A Decentralised Approach to Electronic Consent and Health Information Access Control

Christine M O'Keefe

CSIRO ICT Centre, GPO Box 664, Canberra ACT 2601, Australia
Christine.OKeefe@csiro.au

Paul Greenfield

CSIRO ICT Centre, Locked Bag 17, North Ryde NSW 1670, Australia
Paul.Greenfield@csiro.au

Andrew Goodchild

Distributed Systems Technology Centre Pty Ltd, GPO Box 2434, Brisbane QLD 4001, Australia
andrewg@dstc.edu.au

This paper describes an eConsent *model and demonstrator used to investigate the implementation of patient consent as a means of controlling access to electronic health information shared between healthcare providers. The model and demonstrator described here are designed to operate in an environment of independent cooperating healthcare facilities, such as medical clinics and hospitals, where each facility is responsible for controlling access to the health information in its keeping, according to the patient's expressed conditions as recorded and held by the facility. Novel, privacy-preserving transfer protocols are used to ensure that access to the health information at the receiving facility continues to be governed by the patient's consent. The work was well-received at a symposium where a wide range of stakeholders were offered an opportunity to consider the clinical, legal and technical feasibility of the approach represented by the demonstrator.*

ACM Classifications: H.4 (Information Systems Application); J.3 (Computer Applications: Life and Medical Sciences)

ACKNOWLEDGEMENTS

The authors thank the Australian Government Department of Health and Ageing for funding the research and development for this project. We also thank the other team members for their contributions, particularly Donaugh Austin, Marcus Bowman, Eddy Cheung, Sam Heard, Vincent-Philippe Lauzon, Peter Rickwood, David Rowed, Stephen Shrimpton, Andrew Waugh and Bo Yan.

The work reported in this paper has been funded in part by the Cooperative Research Centres Program through the Department of the Prime Minister and Cabinet of the Commonwealth Government of Australia.

1. INTRODUCTION

The effective coordination of an individual's health care relies on the sharing of personal health information between different health and community care service providers. While the use of

Manuscript received: 2 March 2004
Communicating Editor: Phoebe Chen

computing and communication technologies supports and enhances the ability to share information, it also challenges existing principles and work practices governing information exchange. In particular, patients may not want their personal health information to be transferred without their explicit knowledge and consent, and will want to retain the right to withhold their consent to both transfer and access.

This paper describes a model and corresponding software prototype or demonstrator designed to investigate some of the issues arising from the implementation of patient consent in the sharing of health information between healthcare providers. The research was based on material in the collection of papers by Clarke (2000; 2001; 2003), Coiera (2001) and O'Keefe (2000) which were prepared as background material within the context of the Australian Government Department of Health and Ageing (DoHA) Electronic Consent project (DoHA, 2002). The term *eConsent* was coined to refer to a mechanism through which a patient can record the conditions under which they consent to their personal health information being accessed and shared between health care providers.

In the Electronic Consent project, DoHA commissioned four research and development projects to identify and trial potential mechanisms by which a patient could record the conditions or instructions under which their information may be transmitted from one person to another on an event-by-event basis. The research reported in this paper was funded as one of these projects, and aimed to develop an experimental model and a proof-of-concept demonstrator to explore these ideas. The scope of the project was deliberately limited and, in particular, there was never an intention to develop a production system, nor to address wider issues surrounding the concept of consent and access control over personal health information. DoHA hosted a symposium at the conclusion of the research and implementation phase of the project where a wide range of stakeholders were offered an opportunity to consider their clinical, legal and technical feasibility of the approaches represented by the four demonstrators. These stakeholders were chosen to represent organisations and individuals most likely to be instrumental and influential in any implementations of emerging electronic consent mechanisms in the Australian healthcare sector. Stakeholders included state and territory privacy law officers, health policy and technical managers, patients, clinical craft organisations, and representatives of public and private service delivery organisations.

Our work produced several key outcomes. The most important outcome was the development of the concept of "placeholders", which are used in novel, privacy-preserving, anonymous information transfer protocols. A placeholder essentially presets consent conditions for records such as test results that are expected to be transferred in the future. These were recognised by currently-practising clinicians as "a real step". More generally, we demonstrated the technical feasibility of implementing an electronic consent system compatible with the current healthcare system, where healthcare facilities have control over their locally-held information and patients choose the degree of linking of their own health record. Our model explored the bounds of what is possible in a decentralised and non-standardised healthcare environment and our demonstrator implemented the machine-processing of consent statements.

The remainder of this section describes our model for an eConsent system, including the design principles and goals underpinning our approach. We also give a brief overview of related work. In Section 2, we outline the methods we used in our research, including success measures, the design process and the evaluation process. Section 3 then describes our eConsent model and demonstrator, including details of the eConsent objects and transfer protocols used. This is followed in Section 4 by a discussion, including an evaluation against the success criteria, an extrapolation to expected performance characteristics of a production system based on our demonstrator and a comparison with the demonstrators produced by the other projects.

1.1 An eConsent System

An *eConsent system* is a system which allows the electronic capture and recording of a patient's granting or withholding of consent for access to their personal health information and then uses these consent conditions to manage access to this health information. Accordingly, an eConsent system needs to support mechanisms for:

- capturing a patient's expressed wishes regarding who can access their personal health information. The concept of capture covers the creation, modification and revocation of statements of those expressed wishes.
- applying the statements of patient's expressed wishes to grant or deny access to health information. Application of these mechanisms may be achieved manually by an authorised person or automatically by a computer program.
- communicating a patient's consent between health care providers such as GPs and specialists. In particular, a receiving provider may wish to retain documentary proof of their authority to access particular health information in case of future dispute.

1.1.1 Design Decisions, Principles and Goals

The overarching objective of our research was to design an eConsent system which could be introduced into the existing (Australian) decentralised healthcare environment, with minimal disruption to current infrastructure and work practices. This objective led to the following fundamental design decisions, principles and goals:

Independent, Cooperating Health Care Facilities with Associated Providers

Providers working at independent, autonomous health care facilities cooperate with each other by sharing patients' health information for the benefit of the patient, but only under the constraints imposed by patients' consent conditions. Each health care provider has a relationship with one or more health care facilities, reflecting the fact that the provider provides health care under the legal responsibility and organisational infrastructure of that facility. In our model we assume that each facility has a register of its associated providers. This registry is used to determine whether a given provider is acting in a certain role at any particular time (*provider-role resolution*).

No Centralised Storage of Personal Information

We have based our model on the assumption that patient health and consent information will continue to be decentralised and stored locally at facilities, at least for the foreseeable future.

