

Author:	John Paschoud
Date:	22 September 2004
Programme:	Common Services
Committee Funding:	JCIIE and JCSR
Structure:	3 pages
Document Reference:	CMSS: Paschoud

Core Middleware and Shared Services Studies

UK Assessment of eduPerson and Related Schemas

Contents

Executive Summary	3
1. Introduction	4
2. Background and context	5
3. Report on international scene.....	6
4. Input from the UK education community.....	7
4.1 Data gathering	7
4.2 The Focus Group.....	8
4.3 Mapping the route: significant input from the focus group	9
4.4 Feedback from Further Education Institutions	10
5. Production of the UkeduPerson Schema	11
5.1 UkeduPerson Schema Versions.....	12
6. Vendor consultation	14
7. Looking forward: recommendations to JISC.....	15
8. List of Annexes.....	16
9. Project Team	17
10. References	18

Executive Summary

There is an increasing international awareness of the need to allow access to online resources with reliable user authentication, but without the need to supply privacy-infringing personal information. The challenge lies in being able to assert access rights accurately without revealing the identity of the user. The solution is to move to a system which relies on anonymous role-based authorisation asserted in a standard way by a reliable organisation. One technology which facilitates these secure internet transactions is Shibbolethⁱ. Shibboleth technology originated in the USA, which led the world with the creation of the eduPersonⁱⁱ schema, as this new approach to access rights required a robust method of relating people to their roles.

The UK's interest in the Shibboleth solution led to the consideration of the creation of a standard schema for the description of members of the UK academic communities. The UK however unlike the USA already had a tried and tested access management infrastructure in Athensⁱⁱⁱ. Thus much of the challenge for the UK lies in persuading the academic community to consider new technology. The UKeduPerson consultation exercise illustrated the depth of this problem and emphasised the need to raise national awareness and sponsor a national debate on the issues.

The UKeduPerson project followed three strands: an assessment of the international picture, a "bottom-up" assessment of potential requirements for a UKeduPerson schema and a consultation exercise with Shibboleth aware information vendors. The project resulted in the production of the UKeduPerson schema^{iv} and made recommendations to JISC for future development.

The consultation exercise was based around a questionnaire, but the project team made considerable efforts to extend the knowledge base via one-to-one meetings, email and telephone correspondence. A total of 30 institutions contributed to the research thus the team were confident in being able to report a representative view of large and small institutions from all over the UK. A focus group at the London School of Economics provided much valuable information and directed the project team towards currently available standard definitions (HESA^v and ILR^{vi} categories) and alternative technical solutions.

It had been expected that the UKeduPerson schema would be produced in LDAP^{vii} as with the US eduPerson schema but input from the academic community led the project to consider the advantages of XML^{viii}. The principle impetus for this change of direction was the need to be able to define complex data describing users with multiple roles. Access rights need to be related to roles, not to individuals, and multiple roles present a challenge in LDAP. The technical solution to these challenges was to produce three versions of the schema; an XML version, a description of how to embed the information into the SAML standard^{ix} and an LDAP version for compatibility with legacy directory systems.

The project team recognise that the UKeduPerson schema will need updating and amending and that this should become the responsibility of a standards body. It is recommended that CETIS^x (the centre for educational technology interoperability standards) is offered this role.

1. Introduction

The aim of the project was to develop a usable descriptive schema, which will both fit the requirements of the UK post 16 education community and meet the criteria for access to shared services and for use with future access management technologies like Shibboleth. The approach adopted was to establish the current arrangements for defining roles within the UK Higher and Further Education community with a view to documenting best practice and producing a standard approach for the future.

The project had three strands. The first was intended to provide a context for the UK and involved a review of current practice outside the UK through consultation with Internet2 and other non-UK national bodies. The aim was to assess the strengths and weaknesses of similar 'xeduPerson' schemas already in use, including commonalities and potential incompatibilities^{xi}. The second strand of the project involved detailed data-gathering from UK educational institutions to identify relevant institutional initiatives and to determine common and sector-specific requirements for a potential 'life-long learning schema. This strand represented the major effort of the project. The third strand was a consultation exercise with a small number of 'Shibboleth aware' information resource vendors.

