UK Access Management Federation for Education and Research

Identity Provider Deployment

6 October 2006

Version 0.3

ND/ATG/AAI/DOC/031v0.3

1 Introduction

This is one of a series of documents covering technical, operational, and organisational issues relevant to the management of the UK Access Management Federation for Education and Research.

This document provides advice on the deployment of an identity provider.

Interworking between federation members is based on the willingness of service providers to trust the authentication and authorisation assertions made by identity providers about individual users. This involves two steps:

- a user presents credentials to an identity provider, often acquired by the user as part of a registration process (typically, name and password);
- the identity provider makes assertions about the user, on behalf of the member organisation, to the service provider to which the user has requested access.

Each organisation that wishes its users to have access to Federation services must determine how its registration and identification processes will operate, and how the identity provider service will be arranged. These decisions will be influenced by the licence and other requirements imposed by the providers of the particular services that users need to access and by the organisation's own organisational and technical facilities. In some cases it may be possible and appropriate for either, or both, the Registration and Identity Provider functions to be done by other organisations either as part of an existing relationship or under a new, possibly commercial, agreement.

2 Methods of user registration

Methods of user registration vary in the cost to the organisation and in the quality of the identity assurance guarantees which they provide. The following approaches are in use:

- *Managed registration*. The organisation administers the issue of credentials to individual users. The user's identity can be verified whenever these credentials are presented to the identity provider, and can be linked to each session where assertions about the user are transferred to a service provider.
- *Site registration*. Where an organisation does not provide its users with individual computer accounts, it may arrange for credentials to be issued to any user who makes a registration request on site. These credentials may subsequently be used on or off site. This provides anonymous registration with assurance that the user had access to the organisation's network when the registration was made.
- *Self-registration.* The user presents self-registered credentials obtained from an off-site supplier. This simply provides anonymous registration.

Only the first of these methods provides user accountability. All, however, enable the user to present a persistent identity, albeit an anonymous one in the other two cases.

Methods that do not provide user accountability may be strengthened where the service provider operates an additional out-of-band registration process, typically based on a user's email, telephone, or in-person assertion. This can bind the pseudonym provided by the identity provider to an actual user. A method may be significantly strengthened by an in-person assertion of identity backed up with a photo ID, such as a student card.

In general, each identity provider is regarded as operating a single assurance policy. Where an identity provider offers a range of identity assurance methods (under different product names) each is best represented by a distinct entity in the federation metadata.

3 User accountability

In a number of cases, user accountability is not a service requirement:

- for resources licensed to all members of the organisation (traditionally these have used IP address checking to verify that the user is present on the organisation's secure network);
- for open services, such as many blogs and wikis, where the only requirement is that the user presents a persistent identifier (a 'handle')

Service personalisation can be supported through the use of persistent anonymous identifiers even where the actual identity of the user is unknown.

User accountability is a requirement for many other services:

- where the means must exist to investigate suspected cases of misuse by individuals, such as when commercially sensitive content is involved, or content is supplied on the condition that it is used for educational purposes only;
- where a closed group of collaborating users wishes to preserve the privacy of their work;
- where a teaching resource is part of an integrated product whose operation depends on knowing the actual identity of individual users;
- where an internal service of the organisation is secured by means of a locally assigned entitlement.

In part, the choice of registration method and identity provider type will be determined by the requirements of the external services relevant to the organisation's users.

4 Attributes

Service providers within the federation make their services available to users on the basis of attributes about the user asserted by an identity provider. There are two kinds of attribute:

- *Unscoped.* The attribute is a conventional Directory attribute, often just a simple string value.
- *Scoped.* The attribute is keyed to the user's organisation. It is common for service providers to offer services on the basis that a user is associated with a particular organisation. The organisation with which a user is associated is identified by its *security domain*, referred to as its *scope*, which is usually the same as the organisation's DNS name. The value of a scoped attribute is usually represented as *value@scope* (this should not be confused with an email address, which it resembles). To ensure that service providers can rely on scopes to identify real-world organisations, digitally signed metadata published by the federation shall list the security domains which each identity provider is entitled to use. Standard service provider software discards scoped attributes with scopes not on the approved list for a particular identity provider.