No Federal or State Mandated Patient Identifier

The current Australian health care system does not include a centralised mechanism for uniquely identifying patients, and this is unlikely to change in the near future. Patient identification in our model is therefore based on provider assertion and the underlying trust relationship. Of course, our model would still be applicable if a patient identifier is introduced to the health care system at some time in the future.

eConsent Design Principles (Coiera, 2001)

Briefly, Coeira's design principles for an eConsent system aim to ensure that an eConsent system will grant access to those with consent or authorisation (Principles 1, 2, 4, 5, 6), will deny access to those without consent or authorisation (Principles 3, 7) and will be minimally disruptive and burdensome (Principles 8, 9).

Stateless, Asynchronous Transfers

The system should support stateless, asynchronous transfers of electronic health information. This is important because the information exchanges between facilities that relate to a single episode will often be spread over a period of time. For example, the simple act of a GP requesting a test from a pathology lab and receiving the results will result in multiple transfers that occur over a number of days. It is also vital that transfers do not depend on collaborating system being constantly operational and accessible throughout the entire transfer process.

Anonymous Transfer Protocols

The information transfer protocols should be anonymous, that is, identifying information (such as the patient's name) must not accompany transferred health information, yet the information must be immediately and automatically filed in the correct place in the receiving health information system and must be associated with the appropriate consent conditions as soon as it arrives.

Standing and Single-Use Consent

The system should support *standing* as well as *single-use* consent statements. A standing consent is one which is intended to be long-lasting and cover many requests for access to health information from various providers for various purposes. On the other hand a single-use consent is intended to be applicable for a single request only. It is likely that an eConsent system will need to support both types of consent, for example a patient may set a default consent statement for a health data item but then adjust it on a case-by-case basis with single-use consent statements.

Full Record Transfers

The system should support the transfer of full records, not just tailored health summaries. The system should be capable of supporting the transfer of all types of health information, including health summaries, case notes and medication records, and images such as X-rays and scans. As patients are increasingly mobile, it is becoming desirable to transfer whole practice records between facilities. It would not be sensible to have separate systems for sharing tailored health summaries and full health records.

1.2 Related Work

The work reported in this paper was commissioned as part of the Australian Government Department of Health and Ageing Project "Consumer Consent in Electronic Health Data Exchange – "e-Consent", (DoHA, 2002). Many other countries are addressing the problem of patient consent in the context of making a single, integrated, lifetime patient record available to health care providers, (HealthConnect Program Office, 2002 and 2003; DOH, 2004; National Health Service, 2004; National Research Council, 2004; New Zealand Health, 2004; Office of Health and the Information Highway, 2004).

Privacy and digital rights are closely related concepts, both concerned with letting the 'owner' of information control who can access it and under what circumstances. Digital rights are managed by DRM (Digital Rights Management) technologies. It is possible that DRM technologies could be adapted to express the access control constraints needed by consent-based privacy mechanisms. One significant shortcoming in current DRM languages is that they only express permissions, saying who can do what and when. eConsent also needs to be able to express denials, allowing patients to state who is not allowed access to their personal information, see Ianella (2001).

Privacy and access control have been issues of concern ever since it has been possible to connect repositories of electronic health records over networks. Anderson (1996) and Mandl, Szolovits and Kohane (2001) discuss some of the issues relating to privacy and security of electronic health records.

Much of the background for this project came from the work commissioned by DoHA mentioned above, including the papers Clarke (2000; 2001), Coiera (2001) and O'Keefe (2000). The concepts of eConsent were defined and refined as part of this work. In particular, Coiera introduced a model based on general and specific permissions, such as 'general grant with specific denial', which was taken up and developed as part of this project.

2. METHODS

2.1 Success Measures

The success measures fall into the following two broad categories, *Social and Business Criteria* which reflect an initial evaluation of the likely impact of an eventual production system on the provider-patient interface and on clinical workflow; and *Technical Criteria* which reflect whether the demonstrator has shown that it would be feasible to build an eventual production system based on this model.

The initial indicative relevant Social and Business criteria were determined to be:

- Acceptably short time added to medical consultations to use the eConsent system;
- Enhancement of information flow between providers;
- Simplicity and ease of understanding by both patients and providers;

The measure of success was positive feedback on these Social and Business criteria from the stakeholders at the Symposium.

The relevant Technical criteria were determined to be:

- Implementation of the key features of the model, including:
  – Stateless, asynchronous transfers of electronic health information;
  – Anonymous information transfer protocols;
  – Facility-based information storage, transfer protocols and provider role resolution;
  – Support for full record transfer
  – Support for standing and single-use consent ; and
- Identification of any factors of the model which might negatively impact on the performance, reliability, scalability or security of an eventual production system.

The measure of success was the assessment of the demonstrator by the implementation team and feedback from the stakeholders at the Symposium.

2.2 Design Process

Case studies were used throughout the design process to inform our research and help us define the model and demonstrator. The first two case studies were developed from ones in the Catalogue of Cases (Clarke, 2000), while the third was suggested as a basis for comparing different demonstrators. In the first case study, a patient wishes to have a confidential screening test for HIV, in the second, a patient is in an emergency ward with a suspected cardiac arrest and in the third case study, a patient consults their GP and is referred to a specialist for further investigation.

We first developed formal requirements for an eConsent model and demonstrator, with the goal of satisfying all of the eConsent Design Principles (Coiera, 2001). The requirements analysis led to the design and the development of a software demonstrator at the CSIRO laboratories (see O'Keefe, Goodchild, Greenfield, Waugh, Cheung and Austin, 2002 and Rickwood and Bowman, 2002). In particular, this phase included the development of the model architecture and transfer protocols, classification of requirements as essential or non-essential and simplification of some non-key aspects of the design due to time constraints.

The fundamental requirements were identified as: the health system would continue to be composed of independent, cooperating health care facilities; there was to be no centralised storage of health information and there would be no centralised provider or patient register. These design decisions were embodied in our demonstrator as follows:

- Consent is managed at each facility by a local eConsent system which is tightly integrated to the system that holds and presents the health information
- Default eConsent policies can be defined to minimise the impact on facility workflows.
- eConsent is associated with roles within a facility and each facility maintains a *provider role register* to map the providers to the roles they occupy at any given time.
- Some of the burden of administering the eConsent system can be taken away from providers by defining a separate class of system users called *Administrators*. These users can create and maintain eConsent information while not being able to access any associated health information.
- Health information can be transferred between facilities. Transferred information becomes a normal part of the health record at the destination facility and is managed as such. All transfers are initiated by a request from the destination facility and the transfer and acceptance of consent conditions always precedes any transfer of health information.

Our model and demonstrator provide the functionality proposed by Coiera's design principles (Coiera, 2001), with the following qualifications: our demonstrator supports only explicit consent not implied or inferred consent, although this could be added very easily; and our model supports patient denial of emergency override, however this is not implemented in the demonstrator.

