Department of Health and Aged Care

Acute and Coordinated Care Branch

Consumer Consent in Electronic Health Data Exchange

Implementation Considerations

Final Version

1 December 2002

Roger Clarke
Xamax Consultancy Pty Ltd

Consumer Consent in Electronic Health Data Exchange

Implementation Considerations

EXECUTIVE SUMMARY

This document is one outcome of a project conducted in order to develop a means of implementing patient consent to the electronic transfer of their data. A 'Background Paper' and documents on security and 'Design Principles, Consent Models and Transaction Rules' have established the foundations.

This document examines implementation factors, with particular emphasis on security. It presents a simple, generalised model that encompasses the relevant entities involved in health care. It distinguishes patients, health carers, third parties including diagnostic services, and other service providers. Consent expressed in electronic form (an 'e-Consent') needs to be gathered by a health carer from a patient, for onforwarding to other parties as appropriate.

The facilities that are used by the various entities are also modelled. A particularly important question that is examined more closely is the means that are available to support interaction between the patient and the health carer. Consideration is given to facilities that are currently mainstream and to those that appear reasonably likely to be deployed in the near future.

The nature of information security is reviewed, its elements briefly described, and the process of development of security strategy outlined. These ideas are then applied to e-Consent. Preliminary analyses are undertaken of the data and processes that need to be protected, the threats and the situations in which they arise, and the harm that will result when the safeguards are inadequate. This leads to a general statement of security needs, supplemented by some more specific requirements.

A security strategy is outlined. Central to this is the notion of an e-Consent object. This would contain all of the data necessary to specify consent, and (where appropriate) denial of consent. It would describe the patient information it applies to, the entity or class of entities authorised to send the data, the entity or class of entities authorised to receive it, the purposes for which it can be used, and the authority for and conditions attaching to its further disclosure. It also needs to support a hierarchy of qualifications to the consent. The processes of creation and use of an e-Consent object are described.

A range of technologies are briefly described, and their capacity to support the implementation of an e-Consent object are assessed. No currently available technology appears likely to satisfy the need. On the other hand, a significant number of standards and protocols offer considerable promise. Prototyping and experimentation are necessary in order to establish which of these are most appropriate to the need.

Consumer Consent in Electronic Health Data Exchange
Implementation Considerations

CONTENTS

1. INTRODUCTION **1**

2. SCHEMAS FOR E-CONSENT IMPLEMENTATION **2**

2.1 Indicative Model of Entities and Roles 2

2.2 Indicative Model of Facilities 4

2.3 Facilities for Interactions Between Patient and Health Carers 6

3. INFORMATION SECURITY GENERALLY **8**

4. INFORMATION SECURITY APPLIED TO E-CONSENT **9**

4.1 Information Security Requirements 9

4.2 Information Security Strategy 11

4.3 An e-Consent Digital Object 13

5. IMPLEMENTABILITY OF AN E-CONSENT OBJECT **16**

5.1 X.509v3 Certificates 16

5.2 S/MIME 17

5.3 PGP 17

5.4 SDSI / SPKI 18

5.5 AADS 19

5.6 Brandsian Certificates 19

5.7 Trust Management Systems 20

5.8 P3P 20

6. CONCLUSIONS **21**

REFERENCES **22**

APP. 1: SECURITY SCHEMA FOR MESSAGE MANAGEMENT **23**

APP. 2: DATA-ELEMENTS NEEDED IN AN E-CONSENT OBJECT **24**

APP. 3A: OUTLINE OF THE X.509V3 CERTIFICATE STRUCTURE **25**

APP. 3B: DETAIL OF THE X.509V3 CERTIFICATE STRUCTURE **26**

APP. 3C: X.509V3 REFERENCES **28**

APP. 4: FURTHER INFORMATION RE SPKI / SDSI **30**

APP. 5: FURTHER INFORMATION RE BRANDSIAN CERTIFICATES **40**

Department of Health and Aged Care
Acute and Coordinated Care Branch

Consumer Consent in
Electronic Health Data Exchange

Implementation Considerations

Roger Clarke

1. Introduction

This is the fourth in a sequence of documents designed to develop 'e-consent' from a concept toward being an implementable specification. It is assumed that the reader is already familiar with the 'Background Paper' and the 'Design Principles, Consent Models and Transaction Rules'. This paper also draws heavily on a resource document entitled 'Security Systems Analysis'.

This document commences by presenting a set of models of the context in which consumer consent needs to operate. The first of these is concerned with the entities and roles involved in health care, and the second with the channels, and in particular the computing and communications facilities available to those entities. Because of the complexity of the health care sector, the models are indicative rather than comprehensive.

The next segment of the document considers the security needs of an effective implementation of e-consent. A preliminary section summarises the nature of information security generally. It then considers particular requirements of informaiton security in the context of e-Consent. Drawing on the underlying resource document, an outline specification is provided for a form of digital certificate that could carry the signification of consent.

The final segment of the document provides preliminary assessments of the manner in which these models and the associated security requirements might be able to be implemented using several mainstream and emergent protocols.

2. SCHEMAS FOR E-CONSENT IMPLEMENTATION

This section builds on and extends ideas established in the second paper in the series. The first sub-section provides a simplified but sufficiently rich model of the entities involved in the health care sector, and whose participation in an e-consent process is essential. The second sub-section provides a framework within which the channels and facilities used by those entities can be identified and their essential and desirable characteristics evaluated.

2.1 Indicative Model of Entities and Roles

A wide variety of individuals and organisations are involved in the health care sector. A number of models have been devised, which seek to reflect the entities and relationships that exist in particular sectors. In particular, the Australian Institute of Health and Welfare (http://www.aihw.gov.au/) is establishing the National Health Information Model (NHIM), which is a framework for the National Health Data Dictionary (NHDD).

Available models tend to be very complex, and not workable for purposes such as those being pursued by the e-Consent project. This sub-section proposes a simple model of the relevant entities that is nonetheless intended to be rich enough to enable analysis, design and development of meaningful systems.

The model comprises the following assertions:

- a Patient conducts transactions with, has ongoing associations with, and has data about themselves gathered and stored by, a Health Carer;
- a Health Carer works at a Location and within the context of a Health Care Provider organisation, which also employs other Health Carers and Health Care Administrators;
- a Health Care Provider may pass samples and data to, and receive them from, other Health Care Providers;
- samples and data may also be passed to and received from Diagnostic Service Providers, which employ Diagnostic Professionals and Diagnostic Service Administrators;
- the Patient may or may not conduct transactions with, and may or may not have ongoing associations with, these other Providers;
- transmission of samples and data may be facilitated by third party Transmission Services Providers (including the post, couriers and electronic channels);
- data may also be passed to and received from other organisations and individuals.

The following diagram reflects this simple model of the social context in which e-Consent needs to be managed. It is a generalisation of a Transaction Model analysis provided in the second document in this series.

Exhibit 1: Indicative Model of Entities and Roles

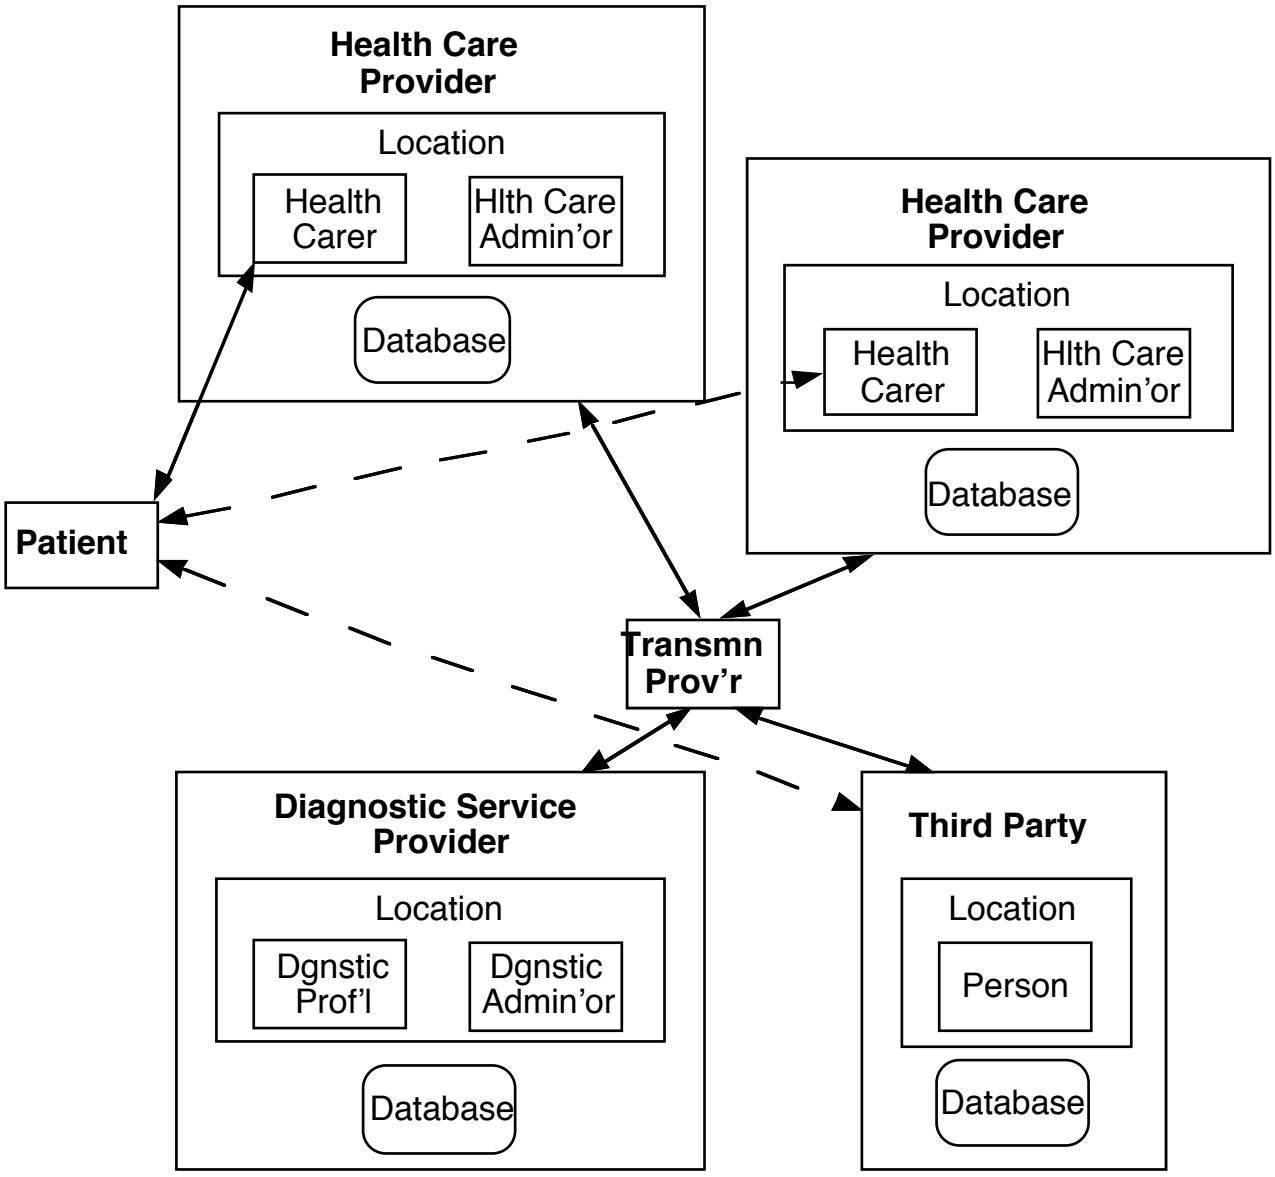

An e-Consent needs to be gathered from the Patient by an appropriate person or organisation (commonly the Health Carer), and communicated to such other parties as is relevant and appropriate. The process is subject to a variety of threats, which need to be secured against.

