



JISC Final Report

Shibboleth Access to Resources on the NGS (SARoNGS) – FINAL Report

Andrew Richards, Claire Devereux, Jens Jensen, Xiaodong Wang, Mike Jones,

Contact: Andrew Richards (Andrew.richards@stfc.ac.uk), 14/09/09
STFC e-Science Centre
Rutherford Appleton Laboratory
Harwell Science and Innovation Campus
Chilton
Didcot
Oxfordshire
Ox11 0QX

Table of Contents

Contents

.....	1
JISC Final Report	1
Table of Contents.....	1
Contents.....	1
Acknowledgements.....	1
Executive Summary.....	2
Background.....	2
Aims and Objectives	2
Methodology	3
Implementation	7
Outputs and Results	8
Outcomes.....	8
Conclusions	9
Implications	10
Recommendations (optional).....	10
References.....	10
Appendix A – SUM submitted to e-Framework for SARoNGS project.	11

Acknowledgements

This project was funded solely under the JISC e-Infrastructure programme. Institutions as part of this collaboration were: STFC (Rutherford Appleton Laboratory and Daresbury Laboratory), University of Oxford, University of Manchester (RCS) and MIMAS (University of Manchester).

Outwith the project, collaborations included University of Reading (David Spence) as part of joint collaboration through University of Oxford, and the VPMAN project which includes the University of Kent. The collaboration with the VPMAN project was to facilitate taking the PERMIS decision engine and incorporating with the overall SARoNGS authentication infrastructure. Overall the technical input

from the VPMAN project and associated collaborators at the University of Kent and the University of Glasgow benefited the project overall. The NGS remains engaged with the VPMAN project to ensure the adoption of PERMIS technologies appropriately within the UK e-infrastructure.

The project also sub-contracted Curtis and Cartwright for the contribution to the e-Framework.

Executive Summary

The aim of the project was to deliver into production a Shibboleth based infrastructure for the NGS, to enable HEI users/researchers to access NGS resources using their institutional identities as provided through membership of the Access Management Federation (AMF) and to show its utility via a domain specific demonstrator. The elements pertaining to this project are encapsulated in four main activities.

- to provide grid authentication tied to the UK AMF (a new service based upon outputs from the ShibGrid project)
- to link this authentication token with VO attributes from the grid computing domain (new)
- to translate attributes within the context of UK AMF into attributes suitable for consumption by grid computing infrastructures (a new service based upon the outputs of the SHEBANGS project)
- to demonstrate these via both subject based and generic demonstrator applications (new).

Background

Building on the work of the SHEBANGS and ShibGrid projects the project aimed to take the outputs from both these demonstrator projects and provide a production ready service for the NGS. To achieve this aim, existing Shibboleth work from the GEMS projects was integrated as part of the development, in order to provide real world examples of where the Shibboleth authentication model can be used and integrated with a production grid service.

The ShibGrid project developed a prototype to enable NGS users with or without UK e-Science certificates to securely access those resources, through the integration of Shibboleth, GSI (Globus Grid Security Infrastructure) and MyProxy. This aimed to bring the confidentiality and privacy aspects of Shibboleth to both grid users and service providers, while making the tools as easy as possible to use. Users are able to access internal and external resources seamlessly using a single institutionally controlled identity.

The SHEBANGS projects aims were divided loosely into three parts, the first part of the project focussed on the development of the basic *Credential Translation Service* (CTS). This CTS creates GSI credentials which are then delegated to a trusted MyProxy server for the consumption by a portal. The second part incorporates VOMS (Virtual Organization Membership Service) assertions created at the CTS into these GSI credentials. The third part incorporates Identity providers from the FAME-Permis project and MIMAS SHIMMER project.

Aims and Objectives

The aim of the project was to deliver into production a Shibboleth based infrastructure for the NGS, to enable HEI users/researchers to access NGS resources using their institutional identities as provided through membership of the Access Management Federation (AMF) and to show its utility via a domain specific demonstrator. The elements pertaining to this project are encapsulated in four main activities.

- to provide grid authentication tied to the UK AMF (a new service based upon outputs from the ShibGrid project)
- to link this authentication token with VO attributes from the grid computing domain (new)
- to translate attributes within the context of UK AMF into attributes suitable for consumption by grid computing infrastructures (a new service based upon the outputs of the SHEBANGS project)
- to demonstrate these via both subject based and generic demonstrator applications (new).

The core SARoNGS infrastructure to achieve this has been successfully developed and deployed as NGS services during this project. Both technical and policy based issues were encountered near the end of the project preventing the CTS from operating with in the UK AMF. These were not resolved until recently and were due to both technical anomalies and restrictive policy which were difficult to resolve with ja.net and the AMF. At the time of writing this report, however, these issues have been resolved sufficiently for the Credential Translation Service to go into production with the corresponding integration with the UK Access Management Federation. A live CTS service can be found here <https://cts.ngs.ac.uk/>. Tests have successfully been completed of the CTS against the IdPs at the University of Manchester, MIMAS, STFC and Glasgow NeSC.

As of the end of the project the MIMAS prototype component, commissioned as part of the project to demonstrate the usefulness of the SARoNGS infrastructure has deployed a test geo-spatial data service based upon the Erdas Apollo Image Manager. This demonstration had been shown at a number of conferences and can be found by following the link from <https://cts.ngs.ac.uk/>.