2. Background and context

Many of the JISC's partner communities are moving towards the adoption of a standard schema for the description of the members of their institutions. The prototype for this is the eduPerson schema devised as part of the Educause/Internet2 middleware programme. One objective of this project is to shift access to online resources away from demanding that the real-world identity of a user (and often other privacy-infringing details) must be revealed (authentication); to a model where access is granted based solely on an anonymous 'role' (such as "student"), asserted in a reliable way by an institution known to the online resource host (role-based authorisation). The principal new Internet technology involved is known as Shibboleth. The production of an eduPerson type schema for the UK is a prerequisite to the implementation of Shibboleth type technology in this country, since the solely American perspective of the US schema would make it difficult to implement here.

Projects in Norway, Australia, Netherlands and the USA are working to adapt and extend the eduPerson schema to fit their requirements. There is therefore a clear need to follow a similar path in the UK. In the longer term such schemas should provide the ability to standardise descriptions across international boundaries.

3. Report on international scene

The Project Director has participated in discussions sponsored by the Internet2 MACE (Middleware Architecture Committee for Education) to explore ways of co-ordinating the trans-national development and use of person schema. The discussions included participants from UNINETT (Norway), CAUDIT (Australia), Columbia (USA), UW-Madison (USA) and U. v Tilburg (Netherlands).

It should be noted that the situation in the USA (and in some European countries) is slightly different from that in the UK, as the US higher education community is creating an access infrastructure from scratch, whereas the challenge in the UK is to move institutions on from existing tried and tested technology to the next phase of developments. Early international discussions proved constructive, but at present they have largely stalled as each country concentrates on the creation of their own schemas.

Discussion of the schema is ongoing, but this has been taken over by the MACE-DIR list^{xii}, where the focus has been on the issues surrounding entitlements and affiliations^{xiii} in a Shibboleth environment. The situation is fluid but UKeduPerson project staff continue to monitor and participate in the discussions.

Proposed deliverables by the international group are at present:

- a Best Current Practice document to cover:
 - schema definitions;
 - date formats;
 - unique identifiers;
 - existing definitions.
- a document on where to find background information on academic community schema definition for schema developers.

In a recent development, the final meeting (6th June 2004) of the TERENA Taskforce for Authentication and Authorisation Coordination for Europe (TF-AACE)^{xiv} raised the subject of the international eduPerson. The general feeling of the meeting was that as more people are starting to get involved in such work, these discussions should continue. It was also felt that the "staff" part of the schema required more discussion. It was agreed that such international schema should be part of the work of the new Taskforce on European Middleware Coordination (TF-EMC)^{xv} in its first year.

4. Input from the UK education community

The team began the work of data gathering by designing a questionnaire to provide an overall structure to the process, whilst encouraging respondents to describe local attributes not defined by the project team. The questionnaire was produced on a spreadsheet (Excel) and was accompanied by a “specimen” return, showing examples of the format of expected returns. The design of the questionnaire was a process of evolution. The project team were keen to consult academic colleagues so requested and received feedback from informal contacts. Members of the project team also gave a presentation to a UCISA^{xvi} meeting during the development of the questionnaire and the feedback received there was fed into the final version. The questionnaire that went out to contacts therefore, included ideas and comments from a range of sources.

The questionnaire was targeted at MIS departmental heads in all HE institutions in England, Wales and Scotland and copied to the heads of library services. It was also sent to a selection of Further Education colleges, London schools and sixth form colleges. The questions were designed to encourage respondents to review and report on local arrangements thus allowing the team to document the full range of key fields.

The results were then fed into a master spreadsheet, showing current arrangements in the old and new university sectors as well as in Further Education and post 16 schooling. The questionnaire also allowed for the collation of data on current software and hardware and other related matters. It was hoped that a picture of common ground would become apparent and would inform the basis for a schema of UKeduPerson attributes. Thus a key challenge of the project revolved around the question of defining the boundaries between nationally common attributes and those of only local significance. This concept of boundary definition also occurs at the national / international level.

4.1 Data gathering

The original intention was that between 30 and 50 institutions would participate in the study and that this would provide a good representative sample of (at least) post-16 education. Considerable efforts were made to ensure inclusion of all UK countries and regions but the data gathering process proved more problematic than expected, resulting in a smaller number of participating institutions than had been anticipated. Completed questionnaires were received from 19 institutions in total, 14 HE, 3 FE and 2 schools. However the project team were able to secure participation in the focus group by some institutions that had not returned questionnaires and the pool of contributors was further extended via telephone interviews and email discussions. These two additional strands of communication widened the consultation and in total 30 institutions took part, covering all UK regions and countries except Northern Ireland.

