

# **Security Research Challenges for e-Science**

Prepared by the UK e-Science Security Task Force

**May, 2005**

## Introduction

An important part of the emerging e-Infrastructure is the security framework and mechanisms. Two aspects of the e-Science programme drive the need for security: the degree and sophistication of connectivity between systems (with the opportunities this creates for abuse), and projects such as those involving health or commercial data where security is a customer concern, or is imposed by ethical or statutory frameworks. Unless security measures are integrated into the programme, it risks failure because the results are not exploitable by their target community. Moreover, future growth of the application domains will be limited if new communities are not convinced that the technology is capable of meeting their security needs.

This report provides a research agenda for Security for e-Science identifying priority needs. These are structured under four headings:

1. Authentication, Authorisation, Accounting (AAA)  
This area has received much attention already, but much remains to be done.
2. Protection of Data  
The goals of privacy, confidentiality, integrity, and so on, form crosscutting concerns. Achievement of these goals is dependent upon AAA, but goes beyond it in many ways.
3. Architectural Components  
Fabric and networking components, platform security, web services issues are all of relevance in designing e-Infrastructure, and deserve attention in their own right, as does the question of how to undertake such design.
4. Operational Characteristics  
Operational characteristics (usability, performance and scalability, manageability, interoperability, assurance) are also given due weight; key operational issues are summarised and their relationship to the priority needs are described.

This document was informed by a data collection exercise that also highlighted a range of important non-technical issues, including the integration and management of security within e-Science projects.

New security technology or understanding will be needed to make project results widely available. Existing technology is sufficient to develop the project results and to deploy them for ongoing operational use, but only within a restricted community. Lack of security may adversely impact future investment in e-Science capabilities.

Finally, this intention of this research agenda is to look just beyond the pressing needs of existing projects. It is by no means intended to circumscribe the range of security technologies, techniques, and insights that may be brought to bear upon e-Science. Rather, its intent is to identify a number of areas that appear promising but so far are unexplored. The focus here is on the research required to solve the stated issues, and whilst there may be a need to develop new technologies and software systems, the focus is not on deployment or production of those systems.

## Priority Technical Needs

This section outlines the priority needs for each security area of concern in turn. The order in which the topics are introduced is of no relevance.

### 1. Authentication, Authorisation, and Auditing (AAA)

There has been a lot of attention to issues in AAA, which have served to both provide mechanisms and solutions to permit prototype e-Infrastructures and to highlight the need of those issues in this area that require further attention.

#### *Authentication*

Authentication is the establishment and propagation of a user's identity in the system.

Technologies are available, and state of the art commercial packages offer many of the required capabilities, although interoperability between authorities and user mobility are still problems. Ongoing JISC projects[5] are investigating large-scale deployment in the UK academic community. Issues concerning authentication include interoperability of authentication systems across different domains, scalability to large communities of users, measures and models of multi-level authentication and support for secure roaming.

#### Priority needs:

- **Levels of trust and responsibility.**

Expensive facilities, secure or sensitive data resources and other something may require strong authentication however other less "valuable" systems may need lesser authentication control. Research is required to understand and provide models and mechanism of multi-level authentication and light-weight systems.

- **Support for User Credential Management.**

Easy to use user credential management is required to make it practical for users to transport their credentials between applications and end-systems. Mechanisms for transporting or managing a user's public key certificate and associated private keys, or a system service to links user-friendly credentials, such as pass phrases, to the authentication infrastructure offer possible solution directions, but more needs to be understood across the requirements of e-Science applications.

- **Support for Secure Roaming.**

Secure roaming will build on the above services to allow users greater freedom of location, and to simplify providing facilities for 'visiting' users. In particular, support for roaming between wireless networks, and stronger logon at remote wireless locations are needed.

#### *Authorisation and Delegation*

Authorisation is concerned with controlling access to services based on policy. A complete authorisation system includes tools to specify and manage policy, mechanisms to distribute or obtain policies, tools to create and manage authorisation tokens, services that use policies to make an access decisions, and mechanisms that request and enforce access decisions.

The management of policies and their distribution, particularly in terms of the flexibility to set up dynamic or short-term groups of users, is a major weakness. The inability to support flexible dynamic delegation policies is a serious impediment to some applications. Other applications have strict privacy concerns about the contents of their authorisation tokens.

