID` attribute values of `uk000200` or higher.

3.4 `<Organization>` Element

The SAML 2.0 Metadata specification defines the `<Organization>` element as specifying "basic information about an organization responsible for a SAML entity or role" ([SAML2Meta], section 2.3.2.1). Its mandatory child elements are:

- `<OrganizationName>`, containing a name that "may or may not be suitable for human consumption"
- `<OrganizationDisplayName>`, containing a name "suitable for human consumption"

- `<OrganizationURL>`, containing a URL specifying "a location to which to direct a user for additional information".

Many SAML federations make use of `<OrganizationDisplayName>` as a convenient location from which to draw a string identifying a particular identity provider. This string is used when selection from a list of identity providers is required: for example this might be done at a central discovery service, often known as a WAYF ("Where Are You From") service.

This convention is unremarkable in an environment where a one-to-one mapping exists between organisations and identity providers, so that the organisation "responsible for" the SAML entity is the same (singular) organisation for which the identity provider speaks. Because the UK federation allows both outsourcing and aggregated identity provision, different conventions are adopted.

Firstly, all entities are provided with an `<Organization>` element in which the `<OrganizationName>` contains a string representing the UK federation's canonical name for the member organisation responsible for the entity. This will normally be the organisation's legal name, as taken for example from the organisation's constitution or from Companies House records.

Secondly, the `<OrganizationDisplayName>` contains a string describing the function of the particular entity, and the `<OrganizationURL>` contains a URL leading to more information as appropriate to the entity's function.

For an identity provider entity:

- The `<OrganizationDisplayName>` should contain the string by which the identity provider is to be known by discovery services.
    - In the case of identity providers representing a single member organisation, this will normally be a simplified form of the canonical name of that member organisation, selected by the federation operator to provide users of discovery services with a coherent selection.
    - In the case of an aggregated identity provider representing multiple member organisations, the `<OrganizationDisplayName>` will be chosen by the federation operator to represent the combined identity community.
- The `<OrganizationURL>` contains a URL leading to either more information about the organisation responsible for the entity, or more information about the identity community served by the entity.

For a service provider entity:

- The `<OrganizationDisplayName>` will be descriptive of the particular service provided. This may include a component representing the organisation offering the particular service.
- The `<OrganizationURL>` contains a URL leading to either more information about the organisation responsible for the entity, or more information about the service provided by the entity.

In the case where member organisation A entrusts the operation of one of its entities to a second member organisation B (or, alternatively, where A purchases services from B):

- The `<OrganizationName>` will refer to member B.
- The `<OrganizationDisplayName>` will refer to member A.
- The `<OrganizationURL>` may refer to either A or B, as appropriate in the particular case.

3.5 Future Directions

3.5.1 SDSS Federation WAYF Namespace

The use of the SDSS federation WAYF namespace will be discontinued at some point. The SDSS-defined `<HideFromWAYF>` marker element will be replaced by a new element in the UK federation label namespace.

3.5.2 `<Organization>` Conventions

The move to the `<Organization>` conventions described here was intended to bring the UK federation metadata into closer conformance with the original SAML 2.0 metadata specification ([SAML2Meta]). In particular, now that the change has been completed metadata consumers have a reliable indication (in the form of the `<OrganizationName>` element) of the organisation responsible for any given entity.

This conformance could be improved still further by making use of the SAML 2.0 `<AttributeConsumingService>` element to describe services. This element specifically includes both `<ServiceName>` and `<ServiceDescription>` child elements, which could be used in place of the `<OrganizationDisplayName>` element for service provider entities.

Such an alternative is unfortunately not available within [SAML2Meta] for identity provider entities. In addition, any move to a UK federation-defined alternative convention for the "WAYF display string" would need to be promulgated well in advance to avoid disruption to any existing WAYF deployments, not all of which can be assumed to be known to the federation operator.

4 Metadata Publication Service

The UK federation makes metadata available to participants and other partners through its Metadata Publication Service, or MPS.

4.1 Service Implementation

The MPS is implemented using a number of distinct physical computers in multiple geographic locations. At present, up to five computers are in use across two locations, but these details are subject to change without notice to allow for service scaling and maintenance.

