Common Metadata for Climate Modelling Digital Repositories

METAFOR Single-Sign-On Evaluation
METAFOR Deliverable 4.1 M15

PROJECT	
Project acronym	METAFOR
Project full title	Common Metadata for Climate Modelling Digital Repositories
Grant agreement no:	211753
Funding Scheme	Combination of Collaborative Projects & Coordination and Support Actions
Call Topic	INFRA-2007-1.2.1 Scientific Digital Repositories
DOCUMENT	
Deliverable	D4.1 Month 15
Deliverable Title	Single-Sign-On Evaluation
Document Identifier	METAFOR-D4.1_M15
Date	June 15th 2009
Work Package	WP4 Deployment of Services
Authors	BADC
Document Status	Final
Document Link	http://metaforclimate.eu/documents

Dissemination Level		
PU	Public	
PP	Restricted to other programmes participants	
RE	Restricted to a group specified by the Consortium	
CO	Confidential	

Document History			
Version	Date	Comment	Author/Partner
0.1	May 20th 2009	First Draft	C. Pascoe and P. Kershaw/BADC
0.2	June 11th 2009	Penultimate Draft	B. N. Lawrence
0.3	June 15th 2009	Final	B.N. Lawrence, following comments

Summary

Single-sign-on is an issue which applies to those developing services which are expected to be deployed in multiple locations yet be accessed via portals or scripts invoked by one user, who may or may not have credentials at all the locations. Even in those cases where the user does have credentials at all sites, there is significant resistance by real users to using services for which repeated access challenges are encountered.

Access control consists of a number of stages: (1) deciding on policies about what access control constraints exist for specific data/metadata objects, (2) deciding who the user is (authentication), and (3) deciding whether that user meet the policy constraints (authorisation).

A prime requirement for Metafor is that the services be deployed in partnership with the Earth System Grid (ESG) team in the United States. It is expected that the IS-ENES team in Europe will also attempt to conform in some way to the Metafor solution, but IS-ENES may also be bound by other European grid initiatives.

In this work, the Metafor team's primary responsibility has been to evaluate access control solutions that will work with ESG to control access to metadata (and data) associated with the fifth Coupled Model Intercomparison Project (CMIP5, an essential part of the forthcoming fifth assessment report for the IPCC).

Metafor has chosen OpenID as an authentication technology, and has already demonstrated (as part of the evaluation) interoperation between BADC (Metafor, in the UK) and NCAR (ESG, in the US).

Metafor has chosen a role based system for authorisation which uses the Security Assertion Markup Language (SAML) interface between the different parties brokering access to resources.

Metafor intends to deploy the access control systems as middleware, which, along with the use of standards, maximises the possibility of using these techniques in future projects, especially, of course the EC funded IS-ENES.

Authentication (authN) **Establishes user identity!**	**Authorisation (authZ)** **Establishes what user can do!**
We have: • Established requirements. • Chosen technology: OpenID • Developed architecture and initial prototypes.	We have: • Established requirements. • Chosen technology: SAML • Developed architecture and implementation principles.

1. Introduction

“Single-sign-on” is required for services which are expected to be deployed in multiple locations yet be accessed via portals or scripts invoked by one user, who may or may not have credentials at all the locations. Even in those cases where the user does have credentials at all sites, there is significant resistance by real users to using services for which repeated access challenges are encountered.

Access control consists of a number of stages: (1) deciding on policies about what access control constraints exist for specific data/metadata objects, (2) user identification (authentication), and (3) deciding whether that user meet the policy constraints (authorisation).

Hence, single-sign-on is the first step to providing the capacity for both project partners and third parties alike to have “one-stop-shop” federated access to services, where users perceive that one set of access credentials can be used for steps 2 and 3 above, and data providers perceive that they still have control over their policies (i.e. step 1 above).

The primary services of interest to Metafor at this stage expose metadata catalogues and descriptions of simulation holdings at the project partners. However, it will be seen that these services need to be integrated within a common environment with other projects.

Advantages of using single-sign-on technology for user authentication:

- No requirement for new user accounts because single-sign on uses existing user account details
- No centralised administration is required
- Single-sign on is scalable which facilitates the inclusion of new institutions.

In the remainder of this document, we outline the requirements driving the solutions chosen with Metafor (section 2.0), and then discuss authentication (section 3.0) and authorisation (4.0), before making a few comments about how we expect to implement these solutions over existing systems.

2. Access Control Requirements

The main access control requirement for Metafor is to know who our users are and what data they have downloaded so that we may contact them if the data they download is ever updated. We also want to make it easy for other institutions to join the federated access control therefore the single-sign-on technology needs to:

- be easy to work with
- be relatively simple to configure
- provide sufficient level of security (but not be too onerous or over-configured).

Given those requirements, the starting point for this work was to think carefully about...

- Who are our users?
- What community are we supporting?
- Minimising the number of systems faced by our users
- Minimise the work required by existing data providers

- What are the key existing systems with which we have to work?
- What threats do we face?
- How valuable is the data we hold?
- What would a malicious person be able to do?

recognising that it is important to work WITH key partners, *within* the project, and importantly *outside!*

The forthcoming fifth assessment report of the IPCC will depend on model simulations organised and archived by the fifth Coupled Model Intercomparison Project (CMIP5) under the auspices of the World Meteorological Organisation (WMO) Working Group on Global Climate Modelling (WGCM).

