Towards fine-grained access control for health grids

Andrew Simpson
(Representing joint work with
David Power, Mark Slaymaker, and Eugenia Politou)

Contents

Background

- The UK is in an exceptional position to support the development of novel technical solutions to support individualised patient-specific healthcare: the UK has a national health service (which is a genuine virtual organisation) and its government is ensuring that IT plays an increasing role in healthcare delivery
- To succeed, it will first be necessary to develop an IT infrastructure capable of interfacing with systems deployed (both now and in the future) within the National Health Service in a secure and ethical fashion
- Such a resource could, for example, support clinical trials, train algorithms, facilitate healthcare monitoring and self-management, and serve as a training resource

Health grids

- There are many definitions of what constitutes a grid
- What commonly runs through each of these definitions is the notion of a virtual organisation—disparate logical and physical entities that span different administrative domains coming together to form a single logical entity
- The term health grid refers to a growing area of joined up healthcare in which medical data is shared between clinical sites in accordance with clearly defined access control policies
- Without such policies it would be neither ethical nor legal to allow access to the data

GIMI

- The GIMI (Generic Infrastructure for Medical Informatics) project is a UK research project funded by the Department for Trade and Industry
- It is led by the University of Oxford, and involves academic collaboration with University College, London and Derby University, and industrial collaboration with IBM and e-San
- The primary focus of the project is to produce a secure, interoperable infrastructure to join together medical data from a wide variety of sources
- The project team is focusing initially on two very different applications with a view to establishing proof of concept

Application 1: long-term conditions

- Within the UK, there is a drive towards self-management as a means of improving the health of patients with diabetes or asthma—the management of diabetes-related complications consuming nearly 10% of the NHS budget in England and Wales, while 1 500 people die from asthma in the UK each year
- Research will focus on the refinement of patient-specific models to provide personalised feedback to patients and the development of robust algorithms for alerting clinicians when the patient's data deviates from the expected pattern
- The development of the models of normality for people with diabetes or asthma must be supported by validation across a large population of users (of the order of tens of thousands of patients)

Application 2: radiologist training

- The PERFORMS (PERsonal perFORmance in Mammographic Screening) system is a self-assessment scheme used by the NHS Breast Screening Programme as part of their quality assurance process
- SMF (Standard Mammographic Format) is a standard to which all mammograms can be converted to with a view to eliminating any differences in, for example, equipment settings
- Within the e-DiaMoND project (which preceded GIMI), the largest collection of annotated mammograms for research was collected
- The second validating application involves is a Web-based self-assessment application based on PERFORMS which takes advantage of the e-DiaMoND archive of mammograms and SMF technology

A requirement for fine-grained access control

- The requirements for an access control model for a system such as GIMI can be simply stated: it should be flexible and fine-grained
- There is no way of stating a priori what access control policies will be implemented at each site: some requirements will be nationally (or internationally) mandated, while others will be due to local policies and requirements
- The best that one can do in such circumstances is to offer a system that is sufficiently flexible to accommodate these different needs
- We intend to offer a model of access control that could support, on the one hand, one logical database with one national DBA and, on the other hand, every patient record being associated with exactly one hospital, which, in turn, has exactly one DBA

Legal and ethical obligations (from the DPA of 1998) include:

- Personal data shall be obtained for one or more specified and lawful purposes, and shall not be further processed in any manner incompatible with those purposes
- Personal data shall be adequate, relevant and not excessive in relation to the purpose or purposes for which they are processed
- Personal data processed for any purpose or purposes shall not be kept for longer than is necessary for that purpose
- Appropriate technical and organisational measures shall be taken against unauthorised or unlawful processing, and against accidental loss or destruction of, or damage to, personal data

Legal and ethical obligations (from the principles of the Caldicott Guardian) include:

- Every proposed use or transfer of patient-identifiable information within or from an organisation should be clearly defined and scrutinised
- Don't use patient-identifiable information unless it is absolutely necessary
- Use the minimum necessary patient-identifiable information
- Access to patient-identifiable information should be on a strict need-to-know basis
- Everyone should be aware of their responsibilities
- Understand and comply with the law