The service is accessed through the DNS name `metadata.ukfederation.org.uk`, which resolves to both IPv4 and IPv6 addresses (A and AAAA records) for each machine. These DNS records have a low time-to-live value (currently 5 minutes) to allow rapid reconfiguration of the service to be performed.

4.2 Service Interface

The MPS makes available a number of defined *aggregates,* or aggregated metadata documents. Each of these aggregates may be retrieved using a standard HTTP GET method, as defined in [RFC2616] section 9.3.

A MIME media type of `application/samlmetadata+xml` is reported for all aggregates, as required by [SAML2Meta] appendix A.[1]

The most important of these aggregates is the *production aggregate,* which is located at the following URL:

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

The production aggregate is intended to be used by all federation participants under normal circumstances.

From time to time, it is necessary to make significant changes to either the format or content of the production aggregate. To allow testing of such changes before they are implemented in the production aggregate, a *test aggregate* is maintained alongside it at the following URL:

http://metadata.ukfederation.org.uk/ukfederation-test.xml

The test aggregate is re-signed and re-published in the same way and at the same times as the production aggregate. This is intended to allow sites wishing to make use of the test aggregate to use it as a direct replacement for the production aggregate without loss of functionality or timeliness. However, as the test aggregate may be used to test experimental features, it is not recommended for long-term use by production deployments.

Although the test aggregate is usually composed of metadata for the same entities as the production aggregate, it may from time to time include additional entities of an experimental nature.

1 Prior to 2011-12-06, all published aggregates used the less specific `application/xml` media type. The test and export aggregates were changed to use the standard media type on 2011-12-06. The production aggregate was changed to use the standard media type on 2012-01-15. At the time of writing, the fallback aggregate is scheduled to be changed to use the standard media type on 2012-02-19.

Features initially introduced for testing purposes in the test aggregate are periodically migrated into the production aggregate. In most cases, because notice is usually given to allow participants to verify these features through the test aggregate, no problems are encountered at this stage. However, the MPS also maintains a *fallback aggregate* to cover transitional problems, located at the following URL:

http://metadata.ukfederation.org.uk/ukfederation-back.xml

The fallback aggregate is composed of metadata for the same entities as the production and test aggregates, but omits features that have been only recently introduced to the production aggregate. The delay in introducing new features, normally of around one month, provides a temporary solution for problems which were not detected through use of the test aggregate.

Like the test aggregate, the fallback aggregate is *not* intended for long-term use by production deployments. Use of the fallback aggregate should always be temporary, and should always be notified to the federation helpdesk.

Use of any other aggregates published by the MPS is not supported.

4.3 Support for Conditional GET

The large aggregate metadata documents provided through the MPS are normally signed and re-published once every working day. Client software accessing the service more frequently then this may therefore end up repeatedly downloading and re-processing large quantities of redundant information.

To allow clients to optimise their behaviour, the service returns both a last modified date and a strong entity tag value, and supports the use of these values with the HTTP conditional GET mechanism described in [RFC2616] section 9.3.

For example, a successful initial fetch of one of the UK federation's published aggregate documents might result in the following HTTP response headers, amongst others:

```
HTTP/1.1 200 OK
Date: Tue, 29 Jun 2010 15:53:36 GMT
Last-Modified: Mon, 28 Jun 2010 17:58:54 GMT
ETag: "9de907-dfb7f380"
Content-Length: 10348807
Content-Type: application/samlmetadata+xml
```

The entity tag and last modified date values returned as part of this initial response could be used as part of a later conditional GET by including the If-None-Match and If-Modified-Since headers in the request:

```
GET /ukfederation-metadata.xml HTTP/1.1
Host: metadata.ukfederation.org.uk
Accept: */*
If-None-Match: "9de907-dfb7f380"
If-Modified-Since: Mon, 28 Jun 2010 17:58:54 GMT
```

Note that as described in [RFC2616] section 13.3.4, both of these headers should always be sent in a conditional GET to the MPS, as both values were provided to the client in the original response. The entity tag value *must* always be sent.

If the requested document has not changed since the initial request, the response headers resulting from this later request might include the following:

```
HTTP/1.1 304 Not Modified
Date: Tue, 29 Jun 2010 15:59:19 GMT
Server: Apache/2.2.3 (Unix) mod_ssl/2.2.3 OpenSSL/0.9.7d
ETag: "9de907-dfb7f380"
```

Here, the 304 status code indicates that the document has not been modified; in this case, the response body will be omitted.