The U.S. Earth Systems Grid (ESG) project is responsible for developing software to manage the archive, and the U.S. Programme for Climate Model Diagnosis and Intercomparison (PCMDI) is responsible for leading a consortium of archives to deliver the storage; that consortium includes three of the Metafor partners.

Metafor has been charged with acquiring the metadata to be used to document the simulations and activities of CMIP5, and the recently initiated EC project: InfraStructure for a European Network for Earth Simulation (IS-ENES) is charged with providing additional support for tools to manipulate both CMIP5 data and other relevant and future climate simulations.

The users (climate modellers), community (initially climate scientists), systems (existing database access, and crucially, CMIP5 systems) and primary issues/threats (interim privacy, accidental overwriting, release of common credentials allowing access to other resources) are very similar for both the CMIP5 data archive and the Metafor CIM repository. We therefore expect to adopt the same security architecture for all three projects (ESG, Metafor and IS-ENES) to ensure compatibility with climate data access control.

It is clearly of prime importance that Metafor develops systems that are as compatible as possible with CMIP5 and IS-ENES, and are compatible as far as possible with either the existing technologies deployed, or failing that, the patterns in which the existing technologies are deployed. Key existing technologies already implemented are NDG security(UK) and c3grid security (Germany) along with existing local solutions at other partner sites (and the Shibboleth efforts in a number of places).

The work carried out in this deliverable was primarily done by Phil Kershaw and Bryan Lawrence, at the British Atmospheric Data Centre (BADC), in conjunction with Stephan Kinderman at the German Climate Centre (DKRZ), and the ESG team (in particular Frank Seibenlist and Rachana Ananthakrishnan at the U.S. Argonne National Laboratory and Luca Cinquini, U.S. National Centre for Atmospheric Research).

3. Authentication

Prior to the collaboration for CMIP5, members of the ESG and NERC DataGrid (NDG) teams had both investigated OpenID and Shibboleth as candidates for a browser based Single Sign On solution. Both groups had already rejected traditional grid based solution using Proxy certificates (alone) on the grounds of usability with the associated technical difficulties for non-technical users managing certificates and embedding certificates in browsers. (That said, it will be seen that certificates do play a role in the access control environment we have constructed.) As a consequence, the main technologies evaluated for single-sign on were OpenID and Shibboleth.

Both Shibboleth and OpenID are implemented as third party services which provide a “form” of identity services, and the “form” of these identify services is a significant distinction between them.

The basic way these services are exploited is that requests for access from an unauthenticated user are redirected by a “Service Provider” (SP) to an “Identity Provider” (IdP) which returns an authentication assertion which an SP can use (generally in conjunction with additional user attributes) to decide what that user can or cannot do.

Of these two choices, we have chosen OpenID for Metafor and ESG authentication (and so, we expect, IS-ENES). In the remainder of this section we briefly describe some points of distinction between OpenID and Shibboleth, and then move to a description of OpenID and how we will deploy it within Metafor/ESG.

Shibboleth v OpenID

Shibboleth is a very credible authentication mechanism, details of how it works can be found in references [1], [2], and [3]. The main reasons why we have not chosen it, now, come down to a few aspects/issues:

1. OpenID is currently simpler to deploy, has Application Programming Interface (API) support in multiple languages, and has considerably more industry support (amongst, for example, Google and Microsoft).
2. There is currently a state of flux in the development of national Shibboleth Federations etc, which means that more individuals and partner organisations (in a wider variety of our user domains) could exploit OpenID “right now” than can exploit Shibboleth.
3. One of the important criterion for Metafor is to know who our users are, and by default, Shibboleth is designed to obscure actual identity (while simultaneously assuring the SP that the user is known and authenticated).
4. At the same time, for a range of Metafor activities, a key issue is to allow users to authenticate via their own IdP, one which is not known to us a priori. This is not normally possible via Shibboleth, see reference [4].

Although OpenID has been selected for SSO, Shibboleth remains an important technology to take into consideration into the future especially given that some national programmes in many of the Metafor partners are expecting to use Shibboleth. In terms of software development, within Metafor, we intend to develop the portal security layers in such a way that at some future time both Shibboleth and OpenID might be supported.

OpenID

As described above, the key advantage perceived for OpenID is its simplicity. There is a concrete self contained specification and as a consequence implementations and software libraries exist for most programming languages. It has developed rapidly and has seen interest and uptake amongst the big industry players. The application programming interface is well supported and because of the multiple implementations, it is easy to integrate into existing site infrastructures.

A detailed descriptions of OpenID can be found in reference [5], in brief, the OpenID system works by

1. All users being identified by their unique Universal Resource Identifier (URI, or or an eXtensible Resource Identifier, XRI[1]), which is associated with an
2. OpenID Provider (equivalent to an IdP).
3. Users authenticate at the OpenID Provider (that is users "log-in" with *their* OpenID provider),
4. and use that authentication status (not the login credentials themselves) at third party locations (known as Relying Parties [equivalent to an SP]).