2.2 Indicative Model of Facilities

The entities that are involved in health care communicate data through various channels, including the post, couriers, telephone, fax and email. They make use of facilities, including computers, and local and wide-area networks. In order to conduct an analysis of security needs in relation to the transmission of e-Consent, a model of these facilities is required. The model provided below focuses on electronic data communications channels, with particular reference to email and email attachments.

The model comprises the following assertions:

- a Patient presents to a Health Carer;
- the Patient may have a card that has some or all of the following capabilities:
  - to identify the card-holder;
  - to authenticate the identity of the card-holder;
  - to securely affix an electronic signature to a message;
  - to securely store a private digital signature key;
  - to securely affix a digital signature to a message;
- a Health Carer may have available to them at the relevant Location a card-reader compatible with the Patient's card, and capable of negotiating with the card a message digitally or otherwise electronically signed by the Patient's card;
- the workstation to which the Health Carer's card-reader is connected may be connected by means of a Local Area Network to a server;
- the workstation or the server may be connected to a wide-area network;
- an organisation with which the Health Carer wishes to communicate may have access to a workstation;
- that workstation may be connected, directly or via a Local Area Network and server, to a wide-area network that is the same as, or connected with, the wide-area network to which the Health Carer's network is connected.

The collection of e-Consent needs to use the available facilities. The infrastructure and the process are subject to a range of threats, which need to be secured against.

Exhibit 2: Indicative Model of Facilities

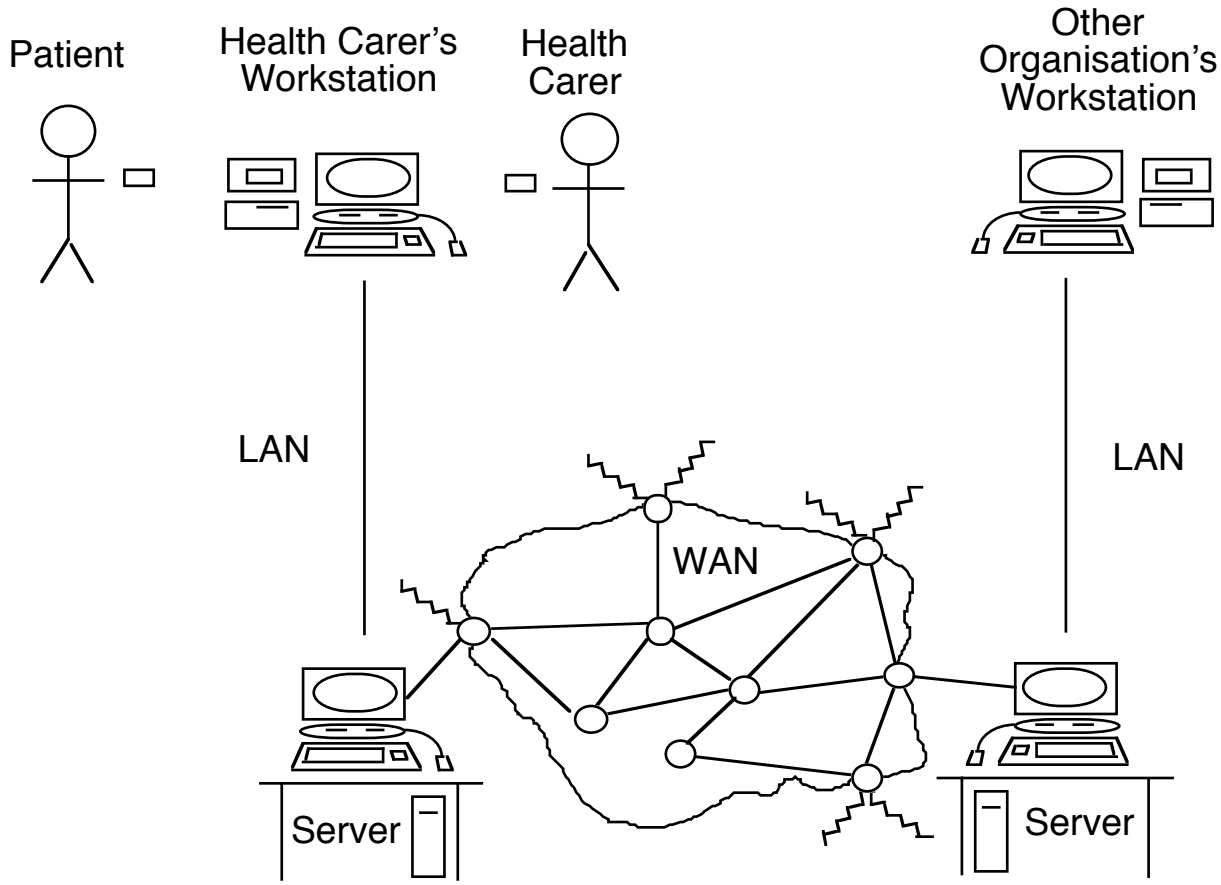

2.3 Facilities for Interactions Between Patient and Health Carers

Many of the facilities and linkages in the indicative model of facilities (Exhibit 2, in the immediately preceding sub-section) are mainstream infrastructure, within the offices of both Health Carers and other organisations that handle health-related data.

Of particular concern, however, are the facilities that support communications between the Patient and the Health-Carer. This sub-section addresses that topic.

(1) Doing Without a Patient-Carried Facility

C could be achieved without any reliance on the Patient having any form of facility, because the necessary functions could be performed using the facilities of the Health Carer. However, this would require Patients to have implicit, uncontrolled trust in each of their Health Carers.

Alternatively, some means could be created to provide a basis for that trust. For example, a Trusted Third Party could operate a register of Patient-ids and authenticators (e.g. passwords, PINs or biometrics). The Health Carer could have an audited application installed on their facility, that enabled the Patient to identify and authenticate themselves by means of the Trusted Third Party, and electronically sign messages arising from actions taken by the Patient using the Health Carer's facility.

(2) Currently Available Infrastructure

An element of infrastructure that may be able to be used is conventional cards, carrying printed data, embossed data, bar-codes and/or data encoded into magnetic stripes. Of particular relevance is the existing Medicare card. It is common for Patients to present these when they deal with Health Carers.

Terminals designed to read the contents of such cards are mainstream devices at points of sale and service, and are generally referred to as EFT/POS terminals. A significant, and possibly still increasing, proportion of locations used by Health Carers have an EFT/POS terminal installed.

Difficulties arise, however, in relation to the reliability of the card for the purposes of e-Consent:

- presentation of the Medicare card is not obligatory, and although greater use can be encouraged (as is likely to occur as a result of the recent changes to the Pharmaceutical Benefits scheme), it is unlikely that it could ever be made obligatory;
- the data currently carried on the Medicare card is limited;
- pseudonymity is not directly supported; and
- Medicare cards are not currently issued to, or in respect of, all Patients. In particular, the card was originally issued on the basis of one card per household.

Moreover, challenges exist in relation to the terminals and networks:

- EFT/POS terminals are commonly installed in an administrative office rather than beside each Health Carer;
- at this stage at least, EFT/POS terminals are mostly not integrated with the Health Carer's workstations or server, but rather are separately connected into a dedicated wide-area network;
- each terminal may or may not have the capacity available to allow an additional Security Access Module (SAM) to be installed to handle the Medicare card; and
- the switches in the dedicated wide-area network may or may not have the ability, and may or may not have the capacity, and their owners may or may not be prepared to provide the service, to switch transactions from Health Carers to other health sector organisations.

(3) Possible Future Infrastructure

It is unrealistic to anticipate that Patients will have a conventional workstation of their own available to them when they are communicating with a Health Carer. On the other hand, micro-processors are increasingly being carried by Patients, in such forms as:

- portable computers;
- hand-held computers;
- mobile phones;
- credit-card-sized chip-cards;
- personal digital assistants (PDAs);
- portable games playstations;
- portable (Walkman-style) music playstations;
- digital cameras; and
- mobile devices arising from convergence of various of the above facilities.

Likely scenarios that would give rise to facilities that could accommodate the needs of e-Consent include:

- an enhanced Medicare card that includes a chip;
- chip-cards issued by other organisations (e.g. telcos, banks, airlines, bus services, local governments, universities and schools, and other governments agencies such as the Australian Taxation Office) which have zones available for hire, and which have features such that a reliable and secure e-Consent service could be established using that infrastructure;
- widespread installation of cost-effective general-purpose card-receiving devices, that can interact securely with a wide range of cards and with a wide range of remote services.

3. INFORMATION SECURITY GENERALLY

The term 'security' is used to refer to both a condition in which harm does not arise, despite the occurrence of threatening events, and a set of safeguards designed to achieve that condition. Threatening events are variously natural, accidental and intentional, and harm can be to persons, property, value, or reputation.

Information security encompasses the whole of an information system, including organisational and individual behaviour, and manual elements of the overall system, as well as computing and communications aspects.

For an information system to be secure, it must have the following properties:

- service integrity;
- data integrity;
- data secrecy;
- authentication, variously of data, identity and attributes; and
- non-repudiation.

Moreover, security is important throughout the information life-cycle, i.e. during the collection, storage, processing, use and disclosure phases, as well as transmission.

A comprehensive security strategy comprises a suite of inter-related safeguards structured in a hierarchical fashion, and dealing with infrastructure, threat management, vulnerability management, and application-specific security. Technical safeguards depend on tools and tool-kits. In telecommunications-based systems, commonality is important, and hence the tools need to be compliant with relevant security standards and protocols.

The process whereby information security is assured comprises a series of phases:

(1) Scope Definition establishes the foundation and framework;

(2) Threat Assessment performs a stocktake of data and processes, and identifies the nature, sources, and situations of threats;

(3) Vulnerability Assessment identifies susceptibilities to those threats;

(4) Risk Assessment evaluates the likelihood and the impact of threatening events;

(5) Risk Management Strategy and Security Plan devices a set of measures that balance the uncertain threats against the plannable costs, and articulates a plan;

(6) Security Plan Implementation ensures the plan is put into effect; and

(7) Security Audit reviews outcomes and identifies adaptations required.

The remainder of the analysis in this document reflects the perspectives on information security that are encapsulated in 'Introduction to Information Security', at:
http://www.anu.edu.au/people/Roger.Clarke/EC/IntroSecy.html
together with the selected works listed in the associated Bibliography, at:
http://www.anu.edu.au/people/Roger.Clarke/EC/IntroSecyBibl.html

4. INFORMATION SECURITY APPLIED TO E-CONSENT

The purpose of this section is to apply the general notions of security outlined in the preceding section to the implementation of e-Consent. It builds on the models developed progressively in the first and second documents of the series and the earlier sections of this third document, and draws directly on a resource document prepared to support this project.

4.1 Information Security Requirements

Information security concerns that arise in the context of e-Consent relate to:

- **personal health data** generally, in particular of Patients, but also of people involved in the provision of services;
- the **consents and denials** given by Patients; and
- the **capacity to create a consent**.