It is recommended that, where possible, client software designed to access the MPS makes use of conditional GET requests as described here in order to minimise both local processing and load on the service.

4.4 Aggregate Specification

All metadata aggregates published through the MPS conform to the profile described by the following sections.

4.4.1 Aggregate Structure

Aggregate documents published by the MPS currently have a simple "flat" structure in which all `<EntityDescriptor>` elements in the aggregate are directly contained within a single `<EntitiesDescriptor>` document element.

Metadata consumers MUST however be capable of processing aggregates containing nested `<EntitiesDescriptor>` elements, as described in [SAML2Meta] section 2.

4.4.2 Aggregate Signature

The `<EntitiesDescriptor>` document element of a UK federation metadata aggregate is digitally signed using a 2048-bit RSA key called the UK federation metadata signing key. The signing key is published in the form of an X.509 certificate referred to as the UK federation metadata signing certificate.

Metadata consumers MUST verify an aggregate's signature against this key and MUST reject an aggregate whose signature cannot be verified. This acts as a protection against attacks in which consumers are provided with fabricated metadata.

Verification of the signature against the signing key SHOULD be performed by direct key comparison as described in [SAML2MIOP]. For the benefit of software which cannot implement [SAML2MIOP] and requires the signing certificate to be taken into consideration, the signing key is *re-certified* from time to time and re-published as a new signing certificate.

The current UK federation signing certificate can be retrieved in Base64-encoded form from the following location:

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

The fingerprints for the version of the signing certificate in use from November 2010 are as follows:

```
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
```

The fingerprints for the version of the signing certificate in use from November 2008 to November 2010 are:

```
MD5:  8E:B3:09:4E:FC:73:83:64:D1:7D:05:74:CA:6A:FF:10
SHA1: D0:E8:40:25:F0:B1:2A:CC:74:22:ED:C3:87:04:BC:29:BB:7B:9A:40
```

The fingerprints for the version of the signing certificate in use from November 2006 to November 2008 are:

```
MD5:  4B:A8:51:42:71:66:76:F7:CD:1B:2D:3F:32:B3:B2:2A
SHA1: BB:F4:CE:85:7A:BC:8C:7F:5B:44:8F:FE:39:4C:25:BE:EC:B9:08:B4
```

4.4.3 Aggregate Validity

The document `<EntitiesDescriptor>` of a UK federation metadata aggregate includes a `validUntil` attribute defining the last instant during which the aggregate should be considered valid. The validUntil attribute's value is set at the time of construction of the aggregate to allow a "validity interval" of a certain number of days after the aggregate's construction. This interval acts as a protection against certain attacks involving replay of old federation metadata containing compromised information.

Metadata consumers SHOULD reject metadata aggregates lacking a `validUntil` attribute and MUST discard aggregates whose `validUntil` instant has passed.

In normal operation, the validity interval used for UK federation metadata aggregates is 14 days. This may be varied in either direction for operational reasons, but until further notice will never be less than 7 days nor more than 28 days.

4.5 Future Directions

4.5.1 Compressed Metadata Service

SAML metadata, as an XML document format, tends to be bulky but repetitive. One result of this is that most large SAML metadata documents are capable of being compressed at roughly a 10:1 ratio.

The MPS will be enhanced to allow metadata clients to request delivery of the compressed form of published metadata. This will allow a large reduction in the amount of data a compatible client needs to transfer. This obviously benefits the individual client while improving the scaleability of the central service.

This enhancement would be provided through use of the HTTP content coding system as described in [RFC2616] section 3.5, with at least "`gzip`" and "`deflate`" compression schemes supported.

It is recommended that client software designed to access the MPS should support at least the "`gzip`" content encoding. Clients indicate which encoding types they support by means of the Accept-Encoding header within the GET request.

4.5.2 Query-Based Metadata Service

The current MPS provides metadata for all entities known to the UK federation within a single, large, aggregate document. This has the advantage of simplicity. However, entities participating in SAML federation do not, in general, require continuous access to metadata for all possible communication partners and in most cases the overwhelming majority of metadata downloaded by clients of the MPS lies unused by the consuming entity.

One way of reducing the burden on both individual MPS clients and on the service itself is to add a second publication method through which an MPS client can request only those individual entity-level metadata documents for which it has an immediate need.

