UK Access Management Federation for Education and Research

Federation Technical Specifications

Ian A. Young
9 September 2010

Version 1.2
ST/AAI/UKF/DOC/004

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).

> *Where appropriate, this document also describes the rationale behind the particular choices made. Paragraphs describing rationale are formatted in this way.*

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 frequently 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 13 September 2010.

1.3 Recent Document Changes

1.3.1 Changes for Edition 1.2

- Add a note about the intended audience for this document.
- Name change: UKERNA has become JANET(UK).
- Complete revision of the chapter on the federation trust fabric to bring it up to date with current practice. This includes previously published notes on how to verify credentials found in UK federation metadata.

- Mention that an XML Schema document now exists for the label namespace.
- Added a conventional prefix for the label namespace.
- The `SDSSPolicy` label is now indicated as having been retired in April 2010.
- Clarify the description of the `EntityDescriptor` element's `ID` attribute now that the transition from the SDSS federation has been completed.
- Remove the description of entities with an `OrganizationURL` value of "`http://www.example.com/`" now that all such entities have been converted to `Organization` convention B.
- `Organization` convention A is no longer in use within the UK federation metadata.
- Added a new chapter describing the Metadata Publication Service.

1.4 Future Directions

Each major section of this document contains 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 `KeyDescriptors`, 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 `KeyDescriptors`, 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://spaces.internet2.edu/display/SHIB2/ExplicitKeyTrustEngine

- Shibboleth 1 implementation:
  https://spaces.internet2.edu/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://spaces.internet2.edu/display/SHIB2/PKIXTrustEngine

- Shibboleth 1 implementation:
  https://spaces.internet2.edu/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.

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 SDSS Policy Label

During the transition from the SDSS federation to the UK federation, entities registered by members of the SDSS federation were temporarily "grandfathered" into the UK federation metadata even though the member's participation was initially under the looser policies devised for the SDSS federation.

Such entities were indicated by the presence of the following label element within the `Extensions` element of their `EntityDescriptor` element:

```
<ukfedlabel:SDSSPolicy/>
```

After a federation member agreed to the Rules of Membership (see [UKROM]), it confirmed to the federation operator those entities which it wished to retain within the UK federation. The federation operator then

added the UKFederationMember label to the metadata for those entities to signal that they now operate under the Rules of Membership.

At the end of the transition period, all entities still labelled as only operating under the SDSS Federation Policy were removed from the UK federation. The SDSSPolicy label itself was removed from the federation's published metadata in April 2010.

3.1.2 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].

3.1.3 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.

3.1.4 Deleted Entity Label

As part of the maintenance of federation metadata, the federation operator may mark an entity as "deleted" by adding the following element to the Extensions element of the entity's EntityDescriptor element:

```
<ukfedlabel:DeletedEntity
    date="deletion date"/>
```

The date attribute should contain the date on which the entity was marked for deletion, in xs:date format (CCYY-MM-DD, for example date="2006-11-30").

The effect of this label is to prevent the metadata for the individual entity from being included in the published metadata for the federation. Consumers of UK federation metadata should never encounter this label, and have no need to check for its presence.

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 Shibboleth 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.

Two different conventions for the information included in `Organization` elements are described here:

- *`Organization` convention A* was used in the precursor SDSS federation.

- *`Organization` convention B* makes explicit the distinction between the responsible organisation and the function of the particular SAML entity.

The metadata published by the UK federation now follows convention B throughout. The description of convention A is included only for purposes of comparison.

3.4.1 `Organization` Convention A

In convention A, all entities are provided with an `Organization` element in which the `OrganizationName` and `OrganizationDisplayName` are identical.

For an identity provider entity:

- The `OrganizationName` and `OrganizationDisplayName` both contain a string describing the identity community on behalf of which the identity provider makes assertions. In many cases, this will be the same as the organisation responsible for the SAML entity, but this will not always be the case when identity provision has been outsourced or aggregated.

- 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 `OrganizationName` and `OrganizationDisplayName` both contain a string describing either the organisation responsible for the entity, or alternatively the service provided by the entity.

- 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.

3.4.2 `Organization` Convention B

In convention B, 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.

In this convention, 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 description of Organization convention A will be removed in a future edition of this document.

The move to the scheme described here as *`Organization` convention B* 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, four 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.

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.

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.[1]

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/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
```

1 This support was first deployed in May 2010.

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 Future Directions

4.4.1 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.

It is very likely that 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.

4.4.2 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 late 2010, at which stage the technology will be evaluated for use within the UK federation.

4.4.3 SAML Metadata MIME Media Type

[SAML2Meta] appendix A defines a MIME media type of `application/samlmetadata+xml` for use with SAML metadata. At present, the MPS uses the less specific `application/xml` media type.

In order to improve standards compliance, the MPS is likely to transition to use of the more specific media type in the future. Client software designed to access the MPS should be prepared to accept metadata under either media type.

4.4.4 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.

5 References

[MDQuery] C. LaJoie, Ed., *Metadata Query Protocol.* Internet Draft, June 30, 2010 (expires January 1, 2011).
Work in progress: this is not a normative reference.
Available as **http://tools.ietf.org/html/draft-lajoie-md-query-00**

[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/**

[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**

[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.* Document ID ST/AAI/UKF/DOC/004; 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.* Document ID ST/AAI/UKF/DOC/003. See **http://www.ukfederation.org.uk/**