The timeframe for the project presented challenges as it necessitated a very rapid turnaround of a fairly complex questionnaire. The target audience were largely unprepared for the underlying concepts of the project and were not always able to put in the time necessary to understand the background. The project team put considerable effort into the data collection process, following up the initial email with telephone calls and personal contacts and where they were able to give a verbal explanation of the

project, this often produced results. However it proved difficult to get an institutional return as many respondents stated that they could only provide a departmental view rather than a holistic view of the institution. Many institutions appear to use different approaches internally and have problems relaying data from one department to another e.g. Registry to the Library.

Complexity of local arrangements raised another issue, as some institutions are set up to store very detailed information within their central directory systems whilst others keep very basic data in a number of different systems. Several respondents reported that they struggled to fill out the questionnaire, which assumed a coherent approach which simply did not exist within their data storage. The project team had attempted to allow for this, by encouraging the respondents to add both categories and values where those offered did not reflect their requirements. Sometimes this appears to be a case of giving the same thing a different name and in others it added to the picture of requirements for the schema. Despite the reservations noted above, the questionnaire returns did allow the project team to establish commonly used fields which enabled the creation of a draft national schema. This initial model was presented to a group of key people at the UKeduPerson focus group where it provided the platform for discussion, constructive criticism and further development.

4.2 The Focus Group

The focus group attracted delegates from a wide range of institutions with representation from all of the English regions and from Scotland. Welsh delegates were invited but unfortunately were unable to attend. The group included people from a good variety of backgrounds in the Higher Education sector. Most participants had technical roles (MIS, CIS) but registry and library staff were also included. The list of institutions covered the full spectrum of potential directory structures ranging from the traditional, relatively simple model of the old universities to the more complex modular approach of the Open University and the majority of institutions in the new university sector. Members of staff from JISC also attended and were able to provide a useful wider context for the project. Unfortunately we were not able to secure the attendance of any representatives from FE institutions at the focus group.

It should be noted that the members of the focus groups were in many ways self-selecting, coming from institutions that had started to consider the issues involved in this project. Most delegates were aware of the potential complexities of the project, although some initially doubted their ability to contribute. The project team were particularly impressed with the level of involvement of the delegates. There was a genuine attempt to drive the project forward, to question and review the available material and to advance their understanding of the underlying reason for a national schema.

The project team became aware that even this relatively knowledgeable group of people fell short of a full understanding of the key issues. This was felt to be significant as it informed the team of the need to raise awareness within the Higher Education community. The group spent some time discussing the concepts of authorisation and authentication without gaining universal support for proposed Shibboleth type model.

Much of the early discussion focussed around the question of the timeliness of the project. Was it appropriate to be planning for a future scenario of complex access decisions when currently resources tended to be bought in as “blanket cover” within individual institutions? The focus group members were initially sceptical about the need to consider Shibboleth type technology in the short to medium term.

Members supported Athens as the current familiar practice and a number suggested that AthensDA was a secure and obvious way forward. Delegates expressed concern about moving towards new technology when the current system and its proprietary successor seemed to cover their requirements. The idea of a globally compatible model was deemed to be unnecessary as long as the UK based Athens was supported internationally and a number of delegates reinforced this idea through examples of cross-national information passing, e.g. the current use of Athens to access US databases in the UK.

The discussion at the workshop explored the legal background, the need for privacy of personal information and the tension between this legal requirement and the increasing attraction for institutions to be able to restrict resource access to particular types of members. Although most institutions were still at an early stage of considering the limiting of electronic resources to particular groups of users there are existing examples of agreements where a high level of granularity is required. Institutions wishing to take advantage of electronically published study packs need to be able to identify individuals as members of a restricted permitted group and will fall foul of copyright law if they fail to meet these requirements. Limited licences for identifiable user roles makes financial sense as the alternative of blanket provision is likely to be costly. The issue of privacy is, however, one of extreme importance so institutions must develop means of targeting resources without revealing personal information to external suppliers.