The OpenID authentication process is illustrated in figure 1. This example shows how a registered BADC user would authenticate themselves at BADC (acting as an OpenID Provider) as one step in gaining access to data at NCAR (acting as a Relying Party). (Note that in this example, the BADC OpenID Provider, only asserts who the user is, nothing is asserted about what the user is able to do, or what data the user is able to access.)

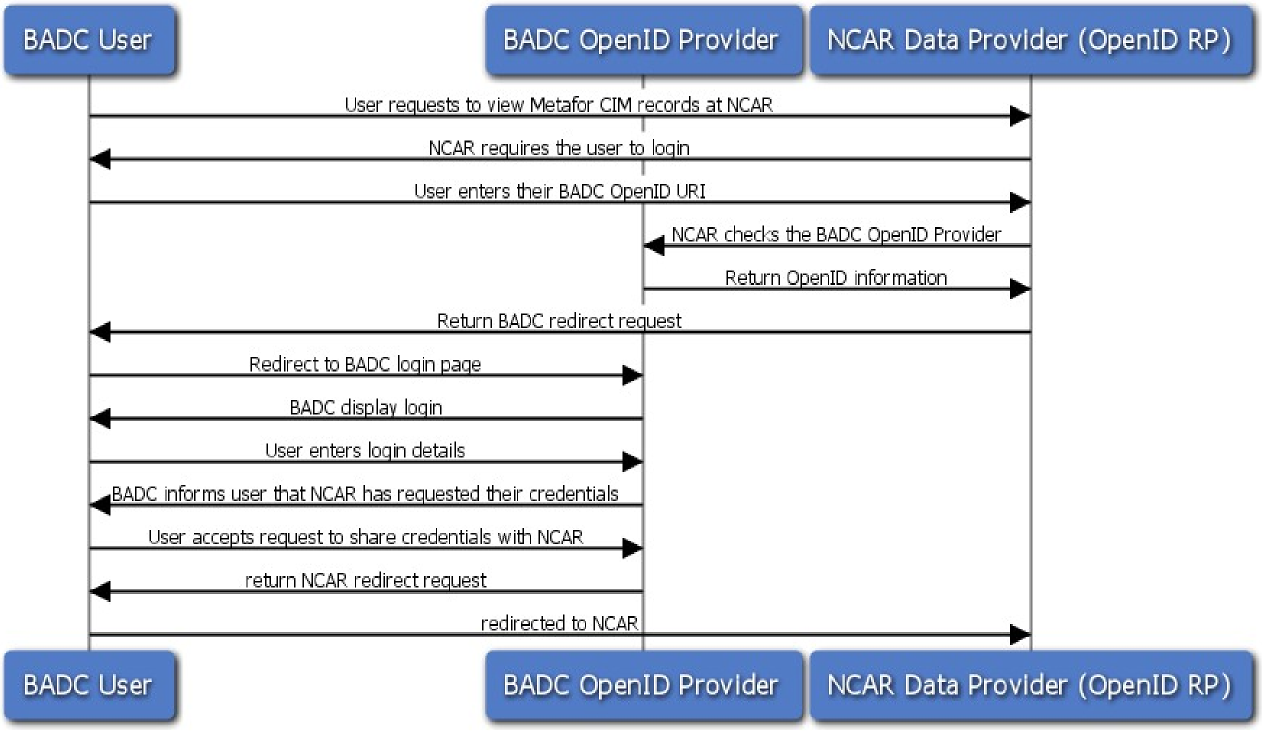
BADC User
BADC OpenID Provider
NCAR Data Provider (OpenID RP)
User requests to view Metafor CIM records at NCAR
NCAR requires the user to login
User enters their BADC OpenID URI
NCAR checks the BADC OpenID Provider
Return OpenID information
Return BADC redirect request
Redirect to BADC login page
BADC display login
User enters login details
BADC informs user that NCAR has requested their credentials
User accepts request to share credentials with NCAR
return NCAR redirect request
redirected to NCAR
BADC User
BADC OpenID Provider
NCAR Data Provider (OpenID RP)

Figure 1. *Sequence diagram illustrating single-sign-on at the British Atmospheric Data Centre using OpenID as part of a data access activity at the U.S. National Centre of Atmospheric Research.*

A key part of Single Sign On is the initial discovery of user identity. With Shibboleth, in most cases a WAYF (Where Are You From) interface is provided to enable the user to select from a static list of trusted IdPs (This restriction was raised as one of the objections to Shibboleth above). By contrast, OpenID IdP discovery is a more open process. The Relying Party interface enables the user to enter their OpenID URI which itself contains the location of their OpenID Provider. This enables a more flexible configuration but carries with it the caveat that the Relying Party must be satisfied with the authentication procedure in place at the user selected Provider organisation.

[1] An XRI is essentially a mechanism for constructing a personal identifier which is independent of a specific domain name (unlike a URI, see http://dev.inames.net/wiki/SRI_and_OpenID for more details).

OpenID is still maturing as a technology. A number of security issues have been addressed with version 2 of the specification but there are still a number of security vulnerabilities. There are various attacks that can be made on the OpenID authentication process (in particular, phishing and man-in-the middle attacks), and so by default, OpenID provides an unacceptable level of security (e.g. reference [6]). However, by securing the connections between the OpenID Provider and the Relying Party (Service Provider) using SSL with mutual authentication, a much more secure system can be delivered. An additional benefit of this sort of approach is that, where desired, sign-on can be restricted to a list of trusted ("white-listed") providers (whether listed in a WAYF page or not).