Service providers vary in the attributes they require:

- *None.* Although this approach is deprecated, some service providers will offer service to any user of a particular identity provider (for example, the identity provider operated by an institution that is a customer of the service provider).
- *eduPersonScopedAffiliation.* It is common to offer services on the basis of the user's relationship (affiliation) with a particular organisation. The organisation is identified by its scope; the relationship may be student, staff, employee or simply "member". So a particular service provider might offer access to any user from either *arbsd.sch.uk* or *ox.ac.uk*, another only to *staff@ncl.ac.uk.* JISC Information Environment services offered by national data centres such as EDINA and MIMAS use this attribute as the main basis for authorisation. Resources are made available to any user with a value of eduPersonScopedAffiliation which has a security domain corresponding to an institution that subscribes to the resource and a value of *member*.
- *eduPersonTargetedID.* Often used in combination with eduPersonScopedAffiliation, this attribute provides a consistent, opaque pseudonym for the user that is different for each service provider, enabling service personalisation without the service provider knowing a user's actual identity.
- *eduPersonPrincipalName.* This scoped attribute provides a consistent name for the user across different services. Since the name is usually based on the same "net ID" used to log in to other services, it will usually be well known to the user, making it easy to create services with authorisation by access control lists that enumerate

eduPersonPrincipalName values of the allowed users (*alice@abdn.ac.uk, erpl99@ed.ac.uk*, …)

- *eduPersonEntitlement.* This attribute allows a service provider to make authorisation decisions based on arbitrary properties of the user, when it trusts an organisation and its identity provider to assert those properties. Multiple values are allowed for the same user, for example "Authorised to access restricted medical content" or "Member of the parking committee".
- *Other attributes.* The attributes listed above are the core attributes recommended for use in the UK Federation. Many other attributes, including surname, given name, e-mail address, and office telephone number are defined by eduPerson and related standards, but at present there is no expectation that these will be required by members of the UK Federation.

5 Choice of identity provider

An organisation may choose whether to operate its own identity provision or to contract identity provision to a third party based on several factors:

- the identity provision requirements of services important to its members;
- the practicality of implementing a solution consistent with the organisation's existing technical resources and administrative procedures;
- the technical and legal constraints in providing information about individual users;
- the balance of cost and benefit for insourced and outsourced identity provision solutions.

Considerations applying to outsourced identity provision are discussed in clause 6.

5.1 Insourced

TBD

5.2 Athens gateway

The Athens Access Management system has for some years provided the H&FE community with its principal means to access licensed resources in the JISC Information Environment, and to users in other sectors.

The Athens to Shibboleth gateway supports managed registration for institutions which have implemented Athens Devolved Authentication and for institutions that supply Athens with uploaded user authentication data for use with classic Athens. Other institutions use classic Athens for site-registration.

The Athens gateway is an important special case of outsourcing both the identity provider and registrar functions together: users continue to use existing classic Athens identities and the gateway proposes to offer three core attributes to service providers: eduPersonScopedAffiliation, scoped to the client organisation; eduPersonPrincipalName, currently proposed as having the scope athensams.net (cf. clauses 6.3 and 6.4); and eduPersonTargetedID. These proposals are provisional and are under active development.

5.3 Other outsourced identity providers

TBD

5.4 Public identity providers

The simplest possible deployment for a client organisation with limited resources is for its users individually to obtain accounts at a public identity provider organisation (e.g., http://openidp.org/). This should be sufficient

for access to small, closed-group service providers prepared to manage individual access control lists (for example, a distributed research group whose members require access to a central data store). The level of assurance in the identities provided is not a major concern if the service provider is willing to add new entries to its access control list on the basis of out of band communications, for example a telephone call from a well-known colleague asking to add his eduPersonPrincipalName of *jim@ openidp.org* to the service provider's access control list. (Jim may work at a large research institution that has not yet deployed its own identity provider.)

This approach, however, does not scale: national data services will not be willing to manage access control lists for thousands of individuals; nor can they accept everyone with an openidp.org scope, because anyone at all can obtain such an identity. Client organisations that need access to the kind of services that have hitherto used Athens for access management cannot use this approach, which will mainly be of interest to smaller organisational sub-units such as research groups, for access to private resources or for initial testing.

6 Outsourced identity provision

A number of considerations apply to an outsourced identity provider concerning its method of identity provision, its representation in federation metadata, its security domain usage, and its attribute usage.