The data, consents and capacity to give consent need to be protected across their entire **life-cycles**. This encompasses creation, storage, processing, use, transmission, disclosure, revocation, archival and destruction.

These data and processes are subject to **natural, accidental and intentional threats**. There are likely to be occasional intentionally threatening events with the potential to create significant harm; but the majority of events that give rise to harm are likely to be of an accidental nature.

Threatening events and vulnerabilities arise in **participating organisations' computing and communications facilities**. These comprise:

- workstations, including hardware, systems software, application software and databases;
- local area networks;
- servers, including hardware, systems software, application software and databases;
- discrete data-storage devices and volumes (e.g. diskettes and CD-ROMs) that are inserted into workstations and servers.

As **Patients** come to use chip-cards or other appliances, their hardware, systems software, application software and databases will embody threatening events and vulnerabilities.

Each participating organisation needs to recognise the scope for threatening events to impinge on vulnerabilities as a result of **incidents in other facilities**, including:

- wide-area networks to which the organisation's facilities are connected; and
- the workstations, local area networks, servers, discrete data-storage devices and volumes, and wide-area network connections of all other organisations connected to common networks, including those with which they communicate, and those they do not, both within and beyond the health care sector.

Threats and vulnerabilities also arise in the context of the **premises** in which facilities are housed, both those of participating organisations, and of other organisations that are electronically connected and may be subject to penetration.

The **supporting infrastructure** is also a source of threatening events and a repository of vulnerabilities, particularly the electricity supply, but also such services as air-conditioning and fire protection.

The **organisational elements** of an information system that harbour threatening events and vulnerabilities include:

- policies;
- manual procedures;
- actual practices; and
- data stored in other than machine-readable form.

The **harm** that arises will most commonly be loss of data, access to and disclosure of data, and inappropriate replication of data.

The first-order impact will be harm to the personal values of affected individuals, generally Patients. The second-order effect will be a loss of reputation by all organisations and health care professionals involved, and of public confidence in them, because of a public perception of a breach of trust by the data-holder.

There will be circumstances in which harm to persons will arise. Examples include:

- incorrect or incomplete data, or missing or inaccessible data, resulting in mis-diagnosis, incorrect treatment, or a failure to address a treatable condition; and
- disclosure of a person's whereabouts, resulting in assault and psychological trauma.

Organisations and professionals will be subject to legal **sanctions** in some of these circumstances. Professional negligence in relation to diagnosis and treatment is one source. Other potential sources include privacy legislation and statutorily enforceable codes of conduct, contracts, and negligence in operational and administrative contexts.

In order to deal with the threats, information systems that handle e-consent in the health care sector need to fulfil the following **general requirements**:

- **service integrity**, to ensure availability, reliability, completeness and promptness;
- **data integrity**, to ensure that data is authentic, reliable, complete, unaltered and useable, and that the processes that operate on data are reliable, legally compliant, comprehensive, systematic, and protective of data;
- **data secrecy**, to ensure that it is only available to those who should have it;
- **authentication**, to ensure that appropriate checks are performed on data, identity and attributes; and
- **non-repudiation**, to ensure that entities cannot convincingly deny important actions they have taken.

In addition, the e-Consent system needs to satisfy the following **more specific requirements**:

- **a facility** is required, to enable a Patient, Health Carer or other participant in the system to:
  - authenticate the identifier and/or attributes of other facilities;
  - provide their identifier and/or attributes in such a manner that they can be authenticated by other facilities;
- **data** (including personal health data, agent identifiers and attributes, consents and denials) needs to be subject to the following:
  - its integrity needs to be assured;
  - its despatch and receipt need to be non-repudiable;
  - it needs to be able to be authenticated;
  - it needs to be encrypted for secret transmission between facilities; and
- **the act of creating a consent** needs to satisfy the conditions expressed in the first paper in the series, viz.
  - be provided by a person with the mental capacity to do so;
  - be provided by a person with the legal capacity to do so in respect of that data;
  - be express, or reasonably implied, or (with great care) reasonably inferred, but in such a manner that an opportunity exists for the implied or inferred consent to be denied;
  - be informed;
  - be freely-given;
  - be specific and bounded;
  - be signified in a manner that provides evidence;
  - be non-repudiable (preventing the denial that the consent was given); and
  - be variable and revocable, by means of a non-repudiable transaction.

4.2 Information Security Strategy

A comprehensive strategy is needed to ensure that the requirements outlined in the previous section are satisfied. Such a strategy needs to comprise:

- an **architecture** that provides a framework within which a cohesive set of safeguards can be devised and implemented;
- a **process**, to ensure that the strategy is articulated into a plan, and the plan implemented;
- **tools and tool-kits** to enable safeguards to be developed, which implement appropriate security standards and protocols;
- **resources**, to enable the plan to be implemented.

Particularly important aspects of the strategy creation process are as follows:

(1) **Scope Definition**. It will be vital to identify stakeholders, and engage credible proxies to consider the initiative on their behalf. Only in this way can their concerns be appreciated and reflected in the design, and their acceptance of and support for the scheme be ensured;

(2) **Threat Assessment**. The preliminary analysis in the preceding sub-section needs to be deepened and further developed in order to reflect the design choices made;

(3) **Vulnerability Assessment**. The preliminary analysis in the preceding sub-section needs to be deepened and further developed in order to reflect the design choices made;

(4) **Risk Assessment**. The implications of the preliminary analysis in the preceding sub-section need to be further examined in the light of the design choices made, and the prevalence of threatening events and their impacts evaluated;

(5) **Risk Management Strategy and Security Plan**. Because public confidence in the health care sector is so critical, strong emphasis is needed on proactive strategies rather than depending on the ability to react to threats that eventuate and the harm that arises. Technical safeguards that appear likely to be critical include access controls, encrypted transmissions, authentication of devices, private-key management, validity checking of digital certificates, auditing to ensure the security of both packaged and custom-built software, transaction logging, analysis of logs, and automated reporting of exceptions.

A detailed plan will be needed. It will be important to implement many features centrally, in order to limit the imposition on, and reliance on, widely dispersed health care organisations, professionals and staff;

(6) **Security Plan Implementation**. The undertaking will demand tight project management, because it will inevitably involve external contractors, who will be required to demonstrate precise compliance with standards, protocols and tightly-specified requirements. Implementation of technical safeguards will need to be complemented by measures to achieve cultural change, and adaptation to the terms of employment and of contract in order to ensure that staff are bound to comply with security requirements;

(7) **Security Audit**. Review and adaptation will be critical to reliability and credibility.

The health care sector comprises vast numbers of legally independent (but functionally inter-dependent) legal entities. There are therefore far more substantial challenges confronting the development and implementation of the security strategy than is the case with a single organisation.

The various security protocols and tools need to be mapped against the entities involved in the health care sector, in order to ensure that the responsibilities can be sheeted home. A provisional mapping is provided in Appendix 1: Security Schema for Message Management.

4.3 An e-Consent Digital Object

An implementation of e-Consent will require secure and reliable communication of the fact of the consent, and its terms, among relevant human and software agents.

It is envisaged that a specially-designed digital object will be used as the basis for that communication. The object would be delivered with the health data to which it relates, in order to provide the recipient with evidence that the Patient has given consent for the provision of the information, and to instruct the recipient as to what they can and cannot do with it.

In some cases, the consent would specify all of:

- the **information**, which might be related to a specific event, encounter, episode, condition, procedure, or medication;
- the identity of the **sending entity**;
- the identity of the **receiving entity**;
- the **purpose** for which the data is to be used; and
- any authority for, and conditions relating to, the **further disclosure** of the data by the recipient.

In many cases, however, the consent would have broader coverage, and would refer to:

- **a general class of information**, e.g. all information, or all information related to a specific event, encounter, episode, condition, procedure, or medication;
- **a general class of sending entity**, e.g. an identified Health Care Professional, or a particular category of Health Care Professional, or Health Care Professionals active in a particular location or practice; and/or
- **a general class of receiving entity**, as for the class of sending entity.

The object needs to support **both consents and denials**. The reason that denials need to be encompassed is that some consent models establish statutory presumptions of consent that are overridable by an explicit denial. (Some European countries do so in relation to organ donation, for example). If such a mechanism were in operation, the e-Consent object would have to be able to record a denial of the presumption of consent.

Moreover, a consent might apply to some broadly-defined class (of information, sending entity, or receiving entity), but the Patient may want to add **qualifications**, in order to deny information about, say, a particular condition or episode. Conversely, the Patient may wish to deny the general presumption of consent, but specifically consent to particular exceptions (e.g. deny the general presumption of consent for organ transplant, but permit use for a specific person or for any family-member).

Hence the e-Consent object needs to support:

- a broadly-expressed consent;
    - qualified by one or more specific denials; and
- a broadly-expressed denial;
    - qualified by one or more specific consents.

Finally, the object needs to support **a hierarchy of qualifications**, e.,g, the Patient:

- consents to all information to all Health Carers; but within that
  - denies all information relating to HIV; but
    - consents to information relating to HIV to STD clinics; and finally
      - denies all information to a specific STD clinic (where Mum works).

There may be multiple qualifications at each level.

In some circumstances, **supplementary data** may need to be supported. In particular:

- in the case of a minor, data about the person *in loco parentis* who provided the consent on the Patient's behalf;
- in the case of a person incapable of giving consent themselves (e.g. because they are comatose, missing or deceased), information about the circumstances and about the person who provided the consent; and
- in the case of a consent given by a formally authorised person on the Patient's behalf, details of the power of attorney or other authority.

Appendix 2 provides a draft of the **data-items** that would need to be contained in an e-Consent object, together with an assessment of the challenges they present.

Possible approaches to the **means of creation** of the e-Consent object include:

- the Patient could carry a card that created and electronically signed an e-Consent object and passed it to the relevant Health Carer;
- the Patient's card could electronically sign an e-Consent object that had been created by a Health Carer's own facility;
- in the absence of a card, a less secure but probably workable arrangement would be for the Patient to indicate agreement to a prepared consent form displayed or printed using the Health Carer's facility, whether by keying a password or PIN, or by signing a printed version. The Health Carer's facility would then prepare the e-Consent object for despatch; or
- a Health Carer might prepare an e-Consent object on the basis of information already held on file.

An e-Consent object could take a variety of forms. One possibility is to use digital certificate standards and associated technologies, supplemented by a Trusted Third Party that administers a set of healthcare consumer ids and associated passwords/PINs.

A possible implementation would involve the following sequence of events:

- the Health Carer uses their computing facility to create the e-Consent object (e-Co);
- the Patient enters their password/PIN, and, if it matches their healthcare consumer id, the transaction is enabled;
- the Health Carer signs the e-Co with their private signing key;
- the Health Carer authenticates the intended recipient's identifier by some challenge-response protocol, including the establishment of a session ID and a secure session key;

- the Health Carer encrypts the signed e-Co with the session key;
- the Health Carer transmits the encrypted, signed message to the recipient.

For revocation purposes, and to resolve disputes, an e-Consent object would need to contain a date-time stamp recording when it was created and/or communicated, and this would need to be reliable and subject to authentication procedures.

The information that is the subject of the e-Co could be included in the same message, or transmitted separately, also subject to channel protection and recipient authentication.

