2007

A Sustainable Approach to Security and Privacy in Health Information Systems

Vicky Liu
Queensland University of Technology, v.liu@qut.edu.au

Lauren May
Queensland University of Technology, l.may@qut.edu.au

William Caelli
Queensland University of Technology, w.caelli@qut.edu.au

Peter Croll
Queensland University of Technology, croll@qut.edu.au

Follow this and additional works at: http://aisel.aisnet.org/acis2007

A Sustainable Approach to Security and Privacy in Health Information Systems

Vicky Liu, Lauren May, William Caelli and Peter Croll
Faculty of Information Technology and Information Security Institute
Queensland University of Technology,
Brisbane, Australia
Email: {v.liu, l.may, w.caelli, croll}@qut.edu.au

Abstract

This paper identifies and discusses recent information privacy violations or weaknesses which have been found in national infrastructure systems in Australia, the United Kingdom (UK) and the United States of America (USA), two of which involve departments of health and social services. The feasibility of health information systems (HIS) based upon intrinsically more secure technological architectures than those in general use in today's marketplace is investigated. We propose a viable and sustainable IT solution which addresses the privacy and security concerns at all levels in HIS with a focus on trustworthy access control mechanisms.

Keywords

Access control, trusted systems, information assurance, health information systems

Introduction

Today's service industries would regard information, computer and telecommunication (ICT) technologies as part of their critical infrastructure. Although some sectors such as healthcare, have been slow in their adoption of ICT, it is evident they are working towards a future where ICT technologies will be both widespread and essential. The use of computer-based information systems and associated telecommunications infrastructure to process, transmit and store health information plays an increasingly significant role in the improvement of quality and productivity in healthcare. Notwithstanding the obvious potential advantages of deploying ICT in healthcare services, there are some concerns associated with integration of and access to electronic health records. Information stored within electronic health systems is highly sensitive by its very nature, therefore health records have clear requirements for confidentiality in order to safeguard personal privacy and to maintain record integrity.

A security violation in a health information system (HIS), such as an unauthorised disclosure or unauthorised alteration of individual health information, has the potential for disaster among healthcare providers and consumers. Although the concept of Electronic Health Records (EHR) has much potential for improving the processing of health data, Goldschmidt (Goldschmidt 2005) warns that electronic health records may also pose new threats for compromising sensitive personal health data if not designed and managed effectively. Goldschmidt also illustrates that malevolent motivations could feasibly disclose confidential personal health information on a more massive scale and at a higher speed than possible with traditional paper-based medical records. Quinn suggests that the key factor to successful implementation of a national health information system is user adoption (Quinn 2004). User acceptability and adoption in e-health relies on the healthcare consumers' willingness to overcome the fear of privacy invasion in relation to their health information. There is also the factor of the healthcare service providers' willingness to accept and adopt a new technology that does not always facilitate efficient working practices. To encourage healthcare service consumers and providers to use electronic health records, it is crucial to instil confidence that the electronic health information is well protected and that consumers' privacy is assured.

Several countries including Australia, the UK, the USA, Canada and New Zealand are actively involved in the development of e-health initiatives. The current approaches to protecting personal privacy and confidentiality of electronic patient records are, in the opinion of the authors, not sustainable. This paper identifies and discusses three scenarios related to information privacy violations or weaknesses which have recently been found in Australia, the UK and the USA. The paper proposes a viable ICT solution to provide appropriate levels of secure access control for the protection of sensitive health data. Increasingly, HIS are being developed and deployed based upon commercial, commodity-level ICT products and systems. Such general-purpose systems have been created over the last 25 years with often only the minimal security functionality and verification. In particular access control, a vital security function in any operating system that forms the basis for application packages, has been founded upon earlier designs based on an access control method known as Discretionary

Access Control (DAC) as described in later sections. DAC systems were defined around an environment where data and program resources were developed and deployed within a single enterprise, assuming implicit trust amongst users. This environmental model is no longer valid for modern HIS. In some commercial systems, for example, even the addition of a simple single printer unit has the capacity to seriously undermine the overall integrity of the information system.

This paper investigates the feasibility of HIS based upon intrinsically more secure technological architectures than those in general use in today's marketplace. Even though such systems are currently commercially available for enterprise system deployment, for example SELinux, they are not in widespread use. The privacy and security issues required of HIS applications are analysed in the context of a new approach to a more trustworthy structure, the Open Trusted Health Informatics Scheme (OTHIS). This scheme consists of a number of trusted models including the Health Informatics Access Control (HIAC) system which is discussed in detail.