#### Priority needs:

- **Access Control**

Many e-Science and e-Health applications require fine-grained access controls, perhaps administered through a number of different policy authorities and distributed decision points. Languages exist for the expression of those policies, but their practical application -- usable policy design, communication in an accessible form, and rigorous validation remain to be thoroughly explored. There is already considerable interest in combination of these languages and their appropriate use in Grid contexts.

- **A policy reference model and supporting protocols.**

The authorisation and policy management system is complex enough to need a reference model that identifies the main components and the protocols by which they communicate. Delivery of such a model would facilitate the development of interoperable components, as opposed to the present situation where each project tends to adopt its own policy language. Emerging standards (SAML[11], XACML[12]) provide a starting point, but more development and consolidation is needed to create a comprehensive model.

- **Policy creation and management**

The manageability of authorisation policies for both users and systems must be addressed if more complex grids are to be established, providing functionality, scale and the capability to deal with the creation and destruction of short-lived virtual organisations. Interoperability and the capability to provide authorisation and delegation across domains as disparate parties come together in virtual organisations are critical.

- **An improved general model for Delegation and Privacy.**

In the long-term, scalability, functional requirements and privacy considerations all require a new, detailed, model of delegation and token privacy. Fine grain delegation and pull as well as push mechanisms for delegation need to be considered.

### *Auditing*

Auditing is the analysis of records of account (e.g. security event logs) to investigate security events, procedures or the records themselves. Logging, intrusion detection and auditing of security in managed computer facilities is well established in theory and practice, although there may still be issues regarding “missing” audit records. Grid computing adds the complication that some of the information required by a local audit system may be distributed elsewhere, or may be obscured by layers of indirection.

#### Priority needs:

- **Models and Tools for the generation of complete diagnostic trails.**

Diagnostic chains depend on the forms of authentication, authorisation, and delegation in use, and there are likely to be several. To avoid the need for each authority to understand the whole infrastructure in the distributed system, it is necessary to create tools that allow some types of record (e.g. user accountability information) to be obtained securely from other parts of the system and interpreted in a common framework.

## **2. Protection of Data**

In many applications the main security requirement relates to the protection of data in various ways: privacy, confidentiality, or integrity. Advanced technologies may also be able to support more elaborate digital rights management, with a wide range of application areas.

### *Privacy*

Privacy requirements relate to the use of data, in the context of consent established by the data owner, or subject. The privacy of the data subject poses special problems when data are amalgamated or copied to third parties. Privacy is therefore distinct from confidentiality, although it may be supported by confidentiality mechanisms including authorisation. Some

users will have privacy concerns about their personal data and their authorisation credentials (attributes/roles) as well as their experimental data.

Privacy is particularly significant for projects processing personal information, or subject to ethical restrictions: projects utilizing health data are particular examples. There is little prior art in privacy grid science, although there is useful UK background in privacy including hospital systems[13]. Web based standards such as P3P[14] may contribute to only a small fraction of the necessary security mechanisms.

#### Priority needs:

- **The generation and promulgation of examples of good practice in both policy and implementation for health systems.**

The exploitation of grid technology in health systems needs a transferable understanding of suitable privacy policies, how they can be applied, and what mechanisms can be used to implement them.

- **A fully worked out model and reference implementation for privacy protection,** including negotiation of attribute and data release by all parties involved in grid transactions and virtual organisations. Additional issues include: how to strongly and inseparably bind privacy policy to the data it is protecting, how to ensure policy is enforced throughout the grid by all systems for the duration of the data's lifetime, and how to ensure that data is properly destroyed when it has reached the end of its life.

- **Prevention of privacy leakage through data mining and merging**

Linking databases together on the grid produces new risks to data privacy. Whereas the data obtained from one source can be anonymised, when data from multiple sources, each of which might be anonymous in itself, is merged together, new inferences can be made which allows the data owners (or sources) to be identified, with a consequent loss of privacy. Furthermore, whereas the retrieval of one data element from a database may in itself be innocuous, the repeated retrieval of millions of data elements from the same source may not be. Ways of providing pan-database search coordination and temporal constraints on data mining need to be explored.

### *Confidentiality*

Confidentiality is concerned with ensuring that information is not made available to unauthorised individuals, services or processes. It is usually supported by access control within systems, and encryption between and within systems.

Confidentiality is generally well understood, but the grid introduces the new problem of transferring or signalling the intended protection policy when data are staged between systems. This is required in support of privacy, and also more generally for sensitive data. Some applications already have the requirement to store encrypted information in their databases, and this brings with it the associated problems of key management and the encryption of query messages.