2.3 Evaluation Process

The evaluation was done in three stages. The model and demonstrator were first evaluated at an internal review workshop, then presented to a DoHA oversight committee and finally, our model and demonstrator were presented at an *eConsent Symposium* in July 2002 hosted by DoHA. This symposium was attended by stakeholders including state and territory privacy law officers and health policy and technical managers, patients, clinical craft organisations, and public and private service delivery representatives. These stakeholders were chosen to represent those organisations and individuals most likely to be instrumental and influential in any implementation of any emerging electronic consent mechanisms in the Australian healthcare sector and were offered an opportunity to consider the clinical, legal and technical feasibility of the models as presented and demonstrated.

The internal evaluation identified advantages and disadvantages of the design decisions that were taken throughout the initial exploration, high level design and implementation phases. During the internal review workshop, currently-practising clinicians were "walked through" the demonstrator in a structured way and were asked to discuss their understanding and views of the model and the demonstrator, although, because it was a demonstrator, it was only possible to hypothesise on how the system would perform in current health care facilities.

The eConsent Symposium started by exploring the eConsent challenge with a group of stakeholders who had had no previous exposure to the issues. We first gave a plenary overview presentation about our eConsent model including the design decisions and goals, and then conducted two demonstrations of the demonstrator software. The first of these demonstrations concentrated on the more technical features and issues, and the second focussed more on cultural features and questions.

3. THE ECONSENT SYSTEM MODEL AND DEMONSTRATOR

The eConsent system model described in this paper is designed to operate in an environment of independent, autonomous health care facilities. These facilities cooperate with each other by sharing patients' health information for the benefit of the patient, but only under the constraints imposed by patients' own stated consent conditions. Each facility is assumed to be running its own independent and autonomous health information system, with no dependencies on each other or on central authorities for the shared storage of health information or for a universal patient registry. Each facility is expected to manage its own database of health information and associated consent conditions, and is responsible for granting or denying all requests to access the health information in its keeping after having evaluated the appropriate eConsent statements. We use the concept of an eConsent object (eCo), see O'Keefe (2000) and Coiera (2001), to capture and store these eConsent statements.

3.1 eConsent and eConsent Objects

Our eConsent system holds patient's formally expressed and agreed access consent conditions in the form of machine-processible eConsent objects (eCos). Each eCo simply holds one or more rules that state who can or cannot access the associated health information under what circumstances. An example of the type of consent and denial statements held inside an eCo is:

> I give access consent only to the GPs at Sollhull Surgery and the mental health specialists and emergency staff at Wall Hospital (except for Dr Smith).

Each eCo can contain a number of eConsent statements, each of which can grant or deny access to the associated health information to individual providers or providers acting in specified roles. Examples of such roles include "primary carer", "treating team member" and "surgeon", and patient-specific roles such as "my treating team". An eCo can be attached to any part of a patient's health information, from an individual health record entry to the whole practice record. Access to any particular health record entry can thus be governed by more than one eCo, each of which can contain more than one eConsent statement. Any request to access a health record entry results in all of the associated eConsent statements being evaluated and reconciled, in the context of the requesting practitioner and the facility holding the information, before access is granted or denied.

eConsent statements do not identify patients in any way and are only associated with them through being attached to some part of the patient's health information. Practitioners wanting to access a particular health record first have to navigate to or otherwise locate the record, and then the eConsent statements associated with the record are used to determine whether or not access is granted.

Our model assumes the existence of a health system-wide registry that can be used to uniquely identify both practitioners and facilities within eConsent statements. Provider roles are also allowed within eConsent statements and these are resolved locally within each facility. It is likely that standardised roles, such as "treating physician", will have to be agreed if role-based consent statements are to be transferable between facilities.

The ability to attach eCos to any part of a patient's health record means that individual entries may be guarded by a number of separate eCos, each of which may contain multiple eConsent statements. For example, an individual health record may be guarded by the practice default eCo, by the patient's default eCo, by a patient-specified eCo applying to all their mental health records and by an eCo specific to this particular record. This possibility raises the question of conflict resolution – what should be done when access is granted by some eCos but denied by others. The demonstrator implementation resolved this problem by assuming that eCos were always applied in a strictly hierarchical or nested manner, as in the previous example, with later (deeper) eCos always overriding earlier (more general) ones. We also assumed that the eConsent statements within a single eCo did not conflict. Further work is still needed on the topics of conflict resolution and eCo ordering to investigate whether this assumption of strict nesting is adequate in practice, and whether we can adequately resolve conflicts in non-hierarchical information stores. The full implementation of provider roles will also introduce the possibility of conflicts within a single eCo that can only be found when roles are mapped to providers at the time when access is requested, and further work is needed to investigate ways of detecting and resolving these conflicts.

3.1.1 eCo Structure

eCos are internal data structures held within the eConsent system. eCos can also be transferred between facilities and in this case they are serialised into XML before transmission. The eCos used in our demonstrator eConsent implementation contained the following information:

- A default consent statement applying to all requests unless overridden by an explicit eConsent statement. This default consent can be:
  - Nothing (no default eConsent statement applies)
  - Deny (deny access unless explicitly granted)
  - Grant in emergency (grant access in emergency and deny otherwise)
  - Grant (grant access unless explicitly denied)
- A set of non-contradictory eConsent statements, each specifying the following:
  - A provider or facility identifier. If a facility identifier is specified, this statement applies to all providers registered with that facility.
  - The explicitly specified access. This may be one of the following: denied, granted or granted in emergency.
  - A link to the eCo it is replacing (if any), used to maintain an audit trail of eCos.
  - Persons to notify when emergency access override is invoked.
  - Time of eCo creation.
  - Author of the eCo.

3.1.2 eCo Creation and Management

eCos are attached to health record entries when they are first created. Newly created records will normally be assigned an appropriate eCo, using the default eCo defined for the patient or the facility, or one previously defined to meet the patient's stated requirements. If none of these eCos meet the patient's current requirements, a new eCo can be created and attached to the newly created records. eCos can also be created to protect health information that does not yet exist but is expected to arrive in the future, such as the results of tests or referrals. The capture of consent conditions would normally take place during a consultation with a health care provider or with an administrator immediately after a consultation.