Access Control

Access control is one of the fundamental security mechanisms used to protect computer resources, in particular in multi-user and resource-sharing computer environments. "Access control" simply refers to a set of rules that specify which users can access what resources with which types of access restrictions. Various operating systems, network control systems, and database management systems (DBMS) can employ a choice of access control mechanisms to allow admission of a user to access protected resources of the system. It should be noted that in any information system a distinction may be made between "security aware applications" and "security ignorant applications". These latter applications usually depend solely upon access control facilities provided by an operating system, DBMS and other like middleware. Controlling appropriate access to data in any information system is a major security issue. Many instances of poor access control management practices leading to security and privacy violations are reported on a regular basis. Recent occurrences include:

Scenario 1: Privacy Invasion Scandal at Australia's Centrelink

Australia's Centrelink, a Commonwealth Government agency, delivers a range of social welfare services and payments to the Australian community including issuing Health Care Cards for concessions on healthcare costs. In carrying out its duties, Centrelink officers may verify information on personal financial and tax records with the Australian Taxation Office (ATO). According to a published media article (Sharanahan & Karvelas 2006), Centrelink conducted a two-year investigation on invasion of privacy by deploying spyware technology to audit and monitor employees' access to client records. The results of this investigation found 790 cases of inappropriate access to client records since 2004. Consequently, 19 IT staff were dismissed, 92 resigned, more than 300 staff faced salary deductions or fines, another 46 were reprimanded and the remainder were demoted or warned. Introduction of the proposed Medical access card in Australia, which may encompass healthcare parameters as well as social security information, is particularly concerning given the findings of this investigation.

Analysis 1: The information collected and stored by Centrelink is of a highly sensitive nature. It is therefore essential that the privacy and integrity of such information is safeguarded from internal and external security threats and attacks. Centrelink deploys spyware software to detect inappropriate access to client records and enforces the penalty for persons convicted in breach, however such steps only deal with occurrences of privacy violations in a reactive manner. It is preferable to adopt a proactive tamper resistant protection approach where such incidents simply cannot occur. The authors propose that this can be achieved by employing the appropriate technological controls to prevent unauthorised access or alteration of the private information ensuring individuals' privacy and integrity of their information.

Scenario 2: A Lack of Adequate Safeguards to Access UK NHS Patient Records

The current UK National Programme for IT (NPfIT) is considered to be the world's largest ICT project providing an HIS for 50 million patients. It has been reported by the media (Leigh & Evans 2006) that a lack of adequate security measures is in place regarding providing access to shared patient records once they are on the national database system. Patient records may contain sensitive information such as mental illness, abortions, pregnancy, HIV status, drug-taking or alcoholism. The article warns that the 50 million patient records may be made accessible by up to 250,000 National Health Services (NHS) staff including police and health managers, counsellor, social workers, private medical practices, ambulance staff and commercial researchers. This has resulted in calls for a boycott of patient records accessible by thousands of authorised NHS staff.

Analysis 2: The confidentiality management approach deployed by the UK NHS to access patient records will be on a "need-to-know" basis. Varied access permissions, based on the role-based access policy, will be granted to access patient records. In its basic form this is a simplistic approach which will not satisfactorily address the primary issue of a lack of adequate safeguards. In particular this approach does not allow patients to selectively

protect particular parts of their uploaded information from being widely accessed. NHS declares that a "sealed envelope" (NHS 2005) mechanism will allow patients to express access restrictions on disclosure of their confidential health information from specific roles. The provision of sealed envelopes however will not be available until the second phase of the release of the NHS Care Record Service.

Scenario 3: Significant IT Security Weaknesses Identified at USA HHS Information Systems

A published security analysis report from the United States Government Accountability Office (GAO) (GAO 2006) assessed the effectiveness of the Department of Health and Human Services (HHS) information security program with emphasis on the Centers for Medicare and Medicaid Services (CMS). The GAO's report reveals numerous significant security weaknesses in the areas of network management, user accounts and passwords, user rights and file permissions, and the auditing and monitoring of security-related events, specifically with HHS unnecessarily granting access rights and permissions to sensitive files and directories.

HHS provides essential health and welfare services to the USA community. CMS, a major operating division within HHS, is responsible for the Medicare and Medicaid programs. HHS is highly reliant on networked information systems to carry out their services including processing medical claims, conducting medical research, managing health and disease prevention, and a food safety program. Because such information systems contain sensitive medical and financial information, it is essential that the security and integrity of such information systems are safeguarded from security threats and vulnerabilities.

Analysis 3: The identified security weaknesses in the HHS information systems increase the very high risk that unauthorised users can gain access to and subvert the systems upon which HHS relies to deliver its vital services. Not surprisingly, this has the potential to expose clients' sensitive information to serious privacy invasions. GAO's recommendation (GAO 2006) to HHS is to implement a complete set of comprehensive information security programs at all operating divisions to address the identified weaknesses.

