

Blueprint for a JISC Production Federation

*Alan Robiette and Terry Morrow, JISC Development Group
Version 1.1: issued 27 May 2005*

1. Purpose of this document

1.1 This document is a consultation paper. It aims to set out the practical and policy issues underpinning the formation of a single production Shibboleth federation for the UK higher and further education communities. Comments – whether on technical, operational or policy issues – are invited from interested parties including universities and colleges, JISC services, and JISC-funded development projects. In so many words the document may be regarded as a “straw man” to stimulate discussion on the basis for a UK federation.

1.2 The document and any responses received during the consultation will be reviewed by the JISC Core Middleware Advisory Board. This board has been appointed to oversee JISC's programme of activities in the area of access management, with particular reference to developing Shibboleth-based services. It contains representation from the majority of JISC committees (JIIE, JCN, JLT, JSR and JCS) and is chaired by Iain Stinson, University of Liverpool. The Advisory Board is expected then to recommend appropriate action by JISC to set up the federation.

1.3 On a similar timescale Becta are managing the setting up of a Shibboleth federation for the UK schools sector, involving schools, local education authorities and regional broadband consortia. The JISC and Becta teams work closely together to ensure that although the two federations are being established as separate entities, they will wherever possible use common standards in such areas as attribute specifications and data formats. Thus if at a future date the decision is taken to move towards a single federation for the whole education sector, there should be no technical obstacle to doing so.

1.4 For ease of reference, each paragraph of the paper is numbered. Comments should preferably use this numbering system to identify the text to which they refer.

2. Background and terminology

2.1 The term “federation” has emerged in recent years to describe groups of organisations which agree to adopt common policies and technical standards to provide a common infrastructure for managing access to resources and services in a uniform way. It is primarily associated with the Shibboleth technology but the term could arguably be used in a more general sense. Examples of Shibboleth-based federations are InCommon (<http://www.incommonfederation.org>), the federation formed by the Internet2 community in the United States, the SWITCHaai federation of the higher education system in Switzerland (<http://www.switch.ch/aai/>), and the HAKA federation (<http://www.csc.fi/suomi/funet/middleware/english/>) developed by the Finnish universities and polytechnics.

2.2 A common nomenclature has also emerged for referring to the various organisations taking part, which all three of the above production federations use. It seems sensible to adopt that terminology also here. It distinguishes the main types of participant into **member** and **partner** organisations. **Member** organisations are primarily the educational institutions who provide and manage the Identity Provider (**IDP**) functions; in other words they represent the communities of users of the federation's resources. (Member organisations may also have a Service Provider (**SP**) function but this is not essential.)

2.3 It is very likely that the member category would need to be defined a little more broadly than simply educational institutions. For example publicly funded research organisations and national libraries would be obvious candidates, and there may well be others. However the basic idea that they represent the educational and research users would remain. It would be possible to classify member organisations into sub-categories should this be useful, but such distinctions are not pursued in the present paper.

2.4 **Partner** organisations on the other hand are bodies who provide resources and services which they wish to offer to the member organisations. Typically they only operate SP functions in the federation.

2.5 It could be useful to define other categories of organisation within a given federation, such as those providing operational support for the infrastructure, but again it will not be necessary to go into any detail for the moment.

3. Governance

3.1 As the strategic infrastructure and services body for ICT in the higher and further education sectors, it is self-evident that JISC should have ownership of the UK federation. JISC is itself not a legal entity: if it were decided that a formal legal framework for the federation should be set up (a question which is deferred until later in this section) it would be necessary to consider how JISC would accomplish this. What is clear, however the possible legal issues are resolved, is that JISC should determine the governance and management structure for the federation and should, through its normal processes, appoint a majority of the members to the resulting committee structure.

3.2 The management structure is for discussion, but a possible structure is as follows.

Board of Management: Responsible for legal and policy issues; financial matters; membership decisions; resolution of disputes (if any). Accountable to the designated point in the JISC committee structure.