The e-Co meets the e-Consent security objectives for this scenario, as follows:

- the Health Carer can authenticate the Patient's identifier by checking that their PIN matches their consumer identifier (entity authentication by something known);
- the Patient's consent is non-repudiable, because the Health Carer has explained to them that by entering their PIN they agree to the transfer as stated in the statement of consent;
- the recipient can authenticate the origin of the e-Co and can be assured of its integrity (data origin authentication and hence also data integrity by digital signatures);
- the e-Co is non-repudiable by the Health Carer (non-repudiability by digital signatures);
- the Health Carer authenticates the recipient's identifier as the recipient of the data (entity authentication by something known);
- the recipient is the only entity able to read the e-Co. This is assured by symmetric-key encryption, with the key established by public-key methods (secrecy by encryption);
- the session ID prevents recording and replay of the session.

The consent is not cryptographically non-repudiable, however, in that it depends on the Patient understanding both the contents of the e-Consent object and the significance of entering their password/PIN.

Coordinated care is facilitated by the process, because the Patient's record accumulates, with their consents associated with them, in locations where the Patient interacts with the healthcare system.

5. IMPLEMENTABILITY OF AN E-CONSENT OBJECT

A great deal of the information security strategy needed in order to support e-Consent applies conventional analyses and techniques. It is important that organisations that handle personal health data, including evidence of consents, design and implement organisational arrangements and processes that protect that data, and deploy technical facilities (including workstations, servers, networks and application software) that embody appropriate technical features that provide comprehensive security protections. In turn, providers of those facilities need to be aware of the needs, and design their offerings to match them.

The particular aspect of the requirements that represents the greatest novelty is the e-Consent object. This section examines a number of available technologies in order to establish whether they are likely to provide a basis for the implementation of an e-Consent object, as described in section 4.3 above and Appendix 2 below.

5.1 X.509v3 Certificates

Appendix 3A provides informal background to the X.509v3 certificate. Appendix 3B provides more formal information, together with the key references.

On the basis of that information, and notwithstanding the fact that v3 was designed to extend the scope of v2, it is still not immediately obvious that the contents of an attribute certificate have the requisite **extensibility**, such that an X.509v3 certificate can carry the additional data that is required by the e-Consent certificate described in section 4.3 above and Appendix 2 below.

Products that apply the X.509v3 certificate format are available. Some products may be restricted in the certificate formats that they support. It is therefore possible that, even if the X.509v3 standard theoretically supports the e-Consent certificate, the **implementation** within some products may be such that the additional data cannot be carried.

In addition, X.509v3 (although originally intended as both an infrastructure and a certificate format) is now only used as a certificate standard, and it therefore depends on **supporting infrastructure**. Although the certificate format specification has been stable since 1996, the supporting infrastructure is less well standardised. The key developments in this area are RFC2459 (1999), RFC2510 (1999), RFC2511 (1999) and RFC2527 (1999). These documents, although they have reached proposed standard stage, are still subject to ongoing amendments (see PKIX 1993-).

A further consideration relates to the availability of an effective and efficient **revocation mechanism**, on which all certificate-based systems are utterly dependent. The adequacy of the batch mechanism called a Certificate Revocation List (CRL), and specified in X.509v3 (1996) and RFC2459, has been called into serious question. The on-line equivalent called Online Certificate Status Protocol (OCSP), and specified in RFC2560 (1999), is still at Proposed Standard level. Moreover, because it was released in June

1999, and longstanding IETF policy and practice has been that Proposed Standards lapse after 12 months, RFC2560 would have to be re-instigated as a Proposed Standard before it could proceed further towards adoption.

Finally, X.509v3 certificates appear to be less than satisfactory from a **privacy** perspective. They include a data-item referred to as the 'name' of the 'subject'. Because this data-item is mandatory, the X.509v3 standard supports identified transactions, but appears to preclude anonymity. It might, however, support pseudonymity, depending on the policies, procedures and practices within Registration Authorities and Certificate Authorities, and legal protections for the records of those organisations.

The considerable range of problems with X.509-based public key infrastructure is documented in a variety of papers, including Winn (1998), Ellison & Schneier (2000), Gutmann (2000), Sneddon (2000), Clarke (2001) and Winn (2001).

5.2 S/MIME

S/MIME is a family of protocols for the transmission of messages with encryption and/or digital signatures. Earlier versions did not proceed to adoption as a formal Internet standard (partly due to patent issues), but version 3 may now be moving towards adoption (FRC2632 1999, RFC2633 1999, IMC 1999, S/MIME 2001).

S/MIME builds on a set of prior standards, including:

- MIME-encapsulation for email attachments (RFC1847 1995);
- X.509v3 certificate formats (X.509 1997); and
- PKIX (X.509v3-based PKI for the Internet – W3C 2000).

It is unclear how readily implementations of S/MIME are available, and how likely they are to become mainstream. If they do, and if the inherent deficiencies of hierarchical PKI are overcome, then S/MIME might provide a considerable amount of the necessary infrastructure for e-Consent. If the e-Consent object were expressed as an X.509v3 certificate, then the problems identified in the previous sub-section would still need to be overcome.

5.3 PGP

The common approach to public key infrastructure depends on hierarchies of Certification Authorities (CAs). An alternative approach is commonly referred to as the 'web of trust'. Under this arrangement, there is little need for commercial CAs (especially if they provide only a limited warranty), because digital certificates can be issued by anyone.

In a web of trust scheme, fault-tolerance is achieved by the participant deciding for themselves how many or how few certificates to rely on. Each certificate can carry a weighting that reflects the degree of trust that the certificate-issuer places in the key-owner and their ability to manage their private key. The approach requires message-

recipients to consider the extent to which they really need assurance, and confront the simple fact that all assurance is relative rather than absolute. There is a reasonable level of both intellectual and commercial support for the 'web of trust' idea (e.g. Maurer 1996, Grossman 2000).

Pretty Good Privacy (PGP) is the earliest and most well-known scheme that implements the 'web of trust' concept (Zimmerman 1995, Garfinkel 1995, Bacard 1995, Stallings 1995. See also PGP/MIME – RFC2015 1996, IMC 1999). The original public domain form is now referred to as PGP 5.x or OpenPGP (RFC 2440 1998, OpenPGP 2001). A proprietary version also exists, marketed and undergoing further development by PGP Inc. Each of the versions carries the signed and encrypted message in a MIME-encapsulated block.

OpenPGP and commercial PGP use an own form of certificates or 'pgp keys'. These contain fewer data-items than X.509 certificates, but can be signed by multiple entities, not just one as is the case with X.509. At some stage, a later version of PGP might support X.509 certificates as an alternative to pgp keys.

The practicality of PGP's specific implementation of the 'web of trust' notion has been criticised. On the other hand, it has been implemented and tested in the field, e.g. by Qualcomm in its popular Eudora email-client. Moreover, there are at least some commercial applications of it. For example, a recently-developed Australian superannuation industry scheme decided on PGP because it was a cheap, straightforward and effective alternative to a CA-based scheme. In addition, the concept has been re-birthed with minor variations, in a scheme sponsored by X.509 providers and referred to as 'mesh architecture'.

PGP could be readily applied in the healthcare context, because many Health Carers would be likely to be confident in certificates issued by other Health Carers. If this was done, then an e-Consent object could be expressed as a message signed by a private key that is attested to by a PGP web of trust. To be practicable, however, email-packages would need to be readily available which make the use of PGP simple and natural for all users involved.

5.4 SDSI / SPKI

The hierarchical and ultimately authoritarian feel of X.509 has drawn criticism and stimulated proposals for alternative schemes that are argued to be more practical. One was Simple Public Key Infrastructure (SPKI) (Ellison 1996, IETF 1997-, Wang 1998, Ellison 2000). The momentum has now shifted to a parallel initiative, the Simple Distributed Security Infrastructure (SDSI) (SDSI 1996-; see also Rivest & Lampson 1996, Ellison 2000). SDSI v2 incorporates the SPKI ideas.

The key element of SDSI is that the X.509 nirvana of a single, global name-space has been abandoned. With it, the presumption has been removed that 'name' (or, better expressed, 'identifier') is reliably bound to a particular entity. The certificate associates a public key (and hence a key-pair) to an entity that only the CA knows, and no warranties are provided by the CA to the recipient of the message as to who the keyholder is. It is up to the relying party to build up an image of the sender based on its successive interactions with the holder of that key.

Attributes are associated with public keys, not with identities of real-world entities. Hence, for example, a recipient can be assured that a particular message was provided by a medical practitioner, or a person over 18, or over 65; but the certificate is silent about the identity of the person who is using the key (Ellison 2000). SPKI and hence SDSI certificates are described in RFC2692 (1999) and RFC2693 (1999).

Although laboratory experimentation has been conducted, at least at MIT, it is not clear whether SDSI has yet given rise to any commercial implementations. If it does, a SDSI certificate could be an appropriate means of implementing an e-Consent object.

Appendix 4 provides further information relating to SDSI/SPKI, which was gathered during this project. This arose from discussions during a visit to the U.S. with the originators of the two schemes, who are based at MIT in Boston, and at Intel in Portland, Oregon.

5.5 AADS

A further alternative approach that can be taken to authentication is to rely on pre-existing relationships within a community. The Account Authority Digital Signature Model (AADS) is a cut-down variant of conventional digital signature processes, applicable to communications among parties that already have 'accounts' with one another, and who have already received and stored one another's public keys (Wheeler 1998). This obviates the need for public keys and key certificates to be transmitted with each message, and thereby avoids some of the problems inherent in X.509.

However, many of the challenges remain, such as how keys are securely generated, stored, backed-up and recovered, and revoked. Furthermore, the system requires that the recipient of a message already have the public key of the sender and the Patient.

If commercial applications of AADS emerge, they could provide some elements of the infrastructure needed to implement e-Consent objects.

5.6 Brandsian Certificates

Brands (2000) proposes a different conception and implementation of digital certificates, such that privacy is protected without sacrificing security. The validity of such certificates and their contents can be checked, but the identity of the certificate-holder cannot be extracted, and different actions by the same person cannot be linked. Certificate holders have control over what information is disclosed, and to whom. Stefan Brands' certificates are expressly anonymous.

It would thus appear that Brands' certificates could provide a mechanism to achieve the objectives of the e-Consent project. Brands is now Senior Cryptographer with Zero-Knowledge Systems, Inc., in Montreal, who are understood to be building applications based on his techniques. Health sector applications are one of their major interests, and they have been and are talking with various key players in the industry, but there is no

large-scale implementation yet. Unfortunately there is no public information available, but Brands has indicated that they are working on a white paper.

Appendix 5 provides further information relating to Brandsian Certificates, which was gathered during this project. This arose from discussions with the originator, who is based Montreal, while on a visit to Washington DC.

5.7 Trust Management Systems

By 'trust-management systems' is meant a generalisation of longstanding access control techniques for achieving security of software processes and data (Blaze et al. 1999). Blaze contends that trust management has five basic components:

- a language for describing **`actions**', which are operations with security consequences that are to be controlled by the system;
- a mechanism for identifying **`principals'**, which are entities that can be authorised to perform actions;
- a language for specifying application **`policies'**, which govern the actions that principals are authorised to perform;
- a language for specifying **`credentials'**, which allow principals to delegate authorisation to other principals; and
- a **`compliance checker'**, which provides a service to applications for determining how an action requested by principals should be handled, given a policy and a set of credentials.

The trust management approach focusses primarily on privileges and restrictions, and is much less concerned about identified individuals; and it can deal with a nym representing a pseudonymous role just as readily as with a name that is associated with an identified human.

Blaze and his team at AT&T have developed a specific prototype trust management system called Keynote (RFC 2704 1999). It is not clear whether and when commercial implementations will become available. If they do, they may represent a highly appropriate framework for the entire infrastructure to support e-Consent objects.

