

# Supporting Security-Oriented, Collaborative nanoCMOS Electronics Research

Richard O. Sinnott, Thomas Doherty, David Martin, Campbell Millar,  
Gordon Stewart, and John Watt

National e-Science Centre, University of Glasgow, Scotland  
{r.sinnott,t.doherty,d.martin,c.millar,g.stewart,  
j.watt}@nesc.gla.ac.uk

**Abstract.** Grid technologies support collaborative e-Research typified by multiple institutions and resources seamlessly shared to tackle common research problems. The rules for collaboration and resource sharing are commonly achieved through establishment and management of virtual organizations (VOs) where policies on access and usage of resources by collaborators are defined and enforced by sites involved in the collaboration. The expression and enforcement of these rules is made through access control systems where roles/privileges are defined and associated with individuals as digitally signed attribute certificates which collaborating sites then use to *authorize* access to resources. Key to this approach is that the roles are assigned to the right individuals in the VO; the attribute certificates are only presented to the appropriate resources in the VO; it is transparent to the end user researchers, and finally that it is manageable for resource providers and administrators in the collaboration. In this paper, we present a security model and implementation improving the overall usability and security of resources used in Grid-based e-Research collaborations through exploitation of the Internet2 Shibboleth technology. This is explored in the context of a major new security focused project at the National e-Science Centre (NeSC) at the University of Glasgow in the nanoCMOS electronics domain.

**Keywords:** Grid computing, e-Research, Security, Virtual Organizations, Shibboleth.

## 1 Introduction

Security and ease of use are critical factors to the success and uptake of Grid technologies in supporting collaborative e-Research. Current end user experience of interacting with large scale computational and data resources such as the National Grid Service (NGS) [1] in United Kingdom typically begins with obtaining an UK e-Science X.509 certificate issued by the trusted UK Certification Authority (CA) [2] at Rutherford Appleton Laboratories (RAL) [3]. This has numerous issues. Firstly, it is off-putting to many potential researchers since they need to deal with unfamiliar security concepts. Furthermore, this *authentication*-based model for Grid security whereby the binding of the user identity to the certificate through the CA is an

extremely limited model of security since it does not restrict what that user can access and use other than at the level of privileges associated with a local user account for example. Instead, to improve the usability and availability of Grid resources to particular individuals or to particular collaborations, finer grained security models are required to ensure that resources are only accessible to appropriate individuals/VOs at the discretion of local resource managers according to their own local policies. That is *authorization* infrastructures are required which allow to define policies on access and usage, which can subsequently be enforced by local resource providers to limit access to their own resources according to appropriate site specific policies.

Critical to the success of any authorization infrastructure are tools to support site administrators in the definition of security policies. End users themselves should also be, as far as possible, shielded from the underlying complexities of authorization policies and associated security attributes or indeed the Grid more generally. In an ideal world end users should be able to access Grid resources in much the same way as they access other Internet resources [4].

In this paper we describe novel solutions which allow system or site administrators to define their own local policies on acceptance of a variety of VO-specific security attributes from potentially remote collaborators which can subsequently be used to make local authorization policy decisions. Through exploitation of the Internet2 Shibboleth technologies, various sources of security attributes - so called attribute authorities (AA), and authorization infrastructures, we are able to provide seamless and transparent access to Grid resources from potentially remote, trusted collaborators. To demonstrate the validity of this approach we show how we have exploited these technologies in the major new security-oriented project: *Meeting the Design Challenges of NanoCMOS Electronics* [5] at the National e-Science Centre (NeSC) at the University of Glasgow. We note that this one example from many projects at the NeSC which have adopted this approach, hence the solutions are generic and widely applicable.

## 2 Collaborative Grid Security Models

Existing Grid security models as typified by X509 Public Key infrastructures [5] underpinning access to resources such as the NGS suffer from several key limitations. Some of these key limitations include the end-user experience; the associated granularity of the security model offered by authentication-only Grid security models, and the trust model underlying the PKI itself. These limitations are described in detail in [6], [7].