#### Priority needs:

- **Mechanisms to transport confidentiality constraints with data**

Data owners need to be able to set restrictions on the use of their data, and these policies must be honoured and enforced by the protection infrastructure. Some existing methods (e.g. MAC policies) utilize universal data labelling schemes; others use encryption methods that only expose data to suitable environments. Both of these approaches have limitations and there is scope for long-term research on improved approaches to this problem.

## *Integrity*

Technical solutions exist to maintain the integrity of data in transit and in storage. The more general question of provenance (maintaining the integrity of chains or groups of related data) is a requirement that is being researched by a number of grid projects.

### Priority needs:

- **Approaches to provenance management.**

When data, particularly medical data, is collected and anonymised for research purposes, the owners of the data would like to be informed when new facts are discovered by the researchers that concern the data owners. There is thus a requirement to be able to rediscover the owners of anonymised data at a later point in time.

## *Digital Rights Management*

DRM has become a popular idea in consumer entertainment products, but has wider applicability to medical record data, mobile code in a grid context, and the distribution of commercially sensitive documents.

### Priority needs:

- **Understanding DRM in e-Science contexts**

The notion of protecting, say, medical data against onward transmission or unwarranted retention is clear enough, but the day-to-day use of such capabilities is not. Nor is it immediately clear how policies covering such applications should be constructed and documented, and integrated into existing authorisation regimes.

- **Technology solutions**

A number of cryptographically-based technologies are on offer which may provide part of the solution but their practical application and integration into real work-flows and application designs remains to be explored.

## **3. Architectural Components**

The large-scale elements of the e-Infrastructure contribute to security in their own right. The networking technologies themselves, and the platforms on which jobs are run, and data stored, have the potential to aid or hinder security design to quite a significant extent.

The trend toward service-oriented architectures, particularly Web Services, is introducing new architectural primitives. These include message-level security, user federation, distributed context management, and the flexible use of claims and assertions in the service authorisation chain. The emerging protocols in this area are more a question of commercial competition than research, but the practical consequences of applying these architectures have considerable scope for investigation and development.

### Priority needs:

- **Security in Service-Oriented Architectures**

Understanding the broad implications of service-oriented architectures and the developing standards in web services within the context of e-Science applications is key to success. Important topics include security design, the balance between complexity and risk in the security infrastructure, and security/performance tradeoffs.

## *Intrusion Detection*

The problem of intrusion detection occurs at all levels in the system, from network activity and low-level resource usage to the interaction between applications. The service-oriented architecture of many grid applications offers new opportunities for intrusion.

### Priority needs:

- **Detection and Classification of Grid Intrusion**

The view that any one administrator has of a distributed system is necessarily constrained, and this may make it harder to distinguish abnormal behaviour. Attacks with a high statistical profile (e.g. denial of service) will not be hard to identify, but those that actively disguise their intent (e.g. phishing) are likely to be harder to identify in highly distributed systems. Research is needed in characterizing and classifying covert attacks, as well as methods of detection.

## *Platform Security*

A significant challenge for the conduct of e-Science is the construction of trustworthy and reliable systems from components with varying degrees of trust (and reliability). From the service provider's perspective, there is the problem of being asked to execute untrusted code on their trusted system. From the client's perspective, the dual problem arises of needing to run trusted code upon an untrusted system.

The first problem is potentially solved by techniques of virtualisation and sandboxing. The Trusted Computing Group is proposing technologies for the second; their adaptation to Grid technologies is under discussion in the Global Grid Forum. Encrypted databases and other advanced encryption techniques may also be of relevance.

### Priority needs:

- **Virtual machine technologies and Trusted Platforms**

Appropriate integration of Virtual machine technologies into grid environments such that they -- and delegated authorities that use them -- are able to undertake just the tasks intended, and no more. Likewise, 'Trusted Platforms' allow secured storage and the measurement of software integrity: open research problems include how to extend those capabilities to distributed collections of heterogeneous systems.

## *System-Level Security Design*

Security cannot be considered at the component level alone, but approaches to design in order to permit the assembly of components - and those for authentication, authorisation, accounting, and so on - into coherent systems is an topic in security requiring further research.

### Priority needs:

- **Policy Development and Management**

System security policies are wider than access control; for example, they may include requirements that limit the behaviour of services, or require additional system functions (e.g. audit) in support of security. It is generally accepted that the only effective decision criterion for security is risk, but risk-based security reasoning has yet to be fully integrated into the system design and engineering process. Research is needed into the process of developing security objectives, determining how system security policies

are designed to meet those objectives and how this process can be integrated into the system design lifecycle.