5.8 P3P

The Platform for Privacy Preferences (P3P) protocol was developed by the World Wide Web Consortium (W3C 2000). Its purpose is to enable:

- web-sites to specify their personal data use and disclosure **practices**;
- web-users to specify their **expectations** concerning personal data disclosure practices; and

- software agents to undertake **negotiation**, on behalf of the parties, in order to reach an agreement concerning the exchange of data between them.

P3P provides means whereby an individual can acquire information from a web-server about its privacy policies, sufficient to make an informed decision on whether or not to accede to the site's conditions. Moreover, it is intended that the decision can be delegated to a software agent acting on behalf of the individual. An overview is provided in Clarke (1998a) and a critique in Clarke (1998b).

P3P enables web-site owners to express their privacy policies (manually, or using purpose-built tools), and have them converted into XML format. The intention is that browsers download the XML statement, interpret it, and compare it with the user's declared preferences.

The P3P standard is close to being finalised, generators have been written to enable the expression of web-site privacy policies in XML, and some server-side tools are available. The capability is not yet embedded in main browsers, nor do any plug-ins appear to be available yet.

One weak point from the perspective of the e-Consent project is that the means whereby the user expresses privacy preferences is not part of the P3P standard. A related standard, A P3P Preference Exchange Language (APPEL), has been drafted to enable preferences to be expressed, "such that a user agent could make automated or semi-automated decisions regarding the acceptability of machine-readable privacy policies from P3P enabled Web sites" (W3C 2000). It is envisaged that APPEL would be used primarily to allow users to import preference rulesets created by other parties and to transport their own rulesets files between multiple user agents.

It does not appear that P3P itself has a great deal to offer the e-Consent project. On the other hand, APPEL could provide a basis for the formatting of the e-Consent object.

6. CONCLUSIONS

The e-Consent project seeks to identify practicable models whereby patient consent can be expressed and communicated in electronic contexts. This series of documents has presented conceptual models that are intended to be rich enough to reflect the diverse array of circumstances that arise in coordinated care, but simple enough to be practical.

Candidate technologies and protocols have been subjected to preliminary evaluation. These three documents provides the basis from which prototypes can be sketched, commissioned and demonstrated, and specifications prepared for products that will meet the needs of the health care sector and the patients it serves.

REFERENCES

Most of the relevant works are listed in the Bibliography at:
http://www.anu.edu.au/people/Roger.Clarke/EC/IntroSecyBibl.html

This list provides additional references.

AIHW 'National Health Information Model', v.2, at http://www.aihw.gov.au/pls/nhik/nhik_info_models.display_model?pModel=NHIM

AIHW 'National Health Data Dictionary', at http://www.aihw.gov.au/pls/nhik/nhik_data_elements.data_element_search?pLinkType=DD&pLink_ID=NHDD

AS 4400 (1995) 'Personal privacy protection in health care information systems' Standards Australia, 1995

Clarke R. (2002) 'e-Consent: A Critical Element of Trust in e-Business' Proc. 15th Bled Electronic Commerce Conf., Bled, Slovenia, 17-19 June 2002, at http://www.anu.edu.au/people/Roger.Clarke/EC/eConsent.html

Gatekeeper (1998) 'Gatekeeper: A strategy for public key technology use in the Government', Office of Government Information Technology, May 1998, at http://www.ogo.gov.au/projects/publickey/Gatekeeper.htm

NHERT (2000) 'A Health Information Network for Australia', National Electronic Health Records Taskforce, July 2000

NHMRC (1993) 'Guidelines for the protection of privacy in the conduct of medical research' National Health & Medical Research Council, 1993

NSW Health (1996) 'Information Privacy Code of Practice' Privacy of Information Committee, NSW Health, May 1996, ISBN 0 7310 0773 5, HP No. (IDU) 96-25, Circular No. CPR 96/34

App. 1: Security Schema for Message Management

<INSERT SECURITY SCHEMA FOR MESSAGE MANAGEMENT>

<John Payne's file Sec Agent Roles JP 010115 .xls>

App. 2: Data-Elements Needed in an e-Consent Object

The following data-items are implied by the description of the e-Consent object:

- a Patient identifier, e.g. a healthcare consumer id; but possibly a pseudonym
- a set of consent/denial statements by the Patient, expressed in a hierarchy.

  The first entry may be either a consent or a denial, and each subsequent entry must be the other category, qualifying the preceding entry. The most common instance would be a single consent entry; and the next most common would be a sequence of a broadly-expressed consent, followed by a specific denial.

  Each entry comprises:
  - an indicator as to whether the entry is a consent or a denial
  - a description of the information to which the entry relates
  - the id of the entity to whom the entry was provided
  - the id of the releasing entity, or a description of the class of entities
  - the id of the receiving entity, or a description of the class of entities
  - a description of the purpose for which the data is to be used
  - a statement as to whether, and under what circumstances, a recipient may further disclose the data
- the means whereby the Patient signified consent
- the date and time when the Patient signified consent
- the date and time of expiry of the consent (optional)
- additional information about the signification of consent (optional, but particularly to cater for circumstances where the Health Carer relies on information already held on file)
- additional information about the person giving the consent (generally optional, but mandatory in such cases as children and a comatose patient, and when the consent is given by another party under a power of attorney or other authority)
- the Health Carer's identifier
- the Health Carer's electronic signature

The following needs are implied, each of which represents a significant challenge:

- a Patient identifier, capable of supporting pseudonymity
- an identifier for health care sector entities
- an identifier for classes of health care entities
- an open-ended coding scheme for classes of information to which the consent or denial relates
- an open-ended coding scheme for the purposes for which the data is to be used
- a coding scheme for the means whereby the Patient signified consent
- means of authenticating the contents of the date-time stamp

App. 3A: Outline of the X.509v3 Certificate Structure

The X.509 standard defines the format for a digital certificate, which is intended to evidence a Certification Authority's certification that an identified identity possesses the private key that matches the public key contained in the certificate. The successive versions are v1 (1988), v2 (1993) and v3 (1996).

Basic Contents (v2)

Certification Authority (CA) Information

- Version / Serial No.
- Signature / Token
- CA / Issuer Name

Subject Information

- Distinguished Name

Key Information

- Algorithm
- Public Key
- Validity Dates
- Identifiers

Extensions (v3)

- Authority Key ID
- Primary Key Attributes
- Certificate Policies
- Key Usage Restriction
- Policy Mapping
- Supported Algorithms
- Subject Attributes
- Issuer Attributes
- Basic Constraints
- Name Constraints
- Policy Constraints
- CRL Distribution Point

The X.509v3 revision enables a parent identity certificate to have any number of child attribute certificates. This has the advantage that the person's credentials do not all have to be visible to everyone that needs access to one of them, e.g. a medical practitioner, in communicating their capacity to prescribe medicine, need not disclose their own health conditions, professional and club memberships, etc.

App. 3B: Detail of the X.509v3 Certificate Structure

The basic format of an X.509v3 certificate is defined as follows (X.509 1997, using ASN.1 notation):

```
Certificate ::= SEQUENCE {
    tbsCertificate       TBSCertificate,
     signatureAlgorithm    AlgorithmIdentifier,
    signature            BIT STRING  }
```

The PKIX Working Group's draft provides the following ASN.1 definition of tbsCertificate:

```
 TBSCertificate  ::=  SEQUENCE  {
    version              [0]    EXPLICIT Version DEFAULT v1,
    serialNumber                CertificateSerialNumber,
    signature                   AlgorithmIdentifier,
    issuer                      Name,
    validity                    Validity,
    subject                     Name,
    subjectPublicKeyInfo        SubjectPublicKeyInfo,
    issuerUniqueID    [1]       IMPLICIT UniqueIdentifier OPTIONAL,
                                -- If present, version must be v2 or v3
    subjectUniqueID   [2]       IMPLICIT UniqueIdentifier OPTIONAL,
                                -- If present, version must be v2 or v3
    extensions           [3]    EXPLICIT Extensions OPTIONAL
                                -- If present, version must be v3
    }
```

From http://www.pdos.lcs.mit.edu/asrg/2000-11-13-verisign.txt. here is a sample [expired] X509 certificate signed by Verisign.

```
Certificate:
   Data:
       Version: 1 (0x0)
       Serial Number:
         48:9c:31:e3:e5:8f:9c:65:f6:b0:92:ff:b0:2d:29:30
        Signature Algorithm: md5WithRSAEncryption
       Issuer: C=US, O=RSA Data Security, Inc., OU=Secure Server Certification Authority
       Validity
          Not Before: Oct  6 00:00:00 1999 GMT
          Not After : Oct  5 23:59:59 2000 GMT
       Subject: C=US, ST=Massachusetts, L=Cambridge, O=Foo and Fu, OU=Development,
CN=snafu.fooworld.org
       Subject Public Key Info:
```

```
        Public Key Algorithm: rsaEncryption
        RSA Public Key: (1024 bit)
          Modulus (1024 bit):
            00:ca:dd:ac:8e:24:b8:b9:28:77:7f:ad:b5:3f:8e:
            6e:5d:cc:ec:ff:80:a7:d9:73:c7:54:77:f8:57:3c:
            3b:d4:85:b5:3c:9e:fc:d5:a6:e0:b9:33:48:86:b6:
            95:f8:1e:88:fe:9c:79:e7:77:99:97:dd:a9:d1:d7:
            6d:60:1c:ce:0c:ff:fc:28:75:d9:f7:e0:74:ca:9a:
            7e:bd:9f:7c:d8:6f:62:6f:8d:be:82:2a:dc:7f:97:
            d1:f7:89:14:a1:c3:17:cd:aa:07:ef:dc:92:aa:40:
            7a:19:8e:fb:6c:01:77:46:4b:85:d4:ad:c5:da:02:
            bb:73:54:3f:bf:a4:29:b8:5f
          Exponent: 65537 (0x10001)
    Signature Algorithm: md5WithRSAEncryption
      48:22:d8:4b:9c:65:ac:52:84:94:40:bc:8d:7c:b7:7e:b8:62:
      00:b3:83:35:b8:5c:4a:2e:bc:4c:67:53:3e:9b:a2:86:de:73:
      ec:26:df:07:b8:1d:18:80:4f:32:16:9d:f7:0e:c9:93:68:e8:
      73:96:ea:fe:07:cf:6f:87:0e:61:35:cb:b8:f6:5f:e0:9b:9e:
      7c:82:d9:29:e4:b0:b9:4f:7b:fc:65:f9:11:31:b4:ce:a4:ae:
      05:0e:61:ba:90:28:81:17:68:2c:95:40:82:39:e5:5b:78:06:
      06:dd:8d:d8:6d:ce:df:1a:ba:b5:55:97:7f:ce:c4:c2:ba
-----BEGIN CERTIFICATE-----
MIICWzCCAcgCEEicMePlj5xl9rCS/7AtKTAwDQYJKoZIhvcNAQEEBQAwXzELMAkG
A1UEBhMCVVMxIDAeBgNVBAoTF1JTQSBEYXRhIFNlY3VyaXR5LCBJbmMuMS4wLAYD
VQQLEyVTZWN1cmUgU2VydmVyIENlcnRpZmljYXRpb24gQXV0aG9yaXR5MB4XDTk5
MTAwNjAwMDAwMFoXDTAwMTAwNTIzNTk1OVowgYExCzAJBgNVBAYTAlVTMRYwFAYD
VQQIEw1NYXNzYWNodXNldHRzMRIwEAYDVQQHFAlDYW1icmlkZ2UxEzARBgNVBAoU
CkZvbyBhbmQgRnUxFDASBgNVBAsUC0RldmVsb3BtZW50MRswGQYDVQQDFBJzbmFm
dS5mb293b3JsZC5vcmcwgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBAMrdrI4k
uLkod3+ttT+Obl3M7P+Ap9lzx1R3+Fc8O9SFtTye/NWm4LkzSIa2lfgeiP6ceed3
mZfdqdHXbWAczgz//Ch12ffgdMqafr2ffNhvYm+NvoIq3H+X0feJFKHDF82qB+/c
kqpAehmO+2wBd0ZLhdStxdoCu3NUP7+kKbhfAgMBAAEwDQYJKoZIhvcNAQEEBQAD
fgBIIthLnGWsUoSUQLyNfLd+uGIAs4M1uFxKLrxMZ1M+m6KG3nPsJt8HuB0YgE8y
Fp33DsmTaOhzlur+B89vhw5hNcu49l/gm558gtkp5LC5T3v8ZfkRMbTOpK4FDmG6
kCiBF2gslUCCOeVbeAYG3Y3Ybc7fGrq1VZd/zsTCug==
-----END CERTIFICATE-----
```