The vision of the Grid is to provide single sign-on access to distributed resources. Through recognizing and trusting a centralized CA in associating the identity of a researcher with a particular digital certificate, single sign-on authentication can be supported. Thus researchers use their X509 certificate (or more often a proxy credential created from that X509 certificate) with a common username given by the distinguished name (DN) associated with that credential and single (strong) password. Through trusting the CA that issued the certificate, the end user is able to access a wide range of resources that recognize that credential without the need for multiple usernames and passwords across those sites. In short, the approach is based upon a public key infrastructure (PKI) supporting user authentication [8].

Knowing the identity of the end user requesting access to a resource is important, but is only the starting point of security however. Finer grained models of security are needed which define precisely what end users are allowed to do on resources across a given inter-organizational collaboration.

Role based access control is one approach that has been advocated for some time to address this issue. In this approach roles are defined and associated with policies describing what a user with that role is allowed to do on a given resource. Attribute certificates capture this information and can be used by resources providers to check the validity of user requests, i.e. that they are in accordance with local authorization policies. Detailed definitions of RBAC based systems and their benefits are given in [9], [10]. RBAC systems are often limited in that they are often complex to administer and use. What are required are simple tools for VO administrators and local system administrators to define and enforce security policies across research collaborations, and user oriented approaches that utilize these information. Examples of some tools for RBAC systems include [11], [12] and experiences in their application are given in [13], [14]. One of the most immediately usable ways to utilize authorization infrastructures is through ensuring that only sites within the VO can access VO-resources. Another way of considering this is scoping of *trust*.

Any useable e-Research collaborative infrastructure has to be aligned with the way in which researchers wish to work. Keeping systems simple from the end user perspective is a key aspect of this, and ideally aligned with the way in which they access resources more generally. The UK academic community and many other countries are rolling out national level federated access control systems, where authentication is devolved to a user's home site utilizing the Internet2 Shibboleth technologies [15], [16]. The UK Access Management Federation [17] was established at the end of November 2006.

The core of Shibboleth is a basic trust relationship between institutions within a federation, where each institute in the federation is expected (*trusted*) to authenticate their users properly. The architecture of Shibboleth defines several entities which are necessary to achieve this seamless integration of separate collaborating institutional authentication systems. The main components of Shibboleth consist of Identity Providers (IdPs, also known as a Shibboleth 'Origin'); a Where-Are-You-From (WAYF) service, and one or more Service Providers (SP, also known as a Shibboleth 'Target'). The IdP is typically the users' home institution and is responsible for authenticating the end users at their institution. Each institution will have their own local systems for authenticating their users, e.g. LDAP or other mechanisms. The WAYF service is generally run by the federation that the institutions are subscribed to. It typically presents a dropdown list to the user that contains all the participating institutions (or projects) that are subscribed to within the federation. Users choose their home institution from this list and are then redirected to the home institution (IdP). The SP provides services or resources for the federation that the end user wishes to access.

A typical scenario of this process is where a user types in the URL of the service or portal (SP) they wish to access. If the SP is protected by Shibboleth, the user will be redirected to the WAYF service where they select their home institution. Once redirected to their IdP they will provide the username and password they would normally use for authentication at their home institution. Once successfully

authenticated, the user will be automatically redirected to the SP they are trying to access. At the same time, the security attributes (privileges) of this user will also be passed to the SP in a secure manner for further authorization from either the IdP or one or more known attribute authorities (AA). What attributes will be released by an institutional IdP or AA and what attributes will be accepted by a given SP needs to be configurable however and targeted towards the needs of particular VOs. It is important that all of this is transparent to the end users (who simply log-in to their home site).