The three illustrated scenarios all have a common security weakness issue which is directly related to access control management. Appropriate computer-based access control schemes can be deployed to address these information security issues. Access control mechanisms, then, are used to restrict users' accesses to resources. Organisations use these controls to grant employees the authority to access only the information the users need to perform their duties. Access controls can limit the activities that an employee can perform on data. Before proposing a viable solution to provide appropriate levels of secure access control for protecting sensitive health data, one must first understand the primary types of computer-based access control. These are examined in the following section.

Access Control models

The two traditional types of access control modes are Discretionary Access Control (DAC) and Mandatory Access Control (MAC). The Role-Based Access Control (RBAC) concept is complementary to both DAC and MAC techniques. RBAC enables easier management by ensuring finer granularity in the access system.

Discretionary Access Control (DAC)

The DAC mechanism is widely implemented for the purpose of managing access control by current commodity software such as Microsoft Corporation's Windows systems, open-source systems such as Linux and the original Unix system. The DAC policy allows the owner of information to grant access permissions to other users or programs at his/her discretion without the system administrator's knowledge. Each user has complete discretion over his/her own objects. Thus, such a policy does not provide the actual owner of the system fully centralised access control over the organisational resources. In fact, the system cannot identify the difference between a legitimate request to modify access control information which originated from the owner of the information and a request issued by a malicious program (Gasser 1988).

DAC mechanisms are fundamentally inadequate for strong system security. One of the major deficiencies with DAC is its vulnerability to some types of Trojan horse attacks. Trojan horses embedded in applications can exploit DAC's vulnerability to cause an illegal flow of information. Applications that rely on DAC mechanisms are vulnerable to tampering and bypassing (Loscocco & Smalley 2001). Malicious or flawed applications can easily cause security violations in the system. This shortcoming of DAC can be overcome by employing MAC policies to prevent information flow from higher to lower security levels.

Mandatory Access Control (MAC)

Gasser (Gasser 1988) states that MAC can be used to prevent some types of Trojan horse attacks by imposing severe access restrictions that cannot be bypassed intentionally or accidentally. MAC can provide the ability to

limit access to only legitimate users. Ferraiolo et al (Ferraiolo, Kuhn & Chandramouli 2003) underscore that MAC is necessary when provision of a truly secure system is required.

With MAC, each user possesses a clearance that is used by the system to determine whether a user can access a particular file. Access permissions are determined by a user's clearance compared with the sensitivity (or security) or classification level label on information stored in the system, not upon the user's discretion. The classification may contain an arbitrary number of categories; for example a conventional hierarchical category set used in military environments might include "top secret", "secret", "confidential" and "unclassified". Each user possesses a clearance that is used by the system to determine whether a user can access a particular file. The access permission to information is determined by the user's clearance compared to the security level of information stored in the system. This is also known as a multi-level security (MLS) policy, which was first introduced by Bell and LaPadula (BLP) (Bell & LaPadula 1973).

With the MLS policy, BLP propose an access control system in the form of a mathematical model for defining and evaluating computer security. This model is designed to address the enforcement of information confidentiality aimed at the prevention of unauthorised information leakage. The BLP model defines two basic rules for making access control decisions: the Simple property and the Star property. The Simple property regulates whether a subject is allowed to read an object (i.e. if the subject's clearance level dominates the security level of the object). It is also known as the "no read up" policy. The Star property determines whether a subject is allowed to write to an object (i.e. if the security level of the object dominates the subject's security clearance level). It is referred to as the "no write down" policy (Ferraiolo, Kuhn & Chandramouli 2003; Gasser 1988).

The traditional MAC policy was originally designed for a military environment based on the MLS hierarchical structure and was quite rigid in its application. More recent research has modernised the traditional MAC approach, overcoming its traditional limitations, in order to better suit contemporary applications such as for the HIS environment.

Role-based Access Control (RBAC)

RBAC is based upon the role concept in managing access control where access permissions are associated with roles. Users are assigned to appropriate roles within the organisation. The user must be assigned as a member of a role in order to perform an operation on an object. Ferraiolo et al (Ferraiolo, Kuhn & Chandramouli 2003) state that the driving force behind the RBAC model is to simplify the management of authorisation. Assigning users' access permissions to each protected object in the system on an individual user basis, particularly in large scale enterprise systems, is an onerous process in security management. With RBAC, users are granted membership into roles according to their responsibilities and competencies. User membership of roles can be included and revoked easily. Updates of assigning privileges can be done to roles rather than updating permission assignments for individual users. RBAC supports users' access rights based on such parameters as job function, enforcement of least privilege for administrators and users, enforcement of static/dynamic separation of duties (SOD) and hierarchical definitions of roles.