4.3 Mapping the route: significant input from the focus group

After discussing the potential use of a national schema, the focus group considered the draft proposal in some detail and made a number of significant points:

- a) There was no current support for a national model for internal directories but there was awareness that a common schema to define access entitlements would be useful. As a result it was not possible to consider the organisational and cost implications of a national standard for directory information as had been intended in the project proposal.
- b) It was felt that personal information should be held in the internal directories and had no place within the UKeduPerson schema, which would be used to support remote role based authorisation.
- c) The concept of a controlled vocabulary of generic terms such as “student” or “member” was felt to be unrealistic but it was suggested that the schema should build upwards by allowing institutions to use combinations of currently understood attributes and values from HESA and ILR. (This superseded the original intention in the project proposal to consider M25 Access Eligibility Tool - now Visit a Library^{xvii} - categories or earlier work done by the Angel project^{xviii}).
- d) The problem of multiple roles was seen as a central challenge to any schema. Where users have multiple roles, it is often important for attributes describing the user to be linked (e.g. student in one department but member of staff in another). It was suggested that LDAP could not cope well with the need to represent multiple roles and that XML would allow such links to be more easily demonstrated.
- e) There was agreement that a first model of a UKeduPerson schema would need to be amended, updated and extended and that it was important to ensure that these changes would be monitored and be compatible with international advances. It was suggested that CETIS would be an appropriate body to overview the standard.

4.4 Feedback from Further Education Institutions

The project team followed up the focus group meeting with visits to GLOSCAT and NESOCOT. These were extremely useful and highlighted the differences and commonalities with the feedback from the focus group.

5. Production of the UkeduPerson Schema

It was originally expected that the UKeduPerson schema would be produced in LDAP to compliment the eduPerson schema. It became clear during the discussions at the focus group that there was a need to be able to express complex data describing users with multiple roles. Where users have multiple roles, it is often important for attributes describing the user to be linked. For example, if a user is a staff member in the drama department and a student in physics, they might have radically different rights in either role, or ones which are different again from a student in the drama department or a staff member in the physics department. It is therefore important to be able to disentangle and match category with departmental affiliation. Producing the schema in XML would provide a mechanism to allow cross referencing in a way that is not possible with an LDAP schema. Three different versions of the schema have therefore been produced, a straight XML version, a description of how to embed the information described in the XML schema into the SAML standard and an LDAP version for compatibility with legacy directory systems.

The discussions leading to the development of the UKeduPerson schema had made it clear that institutional representatives were not happy with the idea of passing personal data about users to external organisations, nor about looking at the schema as a possible standard for internal institutional directory structure at this stage. The main theme of the feedback received on the implementation of the schema was “keep it simple”.

It is essential, however, to record some individualised data (an LDAP record has to be related to an individual to make sense in this context), so a range of individualised attributes have been included in the schema specification. This does not mean that such information needs to be passed to external organisations, merely that it may be available should it be required. Passing attributes should be regulated through the use of one or more attribute release policies, as it is in the Shibboleth architecture. These should be controlled by institutions so that only information that can be released without violating the individual's privacy or legal rights will be made available.

The core of the UKeduPerson schema is the `category` element which describes a group to which an individual belongs, without identifying a specific person directly. This can be used to obtain access to resources without revealing private information. It may be possible in some cases to take a set of category elements describing an individual's group memberships and work out the specific person being described (e.g. a student might be uniquely described by the list of modules they are taking, or there may only be one professor of economics who is also an alumnus) but the effort involved and the access to information required means that this is not considered a serious security issue.

In accordance with the recommendations of the focus group the standard used to record data in the `category` element should be the HESA profiles of staff and students (for higher education) or the Learning + Skills Council Individualised Learner Record (for further education) for the majority of occurrences,. The schema lists the HESA and ILR categories that we would expect to be supported, as a minimum. This is not intended to be prescriptive, either in the sense that *all of these fields must* be supported or in the sense that *only these fields can* be supported. These fields do not need to be passed on to the requester; in many cases the fields would allow an institution to make the authorisation decision but should not be passed on to a resource provider.

The schema also lists some central definitions useful for authorisation for resource access that are unlikely to be covered by other schemes such as “member” and “authorised”.

The `institution` element allows information about the institution asserting data about the individual to be recorded. The data included here is based on the recommendations of the *JISC Scoping Study into Institutional Profiling and Terms & Conditions services*^{xix}.

The simplest possible UKeduPerson document contains a single `category` element, describing an individual as “authorised” (meaning that the institution concerned has checked and determined that the individual is permitted to access a resource).

The UKeduPerson schema, therefore, allows institutions to describe both individuals and categories of users (roles) either for internal use or for the purposes of remote access decisions in a standardised, flexible and simple way, using standards that are already well established within the UK HE and FE community.