Such a metadata publication protocol is currently being standardised (see [MDQuery]), and initial implementations of compatible publication servers and client software are expected to be available on an experimental basis in 2012, at which stage the technology will be evaluated for use within the UK federation.

4.5.3 Export Aggregate

The MPS currently publishes one further aggregate over and above those supported as part of the service interface. This is the *export aggregate,* located at the following URL:

http://metadata.ukfederation.org.uk/ukfederation-export.xml

The export aggregate functions as a testbed for experiments involving the exchange of metadata between the UK federation and other partner federations.

At present, the contents of the export aggregate are derived from a specially selected subset of the entities whose metadata is published as part of the normal aggregates. The format and contents of the export aggregate are subject to change without notice during the experimental phase.

A production service based on inter-federation metadata exchange will be specified should the experimental phase come to a successful conclusion. Such a production service would be likely to be at least initially based on offering entity owners the opportunity of opting in to such an exchange mechanism.

In the longer term, however, the contents of the export aggregate may be based instead on all entities from the normal aggregates which meet appropriate technical eligibility criteria. One likely requirement is that entities included in the export aggregate include embedded key material, so that they can participate in trust fabrics independent of the UK federation's selection of PKIX trust roots.

4.5.4 Aggregate Structure

In order to support future inter-federation metadata exchange, the UK federation metadata aggregates will transition from the “flat” aggregates described above to a “hierarchical” structure. This will allow those entities registered by UK federation members to be separated from those entities imported from other registrars in order to preserve the semantics of attribute release based on relying parties named by the federation URI.

At the time of publication of this document, both the hierarchical aggregate structure and the presentation of a selection of entities imported from partner federations are being evaluated within the test aggregate.

4.5.5 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, provision of the JKS format file is not regarded as being part of the MPS interface defined here, and starting from the next federation signing key recertification in late 2012 the JKS file variant of the signing certificate will no longer be provided.

5 Central Discovery Service

In single sign-on transactions where the user approaches the service provider first, *discovery* is the process by which the appropriate identity provider for the transaction is determined.

Although discovery is best performed by the service provider itself, the UK federation also makes a central discovery service (CDS) implementation available to participants for their use. For historical reasons, this service is often referred to informally as the federation "WAYF", an acronym for "Where Are You From".

5.1 Service Implementation

The CDS is implemented using a number of distinct physical computers in multiple geographic locations. At present, five computers are in use across two locations, but these details are subject to change without notice to allow for service scaling and maintenance.

The service is accessed through the DNS name `wayf.ukfederation.org.uk`, which resolves to both IPv4 and IPv6 addresses (A and AAAA records) for each machine. These DNS records have a low time-to-live value (currently 5 minutes) to allow rapid reconfiguration of the service to be performed.

5.2 Service Interface

5.2.1 Supported Discovery Protocols

The CDS supports two different discovery protocols: a simple "WAYF protocol" based on the Shibboleth authentication request profile described in [ShibProt], and the more modern and functional "DS protocol" as defined in [IdPDisco].

5.2.1.1 WAYF Protocol

The operation of the "WAYF protocol" is defined in section 2.3 of [ShibProt]. In this protocol, a service provider redirects the user agent to a discovery endpoint with query parameters matching those used by the Shibboleth authentication request profile (`urn:mace:shibboleth:1.0:profiles:AuthnRequest`) as described in section 3.1.1 of [ShibProt].

Once the appropriate identity provider has been identified, the WAYF redirects the user agent to an SSO service endpoint derived from the metadata for the selected identity provider. This has the effect of relaying the original authentication request message to the selected identity provider without the service provider's further involvement or knowledge of the selection.

Note that in this protocol the authentication request message contains the assertion consumer service location for the authentication response from the identity provider. This means that the response location (and implicitly the binding or bindings associated with that location in `<AssertionConsumerService>` metadata elements) must be chosen by the service provider before discovery has been performed: that is, before the capabilities of the selected identity provider are known.

To avoid unexpected failures being presented to the user, the `shire` parameter MUST refer to an assertion consumer service location which is bound to the SAML 1.1 Browser/POST profile (`urn:oasis:names:tc:SAML:1.0:profiles:browser-post`).

The WAYF protocol's limitations are sufficient that it is NOT RECOMMENDED for new service provider deployments. Instead, the DS protocol described below SHOULD be used if supported by the service provider software being deployed.