A patient can vary or revoke previously given consents at a facility by varying or revoking the associated eCo. Consistent with our assumption of independent, autonomous, facilities, revocation of consent is handled on a facility-by-facility basis and revocation is never propagated to other facilities even if the associated health information was sent to them. Although this is similar to the situation that exists with paper records today, further work is needed to improve and automate the tracking of transferred eCos and the distribution of revocation requests. Figure 1 shows the

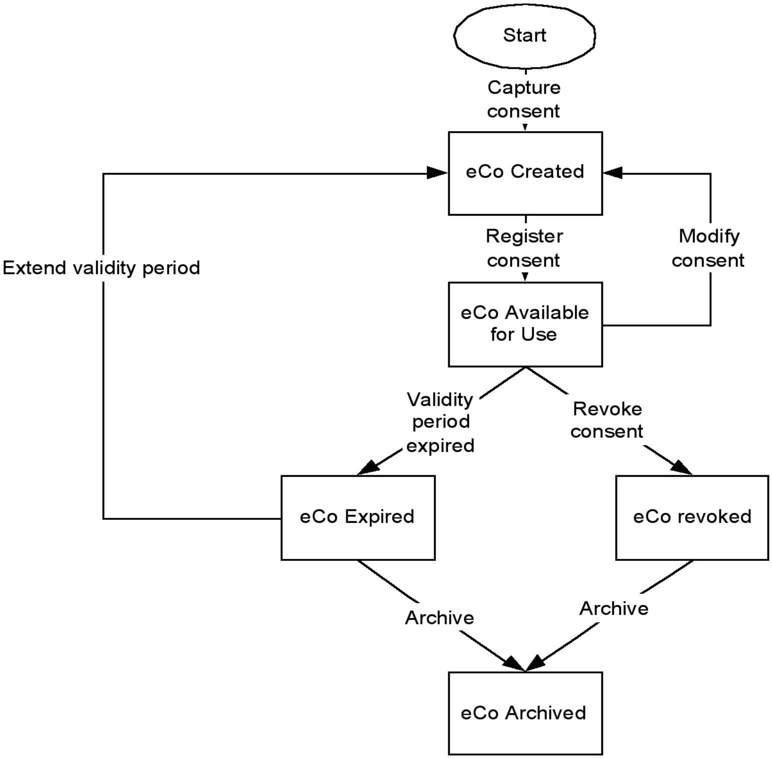

Figure 1: Life Cycle of an eCo

complete lifecycle of an eCo, from its initial creation through to expiry or revocation and subsequent archiving.

3.2 eConsent Systems and Consent Transfer

Each facility is expected to run their own autonomous and independent health information and consent system. Information is shared between these systems through information transfer only, and there is no provision for 'remote access'. Providers only ever access information held at their local facility but they can request that information be transferred from other facilities so that it can be accessed. These transfers will be preceded by consent conditions that will be used to control access to the health information once it arrives.

Facilities are expected to run their own provider-role registries that are used to dynamically map providers onto roles when validating access requests. Each facility holds the details of each of its associated providers and can judge whether a given provider is fulfilling a particular role at a given time for a certain patient. For example, in a GP practice the associated providers would include the GPs and any allied health professionals working at the practice. The practice would thus be able to judge at any time whether a particular GP working there was acting as the "primary carer" for a given patient.

3.2.1 Transfer of health information

There are many situations in coordinated health care in which it is necessary to transfer health information between facilities. For example, health information must be transferred between facilities for referrals for specialist consultation and with requests for pathology tests. The result of all these transfers is that an individual's health information may end up being fragmented, distributed and replicated across several facilities. It is vital that the *same consent conditions* are associated with a given part of a patient's health information record *wherever it is stored*. This implies that the associated consent conditions must always be transferred along with health information to ensure that access continues to be controlled in accordance with the individual's expressed consent conditions.

Our eConsent model and demonstrator uses novel transfer protocols to ensure that transferred health information is automatically protected according to the patient's consent conditions immediately upon arrival at the receiving facility, without needing it to be opened, viewed or otherwise classified. The protocols are anonymous in that no identifying information (such as the individual's name or unique identifier) accompanies the transferred health information. Any requirements in the protocol to identify information for the purposes of filing it away and associating it with the correct eCos are handled through the use of anonymous 'placeholders'.

In our model and demonstrator, *placeholder identifiers* accompany requests for services and information transfer (*response ph* or *info ph* in Figure 2 are placeholders). A placeholder represents a pending entry in the health information system, already filed in the appropriate place and protected by the appropriate consent conditions. When health information arrives at a facility it is immediately and automatically filed and protected using the accompanying placeholder. Placeholders can be used multiple times, and over an indefinite period of time, allowing health information records associated with a single event or request but received at different times and from different sources to be filed together and to be covered by the same eCo.

The use of placeholders is illustrated in Figure 2. There are actually more placeholder identifiers passed between facilities than might at first seem necessary. These additional placeholder identifiers are used to make the protocol "stateless" and more suited to environments where facilities are not constantly connected.

Figure 2 illustrates the health information transfer protocol used in the situation that a GP wishes to refer a patient to a specialist, beginning with an *order request* (the referral) and ending with the *transfer response*, which is specialist's report arriving back at the ordering facility. (Placeholders are

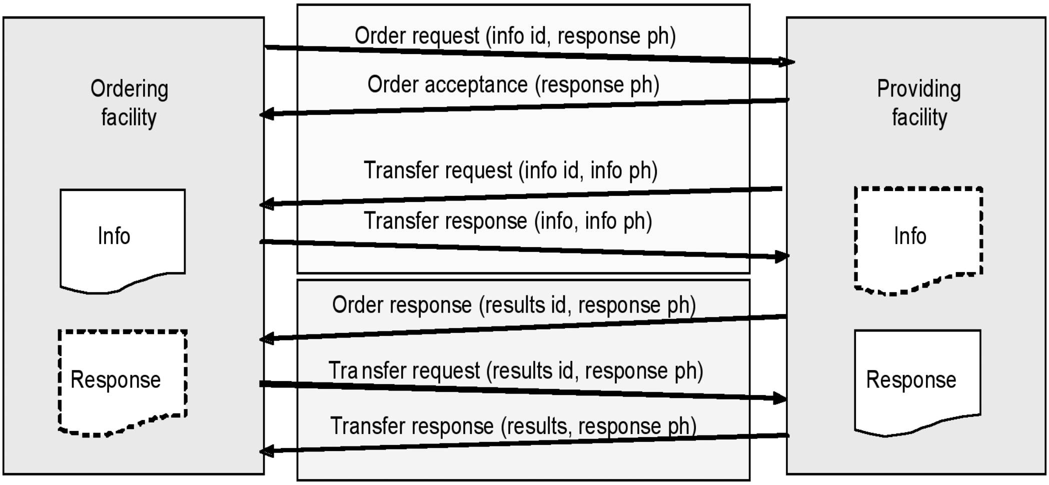
Ordering facility
Info
Response
Order request (info id, response ph)
Order acceptance (response ph)
Transfer request (info id, info ph)
Transfer response (info, info ph)
Order response (results id, response ph)
Tra nsfer request (results id, response ph)
Transfer response (results, response ph)
Providing facility
Info
Response

Figure 2: Referral to a Specialist – Transfer Protocol

represented by *ph* in Figure 2). The ordering facility holds health information relevant to the referral and the identifier (*info id*) is passed over to the providing facility along with the patient's stated consent conditions as part of the *order request*. The providing facility acknowledges their acceptance of the consent conditions by sending back an *order acceptance* message and then fetches this health information by initiating a *transfer request* for the information corresponding to the info id. The ordering facility then uses this identifier to retrieve the requested information and sends a copy over to the providing facility in a *transfer response* message. At the completion of the consultation a similar procedure is used to transfer the report from the providing facility to the ordering facility, triggered by an *order response*.