Methodology

The core SARoNGS infrastructure to achieve this has been successfully developed and deployed as NGS services during this project. Both technical and policy based issues were encountered near the end of the project preventing the CTS from operating with in the UK AMF. These were not resolved until recently and were due to both technical anomalies and restrictive policy which were difficult to resolve with ja.net and the AMF. At the time of writing this report, however, these issues have been resolved sufficiently for the Credential Translation Service to go into production with the corresponding integration with the UK Access Management Federation. A live CTS service can be found here <https://cts.ngs.ac.uk/>. Tests have successfully been completed of the CTS against the IdPs at the University of Manchester, MIMAS, STFC and Glasgow NeSC.

As of the end of the project the MIMAS prototype component, commissioned as part of the project to demonstrate the usefulness of the SARoNGS infrastructure has deployed a test geo-spatial data service based upon the Erdas Apollo Image Manager. This demonstration had been shown at a number of conferences and can be found by following the link from <https://cts.ngs.ac.uk/>.

The architecture of SARoNGS is complex. It revolves around the CTS, with the CTS acting as a hub of internet traffic between and around the user's browser, the Federation, the UK eScience Certificate Authority, virtual organisations, MyProxy services, grid portals, and grid-enabled resources. The SARoNGS project has used the technological exercise described above not only to provide a simple yet secure entry point to grid computing but also to attempt to bring together the conflicting requirements of users and resources. SARoNGS has formed the first tendrils of this process but heavy-weight compromises have had to be made in order to bend to rules of both grid and federated identity management.

The arguments being presented that access to resources should be simple and straightforward are very often opposed to the arguments that access to resources should be safe and secure. Compromises need be made to show the community at large what federated, grid-enabled services can provide. On the other hand rolling out computational systems with such power as the grid provides access to requires a big commitment on the part of the resource providers themselves and thus the trend toward anonymity and unmeasured application of data protection cannot be sustainable if the community wants its choice of more and better resources. We need to establish an environment where risk is assessed properly to avoid situations where we have policies which are either too lax or too paranoid.

Connecting SARoNGS with the federation

Getting CTS as a Shibboleth Service Provider (SP) working with the federation also took some time. We were initially concerned that registering it with the federation would immediately open it to the public; this turns out not to be true. While the service would in principle be accessible by the public, the announcement that the service is available is made as a separate step, once we are sure it is working.

The CTS is a Service Provider as far as the UK AMF is concerned. For the CTS to work as an SP it must have a certificate key pair (as part of a PKI). We first tried a host certificate from the UK e-Science CA, and it took a while to figure out that it didn't work because the federation metadata

published by the UK AMF included the wrong CA certificates for UK e-Science (an obsolete set). On finding this out (a task for which our only option was to analyse and validate the federation metadata ourselves) we informed the Federation. The Federation on receiving this information then informed us that the UK eScience CA was not supported despite the documentation¹ stating otherwise. On learning that UK e-Science CA wasn't "really supported" (because of the need to install CA certificates in browsers – which is not necessarily true²) we were given slightly conflicting advice regarding the CTS' certificate. It was at some point recommended that we try self signed certificates (like a CA root certificate except not able to sign other certificates) for signing the SAML assertions in the SP, and we tried quite a few variations of those with different extensions, although for some reason which we still don't fully understand it never worked. Our final option was getting an SCS certificate from GlobalSign via JANET and maintain the non browser facing component of the CTS with this certificate and the user and grid facing components with the UK eScience certificate. The process of trial and error with the live Federation exacerbated the difficulties we experienced setting up the service provider due to the delay in propagation of the Metadata. Our changes were submitted to the federation, some time later handled by the federation and changes to the Metadata occurred after close of business. A further delay was experienced as we waited for the Identity providers to install their local copy of the metadata.

Getting an SCS certificate (normally required for browser-facing hosts) took longer than expected, mostly because JANET's guidelines for getting one (where we started) were less specific than STFC's. For example, we had to regenerate the request because the name turned out to be wrong according to STFC's guidelines, and STFC's registration authority was not registered to request certificates for the NGS namespace.

Limitations of the SARoNGS Grid credentials

It should be noted here that the certificates are only as good as the material upon which they are based. Ideally, the NGS would have liked to have the SARoNGS CA to become accredited with the [International Grid Trust Federation](#) (IGTF), like the UK e-Science CA is. However, this is currently not possible. Briefly, the issues are:

- **Permitted reuse of eduPersonTargetedId** The UK federation permits reuse of ePTID – see Rules of Membership section 6.4.2 – and this is just for IdPs signed up to section 6 – IdPs not signed up to section 6 can do whatever they like. In grid terms, this means that within 24 months (for institutions signed up to section 6) after a credential has gone out of use, the SARoNGS CA could inadvertently issue credentials with the same name to someone else.
- **Names are not published** It is currently required that all personal certificates have a DN with at least one commonName component containing a "reasonable representation" (i.e. represented in letters from the English alphabet) of the person's real name. Firstly, IdPs are strongly discouraged from publishing attributes other than the four core ones (Technical Recommendations for Participants, section 7.3), particular ones containing personal data such as the commonName. Another minor complication is that the commonName attribute, as defined in the eduPerson schema, does not always contain the common name as required by IGTF. eduPerson says it is "typically the person's full name."
- **Id Management Policies too numerous/ varied** A SLCS (Short Lived Credential Service – up to about 10 days) must carefully explain how a person's real life identity is translated to the Grid. This means getting the information from the IdPs. This is fine, provided the IdPs have signed up to section 6, as they are required by the Rules of Membership item 6.2 to document this process. The trouble arises from the fact that SARoNGS must gather all of these and condense them into a single document. In other words, a baseline sufficiently strong for IGTF must be set, and the SARoNGS CA must then evaluate each IdP to assess whether it meets the baseline requirements for some or all of its users. If the baseline is met only by a subset of users, it must be possible for SARoNGS to filter those out.
- **Revocation vs Lifetime** If long lived certificates are issued, the question about timeliness of revocation arises. Suppose a user leaves the institution – this fact is not communicated to the CA. At least with Classic CAs, the RA will reasonably attempt to revoke the certificate, even if

1 <http://www.ukfederation.org.uk/library/uploads/Documents/technical-recommendations-for-participants.pdf>

2 While the user facing components may choose to use certificates based upon pre-bundled CA sets included with most browsers, these user-facing interfaces do not rely upon the metadata issued by the Federation.

the CA cannot provide an overall guarantee. Conversely, for short lived certificates, the CA needs to ensure that they live long enough for the longest jobs.

The lack of accreditation leads to issues for resources using the IGTF RPM repository³ (and possibly other mechanisms surrounding the IGTF) to manage their CA certificates, NGS *core* sites being no exception to this. This issue arises from the unwillingness of the IGTF to distribute signing_policy files which allow root CAs to sign subordinate CAs which themselves are not accredited. Hence there are two versions of this signing_policy file and automated updating of trust roots is no longer reliably, and work-arounds have had to be put into place.

Linkage with VO Attribute Authorities (VOMS)

Due to the location of the credentials created it was not possible to provide the users with the usual access to the grid VOMS services. Initially we planned to provide a registration facility which would be managed by the CTS. As the CTS was being developed it became apparent that we could reuse some of the code – code designed to perform the communication between the CTS and the VOMS server – to make the CTS behave as a simple, secure web proxy. This mechanism opened up a more direct relationship between the user and the VO allowing them to join a VO, and view and maintain their membership details held by that VO.

This approach removed the requirement to create a local user registration system and hence the need to take the user through such a process. The approach worked until the point where the user's email address needed verification. VOMS requires the email address of the user to be verified, and does so by sending a confirmation URL to the user's email address. The difficulty presents itself when the VOMS server sends the message telling the user to visit the VOMS server directly and further by also requiring that the link be secured via a certificate authenticated HTTPS connection. We needed to provide a mechanism for the user to enter this URL via the CTS. The CTS was consequently enhanced to provide a VO registration form. The form asks for the same information that the VOMS server requires and submits it via the same POST mechanism that the VOMS server expects. This allows the CTS to describe to the user what to do with the email that they will receive from the VOMS server, and duplicates these instructions in an email which will arrive at the users email address at about the same time the VOMS server's email arrives. The extra instruction is simply to cut and paste the URL into a web form on the CTS.

Translation of Federation Attributes

The Federation requires eduPersonScopedAffiliation should be passed via the Shibboleth mechanisms to Service Providers, and recommends that eduPersonTargetedID be used only if persistency (i.e. account management) is required. A successful Shibboleth authentication alone does not guarantee the relative memberships to institutes to satisfy certain license agreements⁴. To avoid issues of direct attribute release the CTS makes use of a policy-based release of attributes which may be added to their grid credentials at the user's request. The aim was to have a call out to a policy decision engine, supplying Shibboleth attributes and requesting authorisation for grid based attribute release. We employed PERMIS to do this given its speed, ease of policy management and its ability to be called via network protocols (removing the need to involve complex API calls from within the CTS Perl scripts). With some extra development work we were able to host the PERMIS in a tomcat container, From the CTS side we developed a SOAP, SAML, XACML client and were able to make authorisation requests to the web service. There were a number of difficulties experienced when trying to use PERMIS in this way:

- **Size of Policy** due to PERMIS only dealing with attributes as complete entities and not allowing policies to parse these attributes, the policy describing the grid attribute creation action based upon Federation attributes was so large that the we had problems uploading the Policy to the LDAP service (this is how PERMIS reads its policies), the Java invocation parameters required special memory settings, and the PERMIS caching process created a substantial delay each time a new request required a new cached policy.
- **Multiple attributes** IdPs provide a number of attributes in this case eduPersonScopedAffiliation may have multiple values e.g. *staff*, *member* and *alum*. We found that in placing all three attributes into the XACML request an exclusive decision based upon

3 <https://dist.eugridpma.info/distribution/igtff/current/>

4 <http://www.ukfederation.org.uk/library/uploads/Documents/technical-recommendations-for-participants.pdf> §7.1.2.1

all the attributes was made i.e. all actions were based on a logical AND of the privileges supplied.