App. 3C: X.509v3 References

Clarke R. (2001) 'The Fundamental Inadequacies of Conventional Public Key Infrastructure' Proc. Conf. ECIS'2001, Bled, Slovenia, 27-29 June 2001, at http://www.anu.edu.au/people/Roger.Clarke/II/ECIS2001.html

Ellison C. & Schneier B. (2000a) 'Risks of PKI: Electronic Commerce' Inside Risks 116, Commun. ACM 43, 2 (February 2000), at http://www.counterpane.com/insiderisks5.html

Gutmann P. (2000) 'X.509 Style Guide', at http://www.cs.auckland.ac.nz/~pgut001/pubs/x509guide.txt

IETF (1997-) 'Simple Public Key Infrastructure (SPKI)', at http://www.ietf.org/html.charters/spki-charter.html

PKIX (1993-) 'PKIX Working Group', at http://www.imc.org/ietf-pkix/

RFC2409 (1998) 'The Internet Key Exchange (IKE)' Internet Engineering Task Force of The Internet Society, November 1998, at ftp://ftp.isi.edu/in-notes/rfc2409.txt

RFC2411 (1998) 'IP Security Document Roadmap', Internet Engineering Task Force of The Internet Society, November 1998, at ftp://ftp.isi.edu/in-notes/rfc2411.txt

RFC2459 (1999) 'Internet X.509 Public Key Infrastructure Certificate and CRL Profile' Internet Engineering Task Force of The Internet Society, January 1999, at http://www.ietf.org/rfc/rfc2459.txt

RFC2510 (1999) 'Internet X.509 Public Key Infrastructure: Certificate Management Protocols', Internet Engineering Task Force of The Internet Society, March 1999, at ftp://ftp.isi.edu/in-notes/rfc2510.txt

RFC2511 (1999) 'Internet X.509 Certificate Request Message Format', Internet Engineering Task Force of The Internet Society, March 1999, at ftp://ftp.isi.edu/in-notes/rfc2511.txt

RFC2527 (1999) 'Internet X.509 Public Key Infrastructure: Certificate Policy and Certification Practices Framework', Internet Engineering Task Force of The Internet Society, March 1999, at ftp://ftp.isi.edu/in-notes/rfc2527.txt

RFC2560 (1999) 'X.509 Internet Public Key Infrastructure: Online Certificate Status Protocol - OCSP' Internet Engineering Task Force of The Internet Society, June 1999, at http://www.ietf.org/rfc/rfc2560.txt

RFC2828 (2000) 'Internet Security Glossary' Internet Engineering Task Force of The Internet Society, 2000, at ftp://ftp.isi.edu/in-notes/rfc2828.txt

Rivest R. 'Cryptography and Security Resource Page', at http://theory.lcs.mit.edu/~rivest/crypto-security.html

Rivest R.L. & Lampson B. (1996) 'SDSI - A Simple Distributed Security Infrastructure', 15 Sep 1996, at http://theory.lcs.mit.edu/~rivest/sdsi10.html

SDSI (1996-) 'A Simple Distributed Security Infrastructure (SDSI)', 1996-, at http://theory.lcs.mit.edu/~cis/sdsi.html

Shaw P.D. (1998) 'Managing Legal and Security Risks in Computing and W3C (2000) 'Public-Key Infrastructure (X.509) (pkix)', at http://www.ietf.org/html.charters/pkix-charter.html

Sneddon M. (2000) 'Legal Liability and e-Transactions` National Electronic Authentication Council' Canberra, Australia, August 2000, at http://www.noie.gov.au/publications/NOIE/NEAC/publication_utz1508.pdf

Wang Y. (1998) 'SPKI' December 1998, at http://www.hut.fi/~yuwang/publications/SPKI/SPKI.html

Wheeler L. (1998) 'Account Authority Digital Signature Model (AADS)', at http://www.garlic.com/~lynn/aadsover.htm

Wheeler A. & Wheeler L. (1998) 'PKI Account Authority Digital Signature Infrastructure', November 1998, at http://www.garlic.com/~lynn/draft-wheeler-ipki-aads-01.txt

Winn J.K. (1998) 'Open Systems, Free Markets,and Regulation of Internet Commerce' 72 Tulane L. Rev. 1177 (1998), at http://www.smu.edu/~jwinn/esig.html

Winn J.K. (2001) 'The Emperor's New Clothes: The Shocking Truth About Ditial Signatures and Internet Commerce` Idaho Law Review, 2001

X.509 (1988, 1997) 'The Directory - Authentication Framework', Volume VIII of CCITT Blue Book, pp. 48-81, CCITT/ITU, 1988, 1997, v1 1988, v2 1993, v3 1996

App. 4: Further Information re SPKI / SDSI

1. Background

The objective of the e-Consent Project is to identify practicable models whereby patient consent can be expressed and communicated in electronic contexts. Several technologies have been identified which may provide an appropriate basis for implementing an e-consent object.

This document reports on discussions concerning SPKI / SDSI, conducted with Prof. Ron Rivest, initially by email, and then in a meeting on 2 October 2001 at MIT in Boston. Subsequent telephone conversations and email interchanges were held with Dr Carl Ellison at Intel in Portland Oregon, and Prof. Peter Szolovits at MIT, and further research was performed on the relevant web-sites.

2. Inadequacies of Conventional PKI

In order to conduct transactions and develop relationships in cyberspace, each party needs to know enough about the other party/ies to engender a threshhold level of confidence. For quite some years now, the expectation has been that parties would provide information to one another that is authenticated by means of digital signatures based on public key cryptography. The tools and processes that enable such digital signatures to be passed around are referred to as public key infrastructure (PKI).

Conventional PKI designs have naively assumed that key-pairs should be associated with (or 'bound to') personal identities. This involves enormous challenges of both a technical and an organisational nature, creates very serious threats to privacy, and is in many cases unnecessary anyway. Clarke (2001) catalogues the problems.

Conventional PKI has been based on a particular certificate format called X.509. But such schemes have not gained widespread acceptance. They have proven to be disappointingly complex, slow to develop and deploy, expensive, and subject to both technical and legal deficiencies. Moreover, claims about the assurances that they offer have frequently been exaggerated, and the terms of contract offered by certificate-providers offer little comfort.

Several attempts have been made to specify a PKI that avoids these problems. Amongst the most important of these are SPKI and SDSI.

3. SPKI / SDSI

Simple Public Key Infrastructure (SPKI) was specified in the mid-1990s by Carl Ellison, previously at CyberCash and now at Intel. See Ellison (1996), IETF (1997-), Wang (1998), Ellison (1998), Ellison et al. (1999), RFC2692 (1999), RFC2693 (1999), Ellison (2000) and Ellison (2001).

In parallel, Ron Rivest had been working on a Simple Distributed Security Infrastructure (SDSI). SDSI v.2 incorporates the SPKI concepts. See SDSI (1996-), Rivest & Lampson (1996), Ellison (2000) and Ellison (2001).

A key element of SDSI is that the X.509 nirvana of a single, global name-space has been abandoned. With it, the presumption has been removed that 'name' (or, better expressed, 'identifier') is reliably bound to a particular entity. The certificate associates a public key (and hence a key-pair) to an entity that only the CA knows, and no warranties are provided by the CA to the recipient of the message as to who the keyholder is. It is up to the relying party to build up an image of the sender based on its successive interactions with the holder of that key.

Attributes are associated with public keys, not with identities of real-world entities. Hence, for example, a recipient can be assured that a particular message was provided by a medical practitioner, or a person over 18, or over 65; but the certificate is silent about the identity of the person who is using the key.

Lay explanations of SPKI's key features are provided in Appendices 1-3. A comparison among the various certificate-types is at http://world.std.com/~cme/html/web.html.

4. Progress in the Application of SPKI / SDSI

Laboratory experimentation has been conducted at MIT, and at some other research institutes including the University of California at Berkeley, and the Helsinki University of Technology.

Two MIT activities are Project Geronimo, comprising two Masters theses (Maywah 2000 and Clarke 2001), in which SPKI/SDSI was integrated into Apache and Netscape; and Project Oxygen, currently in progress, which is targeted at 'pervasive computing', i.e. small computers embedded in other devices, including personal adornments (Burnside et al. 2001). The Oxygen project is relevant, because it enables a challenge-response protocol to be performed by both 'heavyweight' and 'lightweight' devices.

These designs and protoypes are layered over SSL/TLS, TCP and/or UDP, and IP, and hence are capable of being implemented within existing networks.

Projects have been and are being undertaken within industry, e.g.:

- Hewlett-Packard's e-Speak project, at http://www.e-speak.hp.com; and
- Intel's Common Data Security Architecture (CDSA) release, which is open-source, at http://developer.intel.com/ial/security/

5. Applicability of SPKI / SDSI to the e-Consent Object

SPKI / SDSI would provide a very strong basis for prototyping of the e-Consent Object, and indeed for implementing it within the context of Internet protocols. This is because it is oriented towards authorisation rather than identity authentication, and the first example of SPKI application given in Appendix 2 is directly relevant to the project: "If you're running a Web server for medical records and a user asks for a record, you need to decide if the user is authorized to read that record".

As regards health applications, Rivest referred me to Prof. Peter Szolovits, in the Medical Group, also within MIT's Laboratory of Computer Science. Szolovits is well aware of the potential for SPKI/SDSI to be applied within the health care arena. He has not applied it to consent, but is involved in a project to apply it to access controls for a projected personal life-long medical record. This is being undertaken in conjunction with Alberto Riva at the Boston Children's Hospital.

A description is provided under the **Guardian Angel** link from Szolovits' home-page: http://www.ga.org/ga/. Of particular relevance are the following segments of the project:

- **PING (Personal Internetworked Notary and Guardian)**, at http://www.ga.org/ga/#PING, and described in:

  Mandl, K. D., Szolovits, P., Kohane, I. S., et al. (2001) 'Public standards and patients' control: how to keep electronic medical records accessible but private' British Medical Journal 322(7281 283-287, at http://bmj.com/cgi/content/full/322/7281/283.

