



Project Information			
Project Acronym	GFIVO		
Project Title	Grouper to support Federated Identity for Virtual Organisations		
Start Date	01/03/07	Project End Date:	01/03/09
Lead Institution	Newcastle University		
Project Director	Steve Williams (was Paul Hopkins) both director ISS		
Project Manager & contact details	Caleb Racey Caleb.racey@ncl.ac.uk		
Partner Institutions	-		
Project Web URL	http://gfivo.ncl.ac.uk		
Programme Name (and number)	E-infrastructure Call V – Virtual Organisation Management Tools and Services		
Programme Manager	Chris Brown was James Farnhill		

Document Name			
Document Title	Final report		
Reporting Period	-		
Author(s) & project role	Caleb Racey, Project manager		
Date	9 th Jun 2009	Filename	GFIVO_Final_Report_v2
URL			
Access	<input checked="" type="checkbox"/> Project and JISC internal		<input checked="" type="checkbox"/> General dissemination

Document History		
Version	Date	Comments

JISC Final Report –GFIVO project

Grouper to support **Federated Identity for Virtual Organisations** (G-FIV-O)

Final report

Author Caleb Racey

Contact Caleb.Racey@ncl.ac.uk

11/05/2009

Acknowledgements

The project acknowledges the generous funding of JISC e-Infrastructure **Call V – Virtual Organisation Management Tools and Services**.

Executive Summary

The project aimed to develop and deploy an access management infrastructure to support collaborative tools for virtual organisations. This access management infrastructure was based on Grouper and Shibboleth technologies developed by the internet2 (<http://www.internet2.edu>) to support access management. These free open source tools are readily deployable at no cost by other HE institutes; therefore the outputs of this project are of direct value to the HE community at large.

The approach taken was to deploy and populate an access management infrastructure and use it to support the deployment of collaborative tools such as blogs and wikis.

Background

Web application development for HE at present is overcomplicated. A combination of Shibboleth and Grouper with a Platform of web applications deployed according to Service Oriented Architecture SOA principles would go a long way to simplifying development allowing for a much richer set of tools to be deployed. The addition of a virtual identity home to support disenfranchised external collaborators would create a compelling set of solutions that could herald a step change in the value of support that can be deployed to virtual organisations.

Current application provision in many institutes is complex and resource intensive, at present most applications are only developed for large user populations with pressing needs. A large part of the reason for this is that most useful applications require authorisation and authentication. Integration with an institute's existing user name and password stores requires a great deal of technical knowledge of the details of the institute's infrastructure. Therefore only large projects like institutional Virtual Learning Environments (VLEs), Portals, and web based email gateways are developed by institutes. Small focused applications designed to provide a facility to relatively small user populations don't have enough institutional leverage to be granted access to the underlying institutional identity management infrastructures, and the technologists that understand them. These factors work against the development of tools to support virtual organisations, the user base is too small and fluid for it to be worth expending the institutional resource to push development through the heavy weight processes. It is therefore highly desirable that tools are

found that enable an easier and more flexible approach to the development of applications.

Shibboleth can provide a simple light weight authentication interface that enables developers to leverage and extend existing authentication infrastructures. Shibboleth provides the means to authenticate users and make relevant authorisation information available about that user to web applications. Grouper provides an easy means for developers to augment existing identity infrastructures so that they can be used to pass appropriate authorisation information to an application. Grouper is designed to be populated with existing group infrastructure information. It also allows for augmentation of that information, with the control of the augmentation process being delegated to appropriate users. By combining Shibboleth and Grouper existing heavy weight, expensive identity management infrastructure can be made easily available to developers and users, this approach has the benefit that it allows users to enhance existing identity infrastructure as needed for use by applications. In other words the approach that these tools allow democratises the control of identity information, promising a much greater engagement by the user population. The pairing of Shibboleth and Grouper has the additional benefit that developers no longer need to develop their own account management and group management systems for each application, meaning that they can focus their effort directly on creating solutions for users

Aims and Objectives

The project addressed group management issues for groups consisting solely of Newcastle users and also groups consisting of Newcastle users and external collaborators (including external researchers on collaborative research projects). Grouper were used to inform access control decisions for access to web based resources. The project covered all issues related to creating a physical deployment of Grouper as a group management infrastructure. The goal of the project was to develop and document a practical usable group management infrastructure for use by Newcastle University and collaborators. The scope of the project covered technical setup, group population, documentation, user education, usability issues, accessibility, and assessment of fitness for purpose.