Having dealt sufficiently with the size issues, and having provide the CTS with the ability to make independent calls to the decision engine for each attribute, the CTS is able to make a policy decision on what VO credentials it is able to grant to an individual within this VO which we have called *ukfederation.ngs.ac.uk*.

Further issues were noticed when expanding our tests to use other Identity Providers.

- **Case Sensitivity of Attributes** We noticed that some IdPs specified eduPersonScopedAffiliation (ePSA) in upper case and some in lower case. PERMIS is unable to handle this as it matches attributes in a case sensitive manner, and rather than double the size of our policy to deal with the two common forms of these attributes (all capital and all lower case) the CTS translated all ePSA attributes to lowercase before querying the PERMIS service.
- **OIDs vs eduPerson** there is an overlap between attribute spaces e.g. uri:oid and uri:mace:eduPerson. Specifically eduPersonTargetedId and oid:1.3.6.1.4.1.5923.1.1.1.10, both these attributes describe the same entity and should only exist once. In our tests we found one IdP publishing assertions with both these attributes the values of which differing.

MIMAS Geo-spatial demonstrator

The provision of a demonstrator case for the SARoNGS project was through the implementation of the SARoNGS outputs to a new find-view-download service being developed by Mimas-Landmap service. The service utilises a web interface to Erdas Apollo Image Manger (EAIM). The EAIM allows users of the service to locate, view and download datasets of their area of interest. Requests to the Image data repository and database are mostly handled using Open GeoSpatial Consortium requests, such as WebMapService (WMS) providing previews of the data, Web Coverage Service, (WCS) providing download of the data, and Web Feature Service (WFS), providing download of Vector data. The OGC standards also allow for Web Processing Service (WPS) requests to be performed. This is is not currently supported by EAIM of the User Interface (EAIM-UI).

The EAIM-UI allow access to the various OGC functions through local authentication, using either a text file or a database solution.

The Demonstrator needed to prove that users can log into the EAIM-UI select data and then perform a processing request and receive a copy of the results. This was performed in the following manner:

1. The EAIM-UI was extended to integrate the SARoNGS security Solution to allow access to the data.
2. A WPS server was deployed on a workstation to simulate a GRID environment where LANDMAP adapted a Grid-enabled WPS (G-WPS). This was provided by 52north (www.52north.org) G-WPS is a java based program which uses the open source Sextant geo-processing algorithm library and runs in tomcat server The library contains over 250 algorithms that can be performed on raster and vector data including image classification supervised and unsupervised, such as NDVI etc.
3. The EAIM-UI was changed from a find-view-download to a find-view-process-download_results service. Whence the processing request has finished by emailing the user an FTP link to the results of the processing. This is an adaptation of regular WPS which are usually done while waiting to cater for the Asynchronous nature of Grid Processing.

In the Landmap application the selected data for processing are copied to the NGS database before processing. As proof of concept Landmap run a WPS NDVI algorithm on Landsat 7 data successfully. The fact that the obtained results were performed faster than on a desk top, is not interesting, however, the fact that you could schedule such a task for a large number of images is of interest.

Two things should be noted here:

1. the G-WPS are not truly Grid Enabled as most Geo-processing algorithms are not parallelized and the splitting of data sets into smaller subsets that are processed in parallel, then the result are re-assembled is not utilized. The nature of algorithms available for WPS though interesting are not those would lend themselves to huge requirement in processing power.

2. In the context of Remote Sensing data it is not advisable for the data and the processing power to be separated from each other as data volumes are usually huge. Though technically OGC standards cater for this separation, practically the shifting of great amounts of data from one server to other in order to do the processing is not advisable.

EAIM is expected to provide WPS capabilities through their continuous improvement and development of their software. Landmap will be investigating the possible deployment of an experimental service utilising these new capabilities.

Implementation

The project was scoped to take existing work from the Shib-Grid and SHEBANGS projects into production services for the NGS. Given the short timescale of the project and the delay in funding the project as a follow-on from the previous work, it was challenging to obtain access to staff with the relevant skills within the required time. This was most challenging at MIMAS, where the exemplar service using the SARoNGS framework has been developed. It was also initially planned to work with EDINA as well, but time constraints and existing project commitments did not make this feasible and it was requested to the JISC that EDINA did not have to be involved with the project.

The solution developed uses directly the web based (Shibboleth) approach of the UK Access Management Federation. Whilst initially the project aimed to deliver web service based command line tools, these were dropped early in the project as 1) they could not work with the current methodology of the Access Management Federation where users are directly authenticated by their home institute. Such command line client tools would require handling the credentials locally. And, 2) the scope of the rest of the project to deliver SARoNGS services in production required more development effort than initially anticipated and so the project was focussed on ensuring the backend services were delivered properly to provide a platform on which to build web based interfaces for accessing the NGS.

As a service to users SARoNGS has been developed to use the existing NGS Portal as an exemplar service whereby users can register with the NGS and use their institutional username and password via the Access Management Federation to access NGS enabled grid resources.