In spite of several advanced RBAC features, RBAC also brings a number of limitations. Significantly, Reid et al (Reid et al. 2003) point out that RBAC does not efficiently support access policies in the way of general consent qualified by explicit denials. This issue is quite apparent in the privacy vulnerability that occurred in the UK NHS patient record system analysed in Scenario 2. There is also a lack of available products to support the full features of RBAC. A number of research papers discuss the use of the RBAC mechanism for authorisation management in healthcare environments, since role models are suitable for the representation of roles in hospital settings. Ferraiolo et al (Ferraiolo, Kuhn & Chandramouli 2003), the developers of the first model for RBAC and proposers of the RBAC standard, state that RBAC is policy-independent and policy neutral in not enforcing any particular protection policy. Ferraiolo et al also point out that the availability of RBAC does not obviate the need for MAC and DAC policies. MAC is particularly needed when confidentiality and information flow are primary concerns.

Rethink Access Control Models in HIS

Current moves toward Web-based identity and authentication structures present a major challenge where such structures are not based on highly trusted operating systems. All applications and supporting software which necessarily reside atop the untrusted operating systems are also untrusted. We emphasise the need for further research into, and redefinition of, MAC in light of modern information system structures, legislative and regulatory requirements and flexible operational demands in HIS.

Building upon experience with DAC and MAC structures, indications are that a radical re-think is required in the understanding of access control in general in current and future information systems, and in particular in the healthcare environment. One limiting factor in approaches to "hardening" current information systems is the

perceived or expected business requirements to maintain backward compatibility for legacy applications (Microsoft 2006).

Any access control system fundamentally depends upon a trusted base for safe and reliable operation, commonly referred to as a "trusted computing-base (TCB)". Without a TCB, any control structures are subject to compromise. In the past, access control paradigms have been based around fundamentals in operating systems, DBMS and similar IT products. With the ubiquity of information systems, this paper proposes that access control requirements need to be defined against the background of the relevant industries served by such systems.

Information Protection in the Health Sector

A security analysis report published by the USA GAO (GAO 2007) reveals that the USA Department of Health and Human Services (HHS) has initiated actions to identify solutions for protecting personal health information. An overall approach for integrating HHS systems with various privacy related initiatives and for addressing security has not yet been defined. GAO identifies key challenges associated with protecting electronic personal health information in four areas. Two particular areas are relevant to this paper: understanding and resolving legal and policy issues, and implementing adequate security measures for protecting health information. This paper proposes a viable approach which provides the potential for sustainable security measures to protect the privacy and security of health information under an overall trusted health informatics scheme.

Health Information System Architectures

A modern HIS architecture would normally consist of health application services, middleware, database management system (DBMS), data network control system, operating system and hardware, as shown as in Table 1 (c). Many application users wrongly believe that they have sophisticated security at this level since their applications provide role-based access control or equivalent. It should be understood that no matter what security measures are supported at the application level they are only ever going to be superficial to the knowledgeable adversary or malicious insiders. This approach has a significant limitation in that the overall system can be no more secure than the operating system upon which the applications depend. The operating system itself can be no more secure than the firmware and hardware facilities of the computer on which it operates. Likewise, any other software component set, such as "middleware", DBMS, network interface structure or "stack", is constructed above the operating system and so totally depends upon security functions provided by the operating system as well as the robustness of that operating system against attack.

Table 1: (a) OSI Model, (b) TCP/IP Model and (c) General HIS Architecture

	(a)OSI Model	(b)TCP/IP Model	(c) HIS Architecture
Software System Components	**Application**	**Application**	**Health service application** **Middleware** **DBMS**
	Presentation		
	Session	**(not present)**	**Data network management system** **Operating system**
	Transport	**Transport**	
	Network	**Internetwork**	
	Data Link	**Network Access**	
Hardware	**Physical**		**Hardware**

Open Trusted Health Informatics Scheme (OTHIS)

To achieve a high level of information assurance in HIS, our research to date has indicated that an overall trusted HIS involves the definition of structures at a number of levels in computer hardware, operating system, data network control system and health service applications. We propose the Open Trusted Health Informatics Scheme (OTHIS) which is aimed at addressing privacy and security requirements in a holistic manner. OTHIS defines privacy and security requirements at each level within the general HIS architecture to ensure the protection of data from both internal and external threats as well as providing conformance of HIS to meet regulatory and legal requirements.

OTHIS Structure

The OSI reference model (ISO 7289-1) (Table 1(a)) is well known and acknowledged as a baseline for categorisation of network communication functions and assessment. In fact, a fully operational system based on the seven-layer OSI model never attained strong market acceptance. The OSI model envisaged management and control facilities existing at each layer but many of the detailed specifications and activities at each layer were never completed. Instead, TCP/IP (Table 1(b)) is the model used globally for large scale structures in network communications. The TCP/IP model does not exactly match the OSI model (Table 1(a)), however the processes defined in the OSI model are contained in the TCP/IP layers. Normally HIS are based around distributed network systems, therefore it is entirely appropriate to relate the general HIS architecture to the OSI model as well as the TCP/IP model (Table 1). Our research aims to relate and describe the roles and functions performed by each module of the OTHIS architecture, and how they fit into the layers of the OSI and TCP/IP models in a healthcare environment.