Technical Advisory Panel: Responsible for reviewing and advising the Board of Management on technical issues, e.g. recommendations on software and software revisions, attribute definitions, certificate management, security.

Operations Panel: Responsible for reviewing and advising the Board of Management on operational issues, e.g. SLAs of infrastructure services, training and support services, logging and statistics, audit.

3.3 The question of the legal status of a UK HE/FE Shibboleth federation needs careful consideration and may call for professional advice. InCommon, for example, has been set up from the start as a limited liability company with which all members and partners deal as a business entity. In the event of a major incident, it would be possible for an aggrieved party to bring a legal action against InCommon; but since InCommon's limit of liability has been deliberately set very low, there is in practice little point in doing so. In effect, the legal status of InCommon is simply a device to deter all the participating organisations from attempting to resolve disputes through the courts.

3.4 Historically JISC has, to date, used access management technology (in the form of Athens) mainly to provide uniform means of access to electronic resources, in many of which there is of course commercial IPR. Where this is so, the legal aspects are covered by more or less standard forms of contract negotiated for each resource. Thus in the event of a dispute these contracts would in principle provide the basis for litigation. Contracts for JISC content normally specify that the parties agree to use Athens for access management, presumably implying that they are happy that Athens is fit for purpose, and in practice no overall legal framework governing the relationships between the Athens

service (i.e. Eduserv), the universities and colleges, and the IPR owners (e.g. publishers) has ever been deemed necessary.

3.5 The view of the present authors is that in the JISC context, a Shibboleth federation could operate effectively as a JISC service without a special legal status being required. Contracts with publishers could continue on the present basis, and other access management needs – e.g. for inter-institutional learning and teaching or research collaborations – would be best covered by appropriate consortium agreements between the institutions concerned. If however the consensus should be that a limited liability company is a better solution, there are options which could be explored to provide this status.

4. Operational services at federation level

4.1 Shibboleth as a technology requires comparatively little in the way of central infrastructural services. Instead, more responsibility rests on the members and partners to conform to agreed technical standards and operational levels of assurance. This section discusses those federation-level services which are needed and how they may best be provided.

4.2 There will need to be a federation-level **WAYF** (“where are you from?”) service to which all Service Providers can redirect users who are previously unauthenticated, in default of any other method of determining their home institution. The WAYF does not need to be unique – a federation may contain multiple WAYFs if, for example, some service providers choose to offer their own WAYF as part of their front-end web pages – but a default service is a necessity. This should be a straightforward function to provide: its main characteristics are good user interface design and high levels of reliability and resilience (otherwise it could be a key point of failure). It is most likely to be contracted out to an established JISC service provider and the contract should be placed via standard JISC tendering procedures.

4.3 Shibboleth requires a number of **X.509 certificates** to provide security services of varying kinds. Certificates which are (potentially at least) user-visible, in the sense that they are used to establish SSL/TLS connections with users' browsers, are in most current practice obtained from commercial Certificate Authorities (CAs) – i.e. those whose roots are pre-installed in the prevalent browser software – since otherwise the user will receive annoying and possibly confusing pop-up windows during the browser session.

4.4 Certificates which are used to sign Shibboleth messages between participant sites are not subject to the same restriction. It is only necessary to agree amongst the participants which CAs will be trusted at federation level, so that the required information can be installed into the Shibboleth configuration files at each site. Here practices differ: InCommon has taken the decision to run its own CA and to base the federation's trust model wholly on these in-house certificates, whereas the SWITCHaai federation allows both community-generated certificates (i.e. those issued by the SWITCH-operated CA) *and* certificates from a limited number of commercial vendors. SWITCHaai also has a published policy on how the list of allowed CAs is controlled and could potentially be increased.