The uptake and adoption of Shibboleth technologies within a Grid context is not without potential concerns however. Sites need to be sure that collaborating sites have adopted appropriate security policies for authentication. Strength of user passwords and unified institutional account management are needed. Shibboleth is, by its very nature much more static than the true vision of the Grid, where VOs can be dynamically established linking disparate computational and data resources at run time. Instead it is still largely the case that the attributes that are defined and subsequently released from an IdP and how they are used by an SP is an involved and difficult process requiring understanding and pre-agreement on the information exchange between sites. The UK Federation for example has agreed a small set of attributes based upon the *eduPerson* schema [18].

Whilst the combination of Shibboleth and Grid technologies offer numerous direct complementary synergies, few tools currently exist to help facilitate the process of integrating Grid and Shibboleth technologies. For example, on the IdP side, an Attribute Release Policy (ARP) defines which user attributes may be released to the federation for which individuals. Tools such as ShARPE (Shibboleth Attribute Release Policy Editor) [19] provide a user interface to the ARP allowing a user or administrator to interact with the IdP attribute release policy without having to manually edit a raw XML file. At the SP end, the Attribute Acceptance Policy (AAP) component of Shibboleth defines which IdPs will be recognized (the default in the UK federation is that all sites are trusted at the authentication level); which attributes from the set release by any IdP will be recognized to *potentially* gain access to local services; or further, which attributes for specific individuals will be recognized.

Tools are thus required to scope the accepted IdPs and associated attributes. This scoping is likely to be aligned with the particular requirements of different VOs. We note that currently site administrators are required to manually edit the AAP XML file to tighten up the attribute rules. As these rules may change quite frequently (especially in the Grid vision for truly dynamic VOs) it is desirable to provide capabilities similar to ShARPE to allow an administrator to instantly scope attributes for the SP, but also allow a delegated user to dictate the policy for their service through this application. Furthermore there is a risk in deploying a policy which has been edited by hand as any typographic mistake may compromise the whole SP. Services which allow only valid manipulation of the AAP would eliminate this risk.

To improving the usability and uptake of Shibboleth technology in the Grid environment, the SPAM-GP (Security Portlets simplifying Access and Management of Grid Portals) project [20] was proposed to provide tools to support the process of establishing and enforcing fine grained Grid security in a Shibboleth environment. Specifically the project is developing a family of JSR-168 compliant portlets which a Grid portal administrator can use for tailoring access to the resources available behind

the portal, i.e. the Grid services which themselves have authorization requirements that need to be met.

The first such portlet that has been developed is the SCAMP (Scoped Attribute Management Portlet). This portlet allows restricted and syntactically correct manipulation of the AAP of a Shibboleth SP to streamline the subset of IdPs from whom a portal will accept user attributes. The portlet parses the federation metadata for the list of all the IdPs within the federation, and stores the values of the ‘scope’ entry for each IdP.

When the SP is provided with a scoped attribute, the suffix will by definition be one of these scoped values. The list of IdP scopes in the federation is provided to the user/portal administrator in the form of a drop down list, one per user attribute, where the institutions from whom attributes are to be recognized/accepted from may be selected. The first time the portlet runs, the policy will set all attributes to ‘scoped’ but with no scope defined, so the default behavior is not to accept attributes from any institutions – a default common with most security infrastructures, i.e. deny all. Subsequently collaborating sites can be iteratively added to build a VO at the attribute level by the portal (VO) manager. Once defined, these changes can then be added to the AAP file. This policy information will then subsequently be available for the next browser session referencing that resource, i.e. only allowing access to the resources from known and trusted sites with expected attributes.

To understand the benefits of this scoping and how it is used in combination with Shibboleth to tailor access to Grid resources we outline how this has been applied in the nanoCMOS electronics domain.