Starting the project again, it would be recommendable to focus the development activities purely on the core SARoNGS components and to not attempt to integrate with too many service providers and outputs from the VPMAN project. Specifically the planned use of PERMIS for VO membership decisions was envisaged to make the policy setting simple; this turns out not to be the case, policies still need to be manufactured pragmatically and are too large to be processed by humans. A similar policy could be effected in a few lines of Perl within the CTS application itself. For a small short project too many individual sites have been involved. It may be worth recommending to other JISC development projects of this scope that activities are not dispersed over too many project partners.

Applicable Standards

X.509: <<http://www.ietf.org/html.charters/pkix-charter.html>>

RFC 3820 – “Internet X.509 Public Key Infrastructure Proxy Certificate Profile”

RFC 3281 – “An Internet Attribute Certificate profile for Authorization”

SAML: <<http://saml.xml.org/saml-specifications>>

Shibboleth: <<http://shibboleth.internet2.edu>>

Member Integrated X.509 PKI Credential Services (MICS): http://www.tagpma.org/authn_profiles

In addition to the above standards, the published APIs for the following programmes are used:

MyProxy: <<http://grid.ncsa.uiuc.edu/myproxy>>

VOMS: see *Virtual Organisations on the Grid* SUM

Outputs and Results

<Explain the end result of the project work in an objective way. Depending on the project, it might include research results, findings, evaluation results, data, etc. If the project created something

tangible like content, a portal, or software, describe it. Engage the reader, and avoid a long list of deliverables.>

To summarise, the deliverables from this project are;

- Project management work plan and final report
- Interim requirements report after 2 months of project start, final report within 1 month of the end of the project
- A Modified Credential Translation Service (CTS) which works with the “Shibbolised” MyProxy CTS to be enabled to speak the “Shibbolised” MyProxy protocol.
 - CTS to present user interface (VO registration and Credential Management interface).
 - Incorporation of PERMIS into CTS (to allow the VO which hosts the CTS to create policy based decisions for credential release).
 - CTS to provide database to store user preferences
 - CTS to provide WS interface
- A modified MyProxy Server
 - Accepts password change mechanisms based upon SAML assertion authentication.
- One instance of CTS registered as an SP in the UK Federation, to permit every user from every institution with an IdP access to the NGS.
- Software, including various integrated upload and download tools supporting non-portal access.
- Secure website service for use with Shibboleth (to host the CTS).
- Production Deployment of PERMIS from the VPMAN project.
- Exemplar service for accessing MIMAS data sets hosted on NGS Production hardware.
- Project website, documentation and international workshop on usage of VO enabled Shibboleth within a grid environment.
- Outputs submitted to e-Framework.

Outcomes

- The NGS will be able to expand its user base by giving them SARoNGS certificates rather than e-Science certificates;
- For most (ideally, all) services, it is expected that the users will not need to know they have certificates – they log in with their home institution id.
- The SARoNGS infrastructure can easily be adapted to work with other portals. This is advantageous because there is already more than one portal, and because the only way to ensure the user retains the same single identity is by using the same Service Provider (because we translate the eduPersonTargetedId because it is easier to get than eduPersonPrincipalName).
- The close integration with VOMS will help promote VOMS as authorisation management (but will also work with non VOMS-aware services). The fact that we have a simple integrated VOMS server means that we can also manage the roles for the CTS itself internally, without relying on an external VOMS server.
- Integration of the SP with the UK Federation will take a lot more time than expected, at least for non-trivial services – the complexity should not be underestimated.

At the Computing in High Energy Physics (CHEP) conference in Prague, 2009, the need for Shibboleth access to the grid was highlighted in a plenary. It was good to be able to stand up and say we have solved the problem, and more. Likewise, the Terena Grids-and-NRENs working group has expressed an interest, particularly as we were not the only ones who had expressed a need for this. If we can package the software neatly (and make it modular and customisable), we can increase the impact, as others can then adapt it to their own federations which, as a rule, are different from the UK one.

In summary, whenever we have had opportunities for opportunistic dissemination, there has been a lot of interest in this work.

Contributions to the JISC e-Framework

Many of the JISC e-Infrastructure projects are currently either using the current PKI system which some users find difficult and cumbersome or bespoke solutions developed on a per project basis.

This project will contribute a common platform for authentication and authorisation, by building on the existing Shibboleth work of the ShibGrid and SHEBANGS projects for authentication and work to integrate with the VPMAN project for authorisation.

To the e-Framework this project contributed, via the work of Curtis and Cartwright`:

- A common security and authorisation platform that can easily be adopted by other service providers across the local, national and international grid infrastructures. Individual components will be contributed as SERVICES to e-Frameworks.
- Extend the UK Access Management Federation to enable users to access grid enabled resources. For example, enable access to existing services such as MIMAS. Contribution to e-Framework SUMS as use cases on integrating MIMAS and reference documentation provided under the e-Framework GUIDES.
- Recommendations on attribute requirements for the UK Access Management Federation, in order to fully support grid infrastructures. This will be as a contribution to e-Framework GUIDES.

Conclusions

Certificate based access to grid resources has typically been reported as difficult to use by the research community within the UK and internationally. Whilst the need for the UK CA to continue to issue medium level assurance certificates will continue to be needed for the foreseeable future, especially for researchers collaborating at an international level, it is strongly recognised that within the UK HEI community the adoption of a Shibboleth based authentication framework, that enables researchers to access enabled resources using only their institutional ids (username/password) is preferable and more scalable. By utilising the growing HEI base now enabled within the ja.net Access Management Federation, the NGS will be able to support an increasing number of HEIs to access the increasing variety of e-infrastructure enabled by the NGS.