4.5 This paper recommends that the UK federation should adopt the SWITCHaai model. The choice of commercial CAs to be supported requires a consensus to be reached by the interested parties, but it should not be hard to agree on a short list of, say, three or four reputable suppliers. Whether or not a community CA should be added to the list is a moot point. Such a CA would need to operate to quite exacting standards to provide suitable levels of assurance to publishers that their commercial IPR would be adequately protected, and this is not a trivial requirement. The UK Grid Support Centre CA at Rutherford Appleton Laboratory has to conform to internationally agreed levels of assurance (e.g. to

satisfy the requirements of US Federal Government research establishments) and, if it were granted an extended life and wished to provide a server certificate service for Shibboleth sites, could be considered as a possibility.

4.6 A federation-level service to ***maintain and distribute federation metadata*** is also required. This needs to be robust and operated to high standards, particularly in terms of timeliness of updates, but it should be straightforward to set up and run. Like the default WAYF it would be outsourced via standard JISC tendering procedures. There could be advantages if the same service provider operated both the WAYF and the metadata management, but this is not strictly necessary. A policy decision will be needed on which metadata format, or formats, will be supported.

4.7 Finally thought needs to be given to the level of support and training which JISC might see fit to provide to facilitate smooth operation of a UK HE/FE Shibboleth federation. The needs of smaller institutions, and those with limited internal resources to support new IT systems, may need special consideration. These factors should be assessed in the light of the early adopter programmes and the experiences of the Middleware Assisted Take-Up Service (<http://www.matu.ac.uk>) and the SDSS Project (<http://www.sdss.ac.uk>) which are each, in their different ways, funded to provide support to Shibboleth users in UK HE and FE institutions.

5. Technical and policy documentation

5.1 Policy documents are required, and are being drafted, to define a number of important aspects of the federation and the standards to which it will operate. Models for these already exist in operational federations such as InCommon, SWITCHaai and HAKA, but the JISC community has some distinctive characteristics of its own and the UK policies will need to reflect these.

5.2 The document set as currently envisaged includes the following:

- Requirements for Identity Providers (IDPs)
- Requirements for Service Providers (SPs)
- Certificate Authorities and the Federation Trust Fabric
- Approved Software Components and Software Revision Levels
- Federation Attribute Definitions and Associated Management Procedures

5.3 In most of these cases the document content, though it will need to be agreed and signed off, is likely to be uncontroversial. Perhaps the most interesting topic, which merits further discussion here, is the specification of the attributes which participants in the federation are all (in general) expected to support, since these in many ways serve to define what will be possible in the more common applications of the federated access management regime. For this reason the following section is devoted to a preliminary discussion of, and set of proposals for, the attributes which might be placed in a "strongly advised" category.

6. Federation attributes

6.1 Comparison of the mandatory attributes required in InCommon, SWITCHaai and HAKA reveals a surprising lack of unanimity. This no doubt reflects the differing priorities and use cases around which these national federations were designed. Particular local characteristics of the JISC federation are (a) the fact that it may grow to quite a significant size – there are some 180 higher education and 550 further education institutions which might, over time, become members – and (b) that there is a wide range of institution types, missions etc. leading to the expectation that applications could be very varied also.

6.2 These facts argue for a small but carefully chosen set of strongly recommended attributes which all participants should, ideally, be able and be prepared to support. This set should be designed to support a range of basic use cases including teaching and learning collaborations, research collaborations, and access to publishers' resources as is achieved now via the existing Athens access management service. The discussion which follows is based on the eduPerson specification 200312 and the subsequent draft paper *draft-internet2-mace-dir-eduperson-00.pdf* (for copies of these specifications see <http://www.educause.edu/eduperson/>).

6.3 Some applications, notably access to third-party resources, will typically not require the user's real identity to be explicitly known; however many third-party suppliers in practice do prefer to have a persistent identifier for each user, e.g. to allow personalisation of their services. Such an identifier can be opaque, i.e. one which does not identify the real-world individual in an obvious way (and preferably should be, to guard against unnecessary loss of privacy). The eduPerson construct for this situation is ***eduPersonTargetedID*** and this is a clear candidate for inclusion.