Users authenticated with trusted Providers can be attributed much greater access rights to resources than users from an unknown source given that the Provider is known to the Relying Party and transactions have been secured by SSL. Importantly, with OpenID, individuals are represented by a unique identifier. This satisfies the requirement to be able to identify users. However, another criticism levelled at OpenID is that this allows unscrupulous parties to monitor user browsing habits, but in practice we think precisely the opposite: scrupulous parties i.e. recognised IdPs and SPs within the federation, will need to monitor user browsing habits, and users accessing data and information in Metafor and ESG will need to explicitly allow this sort of information to be collected (since the originators of the data and information need this sort of browsing information for metrics). Thus, all parties and individuals will be subject to agreed terms and conditions for data access (and all the relevant parties will of course be subject to relevant data protection acts).

Non-Browser Based Access

While OpenID can be deployed easily in an environment where users are interacting with a browser, it is not suitable for a situation where applications are being run in scripts etc, because the interaction with the Provider involves an indeterminate number of steps which are outside of the scope of the OpenID specification and are usually customised locally at each SP in unpredictable

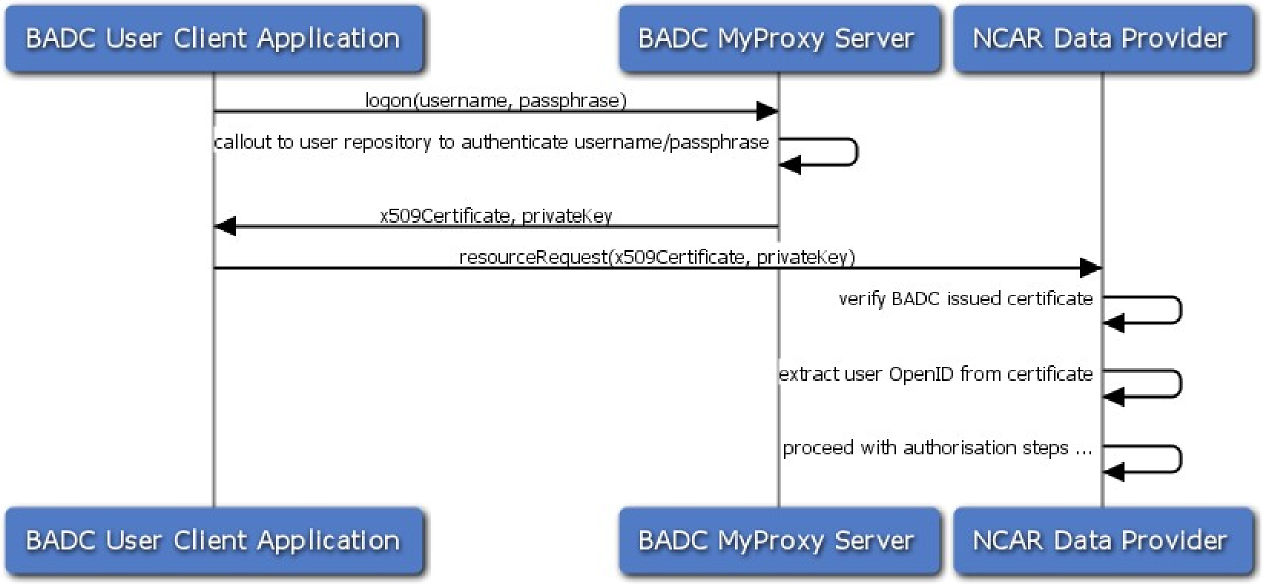
BADC User Client Application
BADC MyProxy Server
NCAR Data Provider
logon(username, passphrase)
callout to user repository to authenticate username/passphrase
x509Certificate, privateKey
resourceRequest(x509Certificate, privateKey)
verify BADC issued certificate
extract user OpenID from certificate
proceed with authorisation steps ...
BADC User Client Application
BADC MyProxy Server
NCAR Data Provider

Figure 2: Sequence diagram showing the actors and interactions for a user client application to sign in at the BADC using MyProxy and use these credentials to authenticate at NCAR and gain authorisation to access a secured dataset.

ways. While scripting is not a high priority requirement for Metafor, it is a key part of the solution that needs to be deployed to work with ESG and IS-ENES, and so we briefly describe it here.

The MyProxy Credential management [7] service fits this use case well since it enables single-sign-on via a discrete programmatic interface MyProxy provides an infrastructure for managing X509[2] security credentials (certificates and private keys), and for generating proxy certificates. Proxy certificates can be used as authentication tokens, and are easily obtained from an online credential repository as needed via the use of a passphrase or via binding to an existing authentication system. An additional advantage of deploying MyProxy is that one can build a certificate authority to construct temporary certificates for either an existing userbase (as is the case for the ESG/CMIP5) or for those users who cannot obtain a certificate from a more authoritative source. These certificates are then used for data and information access.