## 1. Operational Characteristics

Operational characteristics (usability, performance and scalability, manageability, interoperability, assurance) are a significant consideration in the adequacy of existing technical solutions, and are therefore important in the implementation of new technical capabilities. These concerns cut across the functionality requirements, so each concern is cross-referenced to the most relevant functional recommendations.

### *Trust*

Trust is that characteristic allowing one entity to assume that a second entity will behave exactly as the first entity expects [15]. Trusted entities are those for which this expectation is assumed, with the consequence that data they originate are assumed to be correct, and obligations that they promise to undertake will be fulfilled. Contractual or other agreements about what entities are to be trusted to do, and to what extent, are therefore of fundamental importance to virtual organisations.

There is an important distinction between ‘trust management’ systems according to RFC 2704, which implement authorisation, and the wider requirements of trust management. For example, both industrial and health applications require the agreement between users and resources providers of restrictions that cannot be implemented by access control (e.g. restrictions on the export of software, or a guarantee that personal data is deleted after use). There is therefore a need to understand and represent policy and contractual agreements between groups of users and resource providers; such agreements may exist inside or outside the system, and are typically not supported by technical mechanisms today.

### Priority needs:

#### ▪ **Trust Relationships between Collaborating Organisations**

In order to share resources and allow mutual access, independent organisations need to establish a framework of trust, outside the system, that establishes what they each expect of the other. For example, a trust framework may include standards for identifying and managing users, agreements about payments or other boundary settlements, limitations of use, agreed procedures for incident investigation, and escalation procedures in the event of unresolved problems. Although some aspects of this subject have been documented, the trust basis on which large distributed collaborations will operate is still somewhat unexplored.

#### ▪ **A policy framework to allow the establishment of ‘virtual grids’.**

The framework should consider how virtual grids are agreed, implemented and managed. It needs to specify the types of policy that can be supported and the extent to which these policies can be supported by technical security measures. It must also address the question of scalability to ensure that administration of complex grids is feasible.

#### ▪ **A comprehensive trust and contract management infrastructure**

that allows the degree to trust between parties in a virtual organisation to be determined and managed. European and US research projects are starting to address this issue and the results of their research should be available in the medium to long term. Their results should provide a framework and reference implementation, but no doubt will require further development before they are suitable for large scale rollout.

### *Usability*

Usability is concerned with the ease and accuracy with which a system can be used, particularly by end-users who do not have specific security skills or knowledge. In the context of security, ease of use implies that users are able to focus on their main goals, rather than supporting security functions; accuracy implies that users are able to use security functionality correctly, and do not inadvertently introduce security vulnerabilities through ignorance or avoidance of security related actions. Simplicity of use is a critical factor to success.

The critical usability problem in today's infrastructure is the management of private keys by a user. Users are expected to be able to carry out a difficult, technical, process to move private keys between machines; as a result many users will fail to restrict or protect their own private keys, or will avoid the process and be limited to a single terminal.

Other recommendations with a particular usability element include roaming, and health and privacy policies. In the latter, the UK data protection framework [16] requires user interaction and understanding, so the design and implementation of security features in these systems must allow users to understand and configure security policy, as far as necessary to meet these obligations.

Usability can also be extended to the project world, in the sense that security methods and practices must be useable by project practitioners.

### ***Performance and Scalability***

Performance and scalability are concerned with the extent that security services and technology will support large numbers (usually of users and services), or sizes of data, without significantly impacting the performance of the application.

Most grid projects are still in development, so scalability and performance issues have yet to be met in practice. The number of users will rise significantly when current systems move into their exploitation phase. Two areas give immediate cause for concern: the Authentication infrastructure, including the present e-science CA, and the mechanisms for mapping grid users' identities (distinguished names) to end-system usernames.

Other recommendations with particular scalability concerns are the delegation model, which must be careful to retain scalability while introducing more flexible policies, and trust frameworks for virtual grids.

### ***Managability***

Managability relates to setting and changing the policies and parameters that control security, both inside and outside the system. It is essential that they can be flexibly and easily changed, and that they result in intuitive and appropriate controls and agreements. Furthermore, management controls must support individual responsibilities and processes in the context of organisations' overall information security arrangements.