Methodology

The project followed a phased approach to deployment. This enabled the project to produce a series of benefits to the university and wider community during its life cycle rather than producing outputs and benefits in one large lump at the end. This helped to maintain institutional buy in for the project by demonstrating the projects value throughout its 2 year lifecycle.

Methodologies, problems encountered and lessons learned were documented on the project web site (<http://gfivo.ncl.ac.uk>) as they happened.

Implementation

The project was implemented in a phased approach to enable benefits to be extracted as early in the process as possible and to allow feedback generated to be incorporated into the project design.

The phases split as follows

- Phase 1: Deploy a Grouper infrastructure
- Phase 2: Test the use of Grouper for lightweight small flexible internal groups.
- Phase 3: Develop Grouper connectors
- Phase 4: Test the use of Grouper with a complex use case
- Phase 5: Provide a “Virtual Home for Identity”
- Phase 6: Augment Virtual Home for Identity with group membership
- Phase 7: Test complex federated virtual organisations
- Phase 8: Exploit Shibboleth 2.0 functionality

Lesson learned from each phase where fed into the following phases in order to ensure that full benefits were reaped and lesson learned.

A mixed infrastructure of live production servers and development machines were deployed so that the service could be deployed in production while features were developed and trialed on development kit before being integrated into the live production service.

Outputs and Results

A fully supported sustainable Grouper group management service has been deployed and embedded in Newcastle University. This infrastructure consists of a central Grouper database and API which stores and controls group information. This is then interacted with using an infrastructure of Grouper shell scripts for bulk loading of data, Grouper UI user interface for administrators to use and Grouper web services for user interface interaction. This service is in production and is a fully supported official service deployed within the University. Additionally the project proved the feasibility of the use of Mysql (<http://www.mysql.com>) as a deployment database and contributed this information to the community. Data connectors tools developed early in the project for getting data into Grouper have been largely superseded by Grouper loader tool for loading groups from databases. This tool eases the deployment of Grouper significantly making it much easier to deploy and populate with groups source from institutes systems of record.

The project has enabled the deployment of a successful scalable infrastructure of wikis to support collaborative efforts hosted within the university. Prior to the project the institute had a couple of wikis deployed for internal project documentation use cases, however access control was achieved by hand editing of access control lists in files, editing was performed by systems administration staff and requested by logging a helpdesk call. After the infrastructure developed by Grouper was deployed access control could be controlled by the wiki owner resulting in much greater scalability of the service and a much improved user experience. To date this has resulted in the roll out of a wiki infrastructure of 72 wikis with a total of 1040 total users. Wiki roll out has largely focussed on staff use cases such as research group support however a successful pilot deployment in the dental school has proved the value of wikis to support pedagogical use cases. This trial involved dental students contributing to a wiki on course subject matter, student contributions were then assessed and graded.

The project has resulted in the deployment of a proof of concept Shibboleth 2.0 infrastructure. This infrastructure is in the process of being rolled out as the live service (due for go live July 2009). The major benefit of this development is the

integration of Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO) based true SSO capability. This enables Shibboleth to be deployed for internal use cases which benefit from login without user interaction. The most obvious example of this is deployment of a personalised student portal. True single sign on enables Shibboleth to be used to provide user descriptive information to the portal without inconvenience to the user. In this way we will be able to show users their exam paper listing, print credits information, reading lists etc while exploiting the class leading security offered by Shibboleth. This development shows great potential in providing a unified easy to use gateway interface to the disparate web applications that make up the student user experience.

Outcomes

The project proved the power of Grouper and Shibboleth combined as an access management toolset. Grouper provides a flexible role based access control system. The ability to load group information from systems of record and then blend and adapt these groups with hand edited groups is invaluable in a university context. Systems of record often reflect the organisation hierarchy of an institute as seen from a financial aspect and need some adapting to be useful in teaching or research contexts. Grouper fulfils this niche very well providing a variety of powerful and flexible methods for adapting group data.