After the health information has arrived at the ordering facility, the requesting provider is notified that they can now request access to the now locally-held response information. This access is granted or denied in the normal manner according to the associated consent conditions.

Two other novel aspects of the protocol are that information is never transferred until the receiving side has received and accepted the eConsent conditions that are associated with the information; and that information is never 'pushed', only fetched or 'pulled' by a recipient facility after it has been sent a reference that identifies the information available for transfer.

Variants of this protocol can be used for other transfer situations, for example, for emergency access to health information. This protocol variant starts with the emergency facility sending an emergency access request containing information sufficient for the facility holding the information to unambiguously identify the patient, more or less as is current manual practice. The holding facility then responds with a message containing the identifier for the patient's summary health record and its associated eCo. The emergency facility has to reply with an acceptance message and can then initiate a transfer of the required health summary information.

3.3 The eConsent Demonstrator

This section gives an overview of our eConsent system demonstrator, which we have named *MedicClient*. MedicClient is a simplified clinical health information system that uses eConsent to control access to information. The demonstrator supports the creation and access of local health information, and the transfer of information and consent through referrals and test requests.

The purpose of the MedicClient demonstrator was to test ideas from the background papers, experiment with technologies and to provide input to the eConsent research project being undertaken by the Department of Health and Ageing. In particular, it had to demonstrate the processes of capturing a patient's consent, controlling access using that consent, transmitting that consent and supporting subsequent revocation of that consent. The demonstrator was not intended to be a real eConsent or health information system or even a prototype of an eventual real eConsent System.

The scope of the demonstrator was also limited to ensure that we focussed just on the core features of an eConsent system and so were able to complete development in the limited time available for the project. Some of the restrictions present in the demonstrator include:

- Only a restricted form of role-based access is supported. The only provider role allowed is "provider associated with facility A". This role is trivial to resolve since each facility has a register of all the providers that are associated with it.
- Each running copy of the demonstrator, including its local databases, can only be operated by a single user at any one time. This simplification avoids complexities such as the synchronisation of multiple users' changes to the underlying database at each simulated facility
- The problem of conflict resolution was handled by assuming the strict nesting of eCos and forbidding conflicting consent statements within an eCo.

The eConsent model was designed to be technology-independent but the particular technologies chosen for the demonstrator are mentioned here for reference. The development of the demonstrator was carried out using the Microsoft .NET framework using Visual Studio and C#. The database was initially developed using SQL Server but deployed using MSDE to avoid licencing issues.

Messages are transmitted between facilities using TCP/IP and the sockets API. The exchanged information was encoded in XML format, and time-stamped and digitally signed by the sending facility. Each transferred message was encrypted using a symmetric RC2 session key, and the session key was encrypted for transfer using the RSA algorithm.

3.3.1 Information and Consent Views

When a provider logs in and selects a patient, MedicClient displays the *Health Information* view, the normal and familiar view of the patient's practice record. MedicClient shows an index of all the health information to which the provider is granted access and enables viewing of any of this information. MedicClient supports basic health information functions such as creating and filing entries in the patient's practice record.

The patient's practice record is likely to have an organisational structure. Individual entries may be organised into folders, such as "Patient Contact", "Medication List" and "Referrals", and these folders can be nested. The structure of a patient's practice record is not assumed to be uniform across all facilities, nor even across providers or patients in the same facility. It is entirely up to the provider to decide what is most appropriate to a given situation or work practice. Figure 3 shows a patient's practice record, with the provider viewing the contents of a health data item (blood test result).

Clicking on the *Consent View* tab moves the user to a view which shows the consent conditions that apply to the associated health information, as shown in Figure 4. MedicClient supports basic consent management functions, such as the creation, modification and revocation of eCos, as well as the association of consent information with health information.

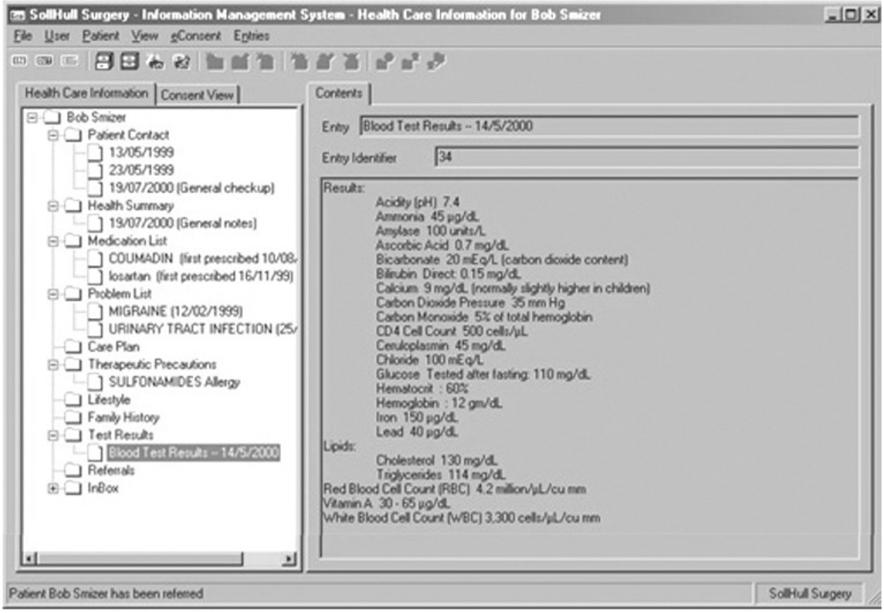
SollHull Surgery - Information Management System - Health Care Information for Bob Smizer
File User Patient View eConsent Entries
Health Care Information
Consent View
Contents

```
Bob Smizer
  Patient Contact
    13/05/1999
    23/05/1999
    19/07/2000 (General checkup)
  Health Summary
    19/07/2000 (General notes)
  Medication List
    COUMADIN (first prescribed 10/08/
    losartan (first prescribed 16/11/99)
  Problem List
    MIGRAINE (12/02/1999)
    URINARY TRACT INFECTION (25/
  Care Plan
  Therapeutic Precautions
    SULFONAMIDES Allergy
  Lifestyle
  Family History
  Test Results
    Blood Test Results -- 14/5/2000
  Referrals
  InBox
```