It should be noted that the OSI model and HIS architecture can also be categorised into software and hardware components. From the point of view of this paper the first group, software system components, will be addressed. The interpretation of the requirements for appropriate levels of data granularity security in healthcare is the basis of this paper and research work performed to date.

Health Informatics Access Control (HIAC) Model

An operating system is a set of software programs that manages the hardware and software resources of a computer between the Physical layer and the Application layer of the OSI model and also forms a platform for other system software and application software. It is an inherently futile exercise to attempt to build an application requiring high levels of trust in security and privacy when the underlying structure within the computer system is a non-trusted operating system. The trusted application relies totally upon the non-trusted operating system to access low level services. The authors contend that ICT is now sufficiently advanced that a MAC-based electronic healthcare management system is feasible. Our research to date has indicated that current operating system structures need to be updated for HIS needs. The Health Informatics Access Control (HIAC) model within the OTHIS architecture is our approach to overcoming many of the privacy and security issues which have plagued previous attempts at electronic health management systems. HIAC is based on the MAC/RBAC type of operating system which primarily satisfies the requirement for confidentiality of records (this is a major impediment in current and previous systems). The HIS is then developed atop the trusted operating system.

Analysis of HIS Access Parameters

Table 2: Analysis of HIS Access Parameters

User role	Capability	DAC	RBAC	MAC	HIAC
Clinicians/ office administrator	User access	Access privileges determined and set by ICT system administrator	Access privileges determined and set (normally) by applications or DBMS/OS	Access determined for each system object (e.g. record) as per set policy	As per MAC
Data custodian	Determine access rights	Tells ICT system administrator who can see what	Tells ICT system administrator who has what role	Specify (possibly create) an appropriate profile for each user (or role with RBAC)	Use suitable profiling language to define HIAC parameters
CEO/CIO	Determine policy	Set organisation general policy	Determines types of roles to suit organisation	Define detailed access policy	Defines organisational policy sets and emergency overrides parameters using natural language
ICT system administrator	Set access rights	Directly program who sees what	As per DAC	Upload (possibly create) policy settings determined by CIO	Upload and manage HIAC profiles
Internal adversary (disgruntled employee)	User access	Can access records inappropriately or feed information to external adversary	As per DAC but more restricted access	Access limited to objects (records) as allowed by relevant policy	HIAC profiles limit violations

User role	Capability	DAC	RBAC	MAC	HIAC
External adversary (e.g. hacker)	Penetrate to obtain user access and/or set access rights	Uses Trojans/viruses, social engineering or other illicit means to gain total access	As per DAC	Cannot gain overall control: limited to social engineering (e.g. gain user password for individual's user access)	Requires infeasible levels of knowledge and covert access (further limited by dynamic risk protection mechanisms)

As indicated in Table 2, the MAC-based system can provide the ability to limit access to only legitimate authorised users. In general, the organisational security policies are defined by the CEO/CIO. Access privileges are determined by the data custodians. The HIAC profiling mechanism allows for the system administrator to configure the organisational access policies defined and determined by the CEO/CIO and the data custodian. With MAC the access privileges of all users are equally bound by the policy, not set by the discretion of the file/program owners as with DAC. The internal adversary or disgruntled employee will not be able to access health information inappropriately or even through giving unauthorised information to an external adversary. The MAC mechanism can protect the system from malicious or flawed applications which can potentially damage or destroy the system and its information. This can prevent an external adversary penetrating the system by exploiting Trojan horse attacks, viruses, malware, social engineering or other illicit means to gain total access control or to tamper with audit systems.

HIAC Implementation

HIAC Platform

For general applications, currently available products that support the MAC principles of trusted operating systems include "Red Hat Enterprise Linux (RHEL) Version 5, "Fedora Core 6", and "Sun Microsystems Solaris 10 with Trusted Extensions Software". The HIAC model exploits the privacy- and security-enhancement features of such trusted operating systems in the healthcare environment. The end result is a dedicated trusted HIS which satisfies all privacy and security requirements. To determine the practical viability of a HIAC model for HIS a demonstrator, based on the Security Enhanced Linux (SELinux) operating system with both the MAC and RBAC approaches, was created (Henricksen, Caelli & Croll 2007). SELinux is based on a flexible, fine-grained MAC architecture named Flask (NSA 2000). The HIAC model is necessarily MAC-based accompanied by RBAC properties for flexibility and a refined level of granularity. This degree of simultaneous control and flexibility is not achievable with DAC, RBAC or MAC individually.