### 3 NanoCMOS Electronics Case Study

The NeSC at the University of Glasgow have successfully demonstrated how single sign-on to a variety of portals across a variety of e-Research domains can be supported enabling inter-disciplinary e-Research combining Shibboleth and Grid technologies. The largest of these projects is in the nanoCMOS electronics domain specifically through the EPSRC-funded *Meeting the Design Challenges of nanoCMOS Electronics* project [5]. This domain is characterized by its heavy dependence and protection of intellectual property. This includes protection of designs, data, processes and the commercial, and often extremely expensive licensed design software that are used. This 4-year project itself began in October 2006 and involves collaboration between the universities of Glasgow, Edinburgh, Southampton, York and Manchester, with many leading industrial partners in the electronics domain including tools providers.

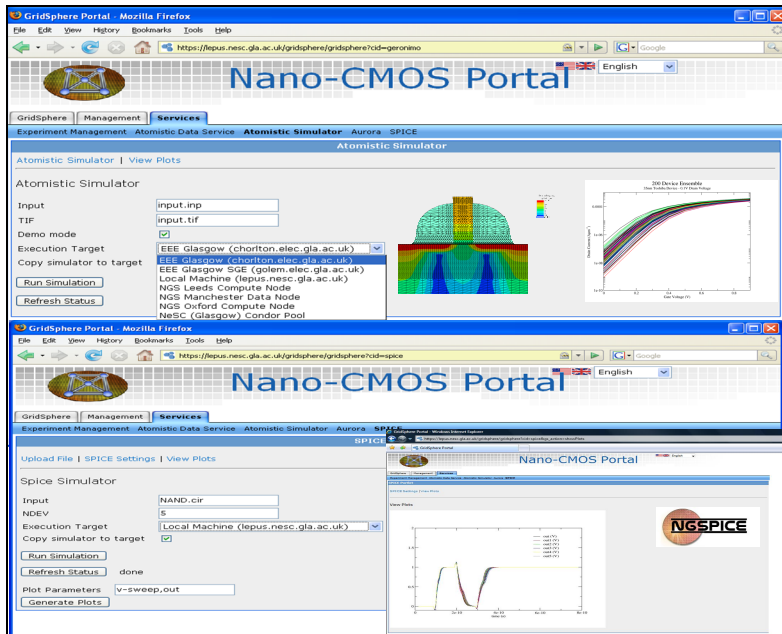
Collaboration in the nanoCMOS domain is essential to overcome the major concerns facing the development of next generation electronic devices. The building blocks of electronics design begins at the transistor level. These transistors are becoming ever smaller and have now reached the nano-scale with 40nm Silicon MOSFETs in mass production at the current 90 nm node and sub-10nm transistors expected at the 22nm technology node, scheduled for production in 2018. 4nm transistors have already been demonstrated experimentally highlighting silicon’s potential for decreasingly small dimensions. These decreasingly small devices are now influenced by atomistic effects resulting in device variability – hence

understanding the influences at the atomic scale and incorporating this into the design process is essential.

At the time of writing, numerous prototypes of the nanoCMOS services have been implemented and made available within a project portal protected by Shibboleth. These technologies have been based upon one of the leading Grid middleware today from the Open Middleware Infrastructure Institute UK ([www.omii.ac.uk](http://www.omii.ac.uk)). These include:

- atomistic device modeling services exploiting transistor designs from commercial device manufacturers and exploiting a range of statistical physics based approaches for atomistic characterization of devices (including modeling of electron mobility, dopant clustering, line edge roughness and exploiting a range of simulation approaches, e.g. Monte Carlo drift diffusion);
- services that supports the generation of compact models from device modeling simulations including exploiting license protected software;
- circuit simulation services incorporating device variability that allow to model the impact of device variability in the circuit/system design process.

The atomistic device modeling service and circuit simulation services are shown in Fig. 1 along with the outputs from the atomistic modeling, namely: a set of I/V curves reflecting the atomistic variability of the dopant concentrations and their distribution, and the output of the circuit simulation of a NAND gate showing the associated variation based on the atomistic device variability.