5.1 UKeduPerson Schema Versions

- *The XML version including the data in the schema.* The human readable version of the UKeduPerson schema is not intended to be used in practice but has been developed for ease of user comprehension. This version and the SAML version allow both cross referencing between attributes and a data structure, where a single `category` element carries information about its value, the encoding used and the URI where more details about the encoding and the field involved are available.
- *A description of how to embed the information described in the XML schema into the SAML standard.* SAML is an XML schema for the transfer of access rights assertions and related information, and is used by Shibboleth and the Liberty Alliance^{xx} for this purpose. In a SAML transaction there are two parties: the client (e.g. a provider wanting to know if a user is authorised to access a resource, or some organisation attempting to find information on an individual) and the SAML authority, which owns the information about the individual and can assert security data about them (e.g. that they are entitled to use the resource). This schema only describes the SAML authority side of the transaction.
- *An LDAP version for compatibility with legacy directory systems.* There are two principal difficulties in rendering the UKeduPerson schema into LDAP. The first is that cross referencing between attributes in LDAP is impossible. The second is that it is impossible to record information where a single `category` element carries information about not just its value but the encoding used and the URI (where more details about the encoding and the field involved are available).

The four options available to resolve these issues are to:

- f) Lose functionality, extendibility and flexibility by creating a list of one-dimensional attributes (which would have the advantage of compatibility with legacy LDAP schemas such as the eduPerson schema);
- g) Create attributes that encode more complex data in a format, which then becomes hard to read for a human (which retains the functionality of the XML version at a high cost);

- h) Create attributes which contain large chunks of the XML from the UKeduPerson schema; this means that applications would need to obtain the attribute values and know how to process them;
- i) Create attributes that are references to fuller data elsewhere (which means that the LDAP directory administrator no longer has complete control over the meaning of the data, and which may require institutions to maintain complex information locally).

Option (a) was chosen. It was considered to be the simplest solution and the most compatible with existing practice. The problems outlined above, which are not unique to the UK scene, are likely to lead to modifications to LDAP or possibly even abandonment in favour of some other technology such as Web Services (which is how SAML is generally deployed).

Model implementations of the SAML and LDAP versions of the UKeduPerson schema have been produced and are available.

6. Vendor consultation

For this strand of the project, a small number of major information resource vendors, were contacted, to cross-check emerging recommendations for a schema against their own perceived requirements for institutional descriptions of end-users who may access their services. Vendors were drawn from those already involved in the Shibboleth Project and therefore already informed about the relevance of this work to future access management models. The Internet2 site lists vendors currently working with Shibboleth.

Most of the existing contracts described simply covered “all staff and students”, although interested was expressed in the targetedId attribute, which allows vendors to personalise their service to each user while allowing the user to remain anonymous. There does not seem to have been any implementation of role based authorisation so far, in fact targetedId seemed the favoured option for many developments in functionality e.g. different skins, better focused resource discovery and storing/tracking of searches.

Some vendors, however, are interested in associating users with a particular school or department (usually medicine, law or business), a particular course or a particular virtual organisation and this would require a role based application.

The impression that most licences are still at the “all staff and students” stage, was reinforced by an LSE Library project^{xxi} to produce an up-to-date record of the electronic resources to which the library subscribes, along with access requirements and access methods. None of the licences restricted access to specifically registered users or user groups. Of the 22 licences recorded, 21 authenticated via IP address, 12 authenticated via Athens, and 3 via username and password.

7. Looking forward: recommendations to JISC

- The Higher Education community remains largely unaware of the importance of the issues surrounding authorisation and authentication. The legal implications of both copyright and privacy law need to be emphasised if institutions are to be persuaded to move forward. The project recommends that JISC considers methods of raising awareness and sponsor a national debate.
- The UKeduPerson schema should be piloted within a variety of different institutions. The project recommends that JISC supports a pilot project to implement the UKeduPerson schema in both old and new universities in Scotland, London and the English regions.
- Further Education needs to be involved but different strategies need to be developed to encourage participation. It is recommended that forward-thinking FE institutions be identified and encouraged to become partners in reviewing the experience of the HE pilot. Future projects should note that resource restrictions make it extremely difficult for FE staff to attend meetings, but they welcome project staff prepared to visit them.
- It would be of benefit to the development of the UKeduPerson schema for JISC to consider the recommendation to develop a XML schema for institutional contact data. This should include the institution URL.
- It is recommended that HESA and ILR documents are maintained as findable targets thus allowing the relevant value to be fully referenced in machine readable form. This would contribute towards the attainment of the Semantic Web^{xxii} in the UK.