6.4 Many contracts for third party resources are on a site licence basis, so that all the supplier needs is an assertion that the user is indeed a *bona fide* member of the home organisation. However it may be useful to classify some resources as available only to staff, or only to students. Thus an attribute conveying information of this type is also a priority. The eduPerson specification provides a choice, the main options being ***eduPersonPrimaryAffiliation*** and ***eduPersonScopedAffiliation***. The first of these takes a single value, indicating the individual's main role in the institution; for instance "staff" or "student". Effectively it links the individual to an IDP via a simple role parameter. The second is more versatile in that the syntax contains a so-called scope parameter (a value is of the type *affiliation@domain*, e.g. *staff@dept.someuni.ac.uk*) so that it links an individual to a specified security domain, a richer concept than just an IDP. Which should be preferred is for debate, but most recent thinking in other federations tends to favour ***eduPersonScopedAffiliation***.

6.5 The catch-all for situations where a basic affiliation role is not sufficient to determine whether or not an individual should be authorised to access a resource or service is ***eduPersonEntitlement***. This can be multi-valued and by appropriate parameterisation can be used to convey information of the general kind "entitled to access resources specified in Contract XYZ123" or even "entitled to place purchase orders on the finance system". It should certainly be supported since it is the key to permitting otherwise complex authorisation criteria to be specified in a simple manner.

6.6 Inter-institutional collaborations involving either research or teaching and learning will frequently require some form of more explicit identifier than the recommendation already given above, i.e. ***eduPersonTargetedID***. The latter type of identifier, just because it is intended to be opaque and not readily linked back to the real user, cannot readily be used for any application which involves for example direct communication back to the user (e.g. via email). Where personal identification at this level is a requirement, the clear choice is to use ***eduPersonPrincipalName***, which has a syntax of the kind *userID@domain* and is guaranteed to be globally unique. [***eduPersonPrincipalName*** is logically distinct from email address, for which a separate attribute of ***mail*** is defined in the iNetOrgPerson schema. However in practice it might be convenient to use an individual's main human-readable email address to serve as ***eduPersonPrincipalName***.]

6.7 Other federations have chosen to make various combinations of the attributes ***cn*** (common name), ***sn*** (surname or family name) and ***givenName*** mandatory. Although these will be held by the home organisation in all cases, it is not so likely that they would be used in inter-organisational access decisions and in the case of JISC, the conclusion is that no attributes carrying identifying information need be recommended for inter-organisational use other than ***eduPersonPrincipalName***.

6.8 Federation policies are likely to be needed on whether and, if so, in what circumstances identifiers previously associated with one individual can be reassigned to another. This issue applies both to opaque identifiers such as eduPersonTargetedID and to personal identifiers such as eduPersonPrincipalName.

6.9 Federation participants with more specific requirements are of course at liberty to use other attributes by mutual consent. For e-learning collaborations an attribute definition allowing courses or modules to be specified may prove particularly useful. An Internet2 working group (see <http://middleware.internet2.edu/courseID/>) is addressing this problem and will be producing a schema to be known as eduCourse. However only a sub-set of the applications within the federation will require this type of attribute and thus it is unnecessary to recommend that this attribute should always be available.

7. Roadmap and timetable

7.1 The top-level project plan for the JISC Core Middleware programme indicates that the groundwork for the production federation should be complete by July 2005. This blueprint, and the consultation exercise it is intended to stimulate, is seen as a significant step towards this goal. In the light of responses received from stakeholders it will be possible to refine the ideas presented here and to incorporate these into the additional (and fuller papers) prefigured in Section 5. A set of drafts for the more detailed policy papers is targeted for early July 2005.

7.2 Assuming sign-off by the Advisory Board, realisation of the plans – for example putting out to tender the infrastructure services – is expected to begin in Autumn 2005. Organisations currently using one of the pilot federations, i.e. those run by SDSS and Eduserv, would be able to transfer to the production federation at their convenience during the 2005/6 academic year. The production federation would then become the main environment for use by institutions of all kinds. The need for a parallel low-assurance, test-bed style federation on the lines of Internet2's InQueue is still being assessed and further announcements will be made about this in due course.