**Fig. 1.** Atomistic Device Modeling Service and I/V Outputs (above) and Circuit Simulation of NAND gate incorporating atomistic variability (below)

Access to these services and importantly to the associated data sets that they generate requires security authorization. This is important both for the commercial value of the licenses, for the intellectual property associated with the designs and data sets themselves.

We note that the atomistic device simulations themselves are especially computationally intensive and the complete characterization of a given device from a commercial supplier can require hundreds of thousands of CPU hours. An atomistic characterization of one such commercial device was undertaken and required >100,000 jobs on the ScotGrid computational resource ([www.scotgrid.ac.uk](http://www.scotgrid.ac.uk)) for its complete atomistic characterization [26].

The front end access to the portal is depicted at the bottom of Fig. 2 below. We note that this portal displays the various attributes that have been released by the identity provider and attribute authority at the University of Glasgow. We note that in this case, the only attributes that are recognised by the portal are those prefixed with *NanoCMOS* from the nanoCMOS partner sites. The top part of Fig. 2 shows another Shibboleth protected portal but this time without scoping of attributes.

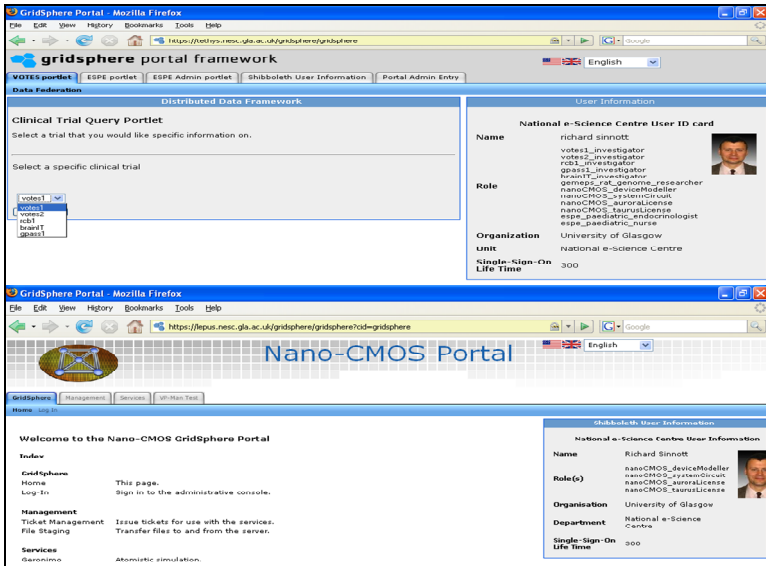


Fig. 2. NanoCMOS Portal with Attribute Scoping (below) and Other Clinical Portal without Attribute Scoping (above)

This scoping allows the portal to be restricted to only accept attributes from known and trusted sources, e.g. the nanoCMOS partner sites or more restrictively, only from specific individuals at those sites. The attributes themselves are then used to restrict access to the associated services that are available within the portal.

The services themselves have been developed to exploit a range of distributed HPC resources such as the National Grid Service, ScotGrid, and Sun Grid Engine-based clusters and Condor pools at Glasgow University. One key use of these attribute

certificates are both to restrict access the specific services but also where appropriate to the back-end computational resources themselves. Thus privileged end users are able to submit jobs, themselves described in Job Submission and Description Language (JSDL) [23] generated through the portlets, via OMII-UK GridSAM instances. This is achieved through providing authorisation capabilities to GridSAM itself, specifically through authorization decisions based on access to the back-end Distributed Resource Management (DRM) connectors of GridSAM. We note that a variety of resource specific DRM connectors are available within GridSAM including connectors for Condor, Sun Grid Engine and Globus.

The focus of the authorization decisions currently supported are through restricting access to the Globus DRM connector for the GRAMSubmissionStage part of the DRM connector sequence. In this model, the authorisation decision is decided before the JSDL document is submitted to the GridSAM instance and converted to a Globus specific Resource Specification Language (RSL) document and ultimately submitted to a GRAM manager. The authorization decisions themselves are made by using policies defined and enforced within the PERMIS RBAC system. The details of how PERMIS can be linked and used to restrict access to Grid services are described in detail in [24], [25].