8. List of Annexes

Annex A	List of Participating Institutions
Annex B	Original questionnaire
Annex C	Summary sheet of returns from questionnaire
Annex D	UKeduPerson Schema version 1.0
Annex E	Licensed electronic material and their access permissions at the LSE Library project overview
Annex F	Spreadsheet of licensed electronic material and their access permissions at the LSE Library

9. Project Team

Project Director John Paschoud, London School of Economics

Project Officers Susan Baker

 Maria Hiscoe

Technical Advisor Simon McLeish

10. References

- ⁱ *Internet2 Shibboleth project*. <<http://shibboleth.internet2.edu/>>.
- ⁱⁱ *eduPerson Object Class*. <<http://www.educause.edu/eduperson/>>.
- ⁱⁱⁱ *Athens DA*. <http://www.athensams.net/development/devolved_authentication/>.
- ^{iv} *UkeduPerson*. <www.angel.ac.uk/UKeduPerson/>.
- ^v *Higher Education Statistics Agency Coding Manuals, 2004*. <<http://www.hesa.ac.uk/datacoll/home.htm>>.
- ^{vi} *Individualised Learner Record 2004/05, Learning+Skills Council, 2004*. <<http://www.lsc.gov.uk/National/Partners/Data/DataCollection/ILR/ILRSpecification/ILRSpec200405.htm>>.
- ^{vii} *Lightweight Directory Access Protocol (v3): Technical Specification, Internet Society Request for Comments 3377, 2002*. <<http://www.ietf.org/rfc/rfc3377.txt>>.
- ^{viii} *Extensible Markup Language (XML), Version 1.0 (Second edition), World Wide Web Consortium, 2000*. <<http://www.w3.org/TR/REC-xml/>>.
- ^{ix} *OASIS Security Services TC, SAML 1.2 definition, 2004*. <http://www.oasisopen.org/committees/tc_home.php?wg_abbrev=security>.
- ^x *CETIS (the centre for educational technology interoperability standards)*. <<http://www.cetis.ac.uk/>>.
- ^{xi} *International Person Schema*. <<http://middleware.internet2.edu/dir/intl-schema/>>.
- ^{xii} *MACE-DIR Discussion list*. <<https://mail.internet2.edu/wws/arc/intl-schema/2004-06/msg00003.html>>.
- ^{xiii} An eduPerson affiliation takes the form “I am a member” “I am a member of staff. An eduPerson scoped affiliation takes the form “ I am a member of the University of XX” I am a member of staff at YY University” An eduPerson entitlement is defined within a Shibboleth Federation where a group of Shibboleth origins and targets agree terms for use by the federation.
- ^{xiv} *TERENA Taskforce for Authentication and Authorisation Coordination for Europe (TF-AACE)*. <<http://www.terena.nl/tech/task-forces/tf-aace/>>.
- ^{xv} *Taskforce on European Middleware Coordination (TF-EMC)*. <<http://www.terena.nl/tech/task-forces/tf-aace/docs/2004/TF-EMC1.pdf>>.
- ^{xvi} *UCISA (Universities and Colleges Information Systems Association)*. <<http://www.ucisa.ac.uk/>>.
- ^{xvii} *M25 Visit a Library*. <<http://www.m25lib.ac.uk/AET/>>.

- ^{xviii} Authentication and Authorisation: A Proposed Namespace for the UK HE and FE Community., McColl, John, 2003. <http://www.angel.ac.uk/public-files/pdf/WP6_d6.5.pdf>.
- ^{xix} *Scoping study into Institutional Profiling and Terms & Conditions services*, JISC 2004. <http://www.jisc.ac.uk/uploaded_documents/InstitutionalProfilingFinalReport%20execsum.doc>
- ^{xx} *Liberty Alliance*. <<http://www.projectliberty.org/>>.
- ^{xxi} *A spreadsheet of licensed electronic material and their access permissions at the LSE Library*. Walker, Joanne, 2004. See Annex E.
- ^{xxii} *Weaving the Web: Origins and Future of the World Wide Web*. Berners-Lee, Tim, Texere Publishing, 2000.