The middleware to be developed for Metafor will allow either OpenID or appropriate MyProxy certificates to be used for authentication. An example of how single-sign-on will work using MyProxy is shown in Figure 2. Here we see that the MyProxy Service links with the same user credentials repository as the OpenID Provider e.g. a user database. Using MyProxy configuration options, and the X.509 certificate extensions capability, it is possible to include the user's OpenID URI embedded with the certificate and if required, any other additional user attributes. Once the user client application has authenticated with the BADC, it can interact with secured data services deployed at any of the other partner organisations within the federation. NCAR is given as the example in the above. The NCAR access control system can validate that the certificate is from a recognised source and proceed with the authorisation process.

[2] X509 Public Key Infrastructure, see http://en.wikipedia.org/wiki/X.509 for a description of key characteristics.

OpenID Deployment Prototype

As part of the development of a single-sign on capability between the ESG federation and the Metafor group, the BADC team have extended the pre-existing NERC DataGrid security layer to support OpenID. Figure 3. shoes how OpenID has been deployed in the prototype, using screen-shots of running software superimposed on a map.

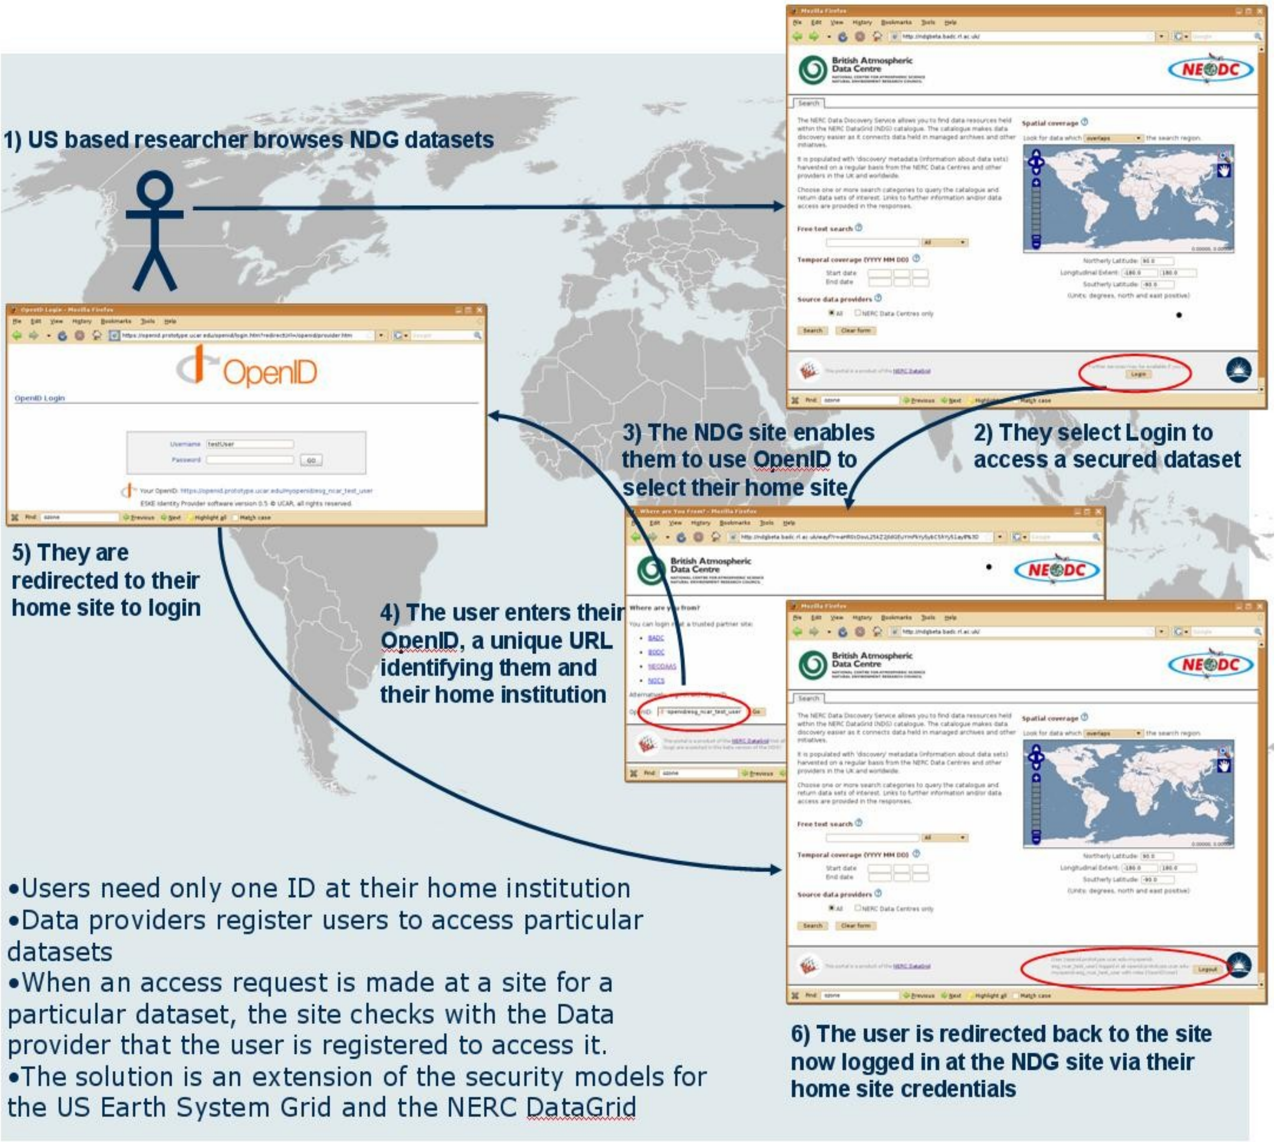

1) US based researcher browses NDG datasets

2) They select Login to access a secured dataset

3) The NDG site enables them to use OpenID to select their home site

4) The user enters their OpenID, a unique URL identifying them and their home institution

