

Security and Privacy System Architecture for an e-Hospital Environment

Kathryn Garson
School of Information Technology and
Engineering (SITE)
University of Ottawa
Ottawa, Ontario, Canada K1N 6N5
kgars062@uottawa.ca

Carlisle Adams
School of Information Technology and
Engineering (SITE)
University of Ottawa
Ottawa, Ontario, Canada K1N 6N5
cadams@site.uottawa.ca

ABSTRACT

Hospitals are now using electronic medical records and computer applications in order to provide more efficient and thorough care for their patients. The Mobile Emergency Triage system provides doctors with decision support for emergency care by pulling information from a patient's health record and a medical literature database. In order to achieve compliance with privacy legislations PIPEDA and PHIPA, security and privacy measures must be put in place. Encryption and access control are necessary for ensuring proper authorization and confidentiality for patient records. Strong authentication and audit logs are required to ensure access only by those allowed. We discuss differences in security technologies and detail the ones used in our MET system. A new encryption technology called policy-based encryption proves to be quite useful within a health care environment for providing both encryption and access control. We propose an extension to an existing scheme which allows for the use of this cryptography in a hospital setting.

Categories and Terms:

D.4.6 [Operating Systems]: Security and protection
K.6.5 [Management of Computing and Information Systems]: Security and protection
E.3 Data Encryption

1. INTRODUCTION

Many hospitals are moving away from paper-based medical records to use electronic health care records. Specialized software and electronic diagnostic tools are offering a new level of patient care. The move towards electronic based systems provides streamlined automated processes and specific applications that can help doctors with diagnosis and treatment of patients. The introduction of these technologies raises privacy risks with regards to patient information. A malicious person trying to compromise many patient records will be able to collect large amounts of data easily if these records are available electronically.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IDtrust '08, March 4-6, 2008 Gaithersburg, MD
Copyright 2008 ACM 978-1-60558-066-1...\$5.00

The Mobile Emergency Triage (MET) project provides doctors with decision support for triage, diagnosis, and treatment of patients in an emergency room setting. The project will be implemented first as a trial and then a full production version in the emergency room of a hospital in Ottawa, Canada. Doctors will be using a tablet PC with the MET software to help them in their tasks. The MET software pulls information from a collection of medical literature and electronic health records (EHRs) to provide doctors with evidence-based decision support. The agent-based system is interactive allowing the doctor to input symptoms and view possible diagnosis and treatment options. The doctor can make treatment decisions based on this information or enter additional data and receive different information.

The MET system is set up on a wireless network giving doctors the ability to travel patient to patient while having full access to the software. This system requires security and privacy technologies to prevent malicious users from having access to sensitive patient data. Patients' EHRs are transmitted over a wireless network to these devices and then stored on the PC. Proper access control needs to be put in place to ensure only authorized users can have access to records. We will discuss in this paper the steps we have taken to ensure privacy of patient's medical records. This project has greatly benefited from the input of a number of disciplines while working on the system. Management, computer science, engineering and medical science professionals have been working on this project together. Getting opinions and ideas these diverse backgrounds has been a great opportunity.

The goal of this paper is to present the technologies to be used with the MET project to add security and privacy functionality. We also discuss our motivations for adding security to this project based on privacy laws in Canada. Different alternatives will be discussed for their advantages or disadvantages in this setting. One of the most promising technologies for use in a health care environment is policy-based cryptography for access control. We discuss this concept and propose an extension to an existing scheme to add more flexibility for use in our environment.

Organization of paper

We first talk about the current environment of the hospital setting and what changes will happen with the introduction of electronic medical records. In section 3 we present our goals for securing our system and providing privacy measures. Then we compare access control and encryption methods available to satisfy these goals in section 4. We discuss a policy-based encryption scheme that works naturally in the health care setting. In section 5 we

compare authentication mechanisms for use with our system. Then we give some details of other security issues we faced and how we solved them in section 6. We give a brief overview of the system architecture and how each privacy module will interact within the MET system in section 7. We discuss related work in section 8 and conclude in section 9.

2. HOSPITAL ENVIRONMENT

The hospital environment is unique in both the fact that it contains such highly sensitive data and this data is required in emergency situations. Emergency situations may overshadow the need for privacy procedures. Doctor's main functions are to treat patients, not to follow complicated steps for accessing data. The workflow in the hospital must be taken into account when designing the system. Our goal should be to provide security for the MET software while minimizing tasks required by staff for using it.

Our trial will run in a hospital that currently uses a paper based system for medical records and is migrating towards electronic records. The paper medical record consists of documents and reports pertaining to one patient. It could include patient admission information, medical history, diagnosis reports, and lab reports from the current stay as well as previous hospital visits. Some parts of the medical record, such as lab reports which are available in electronic form, are printed and stored in the medical record file. These paper based records are considered the real authentic record even if the electronic form is still in existence.