Security architecture

- We create a secure access service to act as a gatekeeper for the web service that we would like to secure
- This access service is itself a web service, and a part of an overall system that has consistent mechanisms for authentication, secure message passing and access control
- At each node the services are grouped into internally and externally facing services and the virtualisation of the data sources is assumed to take place at the service level
- This architecture allows each node to retain full control of its data and to determine who can access it (and when), in accordance with the principles of the Caldicott Guardian
- All user interactions with a site are made via the externally facing services, and it is only these externally facing services that are required to present a consistent interface

Security use cases

- We present, at a relatively high level of abstraction, a number of security use cases that have influenced our thinking
- We do not claim that the use cases are exhaustive, but we would claim that the use cases we present here are representative of the requirements that health grids should satisfy
- From the point of view of security requirements, all of the other use cases that we could provide are specializations of one, or subtle combinations of several, of the use cases detailed: for example, the combination of use cases 2 and 3 are representative of a situation in which a clinician wishes to update the access control policy from a remote hospital
- It should be noted that these use cases are associated with an idealised health **data** grid

Use case 1: distributed queries of patient data

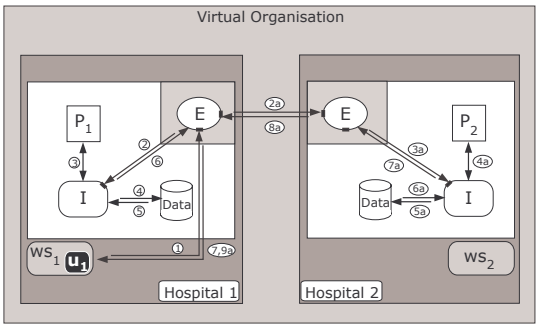

Use case 2: working at a remote hospital

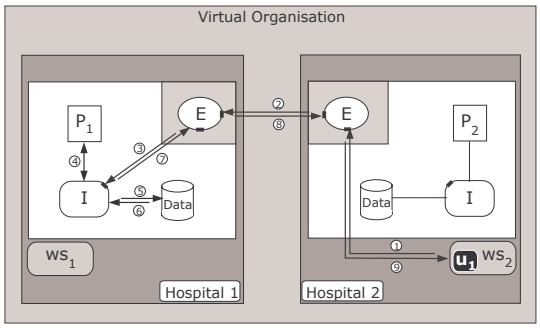

Use case 3: delegation of access permissions

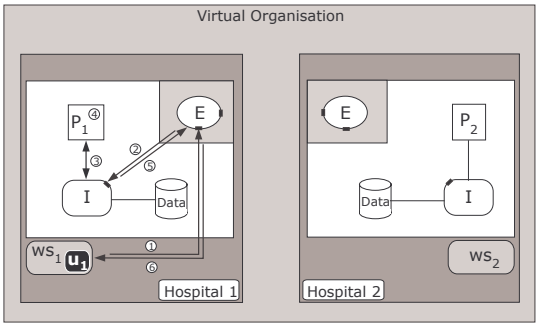

Use case 4: external access

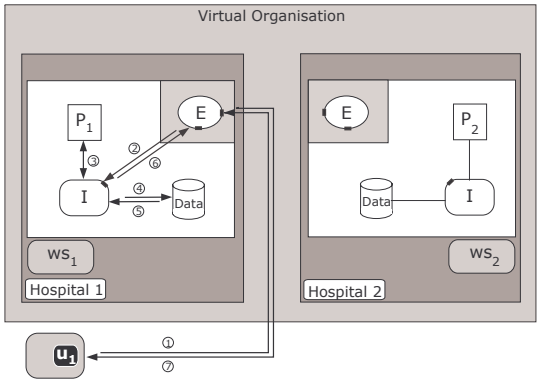

Use case 5: modification of data

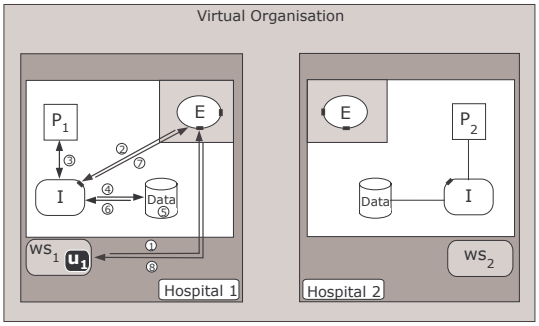

Use case 6: transferring patient records

- Our final use case draws together aspects of each of the previous use cases; we include this as it imposes an additional requirement that is not introduced by the more generic use cases
- To be able to move the data it will first need to be read: this may involve a distributed query as data may already be present at other hospitals
- The data will then need to be deleted from one hospital and copied to another (as the responsibility for it has transferred): this will involve the modification of data
- Finally, the access policies at both of the hospitals may need to be changed to reflect the change of ownership of the data
- If an error were to occur during this transfer process the system could be left in an unstable state: to prevent this, a notion of transaction management must be used