Entry: Blood Test Results -- 14/5/2000

Entry Identifier: 34

```
Results:
          Acidity (pH) 7.4
          Ammonia 45 µg/dL
          Amylase 100 units/L
          Ascorbic Acid 0.7 mg/dL
          Bicarbonate 20 mEq/L (carbon dioxide content)
          Bilirubin Direct: 0.15 mg/dL
          Calcium 9 mg/dL (normally slightly higher in children)
          Carbon Dioxide Pressure 35 mm Hg
          Carbon Monoxide 5% of total hemoglobin
          CD4 Cell Count 500 cells/µL
          Ceruloplasmin 45 mg/dL
          Chloride 100 mEq/L
          Glucose Tested after fasting: 110 mg/dL
          Hematocrit : 60%
          Hemoglobin : 12 gm/dL
          Iron 150 µg/dL
          Lead 40 µg/dL
Lipids:
          Cholesterol 130 mg/dL
          Triglycerides 114 mg/dL
Red Blood Cell Count (RBC) 4.2 million/µL/cu mm
Vitamin A 30 - 65 µg/dL
White Blood Cell Count (WBC) 3,300 cells/µL/cu mm
```

Patient Bob Smizer has been referred | SollHull Surgery

Figure 3: Viewing a Health Data Item

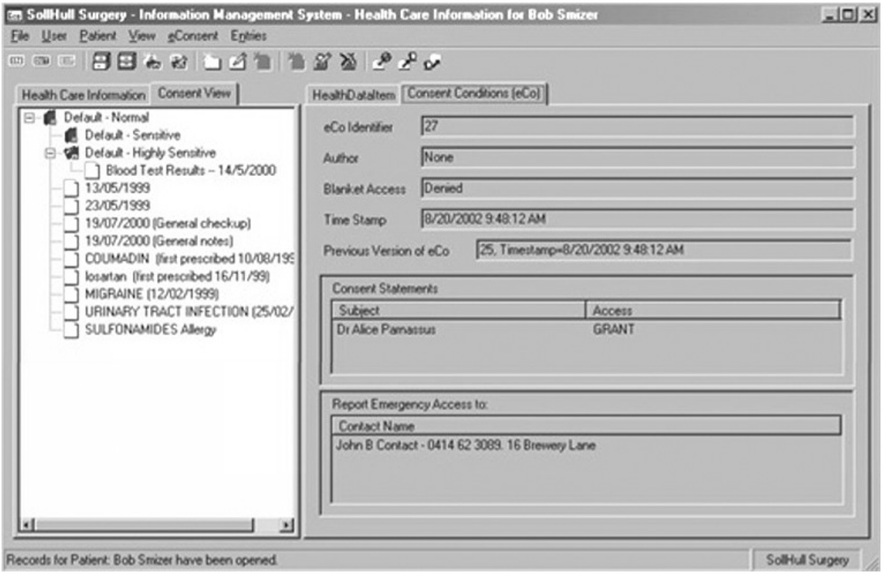
SollHull Surgery - Information Management System - Health Care Information for Bob Smizer
File User Patient View eConsent Entries
Health Care Information
Consent View
HealthDataItem
Consent Conditions (eCo)

Default - Normal
Default - Sensitive
Default - Highly Sensitive
Blood Test Results – 14/5/2000
13/05/1999
23/05/1999
19/07/2000 (General checkup)
19/07/2000 (General notes)
COUMADIN (first prescribed 10/08/199
losartan (first prescribed 16/11/99)
MIGRAINE (12/02/1999)
URINARY TRACT INFECTION (25/02/
SULFONAMIDES Allergy

eCo Identifier: 27
Author: None
Blanket Access: Denied
Time Stamp: 8/20/2002 9:48:12 AM
Previous Version of eCo: 25, Timestamp=8/20/2002 9:48:12 AM

Consent Statements

Subject	Access
Dr Alice Parnassus	GRANT

Report Emergency Access to:

Contact Name
John B Contact - 0414 62 3089. 16 Brewery Lane

Records for Patient: Bob Smizer have been opened.

SollHull Surgery

Figure 4: Consent view

A patient can associate an eCo with any part of their health information. The health information associated with a single eCo is called a consent group and our demonstrator supports nested consent groups. When consent groups are nested the consent conditions in lower groups override consent conditions in higher level groups.

Because consent groups can be nested, they introduce a hierarchical structure over the practice record. This new consent-based structure would normally be completely different from the structure seen in the *Health Information* view. In particular, there would probably be a large default consent group (or groups) and a small number of other groups each with an eCo attached.

The *Health Information* view is the normal view of a patient's practice record. The Consent View is not normally seen and is only used when providers or administrators need to interact with the eConsent system to create or manage eCos.

3.3.2 Transferring Information and Consent

The MedicClient demonstrator supports transfer of referral requests, including attached health information and consent statements, between facilities using the protocols discussed in section 3.2.1. Providers use the Referral Wizard to go through the following steps to send a referral request to another facility.

- Select a specialist facility from the available facilities by consulting the central facility register.
- Select a specialist from the available specialists at that facility by consulting the specialist facility's provider role register.
- Write a referral letter and attach background health records.
- Check that the associated eConsent statements are an accurate representation of the patient's wishes, including granting the specialist consent to access the background health information (see Figure 5).

The referral request and associated health information and eConsent constraints are then transferred to the specialist's facility using the first phase of the transfer protocol shown in Figure 2. The

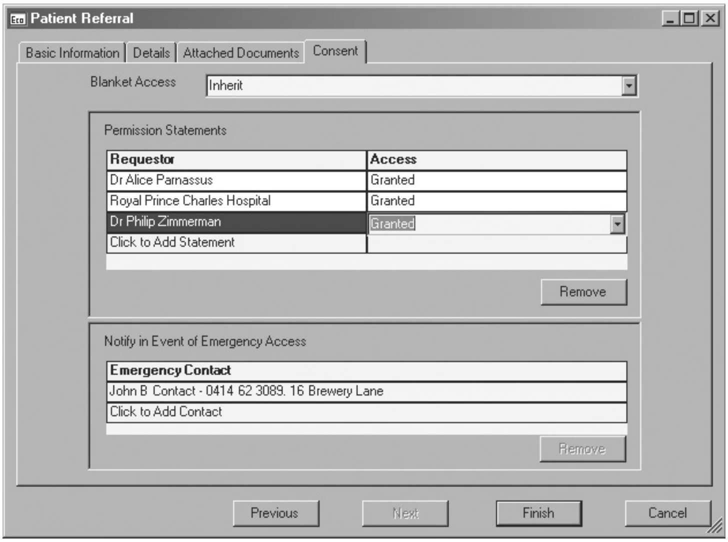

Blanket Access: Inherit

Permission Statements

Requestor	Access
Dr Alice Parnassus	Granted
Royal Prince Charles Hospital	Granted
Dr Philip Zimmerman	Granted
Click to Add Statement	

Notify in Event of Emergency Access

Emergency Contact
John B Contact - 0414 62 3089. 16 Brewery Lane
Click to Add Contact

Figure 5: Granting Access to the Receiving Provider

specialist provider then creates local health information records that contain the findings resulting from the consultation and these are transferred back to the ordering using the second phase of the transfer protocol. All messages exchanged between facilities are signed and encrypted.