We note that since major HPC resources such as the NGS require that X509 certificates are used for job submission, the back end of the portal supports a MyProxy service for creation and management of proxy credentials needed for job submission to major clusters.

## 4 Conclusion

Inter-organizational collaborative e-Research requires tools that simplify access to and usage of distributed resources yet support finer-grained access control. Shibboleth combined with tools that allow management of security attributes offer a suitable model for such collaboration.

Crucial to the success of Shibboleth and the uptake of Grid based e-Infrastructures are tools that support fine grained access to services and data sets. Proof of concept prototypes for definition of attribute acceptance policies have been demonstrated and applied in various e-Research projects. We note that the SCAMP portlet is just one of the several portlets we will produce during the course of this project. Other portlets that will be produced include an Attribute Certificate Portlet (ACP) which will allow users to issue X.509 ACs to other users for use with applications requiring fine-grained highly secure authorization, exploiting results from the recently completed Dynamic Virtual Organizations in e-Science Education (DyVOSE) project, specifically through a portlet enabled version of the Delegating Issuing Service (DIS) [21]; a Content Configuration Portlet (CCP) supporting dynamic configurability of portal content based upon Shibboleth attributes and knowledge of existing available Grid services; and an Attribute Release Policy (ARP) portlet allowing configuration of the attributes released from an IdP.

All of these portlets will be JSR-168 compliant and developed with the intention that a portal based VO administrator can define their own local policies on attribute acceptance, attribute release and how these attributes can configure access to local



Grid resources based upon security authorization policies. We recognize that portlets for administrators are a highly beneficial approach since they overcome the potential syntactic and semantic errors that might be introduced through manual editing of security acceptance policies. Furthermore, through JSR-168 compliance we expect these portlets to be widely applicable and easy to establish and use in other projects (both at NeSC and beyond).

We note that many Grid-based VOs are based upon the Virtual Organisation Management System (VOMS [12]) for definition of the VO-specific attributes. Through the recently funded VPMAN project [22] we are exploring how VOMS attributes can be incorporated into authorization infrastructures such as PERMIS. Thus rather than expecting to aggregate security attributes from one or more IdPs or associated attribute authorities, it might well be the case that we exploit IdPs for authentication and a VOMS server for the attributes that have been agreed upon for that particular VO. These attributes are then used by PERMIS to make an authorization decision. We have demonstrated already how this is supported with a variety of leading Grid middleware including Globus and OMII-UK [24].

One final challenge that remains to be addressed is how to exploit these kinds of tools when defining and enacting workflows comprised of several services where each service in the workflow requires security attributes to be presented to make an authorization decision. To address such kinds of scenarios we are working with OMII-UK to feed them requirements for future security-oriented workflow languages and enactment engines.

## References

1. UK National Grid Service (NGS), <http://www.grid-support.ac.uk/>
2. Jensen, J., The, U.K.: e-Science Certification Authority, Proceedings of the UK e-Science All-Hands Meeting, Nottingham, UK (September 2003)
3. UK Rutherford Appleton Laboratories (RAL), <http://www.grid-support.ac.uk/content/view/23/55/>
4. Sinnott, R.O., Jiang, J., Dr Watt, J., Ajayi, O.: Shibboleth-based Access to and Usage of Grid Resources. In: Proceedings of IEEE International Conference on Grid Computing, Barcelona, Spain (September 2006)
5. Meetings the Design Challenges of nanoCMOS Electronics, <http://www.nanocmos.ac.uk>
6. Sinnott, R.O., Watt, J., Jiang, J., Stell, A.J., Ajayi, O.: Single Sign-on and Authorization for Dynamic Virtual Organizations. In: 7th IFIP Conference on Virtual Enterprises, PRO-VE 2006, Helsinki, Finland (September 2006)
7. Watt, J., Sinnott, R.O., Jiang, J., Ajayi, O., Koetsier, J.: A Shibboleth-Protected Privilege Management Infrastructure for e-Science Education. In: 6th International Symposium on Cluster Computing and the Grid CCGrid 2006, Singapore (May 2006)
8. Housley, R., Polk, T.: Planning for PKI: Best Practices Guide for Deploying Public Key Infrastructures. Wiley Computer Publishing, Chichester (2001)
9. Sandhu, R.S., Coyne, E.J., Feinstein, H.L., Youman, C.E.: Role-Based Access Control Models. IEEE Computer 29, 38–47 (1996)