Our approach to access control for health grids

- The resource which we are trying to provide access to may well have its own access control mechanism: using a user mapping function, it would be possible to use this mechanism via a web service interface provided by the vendor
- However, there are two problems associated with this approach:
  - The existing mechanism may not afford the opportunity to express sufficiently flexible and fine-grained access control policies—which, after all, was one of our initial motivations
  - If the system allows access to a wide range of resources, it would be extremely difficult to coordinate the access control policies across these different resources—again, returning to our original motivation (developing secure health grids) there is a clear need for local autonomy

Our approach to access control for health grids

- Ideally, all access control for the resources at a single node should be determined by a single set of coordinated policies, with all requests for access to the existing web service having to comply with these policies
- As such, the enforcement of access control should take place in the wrapping service
- The use of a generic proxy in this case is an advantage due to the fact that all access control takes place at a single point for each node
- Since we adopt the XACML architecture of defining policies we can safely assume that our secure access service is acting as the Policy Enforcement Point (PEP), and will also handle any obligations

Requirements for deployment

- First and foremost an appropriate service-based infrastructure is needed: the current trend—which appears set to continue—is towards using web services for this purpose
- Many of the differences between the use cases presented are related to the identity of the user and the location of the workstation that they are trying to access data from
  - To provide user credentials, X509 certificates can be used (similarly, it would be possible to use X509 certificates to identify the workstation that the request was sent from)
  - An alternative approach to authentication is to use Kerberos

Further issues: describing denial of access

- They may be situations in which access is denied because a patient does not wish certain aspects of their data to be used for anything but primary care
- In this case it would be relatively simple to deduce much of the data from messages describing the denial of access: there is the potential for information flow
- The suitable description of denial of access—to avoid both the drawing of false conclusions and the potential for information flow—is a key issue
- Within GIMI, we are exploring a query modification approach to get around this problem—this fits well with our gatekeeper approach

Further issues: credentials

- Credentials must be portable if doctors are to be able to access data when at a remote location
- External services need to be trusted to pass users' credentials to other services: proxy certificates require the user to trust the intermediate services to use the proxy as intended; solutions that require the user to sign a more specific request reduce the ability of an intermediary to make useful decisions for the user
- It must be possible to credentials: current systems often rely on the service actively requesting lists of revoked certificates, leaving a window of opportunity for people to misuse the credentials; the use of the Online Certificate Status Protocol (OCSP) may reduce the window to a minimum but this has the drawback of requiring frequent communication

Further issues: secure deletion of data

- One of the requirements of a secure health grid is that it should be possible to give temporary access to data
- It is important that people given temporary access do not make a copy of the data
- In our idealised health grid, data would have a lifetime and would be deleted after use
- The utilisation of Digital Rights Management (DRM) may go some way to solving this problem

Further issues: a formal approach

- As well as the significant benefits it offers, XACML has a number of drawbacks—not least of which is its verbosity
- We are using Z to capture policies: this allows the validation of policies and policy sets against certain requirements
- Translation from the mathematical language of Z to XACML is a relatively mechanical process
- In addition, we have just started investigating the use of Communication Sequential Processes (CSP) to reason about **evolving policies**

Wider applicability

- We acknowledge that there are domains other than healthcare that require flexible, fine-grained models of access control
- The work presented here was originally motivated by e-Health in general, and eDiaMoND and GIMI in particular
- However, given that the services that we would like to secure are often provided by a third party vendor—and, as a result, a key requirement for the approach was that it should be sufficiently generic, extensible and adaptable—we would hope that the approach we are developing will have wider applicability than the healthcare domain

Summary

- Background
- Health grids
- GIMI
- Security architecture
- Security use cases
- Wider applicability