Protection of Health Service Application Data from the Operating System Level

Redhat's SELinux enforces domain separation by 'sandboxes' known as protected zones to prevent processes and applications interfering with each other, such that an unauthorised process cannot gain overall control of the system as with DAC. For example, a sandbox in the application level can be created to protect health service applications accessing health data isolated from another sandbox for general activities allowing a Web browser to access the Internet. Unless explicitly permitted, the Web browser is not allowed to access the health data, nor is the health service application permitted to explore the Internet as the Web browser. Once an adversary attacks a DAC system through the network and manages to obtain super-user access privileges, the entire system is subverted. With SELinux however the adversary would control only a single sandbox, and would need to launch additional exploits, each of which becomes increasingly infeasible with distance from the network.

Creation of SELinux Proxy at the Application Level

A large scale HIS may involve dynamic and frequent changes to the security policies and security servers such as adding/deleting users and applications. Once the request for the change is made, the SELinux policy needs to be modified and the security server is required to be recompiled manually. In order to provide the minimum of disruption to the system operation and avoid creating additional complex interactions between application and operating system level objects, a proxy is suggested. The proxy operates at the application level and is protected in its own sandbox by SELinux. The proxy regulates access by application-level processes to protect data, using its own set of configuration files. This solution can be seen as nested SELinux, whereby the proxy represents a micro-instance of SELinux that deals only with application data. Operating system level processes see only a monolithic object (the proxy) representing application processes, meaning that the number of configuration rules between the two layers is linear rather than exponential.

Proxy Operation

The operation of the proxy mirrors the SELinux mechanism. SELinux separates the policy decision-making logic unit from the policy enforcement logic unit, as shown in Figure 1. For example, a subject X requests to access an object Y in the system. The policy enforcement server unit queries the security server unit for making

an access decision. The security server unit makes the access decision based upon Y's security class and X's security attribute from the security policy database.   The access decision is made by the security server and is then relayed to the policy enforcement server unit.

In the proxy model, a client interacts with the proxy via a pair of Client and Server messages. For each client message received, the proxy sends exactly one server message. In the client message, the client authenticates itself to the proxy with its credentials. Until the next such message is received, the proxy caches the credentials. This mimics the SELinux mechanism, which authenticates a user via a password before transitioning the user into the requested role. The proxy responds to the credentialled message.   The credentials are evaluated whenever the client requests access to a record in the proxy database. The proxy passes the credentials with the record identifier and the policy to the security filter. The security filter assesses the credentials, decides whether the record can be accessed in the way intended and passes this decision to the proxy. Whereas SELinux can protect data to the granularity of the file, the proxy has arbitrary granularity, as determined by tags exchanged between the proxy and its client. The client may wish to retrieve a single word from a database, or an entire collection of files. Our mechanism allows this with as little as a single configuration, although for more complex cases, the number of configuration rules will increase linearly in the number of database items.
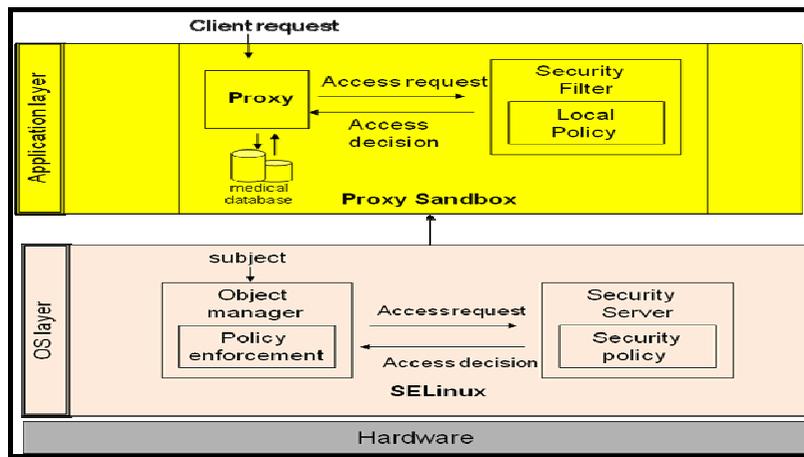

Figure 1: Proxy operation

There are some cases when records must be accessible even in the absence of legitimate credentials. For example, if the authorised viewer of a patient's case file is not present, but the patient requires emergency treatment, then the availability of the information is more important than its privacy. Thus, the proxy is programmed to respond to a special role of 'Emergency', in which case it moves into auditing mode, until a new set of credentials with a differing role is provided. In auditing mode, all records can be retrieved and modified, but each action is recorded and flagged for review by the security administrator. Appropriate punishment for abusing this mode can be meted out at a social level. Our prototype does not handle differential records, whereby the differences between subsequent versions of records are stored, although this would be advantageous for malicious or accidental modification of records in auditing mode.