Implications

The basis SARoNGS framework deployed by the NGS as a production service enables any HEI in the UK to use NGS services such as the NGS Portal to access grid resources easily, or to take the SARoNGS framework and Shibboleth enable their own portal infrastructure to allow access to grid resources. It has been proved as part of this project that it is easy in a short time to Shibboleth enable a portal to be able to access NGS resources. This approach will make it easier for the NGS to work with HEIs to support collaborative research requiring access to e-infrastructure.

Future developments of the SARoNGS project would develop better integrated solutions and tools to enable access to grid resources from a wider variety of client applications, other than just web portals. To achieve this would require not only tool development but service development on the side of the access Management Federation. For example, the NGS would like to enable ssh like access to grid resources where the user can utilise their membership of the Access Management Federation via their institutional ID to assert identity. However, at present the policy of not allowing a third party to handle the users ID, specifically password, would mean that the credential information could not be passed to the Institutional IDP via a ssh server component. This is one example that prevents the current NGS command line tools being able to easily interface with the Access Management Federation.

Recommendations (optional)

1. A review of the Access Management Federation and policies defined on the release of information should be undertaken to enable the broader adoption of e-infrastructure within the UK.
2. Inter country exchange of identity via Shibboleth mechanisms would be beneficial to the long term development of a national grid initiative in the UK and its ability to interact across Europe using similar identity management frameworks. This would require EU policies and research work to be undertaken.
3. Researchers should be supported and encouraged to adopt Shibboleth technology within user facing tool developments. Use of the Access Management Federation and the SARoNGS based infrastructure needs broad scale adoption across institutes to improve the users experience of accessing e-infrastructure.

References

1. SARoNGS: Shibboleth Access to Resources on the NGS. To be presented by Mike Jones on the UK AHM 08
2. Shibboleth Access for Resources on the National Grid Service (SARoNGS). Xiao Dong Wang, Mike Jones, Jens Jensen, Andrew Richards, David Wallom, Tiejun Ma, Robert Frank, David Spence, Steven Young, Claire Devereux, Neil Geddes. Proc. of the 5th International Conference on the Information Assurance and Security
3. SARoNGS: Shibboleth Access for Resources on the National Grid Service. Journal of Information Assurance and Security 5(2010) 293-300.
4. Shibboleth Access to Resources on the NGS. Mike Jones, Robert Frank, Amer Alorichdi, Kamie Kitmitto, Anja Le Blanc, Tiejun Ma, David Spence, Steven Young, Jens Jensen, Andrew Richards, Xiaodong Wang, Andrew Rowley. To be presented by Mike Jones on the UK AHM 09.

Appendix A – SUM submitted to e-Framework for SARoNGS project.

e-Framework Service Usage Model Description

SUM Template v7.2 20070725 © Copyright, e-Framework Partners, 2008
SUM Content © Copyright Curtis+Cartwright Consulting Ltd, 2008 1

e-Framework Service Usage Model Name

Name: Shibboleth/X.509 translation to enable access to Grid resources

Alternative Names: Shibboleth access to Grid resources, SARoNGS

Version

0.2

Version History

Include requested information about all versions of this document.

Version Date Author Description Organization / Project

0.1 1/12/2008 Curtis+Cartwright Consulting Ltd

Initial Draft SARoNGS

0.2 11/12/2008 Curtis+Cartwright Consulting Ltd

Updated following internal review

SARoNGS

Rationale

This SUM provides a mechanism for users with Shibboleth identity credentials to access services which rely on an X.509 authentication infrastructure. Shibboleth is widely used for federated access management, and X.509 is used within Grid computing. Translating between the two allows a much wider range of users to utilise Grid resources.

This SUM does not describe the authorization decisions which relying parties take based on the generated X.509 credential – only the process of generating the credentials and making them available.

Classification

To be provided by the submitter:

SUM Type Domain CORE (a commonly recurring SUM; designation requires e-Framework Integrity Group approval)

Domain(s) Learning & Teaching

Research

Libraries

Administration

IT Services

Common

Maturity Immature Mature

Purpose(s) Exemplar Application Modelling Toolkit

XOR (exclusive “or”) Service Genres Service Expressions

Development Status Proposed Developmental Prototype Production

Deployment Scale Isolated Ubiquitous

State Behaviour Stateful Stateless

Batch Behaviour(s) Individual Batch

Time-Constraint

Behaviour

Hard Real Time Soft Real Time None

Service End Point Provider Requestor Transcoder (both requests and provides)

Authentication/

Authorization

Dependency

Auth-Dependent Auth-Independent

Protocol Binding(s)
(only applies to service
expression-based
SUMs)

Web Service
 SOAP
 REST
 HTTP
 Other

To be determined by the e-Framework:

Status Approved Placeholder
 Unapproved
 Superseded
 Withdrawn
Confidence Level High Medium Low

Notation

UML 2.1 and BPMN 1.2 are used to describe elements of this SUM, and familiarity with them is assumed.

IdP A Shibboleth identity provider. This authenticates a user, and asserts their identity to relying parties. It is typically the user's home institution.