The present authentication and authorisation infrastructure is cause for concern: the mapping of global to local identifiers is difficult to manage and this, together with a lack of management tools, inhibits the creation of short lived or dynamic user groups. Policy management is specifically addressed at, and the mapping issue is described above (scalability).

Policy relationships outside the system are equally important, and the recommendation that a policy and contractual framework for grids is established specifically addresses this issue.

## ***Interoperability***

Interoperability is concerned with the ability to traverse the boundary between two grid environments, in such a way that agreements between the parties, or constraints expressed in the associated protocol, are upheld. Interoperability is dependent on both external policy agreements and technical standards. A different, but practical and desirable, feature of interoperability is to facilitate an open framework for security components and services.

## ***Assurance***

Assurance is concerned with the ability to quantify the reliability with which a particular service upholds a given security policy or function. It contributes to the degree of trust that a user might place in a service, or that might persuade two parties to agree on a particular policy framework. Although it is just one component of trust, it is included here because there are established international frameworks [17] for setting and measuring assurance. Metrics for security both historical and predictive are increasingly important.

There are few assurance issues in the current grid, although the mechanism that maps global to local identities has been subject to an assurance evaluation. In general, any of the recommendations that result in security mechanisms should be careful to use good design practice (such as least privilege, and minimising security critical software) even if a formal assurance evaluation is not required.

## **Summary**

Several themes are evident in this research agenda. The first is the extent to which existing tools for authentication, authorisation, and delegation, and trust can be combined. Although there are several systems that would claim to solve the authorisation problem, there is a need to take this work further to build a framework that is more flexible, and where the elements are better decoupled and interoperable than is currently the case.

A second theme is the consideration of a range of technologies so far only peripherally considered by the e-Science community, but which have potential significantly to extend the kinds of assurance and capability possible. Architectures developed in this way have the potential to broaden considerably the scope of problems that can be successfully addressed by e-Science into the domain of relatively high-assurance requirements.

A third theme is the extent that operational characteristics are a significant issue: many of the recommendations are too often regarded as secondary issues when security infrastructure is built, but it is evident that shortfalls in this area are a significant driver for those developing or deploying security mechanisms. Understanding the impact of or these characteristics within models of the complex security requirements of e-Science applications is a key.

## References

1. *The Globus Project*, <http://www.globus.org/>
2. H Kreger, *Web Services Conceptual Architecture*, IBM, Technical Report WCSA 1. May 2001. <http://www-3.ibm.com/software/solutions/webservices/pdf/WSCA.pdf>
3. I Foster, C Kesselman, J Nick, and S Tuecke, *The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Implementation*, Global Grid Forum, <http://www.gridforum.org/ogsi-wg/>
4. Nataraj Nagaratnam, et al., *The Security Architecture for Open Grid Services*, The Globus Project, July 17, 2002,. <http://www.globus.com/>
5. *Authentication, Authorisation and Accounting (AAA) Programme*, The Joint Information Systems Committee (JISC), 1 October 2002. [http://www.jisc.ac.uk/index.cfm?name=programme\\_aaa](http://www.jisc.ac.uk/index.cfm?name=programme_aaa)
6. L Pearlman, et al., *A Community Authorisation Service for Group Collaboration*, in *Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks*. 2002, IEEE.
7. M Thompson, et al., *Certificate-Based Access Control for Widely Distributed Resources*, in *Proc 8th Usenix Security Symposium*. 1999: Washington, D.C.
8. David W. Chadwick and Alexander Otenko. *The PERMIS X.509 role based privilege management infrastructure*, in *Proceedings of the Seventh ACM Symposium on Access Control Models and Technologies*, Monterey, California, USA. 2002
9. *VOMS Architecture*, European Datagrid Authorization Working group, 5 September 2002.
10. P Hallem-Baker and E Maler, *Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML)*, OASIS, SAML 1.0 Specification. 31 May 2002. <http://www.oasis-open.org/committees/security/#documents>
12. *eXtensible Access Control Markup Language*, OASIS, [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=xacml](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml)
13. Ian Denley and Simon Weston Smith, *Privacy in clinical information systems in secondary care*. *British Medical Journal*, 1999. **318**: p. 1328-1331.
14. *Platform for Privacy Preferences (P3P) Project*, W3C, <http://www.w3.org/P3P/>
15. *Telecom Glossary*, ANSI Standard T1.523-2001.
16. *The Data Protection Act*, in *United Kingdom*. 1998.
17. *The Common Criteria*, August 1999. <http://www.commoncriteria.org/cc/cc.html>