Although the proxy significantly simplifies configuration of application data, it does not address problems at the operating-system level that need to be resolved. Further research in this area needs to focus on simplifying the generic SELinux configuration, to allow realistic deployment of "strict" SELinux, which supports protection of application data. This is indeed happening, as witnessed by the development of modular policy logic in Fedora Core 5, which allows the configuration to be developed and loaded in blocks relating to the processes or daemons being protected. The efficacy of this strategy has yet to be solidly determined.

HIAC Features

HIAC incorporates RBAC which complements contemporary MAC systems by ensuring more flexibility over the more traditional MAC standalone systems. In practice this approach gives more flexibility than in the traditional MAC where accesses are granted to individual persons. The proxy model also includes the extended RBAC model with the function of inheritance of permissions with a role hierarchy, so that the policy configuration can be simplified through the use of role inheritance within hierarchies. The HIAC model includes the principle of least privilege and also enforces domain separation through the use of sandboxes within Redhat's SELinux. These help prevent applications interfering with each other such that an unauthorised user cannot gain overall control of the system as with DAC.

To date Australian privacy laws and health-related privacy legislation prescribe no particular technology to protect personal information. For instance under the Information Privacy Principles (IPP) of the Privacy Act

1988 (Principle 4 – Storage and security of personal information) Principle 4 (a) requires an organisation to take reasonable steps to protect personal information. The National Privacy Principles (NPP) in the Privacy Amendment (Private Sector) Act 2000 also requires a record-keeper to protect personal information by security safeguards as is reasonable. No specific security mechanisms are specified in both the IPP and NPP, thus any reasonable and adequate security measures are allowed for protecting personal information. The HIAC structure enables an effective safeguard strategy for the protection of the confidentiality of individual health information to assist the healthcare industry to comply with Australian privacy legislative and regulatory requirements. Australia's privacy regime is currently under review. This research will continue to observe the update of privacy and e-health privacy legislation in Australia, in order to design the OTHIS architecture for legal compliance.

In general HIAC provides for maximum flexibility within a strongly secure environment. This means that it provides the potential for achieving a balance between security needs and flexibility of implementation, which is primarily determined from a privacy risk assessment. For example HIAC provides the flexibility of having timely access control to assist information resources with an emergency override function by switching to the emergency policy in emergency circumstances. Full auditing of the system deters potential abuses of this flexibility. A major area for future research concerns the simplification of the MAC profile definition. At present the methods and processes needed to define and deploy a mandatory security policy within an overall HIS are complex and could be considered to be beyond the expertise level of many CIO in health related organisations. Integration of such security profiling structures is required in relation to such other enterprise systems as overall human resource management systems and the like. This allows for definition and deployment of security policies that represent legal, regulatory, policy and enterprise level requirements for reliable and consistent enforcement at the computer system level. This future research requires the definition and implementation of appropriate interfaces between such large scale enterprise systems and the proposed HIAC structure.

Protection and Enforcement using Cryptography in OTHIS

Cryptographic technologies have long been used for integrity and confidentiality purposes. Large numbers of security-related tools use encryption to protect sensitive information, particularly to maintain privacy. It is important to understand that the principle role of cryptography is to ensure the quality of service of the technology, and thus ensure that the technology satisfies the business requirements of the system. Cryptography then is primarily an enabler of services. Detection and prevention of security breaches is a subset of this primary function. For integrity, a "keyed hash function" may be applied to each relevant data record to prevent unauthorised insertion of records as well as unauthorised alteration of existing records. An unauthorised third party (or an authorised party extending beyond their authorisation) would need to possess the necessary key to either create or recreate the integrity enforcing checksum, commonly referred as a "message authentication code (MAC)". Confidentiality can be enforced using a single-key cipher, but key management structures to allow for multiple roles to have access to a healthcare record would be necessarily complex. As such, maintaining record confidentiality using public key cipher schemes may be advantageous. Historically with this approach, a performance penalty may have been involved, but with current hardware bases for the implementation of these ciphers, such performance problems are normally minimal.

Our research intends to investigate the use of suitable cryptographic techniques embedded into the OTHIS architecture for protecting confidentiality and security of personal health data. Encryption should be used, and normally is used, to protect data in transit for complete end-to-end protection, including within the node systems at each end of a connection. Data in storage should also be encrypted for end-point security against unauthorised or accidental access or eavesdropping. Identity-based encryption mechanisms may be used for identity and/or role management in the healthcare environment. An assessment of suitable cryptographic services and mechanisms for the healthcare sector will be undertaken. Cryptographic integration in the UK, USA and New Zealand healthcare sectors will be investigated. A particularly relevant contribution envisaged from protection and enforcement using cryptography in OTHIS is the elucidation of the requirements for the integration and management of such cryptographic systems in OTHIS for enforcement of privacy and security of electronic health data.