SP A Shibboleth service provider. This provides a resource, and relies on the identity asserted by an IdP.

VO Virtual Organisation. VOs essentially provide group membership information which can be used to make authorisation decisions in a Role-Based Access Control system (RBAC).

AC Attribute Certificate (see RFC 3281)

PC Proxy Certificate (see RFC 3820)

Description

This SUM sets out a process to allow users with Shibboleth identity credentials to access resources which rely on X.509 identity certificates, through a portal. Other grid access mechanisms are in common use (primarily UNIX terminal sessions), but this SUM does not support those cases. Together, the functionalities of this SUM are provided as a Credential Translation Service (CTS), which orchestrates the services required to authenticate to grid resources.

This SUM generates X.509 Proxy Certificates (PC) for users, and permits users to assert their membership of, and roles within Virtual Organisations (VO). It mediates interactions between the user, the server or servers which are responsible for their VOs, and the Proxy CA and server which generate and host certificates for them to use on the Grid.

Business Process Modelling

Due to the nature of this SUM, many of the business processes are described within the *Virtual Organisations on the Grid*, *Proxy access to Grid resources* and *UK Federation: access resources* SUMs. This SUM orchestrates those nested SUMS.

Authenticate User: utilises a Shibboleth identity provider (IdP) to authenticate and identify a user, then generates an X.509 Proxy Certificate (PC) for that user.

Enrol for VO: allows a user to request membership of a role and group within a VO.

Create Credential: allows a user to assert membership of groups and roles across one or more VOs, confirms the validity of these assertions, and generates an AC for the user.

Store Credential: delegates the complete PC and AC to a proxy server from where the user can access it.

Access Resources: a user can access resources using their generated X.509 PC.

[activitySARoNGS Overall Flow {1/1}](#)

AuthenticateUser
CreateCredential
StoreCredential
AccessResources
EnrolForVO

SUM Diagram

This SUM uses no services directly (at least insofar as services are defined within the e-Framework), rather it is composed entirely of nested SUMs, which are depicted as a box with an internal border. The functions of the CTS do not stand alone from this SUM.

Usage Scenarios

The use case for this SUM is for users who have a Shibboleth identity credential to authenticate to a portal which relies on X.509 identity credentials. Through this portal, they can access resources on the Grid.

Applicability

This SUM is used to allow Shibboleth users to access Grid resources through a portal. Other grid access mechanisms are in common use (primarily UNIX terminal sessions), but this SUM does not support those cases.

Services which allow access to users authenticated using this mechanism **MUST** consider the different user enrolment requirements, and the consequent level of trust in the asserted identities.

Functionality

Authenticate User

The user is authenticated using Shibboleth (see *UK Federation: access resources* SUM), and this authentication information is passed to a Grid proxy server which decides whether to issue a grid proxy certificate (PC) for them. If so, the certificate is issued.

Enrol for VO

This allows an authenticated user to request access to a VO, and a role and group within that VO. As the user's PC only identifies them pseudonymously, they must provide a real-world identity at this point, for the VO manager to make access decisions. In this SUM, the user's email address is used, which is verified before the access request is made. The VO manager can then grant the user access (using functionality of the *Virtual Organisations on the Grid* SUM).

Create Credential

Users assert membership of none or more roles/groups within VOs (which they have previously enrolled), their authorisation is checked, and ACs are issued to them as appropriate, associated with their PC. The ACs are used to generate a VOMS-compatible PC.

Store Credential

The VOMS PC is stored on (*ie* delegated to) a proxy server. The user is then redirected to the portal they initially accessed, with details of how to obtain their PC.

Access Resources

Following the storage of their generated VOMS PC on the proxy server, users can access resources on the Grid via the portal. This uses functionality of the *Proxy Access to Grid Resources* SUM.

Structure & Arrangement

Authenticate user

This SUM depends on functionality described within the *UK Federation: access resources* SUM. The authentication and authorisation are undertaken by the proxy CA, using Shibboleth. Although the CTS is the SP from the perspective of the Federated Access Management system, the assertions from the user's IdP are passed to the proxy CA directly for interpretation and validation. The proxy CA undertakes this validation, and returns a PC for the user.

Following authentication (using the functionality of *UK Federation: access resources* SUM), the user is associated with a persistent identifier (eduPersonTargetedID). The signed (see *Crypto core* SUM) attribute set is used to generate a proxy certificate for the user (using functionality of the *Proxy Access to Grid Resources* SUM). In this proxy certificate, the Distinguished Name is an aggregate of an identifier for the service + the hashed value of eduPersonTargetedID. The hash value is essentially used as a pseudonym.¹ The entire Shibboleth assertion from the user's IdP is passed to the proxy CA, which undertakes the authentication and authorisation decisions.

¹ Although eduPersonTargetedID is itself opaque, the rules of membership of the federation where this SUM has been developed forbid sharing of this attribute beyond the IdP.