3.3.3 Limitations

The implementation of eCos and eConsent in the demonstrator did not fully support provider roles due to lack of time and resources. Supporting roles would just require the addition of provider role identifiers as possible eConsent statement principals, alongside providers and facilities, and mechanisms to resolve these roles when access requests are being evaluated.

4. DISCUSSION

4.1 Evaluation Against Success Criteria

Our success measures fell into two broad categories, Social and Business Criteria and Technical Criteria. Success on the Social and Business criteria was judged through evaluating the feedback received from the stakeholders at the Symposium. Success on the Technical Criteria was evaluated through an internal assessment of the demonstrator conducted by the implementation team and from feedback from the stakeholders. A full report on the feedback received at the internal workshops and at the Symposium can be found in O’Keefe (2002).

Feedback at the Symposium was positive overall, with some questioning of the practicality of eConsent in clinical practice and of the consequences of the replication of health and consent information. Some of the comments received were:

- “This is a very good system for the [health care] system as it currently works and adding to that. Most of us in fact do work at a facility-to-facility level.”
- In the clinical demonstration, there was a comment supportive of our decision to make the facility ultimately responsible, in line with answering the question “Who does the Coroner call?” in the event of something going wrong.

There was concern at the Symposium that the use of eConsent would increase consultation times by adding new tasks to consultation practices. However, with well-considered default consent policies and appropriate workflow, the the burden of eConsent may not be so great. For example, patients can be introduced to the default consent conditions of the facility on initial consultation and the majority may well agree to these policies. Providers and patients may still have to reflect on the sensitivity of information that is being recorded within consultations and modify the consent conditions, but it is arguable how great an imposition this will be.

The questions about the impacts of replication echo our own view of one of the disadvantages inherent in decentralised models. In answering a question about consent revocation we talked about how a patient might need to visit multiple facilities in order to change and/or revoke consent. It is worthwhile to remember that our current health care system is decentralised and patients' health information is already replicated in multiple places and kept in multiple distributed storage sites. It would be a worthwhile research project to build a distributed protocol that would address this issue by improving the tracking of information as it flows around the network.

Technically, we successfully implemented all of the key features of the model, including:

- Stateless, asynchronous transfers of electronic health information;
- Anonymous information transfer protocols;
- Facility-based information storage, transfer protocols and (simplified) provider role resolution;
- Support for full record transfer
- Support for standing and single-use consent

During the evaluations, clinicians made comments that reinforced the soundness of our design choices, such as:

- "Clinicians are becoming unhappy with emailing health information, they want health information to go straight into their databases";
- "The idea of using placeholders to preset consent conditions for expected transferred records such as test results is a real step"; and
- "The fact that a facility needs the ID of a record before it is transferred is good; it implies a relationship between the facilities before health information is transferred."

4.2 Expected Performance of a Production System

Our model, based on facility-based access and facility-level transfers, is potentially much simpler than alternatives that directly transfer information between providers. For example, a register of facilities is smaller and easier to maintain than a register of providers and roles, and each facility only needs to be able to authenticate the identity of other facilities, instead of each provider needing to authenticate potentially the whole community of providers. The facility-based provider register and the associated task of role resolution, is efficient and can reflect dynamic changes in provider roles. The use of distributed, facility-based storage of patient health and consent information also avoids the potential bottlenecks and information ownership disputes that might arise if all information was to be held in a centralised system. We emphasise that, although transfers occur at the facility-level, no user is able to access health information without being granted permission by the patient through the relevant eCos.

The lack of centralised identifier and storage systems means that the linking of health information must be managed by the patient with assistance from the relevant facilities. As soon as a health information transfer takes place, a patient will end up having copies of health and consent information being held at different facilities. When the patient wishes to change or revoke consent information, the problem is to then associate the new consent information with all copies of the related health information. There are several possible technical solutions to this problem, including modifying the transfer protocol to support the tracking of transfers so that modifications and revocations of consent information can be propagated.

4.3 Comparison with the other Approaches

There were significant areas of both commonality and difference between the four demonstrators developed within the DoHA eConsent project.

- Each of the demonstrated approaches granted or denied access to health information through the use of eCos to express the patients' stated access control policies – a Gatekeeper model.
- All of the projects started with an assumed trust relationship between a patient and provider, and some of the other projects also required a trust relationship between facilities or between a facility and a centralised system or data server to appropriately handle access to shared information.
- Each of the projects grappled with the question of how much centralisation was acceptable to stakeholders.
- All the projects demonstrated the use of protocols for exchanging eCos, and three of the four teams demonstrated the use of provider roles in eCos and recognised the need for standardisation of these roles.
- Each of the projects experimented with some form of facility default policy, patient profile or overarching default policy.

Our model differed from the other approaches through its focus on flexibility and working within the current decentralised health care system, although it could also be used in more centralised and standardised systems. Two unique features in our model and demonstrator are the low-risk transfer protocol and the use of placeholders.

Under our transfer protocols, the health information is automatically protected according to the patient's consent conditions immediately upon arrival in the receiving facility without anyone needing to open or view it. In fact the health records and consent information are stored directly into the facility's databases even before anyone is alerted that information has arrived. The sending facility first sends the consent conditions associated with the health information to the requesting facility. If the receiving facility accepts the consent conditions (this would normally be an automated step) then these are stored in the eConsent database on the receiving system and the facility then requests the transfer of the health information. The health information is then transferred to the receiving facility's health information database and is immediately automatically protected by the previously stored consent conditions. This transfer sequence also protects against the accidental sending of health information to the wrong address, since health information is only sent after the establishment of a relationship between the sending and receiving facilities.

A placeholder corresponds to an empty consent group or entry in the patient's consent view or health information record. When the expected health information arrives, it is stored in the spot defined by the placeholder and automatically protected by the appropriate consent information before anyone at the receiving facility even knows that it has arrived.

5. CONCLUSION

We believe that our model supports the degree of flexibility demanded by the current range of different facilities, contexts and social and cultural preferences that exist in the health care environment today. We explored the bounds of what is possible in a decentralised and non-standardised health care environment, although our model would also be suited for more centralised and standardised environments.

As part of our design, we developed novel and secure transfer protocols for transferring health and consent information. Under our transfer protocols, the health information is automatically protected according to the patient's consent conditions immediately upon arrival in the receiving facility without anyone needing to open, view or classify it. In fact the health and consent information are stored directly into the facility's databases even before anyone is alerted that information has arrived. This is much safer than sending health and consent information by email, say, where it has to be received and opened before it can be filed.