5) They are redirected to their home site to login

6) The user is redirected back to the site now logged in at the NDG site via their home site credentials

- Users need only one ID at their home institution
- Data providers register users to access particular datasets
- When an access request is made at a site for a particular dataset, the site checks with the Data provider that the user is registered to access it.
- The solution is an extension of the security models for the US Earth System Grid and the NERC DataGrid

Figure 3. *OpenID is being used at the BADC to provide an authentication service. The steps show a user from NCAR in the U.S. using their OpenID provider to gain credentials within the NERC DataGrid.*

4. Authorization

Once a user has been authenticated, it is then desirable to decide what that user can do; we know who they are, and we can then make choices about what they can do. As discussed above, the authentication framework was constructed in partnership with NCAR, ANL, and PCMDI, and so too was the authorisation framework. Together we have decided to implement a role based access control system which embodies:

1. Authentication
2. Acquisition of user attributes
3. Attributes/constraints that apply to a given resource

Within this role based access control framework:

4. Resources to be secured have attributes (“required roles”) associated with them.
5. Users have attributes associated with them (“roles”) determining what resources they can access.
6. User roles are held in Attribute Authorities
7. Access control decision are handled by Policy Enforcement Points (PEPs) and Policy Decision Points (PDPs). There may be any number of PEPs deployed, securing different resources at different service providers. Any givenPEP, secures the resource it protects by referring access control decisions to a PDP and enforcing this decision allowing or denying access.

Attribute Authorities are resources which hold mappings of user roles against identifiers and we might expect that these are associated with data and information providers (although many providers can share one attribute authority, and indeed, providers can exploit multiple attribute authorities). In bringing together a heterogeneous environment of multiple data providers protecting multiple datasets each with their own pre-existing access policies, a conscious decision was made to support the capability for multiple Attribute Authorities rather than a single Attribute Authority within the federation as for example with VOMS (reference [10]).

In practice, the authorisation protocol is implemented by a Service Provider's access control system maintaining metadata associating the role(s) protecting a resource with an Attribute Authority that has responsibility for issuing role(s) to users. To make an access control decision, user attributes are pulled from the relevant Attribute Authority: the user's (OpenID) identifier, is passed so as to query what roles that user is registered for. The PDP makes an access control decision by comparing the user roles retrieved against the constraints/roles required by resources.

This sequence of events (see fig 4) is relatively standard, and as a consequence there are standards available to request and encode security assertions. We have chosen to use the Security Assertion Mark-up Language (SAML, see reference [8]) for this task. SAML is a standard protocol that defines the interface between different parties brokering access to resources.

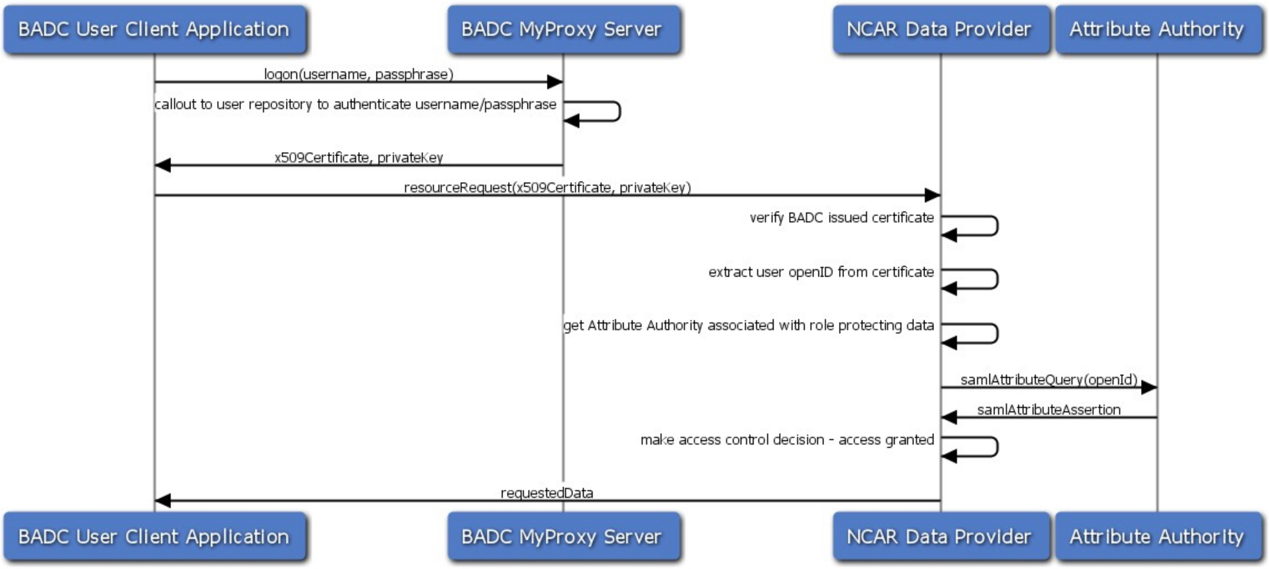
BADC User Client Application
BADC MyProxy Server
NCAR Data Provider
Attribute Authority
logon(username, passphrase)
callout to user repository to authenticate username/passphrase
x509Certificate, privateKey
resourceRequest(x509Certificate, privateKey)
verify BADC issued certificate
extract user openID from certificate
get Attribute Authority associated with role protecting data
samlAttributeQuery(openId)
samlAttributeAssertion
make access control decision - access granted
requestedData
BADC User Client Application
BADC MyProxy Server
NCAR Data Provider
Attribute Authority