10. Ninghui, L., Mitchell, J.C., Winsborough, W.H.: Design of a Role-based Trust-management Framework. In: Proceedings of the 2002 IEEE Symposium on Security and Privacy (2002)
11. Chadwick, D.W., Otenko, A.: The PERMIS X.509 Role Based Privilege Management Infrastructure, Future Generation Computer Systems, vol. 936, pp. 1–13. Elsevier Science BV, Amsterdam (2002)
12. Virtual Organization Membership Service (VOMS), <http://hep-project-grid-scg.web.cern.ch/hep-project-grid-scg/voms.html>
13. Sinnott, R.O., Stell, A.J., Chadwick, D.W., Otenko, O.: Experiences of Applying Advanced Grid Authorisation Infrastructures. In: Sloot, P.M.A., Hoekstra, A.G., Priol, T., Reinefeld, A., Bubak, M. (eds.) EGC 2005. LNCS, vol. 3470, pp. 265–275. Springer, Heidelberg (2005)
14. Sinnott, R.O., Stell, A.J., Watt, J.: Comparison of Advanced Authorisation Infrastructures for Grid Computing. In: Proceedings of International Conference on High Performance Computing Systems and Applications, Guelph, Canada (May 2005)
15. Shibboleth, <http://shibboleth.internet2.edu/>
16. Shibboleth Architecture Technical Overview, <http://shibboleth.internet2.edu/docs/draft-maceshibboleth-tech-oberview-latest.pdf>
17. UK Access Management Federation, <http://www.ukfederation.org.uk/>
18. eduPerson Specification, <http://www.educause.edu/eduperson/>
19. Shibboleth Attribute Release Policy Editor, <http://federation.org.au/twiki/bin/view/Federation/ShARPE>
20. OMII SPAM-GP project, <http://www.nesc.ac.uk/hub/projects/omii-sp>
21. Delegation Issuing Service (DIS), <http://sec.cs.kent.ac.uk/permis/downloads/Level3/DIS.shtml>
22. Integrating VOMS and PERMIS for Superior Secure Grid Management (VPMan), <http://sec.cs.kent.ac.uk/vpman/>
23. JSDL, <http://www.gridforum.org/documents/GFD.56.pdf>
24. Sinnott, R.O., Chadwick, D.W., Doherty, T., Martin, D., Stell, A., Stewart, G., Su, L., Watt, J.: Advanced Security for Virtual Organizations: Exploring the Pros and Cons of Centralized vs Decentralized Security Models. In: 8th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2008), Lyon, France (May 2008)
25. Sinnott, R.O., Watt, J., Chadwick, D.W., Koetsier, J., Otenko, O., Nguyen, T.A.: Supporting Decentralized, Security focused Dynamic Virtual Organizations across the Grid. In: 2nd IEEE International Conference on e-Science and Grid Computing, Amsterdam (December 2006)
26. Reid, D., Millar, C., Roy, G., Roy, S., Sinnott, R.O., Stewart, G., Asenov, A.: Supporting Statistical Semiconductor Device Analysis using EGEE and OMII-UK Middleware. In: 3rd EGEE User Conference, Clermont-Ferrand, France (February 2008)