Title	**SSO Discussion document (D5.6)**
Work Package	WP5
Authors	Atle Alvheim, Ken Miller
Source	Collective experience and work undertaken at the UK Data Archive and NSD
Dissemination Level	PU (Public)

Summary/abstract

Summary including use-cases and issues associated with Single Sign-On (SSO) for CESSDA.

Single Sign On – process summary

Users enter the system via the portal. The portal leads researchers to datasets, stored in (sets of) national servers (services). Datasets may represent open metadata and restricted access data. Different datasets may have different restrictions. Users are defined as being academic if they belong to an academic institution.

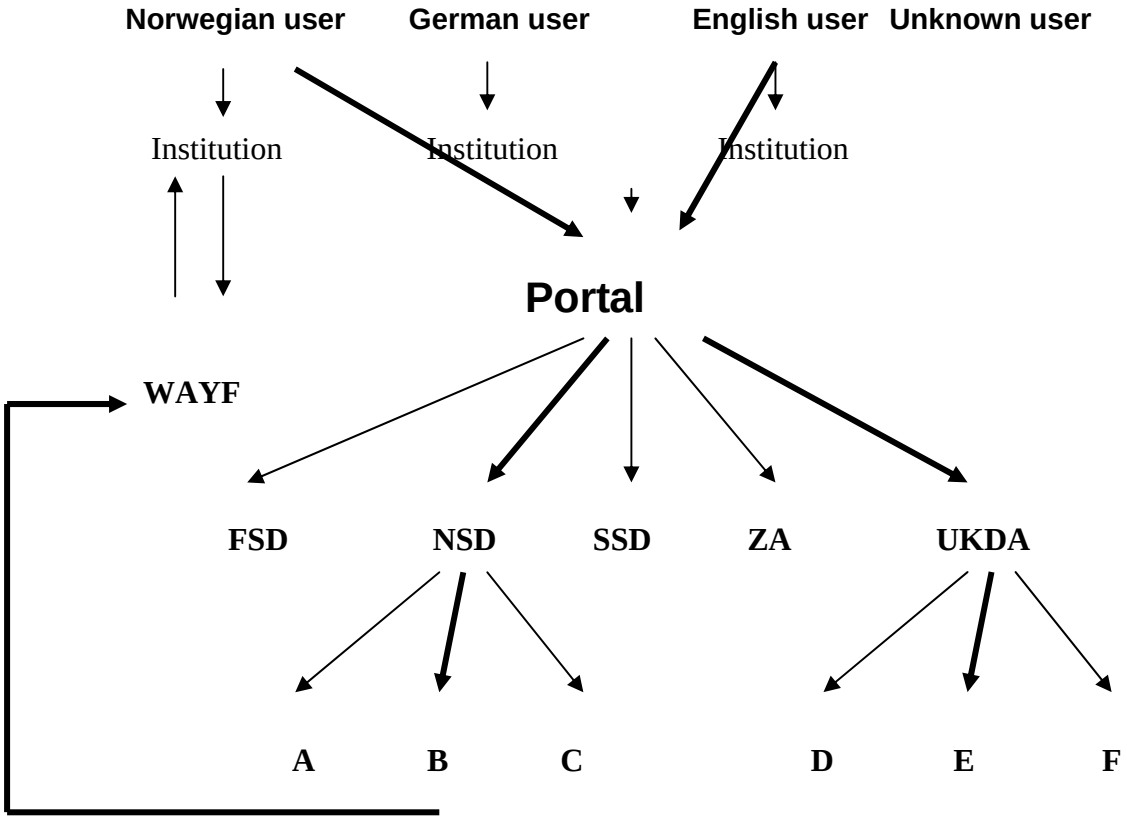

If a Norwegian user, via the portal goes to the NSD server and tries to access dataset B, he needs to be authenticated and authorised. The metadata part is freely available but the access to dataset B (or datasets in the group classified as restriction type B) are guarded by specific conditions. The most simple condition is that the user has to be affiliated to an academic institution.