Our work in developing a model and a corresponding demonstrator for an eConsent system led us to the following observations. First, it is clear that there are many benefits to sharing health information and eConsent is a key enabler for ensuring that sensitive information can be shared in a way that respects the patient's wishes. The keys to a successful eConsent system is to make it flexible so that it can be tailored to different environments, to make it technology independent, and to ensure it can be used with existing systems. Further, the system should be lightweight in its use so that it has minimal impact on business processes. A good set of default consent statements or policies is vital to ensuring that eConsent has a minimal impact on clinical workflow.

In terms of taking eConsent forward, we see the most important step is to implement one or more trials that integrate eConsent into existing health information management systems and demonstrate the business value of eConsent. This trial in turn should be supported by the appropriate processes to achieve an agreement on the underlying principles for eConsent and standards for eConsent models and communication protocols.

REFERENCES

ANDERSON, R.J. (1996): Security in clinical information systems, British Medical Association, 1996 http://www.ftp.cl.cam.ac.uk/ftp/users/rja14/policy11.pdf Accessed 23-Feb-2004.

CLARKE, R. (2000): Consumer consent in electronic health data exchange: Catalogue of Cases. http://www.health.gov.au/hsdd/primcare/it/pdf/testcase.pdf. Accessed 28-Jan-2003.

CLARKE, R. (2001): Consumer consent in electronic health data exchange: Implementation Considerations. http://www.health.gov.au/hsdd/primcare/it/pdf/implement.pdf. Accessed 28-Jan-2003.

CLARKE, R. (2003): Consumer consent in electronic health data exchange: Background Paper. http://www.health.gov.au/hsdd/primcare/it/pdf/e_bground.pdf. Accessed 28-Jan-2003.

COIERA, E. (2001): e-Consent consumer consent in electronic health data exchange. http://www.health.gov.au/hsdd/primcare/it/pdf/coiera.pdf. Accessed 28-Jan-2003.

DoHA (2002): Australian Government Department of Health and Ageing Project: Consumer consent in electronic health data exchange – e-Consent, see http://www.health.gov.au/hsdd/primcare/it/econsent.htm. Accessed 16-Feb-2004

DOH (2004): UK Department of Health http://www.doh.gov.uk

HEALTHCONNECT PROGRAM OFFICE (2002): Consent and electronic health records – A discussion paper. http://www.health.gov.au/healthconnect/pdf_docs/cons_dp.pdf Accessed 16-Feb-2004.

HEALTHCONNECT PROGRAM OFFICE (2003): Research Report 5: What will be required to manage privacy, http://www.healthconnect.gov.au/pdf_docs/v2-5.pdf Accessed 16-Feb-2004.

IANELLA, R. (2001): Digital rights management (DRM) Architectures. *D-Lib Magazine*, 2001.
http://www.dlib.org/dlib/june01/iannella/06iannella.html Accessed 23-Feb-2004.

MANDL, K.D., SZOLOVITS, P. and KOHANE, I.S. (2001): Public standards and patient control: how to keep electronic medical records accessible but private. BMJ 2001; 322:2837, http://bmj.bmjjournals.com/cgi/content/full/322/7281/283 Accessed 23-Feb-2004.

NATIONAL HEALTH SERVICE (2004): UK NHS Information Authority, http://www.nhsia.nhs.uk Accessed 23-Feb-2004.

NATIONAL RESEARCH COUNCIL (2004): For the record: Protecting electronic health information, US National Research Council. http://www.nap.edu/readingroom/books/for/ Accessed 23-Feb-2004.

NEW ZEALAND HEALTH (2004): NZH Information Service, http://www.nzhis.govt.nz Accessed 23-Feb-2004.

OFFICE OF HEALTH AND THE INFORMATION HIGHWAY (2004): Canada Office of Health and the Information Highway – eHealth Resource Centre. http://www.hc-sc.gc.ca/ohih-bsi/menu_e.html Accessed 16-Feb-2004.

O'KEEFE, C.M. (2000): Consumer consent in electronic health data exchange – eConsent – Preliminary security systems analysis discussion paper. http://www.health.gov.au/hsdd/primcare/it/docs/security.doc Accessed 16-Feb-2004.

O'KEEFE, C.M., GOODCHILD, A., GREENFIELD, P., WAUGH, A., CHEUNG, E. and AUSTIN, D. (2002): Implementation of electronic consent mechanisms – Final analysis paper. http://www.health.gov.au/hsdd/primcare/it/docs/paper.doc Accessed 23-Feb-2004.

RICKWOOD, P. and BOWMAN, M. (2002): MedicClient software user guide. http://www.health.gov.au/hsdd/primcare/it/csirop.htm Accessed 23-Feb-2004.

BIOGRAPHICAL NOTES

Christine O'Keefe

Dr Christine O'Keefe is a research leader in the CSIRO ICT Centre, and currently leads projects in Health Data Integration and Information Security and Privacy Technologies. Christine's recent research in consent and other privacy-enhancing technologies has focused on applications in health informatics. Christine has more than 65 publications and is on the editorial board of two international scientific journals.

In 2000 Christine was awarded the Australian Mathematical Society Medal *for distinguished research in the Mathematical Sciences and in 1996 she was awarded the* Hall Medal of the Institute for Combinatorics and its Applications *in 1996 for outstanding contributions to the field.*

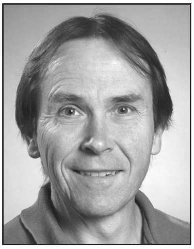

Paul Greenfield

Paul Greenfield is a senior researcher in the CSIRO ICT Centre and runs the Distributed Systems Research stream. This group undertakes research and consulting into enterprise-level distributed computing technologies, trusted storage and service-based architectures. Paul has also consulted on Internet security, content filtering and e-commerce for the Australian government.

Paul has worked in software research and development for 18 years. Since moving to CSIRO, Paul has been the co-author of 10 papers on enterprise computing and information privacy. His current research interests are consistency management for Web Services applications and the impacts that the emerging trusted computing technologies will have on enterprise-level computing.

Andrew Goodchild

Dr Andrew Goodchild is a senior research scientist at the DSTC, and works in a wide variety of roles, including: project management, research or software architecture. For the last five years, Andrew has focussed on research in semantic interoperability and whole of sector information sharing in the defence and health industries. Recent research has focussed on electronic health records. In particular, how a scalable and extensible national electronic health record can be built for Australia. Recently this work has culminated in the development of the Brisbane Southside HealthConnect Trial. Future research topics of interest to Andrew include basic building blocks for a health knowledge network, such as security, health terminologies and computer interpretable guidelines and their application to chronic disease management and preventative healthcare.