Conclusion

Our research indicates that an overall trusted HIS should implement security at all levels of its architecture to ensure the protection of personal privacy and security of electronic health information. From an information security perspective, we propose OTHIS for the overall HIS architecture. This paper relates an HIS architecture to two internationally recognised standards, the OSI reference model and the TCP/IP model, to describe how OTHIS fits into these reference models in a HIS. A development of the OTHIS architecture comprises a number of modules with viable and suitable security mechanisms to achieve a high level of security, including the HIAC

model. HIAC is a trustworthy access control mechanism to provide the privacy and security of personal health data at the levels of health service application, DBMS, middleware, network control system and operating systems in HIS. HIAC is proposed as a viable solution which has the potential to address the common types of information privacy violations and weaknesses illustrated by the recent access control management scenarios from Australia, the UK and the USA.

This paper contends that it is both timely and desirable to move electronic HIS towards privacy- and security-aware applications that reside atop a trusted computing-based operating system. Such systems have the real-world potential to satisfy all stakeholder requirements including modern information structures, organisational policies, legislative and regulatory requirements for both healthcare providers and healthcare consumers (privacy and security), and flexible operational demands in HIS. This paper emphasises the need for well-directed research into the application of inherent privacy- and security-enhanced operating systems to provide viable, real-world trusted HIS. The authors propose an HIAC model which has the potential to fulfil these requirements. Future work will be continuing on the development of the other modules within the proposed OTHIS structure with the ultimate goals of maximum sustainability, flexibility, performance, manageability, ease of use and understanding, scalability and legal compliance included in the healthcare environment.

References

Bell, D.E. & LaPadula, L.J. 1973, *Secure Computer Systems: Mathematical Foundations and Model*, The Mitre Corporation.

Ferraiolo, D.F., Kuhn, D.R. & Chandramouli, R. 2003, *Role-Based Access Control*, Artech House, Boston.London.

GAO 2006, *Information Security: Department of Health and Human Services Needs to Fully Implement Its Program*, United States Government Accountability Office, viewed 20/11/2006 <http://www.gao.gov/new.items/d06267.pdf>.

---- 2007, *Health Information Technology Early Efforts Initiated but Comprehensive Privacy Approach Needed for National Strategy*, United States Government Accountability Office, viewed 10/05/2007 <http://www.gao.gov/new.items/d07237.pdf>.

Gasser, M. 1988, *Building a Secure Computer System*, Van Nostrand Reinhold, New York.

Goldschmidt, P.G. 2005, 'HIT and MIS: Implications of Health Information Technology and Medical Information Systems', *Communications of the ACM*, vol. 48, no. 10, pp. 69-74.

Henricksen, M., Caelli, W. & Croll, P.R. 2007, 'Securing Grid Data Using Mandatory Access Controls', paper presented to 5th Australian Symposium on Grid Computing and e-Research (AusGrid 2007), Ballarat Australia.

Leigh, D. & Evans, R. 2006, 'Warning over privacy of 50m patient files', *Guardian News and Media Limited*, 01/11/2006, <http://society.guardian.co.uk/health/news/0,,1936403,00.html>.

Loscocco, P. & Smalley, S. 2001, 'Integrating Flexible Support for Security Policies into the Linux Operating System', paper presented to Proceedings of the FREENIX Track: 2001 USENIX Annual Technical Conference(FREENIX '01).

Microsoft 2006, *The Road to Security*, Microsoft Corporation, viewed 28/11/2006 <http://www.microsoft.com/resources/ngscb/default.mspx>.

NHS 2005, *"Sealed Envelopes" Briefing Paper Draft*, National Health Services, viewed 03/11/2006 <http://www.ardenhoe.demon.co.uk/privacy/Sealed%20Envelopes%20briefing%20paper.pdf>.

NSA 2000, *Security Enhanced Linux*, National Security Agency, viewed 20/1/2007 <http://www.nsa.gov/selinux/>.

Quinn, J. 2004, *Lessons from the UK EMR: Not Exactly Apples to Apples*, HealthLeaders Inc.,, viewed 17/08/2005 <http://www.healthleaders.com/news/print.php?contentid=60316>.

Reid, J., Cheong, I., Henricksen, M. & Smith, J. 2003, 'A Novel Use of RBAC to Protect Privacy in Distributed Health Care Information Systems', paper presented to Information Security and Privacy, 8th Australasian Conference, ACISP, Wollongong, Australia.

Sharanahan, D. & Karvelas, P. 2006, 'Welfare workers axed for spying', *The Australian*, 23/08/2006, <http://www.theaustralian.news.com.au/story/0,20867,20223075-601,00.html>.

Copyright