6.1 Available registration methods

An outsourced identity provider may support any of the methods of identity provision described above (managed, site, or open registration) for a client organisation.

a) It can support managed registration, and hence user accountability, if it is linked to the organisation's staff and student databases. This may be accomplished by periodic upload of user information to the identity provider, or by some other form of integration. Note that any disclosure of personal information must be covered by a data processor contract in order to comply with the *Data Protection Act 1998*.

b) Where no such integration is available, the outsourced identity provider can support site registration by enabling the organisation's users to self-register using equipment known to be present on the organisation's secure network. The identity provider performs an IP address check on each registration request. Such registrations should be regarded as valid for a period not exceeding one year.

c) Otherwise, the organisation can simply recommend one of several available open registration identity providers to its users.

6.2 Representation in federation metadata

An outsourced identity provider acting on behalf of a number of client organisations may be represented in federation metadata in two ways.

a) *Multiple entities.* A distinct entity for each of the client organisations may be present in the federation metadata. Each entity has its own security domain (corresponding to the security domain of the client organisation), but otherwise, most of the configuration information is identical for each entity.

b) *Composite entity.* The identity provider may be represented by a single entity in the federation metadata which comprises all its client organisations (i.e., the single physical identity provider machine is represented by a single metadata entity). The identity provider is permitted to assert any of the security domains (scopes) belonging to its client organisations and is trusted to assert the appropriate security domain in each case.

The only practical distinction between these is the ease of handling potentially large numbers of entities in the federation metadata. For this reason, the use of a composite entity is recommended.

6.3 Security domain usage

The security domain present in an attribute assertion (for scoped attributes) is often the key information relevant to a service provider's authorisation decision. In the case of an organisation which outsources its identity provision, the normal requirements apply regarding its claimed security domain or domains:

- the security domain must correspond to an existing DNS name;
- it must be readily distinguishable from other security domain names;
- normally, it must be owned by the organisation, either through formal registration or as part of a systematic name allocation scheme, such as *.sch.uk* assignment in the Schools sector;
- where it is not owned by the organisation, the actual owner of the DNS domain must approve the proposed use, except where authority has been explicitly delegated by the owner to an outsourced identity provider to assign a set of security domains corresponding to a group of organisations under the owner's common management.

This last case is common in the Schools sector, where a Local Authority (LA) which owns a set of security domains delegates to a third party, such as a Regional Broadband Consortium (RBC), the management of identity provision for a collection of schools under the LA's responsibility.

6.4 Attribute usage

An organisation that employs an outsourced identity provider should obtain a guarantee that the information presented to service providers on its behalf will contain as little information as possible that is specific to the outsourced identity provider. A failure to ensure this may lead to effective lock-in to a particular identity provider where provider-specific information is recorded in licence registration information and other configuration data used by service providers. If identity provider-specific information is used then stored user information, accounting data and other licence permissions are likely to be lost when the organisation chooses either to change identity provider or to operate its own identity provider system. If the same information is provided by an outsourced identity provider for a number of organisations then service providers may find it difficult or impossible to provide tailored services to individual organisations.

In particular, all security domains present in attribute assertions should belong to the organisation, or be registered to an agent acting under the instruction of the responsible authority (e.g. the LA). They should not be values relative to the outsourced identity provider.

The four core attributes described in [UKTRP] may be handled as follows:

- *eduPersonScopedAffiliation*. Unless the identity provider has access to information supplied by the organisation which indicates the user's status, it will be able to assert only the value *member*. Otherwise, it may

assert the value appropriate to the user's status where permitted by the governing attribute release policy.

- *eduPersonTargetedID*. The identity provider should be able to supply this attribute in the normal way. Where the new, unscoped, form of eduPersonTargetedID is used, the value presented contains the name of the identity provider rather than the name of the security domain of the client organisation.
- *eduPersonPrincipalName*. Where the identity provider is unable to guarantee user accountability, it should not assert eduPersonPrincipalName.
- *eduPersonEntitlement*. This attribute may be supported where suitable procedures exist for its administration by the organisation.

6.5 Attribute release policy

Ideally, each client organisation will instruct its outsourced identity provider to release attributes only in accordance with the organisation's stated attribute release policy. This policy should make stipulations for each service provider, and may preclude the release of specific attribute values as well as all values for a given attribute type. The identity provider should be able to apply a different attribute release policy for each organisation it represents.

Where an authority is responsible for the common management of a group of organisations, it may specify a common attribute release policy which the outsourced identity provider is instructed to observe for all client organisations for which no individual policy is available. In no cases should attributes be released which disclose personal information without the explicit instruction of the organisation or the organisation's managing authority.

The Federation Guide to Use of Personal Data contains more information on attribute release policies and the protection of individual privacy.