  Riva A., Mandl K. D., Oh D. H., Szolovits P., Kohane I. S. (2001) 'The Personal Internetworked Notary and Guardian' International Journal of Medical Informatics. 62: 27-40, at http://www.ga.org/ga/Riva2001.pdf;
- to some extent also **HealthConnect**, at http://www.ga.org/ga/#HC; and
- to some extent also **W3-EMRS: World Wide Web based Electronic Medical Record System**, at http://www.ga.org/ga/#W3EMRS (although at a brief glance this seems to be strikingly unconcerned about consent).

The PING sub-project in particular appears to be closely complementary to the e-Consent project, and Szolovits said he would be pleased to interact further on this matter.

6. Availability of SPKI / SDSI Code Libraries

A list of sites offering source-code and products is at:
http://world.std.com/~cme/html/spki.html#Code.

A Java implementation is available from MIT at:
http://theory.lcs.mit.edu/~cis/sdsi/sdsi2/java/SDSI_Java_Intro.html.
It is still marked as being for download only in the U.S.A. and Canada. But this is a hangover from the days when cryptographic products were restricted for export, and

there are no longer any constraints on licensing the code-library to an Australian person or organisation. Rivest pointed out, however, that this was developed by postgraduate students, as a research resource, and that it would probably be more appropriate to use a commercial code-library, given that one was available.

Rivest referred me to the CDSA code-library, available from Intel, and provided contact details for Dr Carl Ellison at Intel. I have previously met Ellison, called him in Portland while I was in Boston, on 2 October, and exchanged emails with him subsequently.

Intel describes its **Common Data Security Architecture (CDSA)** at:
http://developer.intel.com/ial/security/
http://developer.intel.com/ial/security/tech.htm
as "a security middleware specification and reference implementation that is open source, cross-platform, interoperable, extensible, and freely exportable. It is a set of layered security services".

Ellison confirmed that a particular component of the CDSA code library was relevant to the e-consent project. The **Authorization Computation, AuthCompute or AC module** implements SDSI v.2. It can be used independently of the remainder of CDSA. One small part of the process (the signing of certificates) is performed by a separate module; but this can be hand-written, or alternatively Ellison offered to send some C code that implements it. He also expressed interest in the e-consent project more generally.

CDSA/AC source code is available for download from:
http://sourceforge.net/projects/cdsa.

7. Application in Australia

None of the interviewees were aware of any downloads of the code-libraries, or other SPKI / SDSI-related activity, within Australia. Ron Rivest was aware of Jenny Seberry in Wollongong, and Peter Szolovits knows Enrico Coiera at UNSW.

Participation by any Australians might be uncovered by a search of the mailing list archives relating to CDSA, at http://sourceforge.net/mail/?group_id=13228.

8. Conclusions and Recommendations

A design using the CSDA/AC tool-set would very likely be an appropriate and effective means firstly of prototyping the e-consent object, but also secondly of implementing it.

In the event that no tender is submitted that involves use of the CDSA/AC tool-set, it would be valuable to withhold a proportion of the available funding from the initial round of grants, and distribute a second round of invitations, specifically requiring proposals to apply SPKI/SDSI by means of that code-library.

References

Barlow L. (2000) 'CDSA Now Includes Biometric Authentication, Authorization' Intel Developer UPDATE Magazine, September 2000, at http://developer.intel.com/update/departments/initech/it09003.pdf

Burnside M., Clarke D., Devadas S. & Rivest R. (2001) 'Distributed SPKI/SDSI-Based security for Networks of Devices' Working Paper, August 2001

Clarke D. (2001) 'SPKI/SDSI HTTP Server /Certifcate Chain Discovery in SPKI/SDSI' Master 's thesis,Massachusetts Institute of Technology, 2001

Clarke D., Elien J.-E., Ellison C., Fredette M., Morcos A. & Rivest R. (2001) 'Certificate Chain Discovery in SPKI/SDSI' Forthcoming in Journal of Computer Security, 2001

Clarke R. (2001) 'The Fundamental Inadequacies of Conventional Public Key Infrastructure' Proc. Conf. ECIS'2001, Bled, Slovenia, 27-29 June 2001, at http://www.anu.edu.au/people/Roger.Clarke/II/ECIS2001.html

Ellison C. (1996) 'Establishing Identity Without Certification Authorities', Proc. 6th USENIX Security Symposium, San Jose CA, July 22-25, 1996, at http://world.std.com/~cme/usenix.html

Ellison C. (1998) 'SPKI Examples', March 1998, at http://world.std.com/~cme/examples.txt

Ellison C. (2000) 'SPKI/SDSI and the Web of Trust' September 2000, at http://world.std.com/~cme/html/web.html

Ellison C. (2001) 'SPKI/SDSI Certificates' 21 September 2001, at http://world.std.com/~cme/html/spki.html

Ellison C., Frantz B., Lampson B, Rivest R., Thomas B. & Ylonen T. (1999) 'Simple Public Key Certificate' The Internet Society, July 1999, at http://world.std.com/~cme/spki.txt

IETF (1997-) 'Simple Public Key Infrastructure (SPKI)', at http://www.ietf.org/html.charters/spki-charter.html

Maywah A. (2000) 'An Implementation of a Secure Web Client Using SPKI/SDSI Certificates' Master 's thesis, Massachusetts Institute of Technology, 2000

RFC2692 (1999) 'SPKI Requirements' Internet Engineering Task Force of The Internet Society, September 1999, at ftp://ftp.isi.edu/in-notes/rfc2692.txt

RFC2693 (1999) 'SPKI Certificate Theory' Internet Engineering Task Force of The Internet Society, September 1999, at ftp://ftp.isi.edu/in-notes/rfc2693.txt

Rivest R. 'Cryptography and Security Resource Page', at http://theory.lcs.mit.edu/~rivest/crypto-security.html

Rivest R.L. & Lampson B. (1996) 'SDSI - A Simple Distributed Security Infrastructure', 15 Sep 1996, at http://theory.lcs.mit.edu/~rivest/sdsi10.html

SDSI (1996-) 'A Simple Distributed Security Infrastructure (SDSI)', 1996-, at http://theory.lcs.mit.edu/~cis/sdsi.html

Wang Y. (1998) 'SPKI' December 1998, at http://www.hut.fi/~yuwang/publications/SPKI/SPKI.html

Attachment 1: Description of SPKI / SDSI

Extracts from Burnside M., Clarke D., Devadas S. & Rivest R. (2001) 'Distributed SPKI/SDSI-Based security for Networks of Devices' Working Paper, August 2001

…
The SPKI/SDSI framework provides mechanisms for easy maintenance of access control lists (ACLs). It also features an elegant model for forming groups and delegating authority.
…
Using the SPKI/SDSI framework, access control lists (ACLs) associated with resources can be created once and rarely need to be modified. User access rights are modified by issuing certificates based on group membership. SPKI/SDSI also facilitates short certificate validity periods to assist in the problem of certificate revocation. In addition, **SPKI/SDSI features an elegant model for delegation of authority, allowing for the partitioning of responsibilities**. The principal maintaining the ACL could, but need not be, the same principal that authorizes users to use a resource. This significantly eases the burden of system administration.
…
SPKI/SDSI provides fine-grained access control using a local name space architecture and a simple, flexible,trust policy model. SPKI/SDSI is a public key infrastructure with an egalitarian design. **The principals are the public keys and each public key is a certificate authority. Each principal can issue certificates on the same basis as any other principal. There is no hierarchical global infrastructure. SPKI/SDSI communities are built from the bottom-up,in a distributed manner, and do not require a trusted "root."**
…
The client proxy generates a chain of certificates using the SPKI/SDSI certificate chain discovery algorithm [4,3]. This certificate chain provides a proof of authorization that the user 's key is authorized to perform its request. The certificate chain discovery algorithm takes as input the ACL and tag from the server, the user's public key (principal), the user 's set of certificates, and a timestamp. If it exists, the algorithm returns a chain of user certificates which provides proof that the user 's public key is authorized to perform the operation(s) specified in the tag, at the time specified in the timestamp.
…
Since each principal can issue name certificates, each principal has its own local name space consisting of the names it defines. **SPKI/SDSI, thus, has a local name space architecture which helps to make the infrastructure scalable**. A user does not have to ensure that the names he defines are unique in a global name space; he can define names which are meaningful to him,which he can easily remember and recognize.
…
Note that **it is easy for a person to belong to multiple groups**. … The ability to define groups is one of the principal notions of SPKI/SDSI. This feature facilitates easy management of ACLs,as will be described in Section 5.2. It also makes it easier and more intuitive to define security policies. Because the names of the groups are at the discretion of the owners, the groups' names can be meaningful and intuitive. Security policies can be defined in terms of these groups, simplifying the auditing of group definitions and ACLs.

…
SPKI/SDSI is primarily concerned with authorizing principal perform particular operations on protected resources. An administrator controls access to a resource by setting up an ACL to protect it. A principal (public key) makes a request to perform a particular operation on the resource. Examples of requests are a request to read a file, a request to login to an account, or a request to turn on an appliance. In these examples,the protected resources are the file, account and appliance, respectively.
…
If an administrator wishes to grant a set of principals access to a number of resources, each protected by a separate ACL, he can simply define a group and place the group name on each ACL. The ACLs need only be updated once. From then on, as new members join the group, they are issued the relevant certificates and are automatically authorized to access the protected resources without the administrator having to update the ACLs again. It is clear that using groups makes it easier and more efficient to maintain and update ACLs, since an explicit list of all the principals does not have to be stored on each ACL.
…
…there can be a delayed definition of the group. An administrator can add a group to one or more ACLs without knowing the group 's members beforehand. He can update the ACLs with an entry for the group, and at some later time that is convenient and appropriate, he can issue name certificates adding principals to the group. He does not have to know all the members of a group when he is setting up his ACLs. Note that an administrator is free to add any principal's group to his ACL. He is not restricted to just adding his own groups …
…
…sometimes the entity responsible for protecting a resource may prefer to delegate the responsibility for making this determination. For example, in a research laboratory, it may be the sys-admin who is responsible for setting up and maintaining ACLs on the laboratory 's color printers. Instead of having every new graduate student come to him for access credentials, he may want to delegate this responsibility to one or more floor managers. With SPKI/SDSI, the sys-admin can delegate to the floor managers the authority to determine who is allowed to access the color printers. The sys-admin must trust the floor managers to correctly identify new graduate students before issuing them certificates.
…
SPKI/SDSI treats key compromise as a separate issue from certificate revocation. It argues that "certificates should not be revoked merely because the key is compromised. Rather, the signer should present separate evidence to the acceptor that the key has not been compromised. Since, in this framework, the no-compromise evidence is separate, the ordinary certificates can continue to be 'valid 'even though the key has been compromised." [19] SPKI/SDSI suggests using a new kind of agent, called a key compromise agent (KCA), or a suicide bureau (SB). Multiple SBs cooperate to serve SPKI/SDSI communities. When Alice creates her key pair, she also signs a personal suicide note which she protects in a private place, and also registers her public key with an SB. In the unlikely event that her key is compromised or lost, she sends her suicide note to the SB. The SB broadcasts this note on the SB network so that other SBs are made aware of the compromised key.
...

Attachment 2: Intel's Description of SPKI

Extract from Barlow (2000) – emphasis added:

SPKI Authorization

SPKI—the next step beyond digital certificates—is the simplest **way to determine whether a user is authorized to perform an action or use a resource**. This mechanism can also allow the authorized user to delegate some of his or her authority to other users, by binding a public key to a set of permissions.