Figure 4: T*he sequence of events for MyProxy based Single Sign On and authorisation via a call-out to an Attribute Authority. The NCAR access control system has access to metadata associating the roles protecting access to the data with the Attribute Authority (or Authorities) with responsibility for issuing those roles to users.*

In addition to this role querying interface, Attribute Authorities will also support a web based interface to enable new users to register for roles. This will be an asynchronous process, with the user registration information submitted to an approval pool to assess the users suitability for the access rights in question.

We have yet to build a prototype that tests the proposed authorisation configuration, but the current expectation is that we will build and deploy a working system for CMIP5, in partnership with ESG during the second half of 2009. The agreed federated access control for CMIP5 will only enforce protocols for single-sign-on, attribute management and attribute release. In effect, only the interfaces between organisations to broker access need be defined. It will be up to individual organisations to decide how to apply access control decisions. Therefore Metafor will be able to develop CMIP5 compliant access control that will be able to be integrated into existing infrastructures, see fig 5.

Under the proposed CMIP5 access control system (reference [9]), in addition to vanilla OpenID and SAML, we will exploit an additional property of OpenID, the ability to request additional user attributes beyond the URI alone. In this configuration, we will also request the IdP to provide the SP with a users first and last name, and email address. In practice these user attributes become "federation-level" attributes (and these are necessary primarily to meet data provider access logging requirements). In practice, CMIP5 federation Identity Providers (including those deployed in Metafor) need only collect this minimal number of user attributes, making this an easy federation to join. However, to get access to CMIP5 data, or write access to Metafor information, users will need to establish with Attribute Authorities their "fitness" for specific roles, and this may involve the collection by the organisation responsible for those authorities of additional metadata ("research proposals", "job-function", "project-status" - whether "commercial" or not etc). However, that metadata is not exposed, all attribute authorities will do is expose roles that users have; this other metadata is simply used to establish those roles, and will be private to the organisations collecting it.

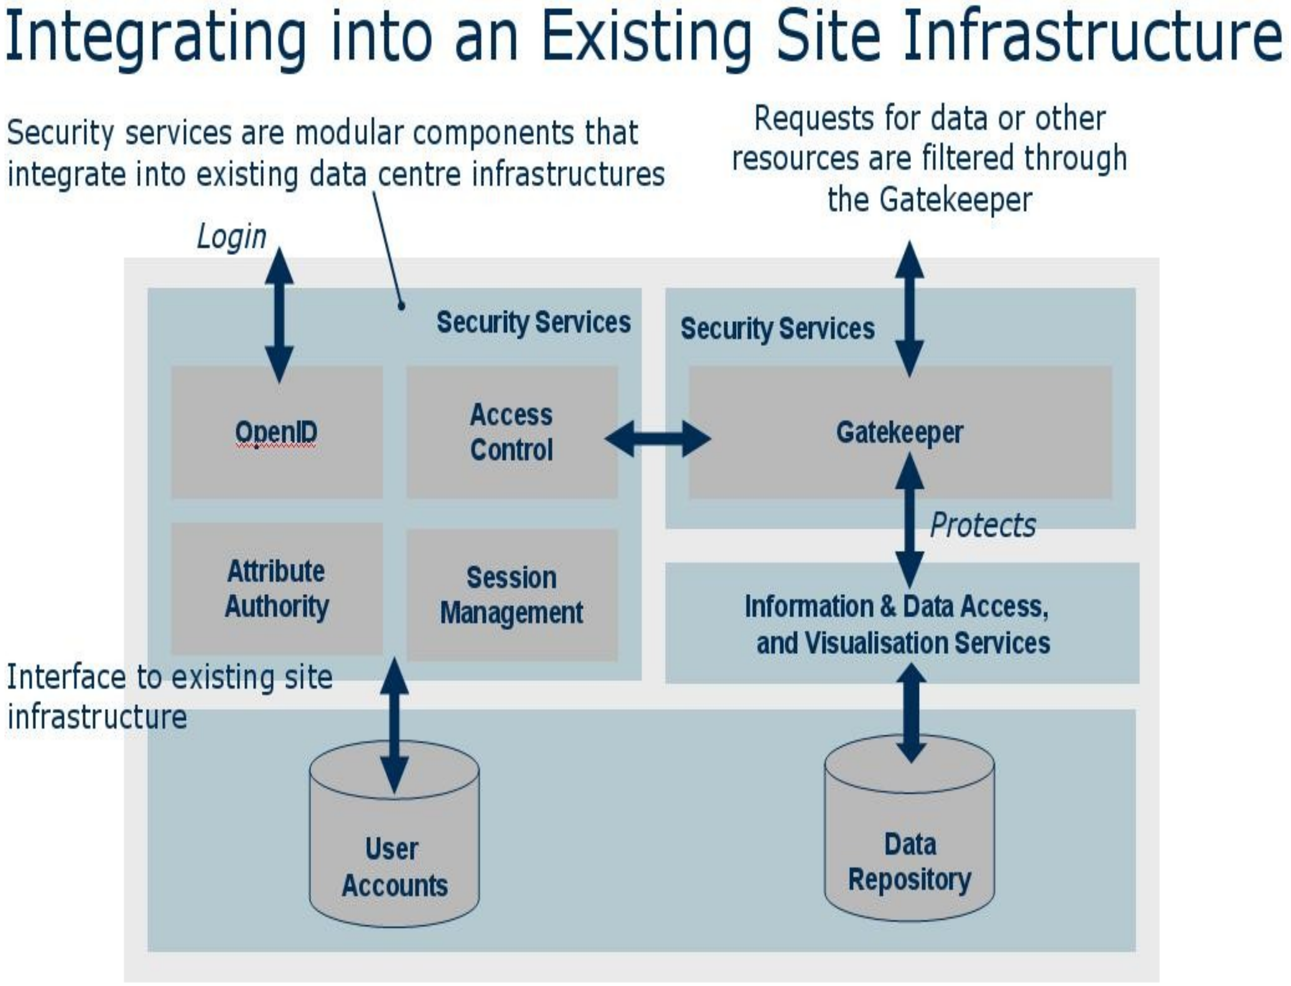
Integrating into an Existing Site Infrastructure
Security services are modular components that integrate into existing data centre infrastructures
Requests for data or other resources are filtered through the Gatekeeper
Login
Security Services
Security Services
OpenID
Access Control
Gatekeeper
Protects
Attribute Authority
Session Management
Information & Data Access, and Visualisation Services
Interface to existing site infrastructure
User Accounts
Data Repository