This flow is set out in the interaction diagram below (including Shibboleth elements). Note that the calls are "pseudo-code" – they do not map directly to any implemented methods.

```

Browser Portal CTS WAYF IdP MyProxyCA
alt ["Browser/POST"]
redirect()
InvokeBrowserPOST(SignedAssertion)
authenticate(Credentials)
redirect()
request()
redirect()
request()
request()
["Browser/Artifact"]
redirect()
redirect()
issue(SignedCert)
request(AssertionReference)
authenticate(Credentials)
request(AssertionReference)
request(SignedAssertion)
MyProxyCA::get_delegation(SignedAssertion,Key)
identify(SignedAssertion)

```

e-Framework Service Usage Model Description

SUM Template v7.2 20070725 © Copyright, e-Framework Partners, 2008
 SUM Content © Copyright Curtis+Cartwright Consulting Ltd, 2008 6

The data flow involved in issuing the PC is:

```

Assert
attributes
Request PC
Authorise Issue PC
Check SAML assertion
1) is signed by IdP
2) contains eduPersonTargetedID
3) is valid
4) is issued by Federation member
5) was passed by CTS authenticated through SSL
OK
SAML
Cert info Assertion
PC
SAML
Assertion
Federation
metadata
IdP cert User info

```

Enrol for VO

This business process is provided by the *Virtual Organisations on the Grid* SUM, but with the CTS mediating the flow between user agent and VOMS server.

Create Credentials

In this SUM, the user agent is a browser which can execute an AJAX programme. This programme interacts with the CTS, downloading data as required. The CTS contacts one or more VOMS servers to fulfil these requests. When the user has selected the set of ACs which they wish to assert, they submit the set to the CTS. The CTS then generates a VOMS-compatible proxy certificate.

The overall flow is described in the interaction diagram below (note that although one VOMS lifeline is drawn there could be multiple VOMS servers, each contacted for a specific AC):

```

interaction CreateCredentials CreateCredentials {1/1}
Browser CTS VOMS
loop (0,*)
opt["Other VO server"]
request(VO)
request(VO)
request()
respond(AJAXPage)
request(Role)
VOMS::issue_AC()
//
This provides an ability to
specify a VO which is not
known to the AJAX app
request()
request(VO)
respond(VOlist)
request(VO)
respond(VOConfig)
request(VO)
respond(GroupsList,RolesList)
respond(GroupsList,RolesList)
respond(AttrCert)
respond(AttrCert)
respond(VOlist)
respond(VOConfig)
submit(UserCert)

```

The optional element is invoked if the user selects a VO server which is unknown to the application.

Store Credentials

After the CTS has created a VOMS-compatible PC, it delegates this to the proxy server. The user agent is then redirected to the portal which they originally visited, with details of the proxy server which is holding their credential.

[interaction StoreCredentials StoreCredentials {1/1}](#)

```
Browser CTS MyProxy
redirect()
MyProxy::store(VomsProxyCert)
//
```

Completed VOMS proxy
certificate generated by CTS

Applicable Standards

X.509: <<http://www.ietf.org/html.charters/pkix-charter.html>>

RFC 3820 – “Internet X.509 Public Key Infrastructure Proxy Certificate Profile”

RFC 3281 – “An Internet Attribute Certificate profile for Authorization”

SAML: <<http://saml.xml.org/saml-specifications>>

Shibboleth: <<http://shibboleth.internet2.edu>>

In addition to the above standards, the published APIs for the following programmes are used:

MyProxy: <<http://grid.ncsa.uiuc.edu/myproxy>>

VOMS: see *Virtual Organisations on the Grid* SUM

Design Decisions & Tradeoffs

In this SUM, the CTS is “dumb” – it mediates between a user and a range of services which can provide the required functionality. It does not make authorisation decisions, nor have any state information regarding the credentials which have been issued through it.

Implementation Guidance & Dependencies

Systems which depend on X.509 certificates for identification of users frequently operate in medium or high-assurance environments. Typical Shibboleth authentication systems depend on only a username and password at the IdP for authentication. Additionally, the Shibboleth protocols and software have undergone no formal security audit. The assurance in a user’s identity, both in terms of linking their real-world identity to their online identity (*ie* enrolment), and ensuring that the user is the rightful holder of the electronic identity which they present must be assumed to be lower in a Shibboleth system than in an X.509 system with strong authentication at enrolment. Operators of services on the Grid must decide whether to accept the proxy certificates issued by a Shibboleth/Grid CTS.

Known Uses

This SUM was developed by the SARoNGS project, and at the time of writing is being taken to production within the UK National Grid Service.

Data Sources Used

Configuration information: information required to provide the service, including the endpoints of the

MyProxy server and CA, and the VOMS servers.

PKI information: the set of information required to operate within the public key infrastructure. This includes public and private keys for the CTS, public keys and certificate revocation information for appropriate certificate authorities.

Federation metadata: the set of information regarding members of the federation. See *UK Federation: top level view* SUM.

Related SUMs

Proxy access to grid resources

Virtual Organisations on the Grid

UK Federation: use services

CORE SUMs Used

Crypto

References

<<http://www.jisc.ac.uk/whatwedo/programmes/einfrastructure/sarongs>>
See applicable standards above.

Terms

The terms used within this SUM are explained within the nested SUMs from which it is constructed.

This SUM is licensed under:

Creative Commons Attribution-NonCommercial-ShareAlike 2.5 licence

<http://creativecommons.org/licenses/by-nc-sa/2.5/au/>