5.2.1.2 DS Protocol

The *Identity Provider Discovery Service Protocol and Profile* ("DS protocol") is defined in [IdpDisco]. Use of this protocol is RECOMMENDED for all new service provider deployments.

Whereas in the WAYF protocol the result of the discovery process is a message relayed to the selected identity provider, in the DS protocol the result of the discovery process is a message returned to the service provider indicating the selected identity provider in terms of its entity ID. This means that the service provider can select the appropriate protocol and profile to use with the particular identity provider rather than being forced to take a "lowest common denominator" approach. In particular, the DS protocol is SSO protocol agnostic and therefore allows the use of both SAML 1.1 and SAML 2.0 profiles rather than being limited to SAML 1.1.

A secondary advantage of this protocol is that problems arising from any mismatch between the profiles supported by the identity provider and the service provider are detected at the service provider. This allows more suitable error messages to be generated than is the case when the CDS is responsible for error reporting.

Note that any service provider making use of the CDS with the DS protocol MUST declare appropriate `<idpdisc:DiscoveryResponse>` elements in its metadata.

5.2.2 Supported Service Endpoints

The following sections describe the service endpoints supported by the CDS. Service providers MUST NOT use any endpoints at the CDS which are not listed below. 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.

5.2.2.1 Production Endpoints

Service providers capable of implementing the DS protocol SHOULD use the following discovery endpoint with the DS protocol:

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

Service providers not capable of implementing the DS protocol MUST use the following discovery endpoint with the WAYF protocol:

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

5.2.2.2 Test Endpoints

The following endpoints are maintained as alternative discovery endpoints:

https://wayf.ukfederation.org.uk/DS-test

https://wayf.ukfederation.org.uk/WAYF-test

In normal operation, they have the same functionality as defined above for the similarly named production endpoints. From time to time, however, they will be used as ways to expose the next generation of CDS implementation for testing purposes.

The test endpoints SHOULD NOT be used by production service providers except when actively testing next-generation discovery systems.

5.2.2.3 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.

5.3 Future Directions

5.3.1 Deprecated Endpoints

Discovery service endpoints listed above as deprecated may be removed from the service definition at some point in the future.

6 SAML V2.0 Browser SSO Implementation Profile

This profile specifies behaviour and options that implementations of the SAML V2.0 Web Browser SSO Profile [SAML2Prof] are required to support. It is layered on, and supplements, the InCommon SAML V2.0 Browser SSO Deployment Profile [ICSAML2].

Compliance with this profile is RECOMMENDED for all SAML products intended for use within the UK federation.

Although the UK federation does not mandate compliance with this profile as a requirement for deployment, software which does not comply with this profile may not interoperate with a significant proportion of other entities and deployment of such software is therefore NOT RECOMMENDED.

Implementations MUST comply with all normative requirements of [SAML2Prof], as modified by the Approved Errata [SAML2Err].

Implementations MUST comply with all normative requirements of the InCommon SAML V2.0 Browser SSO Implementation Profile [ICSAML2], except that for the time being the following requirements are relaxed:

- support of the use of the "`ETag`" header for metadata cache management is strongly RECOMMENDED
- support of the Identity Provider Discovery Service Protocol Profile in conformance with section 2.4.1 of [IdPDisco] is strongly RECOMMENDED

Implementations SHOULD include support for all non-normative recommendations of [ICSAML2].

7 SAML V2.0 Browser SSO Deployment Profile

This profile provides requirements and recommendations to deployers of the SAML V2.0 Web Browser SSO Profile [SAML2Prof]. It is layered on, and supplements, the following profiles:

1. InCommon SAML V2.0 Browser SSO Deployment Profile [ICSAML2]
2. Interoperable SAML 2.0 Web Browser SSO Deployment Profile [SAML2Int]

Deployments SHOULD make use of the recommendations contained in [ICSAML2] and [SAML2Int] except where they conflict with this profile. In such cases, this profile MUST be regarded as taking precedence.

Normative requirements of this profile are enforced by the UK federation registrar; metadata not meeting these requirements will not be registered.