The server needs to know if this person is the one he claims to be, and if he is authorised to access data at this level of security. Exploring metadata is not restricted and at this level a prospective user does not need to be logged in. When there is an action to access the actual data, the server tells the Discovery Service (WAYF) that user N is trying to access resource B. The Discovery Service asks the user: 'Where Are You From?' and expects to identify the institution the user is affiliated with. It then tells the institution that it has a user that the institution needs to authenticate. The institution, via its identity provider (in Norway FEIDE, and in the UK, the UK Access Management Federation for Education and Research – which uses Shibboleth middleware) responds by asking for a user id and password. If that is delivered and accepted, the message is passed back via the Discovery Service that the user is authenticated.

This does not necessarily mean that the user is authorised to access resource B. When the server knows that the user is authenticated, it can go to the next step, which is to authorize him for use of the data he has requested. In this case we could postulate that the data requested require that users should be affiliated to an academic institution **AND** that they are enrolled in a relevant project.

All the data archives need to have an explicit data access policy. Across all servers (services) there will be a need for a central authorization server, where every dataset is classified into some class / type of access. The server sends a request to this central authorization server and asks whether the user has the appropriate authorizations to view the data he has requested.

If the information is held in structured xml format of some kind, then it can be indexed by Lucene and accessed by an authorization server.

The authorisation server

An authorization server must provide the service providers with information about which authorizations a given identity has been granted.

Depending on the dataset, user may have to be bound to an institution, be enrolled in a project, or have signed an agreement of some kind. These restrictions should be bound to the data itself. Regardless of the restrictions, any attempt to access restricted data or to perform restricted actions (e.g. publishing data) will be checked against the authorization server.

This server should hold information about what conditions (i.e. projects) a given identity is connected to, what agreements the user has signed, etc, and also which other authorizations the user has been given explicitly as well as, possibly, an explanation of why the user has been granted this authorization).

All authorizations should be time-limited, with a start- and end-date specified for each. The server should regularly check that these time-limits have not been passed. When there are **x** days left of an authorization, the user should be notified and be presented with either a warning that his authorization is about to expire, or offered the possibility of extending the time-limit; either by signing a new agreement or by other means.

Issues to be resolved

Currently, not all CESSDA partners are part of a national authentication federation although the membership requirements set out in the Statutes require this.

Therefore it is **recommended** that authorization is dealt with at the national level. However, it is possible, as is the case in the UK, for one service provider to act as an issuing authority for others. Although highly unsatisfactory, as it would prove difficult for any organization to approve and authenticate users from outside their country, especially in the case of sensitive data, if this route were chosen, the **recommendation** is that non-issuing countries should be expected to reimburse the

issuing service provider as there are significant overheads in managing, maintaining and administering authorization systems, server and supporting database.

Technically, all identity- and service-providers should use technologies that communicate using SAML 2.0, or standards that are compliant with SAML 2.0. The **recommendation** is that all service providers agree on using the same technology (e.g. Shibboleth).

SSO cannot be fully implemented internationally until there is consistency of information between service providers. Specifically, all legal agreements will have to be directly comparable, the definitions of projects and teams have to be standardized, and the authorization requirements will have to be equal in all participating countries.

These legal and administrative matters are addressed in detail in Work Package 10 which considers the need for a European-wide categorization of users and a common system of licences. See WP10 documents and deliverables.

Similarly, the prototype developed is based on a simple but efficient architecture that uses a single authorization server (as in the UK's federated system). On a small scale this is workable but it would need to be scaled up significantly for it to function effectively across more countries. There will be significant financial overheads in its provision and management of a European-wide system in terms of: technical infrastructure; the manpower needed to maintain a running system; security and; back up systems (some service providers are contracted to provide 24/7 access). In addition the authorization process is the source of reporting information for national services. So, for example, the funders of national services require consistent, regular information on data use and dissemination and data providers sometimes expect to detailed information about the use and dissemination of their data to be available if requested. Agreement would need to be reached as to precisely what information must be collected for such reporting. Investment would be needed to set up and then maintain the supporting database and interface for service providers to create their reports. The **recommendation** is that the ERIC must treat SSO as a high priority; work to resolve all the outstanding issues and; ensure that sufficient funding is available to provide a 24/7 service that is acceptable to all members. Equally, it is critically important to realize that these problems are not isolated to CESSDA and a potential future ERIC. They are common across all similar data sharing and exchange research infrstructures. Thus, in an ideal case, common solutions across research infrastructures should be found and implemented with multiple ERICs and related research partners working together to pool expertise and build integrated and interoperable solutions.