The project also demonstrated that scalable access management is the key factor in achieving scalable sustainable roll out of collaborative web based tools. Scalable access management in particular greatly enhanced the value of the community of wikis, allowing a successful sustainable service to be deployed to support research and pedagogical use cases.

Further the project has demonstrated the value of being able to represent organisational hierarchies from multiple different sources in one system that enables them to be combined and used. This is particularly valuable in an HE context as universities have many different structures and hierarchies that represent the different ways members are structured (e.g. by location, by pay grade, by responsibility etc). An example of this is the use of Grouper to provide access control to a room booking system. By making available the staff organisational structure and the meeting room structure as groups we are now able to blend these two structures together to say which groups of people can book which rooms, previously room booking was open access to all with an administrator checking to see the right people are booking the right rooms.

Lessons learned in the project, particularly around the area of acquiring high quality institutional data, informed a successful further JISC funded project in the institutional exemplar programme called IDMAPS (<http://research.ncl.ac.uk/idmaps>). This project will look at the policy and technical requirements of systems and services required to transfer institutional data between disparate institutional systems, additionally it will look at how improved data flow can be leveraged to improve the cohesion and value of the student experience presented by institutional systems. As a part of this effort the project will look at improving data availability for use in access control systems like Shibboleth and Grouper.

Conclusions

The combination of Shibboleth and Grouper provide a powerful and flexible toolset to support access control for web applications in HE institutes. Shibboleth provides the framework for single sign on systems based on user information or attributes. Grouper provides a compelling tool for expressing controlling and augmenting those attributes. The combination of the pair provides an access management solution that is compelling and is architected with the complexity of HE identity stores in mind.

The deployment of simple collaborative tools such as wikis provides a simple and compelling communication platform that is highly regarded by a cross section of user groups. Simple scalable access control architectures greatly enhance the ability of institutes to deploy services like wikis to support their users.

Widespread use of Federated Identity to support Virtual organisation deployment that crosses institutional boundaries is currently hampered by the lack of maturity in the community. In order to achieve easy deployment of tools like wikis and VLEs to support federated virtual organisations each user needs to have an identity provider that is registered into the federation, they also need to know this and have been educated about how to use federated identity by their home institute. At the time of the project federated identity was still in its embryonic form resulting in patchy levels of user education, this hampered user acceptance. These issues will not be solved by the deployment of new technology such as OpenID (<http://openid.net/>) CardSpace, (<http://www.microsoft.com/net/windowcardspace.aspx>) etc, and are problems that would be encountered regardless of technology. The project expects these issues to diminish with time as the use of federated approaches grow and users become more familiar with its use and the value it offers. The need for federated identity continues to grow and all effort should be made to support its maturation.

Implications

True single sign on for web applications would greatly enhance Shibboleth deployment for internal use cases and would be relatively easy to develop. It is on the roadmap for Shibboleth 2.2 as “desirable” and the project recommends that JISC looks into funding development work. While SPNEGO based true single sign on is possible utilising a combination of Shibboleth and CAS much better functionality would be possible with an integrated Shibboleth login handler that implemented SPNEGO. Browsers different handling of SPNEGO requests is a major barrier to widespread deployment with Internet Explorer displaying damaging behaviour in off campus contexts. It would therefore be beneficial to be able to finely control the login behaviour allowing SPNEGO when appropriate and failing back to normal form based login when needed. The possibility of using SPNEGO based single sign on and Shibboleth’s “isPassive” login flow to achieve web application sign on without user interaction would greatly enhance the use of Shibboleth in Portal use cases. The project published an initial scoping study for integrating SPNEGO functionality directly into a Shibboleth login handler, we recommend that this work is properly scoped and funded. While the project team can contribute to the scoping exercise we do not possess the skill set to develop a SPNEGO based login handler, we therefore recommend that a call is put out for expressions of interest. Our initial assessment is that much of the work is already done in the CAS SPNEGO support and much of that could be transferred to a bespoke *Shibboleth* login handler.