7.1 Metadata and Trust Management

It is the responsibility of each deployment to incorporate the metadata supplied by the UK federation into its trust management infrastructure. It is RECOMMENDED that use of the metadata conforms to the SAML V2.0 Metadata Interoperability Profile Version 1.0 [MetaIOP] and that metadata be updated at least daily. Metadata update with a higher frequency than once every six hours is NOT RECOMMENDED unless constrained by use of the "`ETag`" header for cache management. Metadata update with a higher frequency than once every hour is NOT RECOMMENDED.

The use of TLS for Assertion Consumer Service endpoints is REQUIRED.

Provision of metadata supporting the Identity Provider Discovery Service Protocol Profile [IdPDisco] is RECOMMENDED.

7.2 Attributes

It is RECOMMENDED that any `<saml2:Attribute>` elements exchanged via any SAML 2.0 messages, assertions, or metadata conform to the MACE-Dir Attribute Profile for SAML 2.0 [MACEAttr]. This includes any use of `<md:RequestedAttribute>` elements in entity metadata.

7.3 Authentication Requests

7.3.1 Binding and Security Requirements

The use of TLS on endpoints at which an Identity Provider receives a `<saml2p:AuthnRequest>` message, and for all all subsequent exchanges with the user agent, is REQUIRED.

7.4 Responses

7.4.1 Binding and Security Requirements

The use of TLS on endpoints at which a Service Provider receives a `<saml2p:Response>` message is REQUIRED.

7.5 Future Directions

7.5.1 [SAML2Int] Move to Kantara

The [SAML2Int] specification was developed independently rather than within a formal standards body. It is anticipated that at some point during 2012 this specification will be migrated to the Kantara initiative and brought under that organisation's change control.

Once the migration process has been completed, this specification will be modified to refer to the stable Kantara-based version of [SAML2Int].

8 References

[ICSAML2] *InCommon Federation SAML 2.0 Profiles; Working Draft 03.* InCommon Federation, February 18, 2010.
See **https://spaces.internet2.edu/display/InCCollaborate/SAML+2.0+Profiles**

[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**

[MDQuery] C. LaJoie, Ed., *Metadata Query Protocol.* Internet Draft, December 31, 2010 (expired July 4, 2011).
Work in progress: this is not a normative reference.
Available as **http://tools.ietf.org/html/draft-lajoie-md-query-01**

[RFC 2119] IETF RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels,* March 1997. See **http://www.ietf.org/rfc/rfc2119.txt**

[RFC2616] R. Fielding et al., *Hypertext Transfer Protocol – HTTP/1.1.* IETF Request for Comments 2616, June 1999.
Available as **http://tools.ietf.org/html/rfc2616**

[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/**

[SAML2Err] OASIS Approved Errata, *SAML V2.0 Errata.*
See **http://docs.oasis-open.org/security/saml/v2.0/sstc-saml-approved-errata-2.0.pdf**

[SAML2Int] A. Solberg et. al., *Interoperable SAML 2.0 Web Browser SSO Deployment Profile, Draft.*
See **http://saml2int.org/profile/draft**

[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/**

[SAML2MIOP] S. Cantor, *SAML V2.0 Metadata Interoperability Profile Version 1.0. Committee Specification 01.* OASIS SSTC, 4 August 2009.
See **http://wiki.oasis-open.org/security/SAML2MetadataIOP**

[SAML2Prof] OASIS Standard, *Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0,* March 2005.
See **http://docs.oasis-open.org/security/saml/v2.0/saml-profiles-2.0-os.pdf**

[ShibProt] S. Cantor et al. *Shibboleth Architecture: Protocols and Profiles.* Internet2-MACE, September 2005. Document ID internet2-mace-shibboleth-arch-protocols-200509.
See **http://shibboleth.internet2.edu/docs/internet2-mace-shibboleth-arch-protocols-200509.pdf**

[UKFTS] *UK Access Management Federation for Education and Research: Federation Technical Specifications.* This document.
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/**

[UKROM] *UK Access Management Federation for Education and Research: Rules of Membership.* Document ID ST/AAI/UKF/DOC/001.
See **http://www.ukfederation.org.uk/**

[UKTRP] *UK Access Management Federation for Education and Research: Technical Recommendations for Participants*
See **http://www.ukfederation.org.uk/**

[XMLSig] D. Eastlake et al., *XML-Signature Syntax and Processing. World Wide Web Consortium.*
See **http://www.w3.org/TR/xmldsig-core/**