Figure 5. *How federated access control fits into existing infrastructure. The key requirements are that modular components are deployed as services interfacing with existing infrastructure. Gatekeepers (Policy Enforcement Points, PEPs) intercept all transactions that need to be secured, and refer decisions on access to Access Control modules (Policy Decision Points). Session Management services hold access control state, and OpenID and Attribute Authority services can interact directly with existing user account structures via whatever interfaces exist to those established systems.*

5. Conclusions

1. Metafor single-sign-on technology needs to be deployed at existing sites with thousands of users, and in partnership with other major international activities.
2. Given that the most important use-case for Metafor is currently to work with climate data associated with the CMIP5 project, the agreement of a common access control paradigm with the Earth System Grid is crucial for the success of Metafor.
3. A common access control paradigm has been agreed for Metafor and CMIP5, and consists of deploying:

   a) OpenID as the authentication framework for browser based usage, and MyProxy based certificates for non-browser based usage.
   b) Attribute Authorities which expose user roles over a SAML interface and enable users to register to access secured resources with those roles.
   c) Policy Enforcement Points which constrainn access to resources
   d) Policy Decision Points which make access control decisions.

4. We have deployed prototype authentication services which conform to the agreed paradigm,and are working on the other services. The code developed thus far could be extensible to support Shibboleth for browser based single-sign-on if necessary.
5. Software and protocol selection is based on available widely-deployed standards, which lowers entry costs (in terms of software development and deployment costs), improves functionality (in terms of exploiting existing experience to get functionality “out-of-the-box”) and since the standards are peer reviewed it reduces the risk of security vulnerabilities.

6. Bibliography

(all weblinks retrieved June 10th to check current validity)

[1] Shibboleth Homepage: http://shibboleth.internet2.edu/
[2] Shibboleth Protocol Document: http://shibboleth.internet2.edu/docs/internet2-mace-shibboleth-arch-protocols-latest.pdf
[3] Shibboleth Technical Overview: http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-tech-overview-latest.pdf
[4] Evaluation of Shibboleth for ESG Web SSO, ANL, r9, 3 April 2008, http://www.ci.uchicago.edu/wiki/bin/view/FranksProjects/ShibbolethEvalESG
[5] OpenID Technical Specification: http://openid.net/specs/openid-authentication-2_0.html
[6] OpenID Security Research/Evaluation Page, ANL, http://www.ci.uchicago.edu/wiki/bin/view/FranksProjects/OpenIdSecurity
[7] MyProxy Credential management Service: http://grid.ncsa.uiuc.edu/myproxy/doc.html
[8] Security Assertion Markup Language (SAML): http://saml.xml.org/
[9] ESG Security Architecture for IPCC AR5 Interoperability, v0.2, P J Kershaw, 24 October 2008, http://proj.badc.rl.ac.uk/ndg/export/5003/TI12-security/trunk/documentation/esgInteroperabilityForIPCCar5/ESG%20Security%20Architecture%20for%20IPCC%20AR5%20Interoperability.doc
[0] Alfieri R, et al. VOMS: an authorization system for virtual organizations, 1st European across Grids conference, Santiago de Compostela. (http://grid-auth.infn.it/docs/VOMS-Santiago.pdf)