7 Institution types

It is likely that different solutions will be prevalent in different communities. Some of the factors relevant to different types of institution are described below.

7.1 HE institutions

Many large HE institutions already operate a single-signon scheme for their users and are likely to opt for an in-house identity provider solution. The main ongoing costs, following initial deployment, are likely to be in the software maintenance of service equipment and the management of attribute assignments for individual users (in particular for an attribute such as eduPersonEntitlement).

Other HE institutions will have well-developed registrar functions for assigning computer identities to students, but still may not wish to spend time acquiring the required skills to deploy their own in-house identity provider. Outsourcing to an identity provider organisation that asserts attributes based on an institution's own user database should be possible, but is likely to incur significant up-front costs, both in the inevitable customisation required for the identity provider organisation's systems to access this database and in creating a suitable contract in the absence of existing models. Bear in mind that both the identity provider organisation and the client organisation will have to commit to observing the UK Federation's terms and conditions, and therefore they must divide the obligations which this entails between them clearly.

7.2 FE institutions

Many FE institutions do not give their students individual computer identities and therefore do not maintain a large-scale registrar function. In those cases, an outsourced identity provider organisation may be attractive if it is willing to undertake either site registration or managed registration on behalf of the client organisation. The identity provider must be able to assert a scope associated with the client organisation rather than the identity provider organisation, e.g., for scoped affiliation *member@jevc.ac.uk* for Jewel and Esk Valley College, rather than *member@outsourcer.com*. (The alternative of *member@jevc.outsourcer.com* is feasible for identifying the organisation to service providers but deprecated since the client organisation is obviously tied to the identity provider organisation.) Note that the identity provider organisation should offer some convenient mechanism to give the client organisation control over which user attributes may be released to which service providers. Minimising the attributes released to each service provider to only those genuinely necessary helps preserve user privacy.

One issue arises when the outsourcer also owns the user registration data. In this case, there is substantial risk of vendor lock-in. For example, switching vendors or setting up an in-house identity provider may require

changing the eduPersonTargetedID values associated with users, resulting in personalisation information being lost, including saved searches or stored user preferences. At worst, depending on the contract negotiated with the old vendor, users might need to be issued with completely new identities (if the user database is not available to the client organisation). Since drawing up effective contracts will require some of the technical knowledge that the client institution is trying to outsource, it would be wise to assume, at least until model contracts emerge over time, that the client organisation's relationship with an outsourcer is likely to be a long one, and hard to change.

At least one service within the UK Federation currently bases its authorisation decisions on whether the identity provider is in its list of customer identity providers, rather than on attributes asserted by the identity provider about a user. Institutions wishing to make use of such a service cannot use an identity provider organisation that presents itself to service providers as a single identity provider with different scopes distinguishing users from its several client organisations. Only identity provider organisations offering a distinct identity provider entity for each client organisation can support such service providers. However, service providers are being strongly recommended to base their authorisation decisions on user attributes rather than maintaining lists of approved identity providers, so this should not be seen as constraining the choice of identity provider organisation in the longer term.

7.3 The Schools sector

The Schools sector is in many ways the most complicated case, both in its scale and in its administrative arrangements. The parties involved are:

- *Client organisation.* The school. Note that service providers rely on scopes to *distinguish* organisations, so in order for a service provider to be able to render its services to some schools in an LA area but not others, users from a particular school should have a scope associated with the school, not with the LA.
- *Identity provider organisation.* May be a commercial organisation operating under contract to an LA or RBC. Alternatively, it may be the RBC itself.
- *Registrar.* May be an LA, an RBC, in some cases the school itself, or a commercial organisation or consultant operating under contract to any of them.

When an identity provider organisation is not performing the registrar function, there is less risk of vendor lock-in, provided that scopes are associated with the client organisation rather than the identity provider organisation, as described previously (though retaining personalisation data when changing identity provider is still a problem).

Both identity provider organisations and client organisations must agree to the UK Federation's terms and conditions. For schools, an LA or RBC is allowed to agree on behalf of the schools for which it is responsible, but even with this simplification, ensuring that suitable contracts are in place between all the multiple parties, particularly commercial organisations, to assure

service providers and the federation that an identity provider asserting attributes about a school pupil will make only accurate assertions, which everyone involved will stand behind, is one of the less obvious costs of outsourcing. Creating good contracts is likely to require considerable effort in the early days before model contracts become available.