SPKI offers a different kind of certificate format. The original certificate, as designed in 1978 [i.e. X.509] , bound names to public keys. **SPKI was designed to satisfy a different problem: authorization**.

Developers with security needs often must make authorization decisions. **If you're running a Web server for medical records and a user asks for a record, you need to decide if the user is authorized to read that record**. If you're sending a classified document by e-mail, encrypting the document isn't enough; you need to decide whether the recipient is authorized to read that document. If you are running a Web browser and connecting to an e-commerce page, you need to decide whether the machine at the other end is authorized to handle your payment information (e.g., your credit card number).

When a certificate gives only the name of a keyholder, that's enough information if you know the keyholder and what he or she is authorized to do. SPKI was designed for the more modern case, in which you do not know the keyholder and what he or she is authorized to do.

SPKI certificates bind authorizations directly to public keys. They can also be used to name people you do know and bind authorizations to those names, if you prefer to deal with names for your local community, either way **the end result is to communicate an authorization to some code that needs to make a security decision and to allow that decision to be made on the basis of a strongly validated (cryptographically proved) authorization, without requiring that code to know things by some other channel**. SPKI certificates simplify the job of making security decisions because they were designed specifically for that purpose, by developers who had that need.

Attachment 3: Ellison's Description of SPKI

Most of Ellison's work is written for technical audiences, and needs to be interpreted in order to speak to management needs. The following, however, is an extract from a recent email interchange:

SPKI/SDSI does not claim to bind between key pair and user. The RFC specifically states that SDSI names for a user are defined by the keyholder of that namespace (an individual) and are meant to be meaningful to that keyholder and to no one else. The original SPKI certificates didn't use names at all -- just raw keys.

{If you've seen my certificate triangle, SPKI supplied the certificates along the base of the triangle, while SDSI supplied the name definition (the top vertex) and the certificates on the right edge (name to key). When they were merged, the use of a name as a subject in an SPKI certificate provided the left edge of the triangle.}

None of this is designed to indicate to a human relying party a name for the subject keyholder that would be meaningful to that relying party. That is the function of X.509 and one that SPKI/SDSI discounts as meaningless in practice. SPKI/SDSI concentrates instead on specifying what permissions a given key has been granted.

Attachment 4: Contact Points

Prof. Ron Rivest
Laboratory of Computer Science, MIT
200 Technology Square, Room 324
mailto:rivest@mit.edu (Ron Rivest)
http://theory.lcs.mit.edu/~rivest/
Tel: +1 617 253 5880
Dwaine Clarke
200 Technology Square, Room 226
mailto:declarke@mit.edu
Tel: +1 617 253 0702
Matthew Burnside
200 Technology Square, Room 226
Tel: +1 617 253 5970
mailto:event@mit.edu

Carl Ellison
Intel, Portland OR
mailto:cme@acm.org (Carl Ellison)
http://world.std.com/~cme/
Tel: +1 503 264 2900
Mob: +1 503 819 6618

Prof. Peter Szolovits
Laboratory of Computer Science, MIT
200 Technology Square, Room 416
mailto:psz@medg.lcs.mit.edu (Peter Szolovits)
http://www.medg.lcs.mit.edu/
Tel: +1 617 253 3476
Mob: +1 617 686 4932

App. 5: Further Information re Brandsian Certificates

1. Background

The objective of the e-Consent Project is to identify practicable models whereby patient consent can be expressed and communicated in electronic contexts. Several technologies have been identified which may provide an appropriate basis for implementing an e-consent object.

This document reports on discussions concerning Brandsian Private Credentials, conducted with Dr. Stefan Brands, by email, and in meetings on 3-4 October in Washington DC.

2. Inadequacies of Conventional PKI

In order to conduct transactions and develop relationships in cyberspace, each party needs to know enough about the other party/ies to engender a threshhold level of confidence. For quite some years now, the expectation has been that parties would provide information to one another that is authenticated by means of digital signatures based on public key cryptography. The tools and processes that enable such digital signatures to be passed around are referred to as public key infrastructure (PKI).

Conventional PKI designs have naively assumed that key-pairs should be associated with (or 'bound to') personal identities. This involves enormous challenges of both a technical and an organisational nature, creates very serious threats to privacy, and is in many cases unnecessary anyway. Clarke (2001) catalogues the problems.

Conventional PKI has been based on a particular certificate format called X.509. But such schemes have not gained widespread acceptance. They have proven to be disappointingly complex, slow to develop and deploy, expensive, and subject to both technical and legal deficiencies. Moreover, claims about the assurances that they offer have frequently been exaggerated, and the terms of contract offered by certificate-providers offer little comfort.

Several attempts have been made to specify a PKI that avoids these problems. Amongst the most important of these are Brandsian Private Credentials.

3. Brandsian Private Credentials

Brands worked through the 1990s developing his theory and articulating it into a complete proposal. His motivation was to produce a privacy-protective alternative, and his articles and book are written from that persepctive. The Appendix contains a series of linked quotations from his works, summarising the manner in which he has been presenting his approach. This section provides a brief overview of his technology, observed from a conventional perspective rather than from a specifically privacy protection viewpoint.

Brandsian Private Credentials represent a superior alternative to conventional approaches to PKI. Their cryptographic design provides the following significant advantages:

- they typically involve short-life certificates, which promotes a different approach to the revocation of certificates and thereby removes one of the significant risk-exposures in conventional schemes;
- they support all of the following:
  - identity certificates (in much the same manner as X.509);
  - attribute certificates associated with identity certificates (but in a manner more flexible than that provided by X.509); and
  - attribute certificates independent of identity (which is a highly desirable service in many circumstances, but which is not supported by X.509);
- they enable significantly greater privacy protection;
- they enable significantly greater protection against fraud, and in particular against identity theft and unauthorized lending of certificates;
- they enable the use of simpler smartcards than is possible with other cryptographic approaches, with the result that the total costs of implementation and operation are potentially considerably lower, especially in large-scale schemes.

4. Progress in the Application of Brandsian Private Credentials

The theory underlying Private Credentials was laid in Brands doctoral work, conducted in The Netherlands between 1992 and 1999. He worked initially with David Chaum, but that research team broke up when Chaum moved to the U.S., and it was eventually completed at Eindhoven. Publication of the thesis was preceded by the issue of multiple patents and a series of technical papers.

Various aspects of Private Credentials have been applied in multiple prototypes. The CAFE and OPERA payments projects for the European Commission involved close to twenty companies, including Gemplus and Siemens. Several other e-cash implementations have used it.

During 2000-2001, Zero-Knowledge Systems (ZKS), based in Montreal have developed e-cash for PCs and RIM's Blackberry (a handheld device specially suited for handling e-mail), and a modular toolkit in C. As a result of the downturn in the I.T. industry, ZKS have recently suspended their Private Credentials project (and their long-running

Freedom service), in order to focus on short-term revenue streams. Brands left ZKS in August 2001, and is currently seeking an appropriate host-corporation for his work.

Brands owns 9 international patents. He has twice licensed it, to Chaum's company DigiCash, and to ZKS; but both licences have been terminated. (This was not because of problems with Private Credentials, but because DigiCash went bankrupt and ZKS has withdrawn from the area).

Demonstrators exist, but are owned by ZKS. A toolkit exists, but is also owned by ZKS. Brands is currently negotiating to acquire the rights to them, and expects little difficulty in achieving that end.

5. Applicability of Brandsian Private Credentials to the e-Consent Object

Brandsian Private Credentials are capable of being implemented in a variety of forms, supporting functions ranging from attribute authentication without identity, to identity certificates.

Their advantages in the context of e-consent in the health care sector would include:

- reduced dependence on the complex but ultimately ineffectual bureaucracy inherent in conventional PKI;
- lower infrastructure investment than would be necessary using conventional PKI; and
- greater privacy-protection than is possible using the fixed structures of X.509 certificates.

In order to construct a prototype, a laboratory would need to:

- acquire the available documentation from Brands (or ZKS);
- acquire demonstrators from Brands (or ZKS);
- acquire a licence from Brands;
- acquire the code-library from Brands (or ZKS); and
- extend the code-library as necessary.

It is recommended that the following steps be taken:

- an amount of money be reserved from the current round of grants, with the intention that it be used for prototyping using Brands Private Credentials;
- the information in this document be disseminated to the laboratories that were invited to tender for grants;
- proposals be invited for prototype development using the technology; and
- contact be sustained with Brands, in order to facilitate access by the selected laboratory to licences and the code-library.

References

Brands S. (2000) 'Rethinking Public Key Infrastructures and Digital Certificates – Building in Privacy' MIT Press, ISBN 0-262-02491-8

Brands S. (2001) 'A Technical Introduction to Private Credentials', 1 October 2001

Clarke R. (2001) 'The Fundamental Inadequacies of Conventional Public Key Infrastructure' Proc. Conf. ECIS'2001, Bled, Slovenia, 27-29 June 2001, at http://www.anu.edu.au/people/Roger.Clarke/II/ECIS2001.html

ZKS (2000) 'Private Credentials' White Paper, Zero-Knowledge Systems Inc., November 2000

Contact Points

Dr Stefan Brands
3781 av. Laval
Montreal, Quebec
H2W 2H8 Canada

mailto:brands@xs4all.nl (Stefan Brands)
Tel: +1 514 281 1417
Mob: +1 514 583 2726

Attachment: Abbreviated Rendition of Brands' Original Presentation

1. Inadequacies of Conventional PKI

Brands identifies serious problems with conventional digital certificates and public key infrastructure: "each digital certificate can be traced uniquely to the person to whom it has been issued (or the device in which it has been incorporated), and can be followed around instantaneously and automatically as it moves through the system. Even digital certificates that do not specify the identity of their holder can be traced in a trivial manner ... On the basis of these unique serial numbers, which will travel along whenever an individual engages in a communication or a transaction, organisations and even individuals can compile detailed personal dossiers" (2000, p. vii-viii).

He identifies the privacy dangers in (2000) pp. 16-20, examines previous privacy-protection efforts and their shortcomings in (2000) pp. 20-24, and defines desirable privacy properties of certificates in (2000) pp. 24-26.

2. Brandsian Private Credentials

Brands has designed an alternative form of digital certificates "that fully preserve privacy, without sacrificing security. The new certificates function in much the same way as do cash, stamps, cinema tickets, subway tokens, and so on: anyone can establish the validity of these certificates and the data they specify, but no more than that. Furthermore, different actions by the same person cannot be linked [except where that is expressly designed into the application]" (2000, p. viii).

The basic definition of the Brandsian 'secret-key certificate' is in (2000) at pp. 74-77. Details are in a patent filed in 1994, and are explained through the remainder of (2000).

Among other features of Brandsian certificates, "certificate holders can decide for themselves, depending on the circumstances, which property they disclose of the data encoded into their certificates. Also, a certificate can be presented in such a manner that the organisation is left with no evidence at all of the transaction, or only with self-authenticating evidence of a message or a part of the disclosed property. Furthermore, the self-authenticating evidence can be limited to designated parties" (2000, p. ix).

Brandsian certificates may attest to the holder's identity (where that is what the application calls for), but more generally attest to the (unidentified) holder having one or more specific attributes that are relevant to the transaction. They encompass conventional X.509v3 digital certificates as a special case (2000, p.29). They may be used over an extended period (in which case the scope for a dossier to be constructed increases, because each use of the same certificate carries the same public key). Alternatively, each certificate may be used only a few times, or just once, and the holder can obtain replacements for them (2000, p. 166).

The technique can be implemented entirely in software, or in smartcards.