The electronic system for accessing the available electronic documents is currently a read-only system since reports are not modified electronically. Employees sign in with their username and password to the software system on shared computers. The shared computers are in areas of high traffic and usually are already logged in to a general account. Each staff member who wants to access software or records will sign in to that software with their own information. Access rights are basic in that everyone who uses the system can read all records. There are groupings of staff and these groupings determine read and write access rights to documents in medical records. When the hospital moves to a full electronic system, these groupings will be defined formally in policies. Currently staff members only have access to the network from inside the building, but are moving towards remote login. This will not be a part of our initial trial but may be something to consider for the long term.

The procedure for obtaining a paper medical record is straightforward. A doctor can place a call to the records department requesting a document. Porters or other employees will bring the requested documents to their department. Records are assigned either to one doctor or one area of the hospital specifically, and are returned by the end of the day. If moved or passed on, medical records should be notified to be able to track the documents. Generally if a medical record stays in one area of the hospital is not necessary to notify them, since many doctors may see one patient during their time in that area.

3. PRIVACY GOALS

In Canada, the Personal Information Protection and Electronic Documents Act (PIPEDA) describes guidelines for the protection of personal electronic information^[15]. PIPEDA provides guidelines for handling personal information in electronic form.

The Personal Health Information Protection Act (PHIPA) is privacy legislation in the province of Ontario which builds on PIPEDA^[14].

Principle 7 of PIPEDA describes the safeguards that should be in place to protect sensitive data. Here we highlight the aspects that pertain to our project:

1. **Sensitive information should be protected with a higher level of security:** Medical records are always considered to contain highly sensitive data. For our project we need to ensure that all patient data is secured by the best available methods. The sensitive information needs to be protected from unauthorized users during both transmission and storage.
2. **Methods of protection should include technological methods such as the use of passwords and encryption:** We should investigate different encryption and authentication methods to find which would be most suitable for a health care environment such as the one for the MET project. The most secure password scheme may be great for privacy but may not be feasible in a health care emergency environment. We need to find technological solutions that are practical.
3. **Limit access to a 'need to know' basis:** Access control methods need to be used. We can limit access to individual documents of a medical record for both read and write permissions. Employee roles within the hospital and permissions associated to those roles can be used in determining a good access control model.
4. Both PHIPA and PIPEDA specify that medical records should be **protected from loss, theft, unauthorized access, disclosure, copying, use, or modification** regardless of what format they are in. Electronic records must be protected by access control and include an integrity mechanism. Furthermore, we must have an audit functionality to ensure that no malicious use of the system goes unnoticed.

These privacy guidelines are not restricted to Canada. In the US, title II of the Health Insurance Portability and Accountability Act (HIPAA 1996) outlines standards for security and privacy of patient health records. They encourage the use of electronic data interchange in the health care system while protecting information. Physical and technical safeguards similar to the PIPEDA safeguards are outlined. Similarly in Europe they have the EU Data Protection Directive (1995).

Our goals for the MET project are to satisfy these privacy guidelines using the appropriate technology. We plan to use encryption to secure all data transmission and storage of patient records. We want to enforce read and write access control not only for a medical record but for individual documents within that record. The security measures put in place should include end-to-end encryption of patient records, strong authentication, authorization, data integrity, and audit logs.

Another privacy goal specified by the hospital staff is to make sure none of these records leave the hospital premises. With the paper based system, paper records are not to be removed from the

hospital. Because in our case the tablet PCs may leave the hospital, we don't want any records that are currently on the PC to be able to leave with it. We will investigate the options we decided on for dealing with this privacy issue.

4. ENCRYPTION AND ACCESS CONTROL

Encryption is necessary in our system for protecting sensitive data such as patient records. We want this data to be stored encrypted and transmitted encrypted so that no one sniffing the network can get access to the data (particularly since a portion of the network is wireless). There are many options for securing data using encryption such as network encryption using SSL and data encryption using public key cryptography. We will discuss the options we considered as well as propose a policy-based encryption solution.

4.1 Network Encryption & Access Control

In complying with PIPEDA we want to ensure all transmissions with sensitive data are encrypted. The Wired Equivalent Privacy (WEP) protocol is intended to provide confidentiality on wireless networks however many weaknesses have been found that could lead to a relatively easy attack^[4]. It is not considered to be secure against any more than a casual eavesdropper. For our purposes we need more than this in order to comply with our first privacy goal of protecting sensitive data with a high level of security.

A Virtual Private Network (VPN) is another consideration for securing our wireless network. IPsec and SSL can each be used to implement a VPN. IPsec has been notoriously complex to configure and manage and we will not consider it for this implementation^[2]. SSL meanwhile has been used as a much easier alternative and has implementations that have proven to be secure.

SSL/TLS uses a handshake protocol to establish shared keys between a client and server. All communications between client and server use the shared key to encrypt data. This creates an 'encrypted channel' between the client and server. SSL provides confidentiality and data integrity through cryptography. The handshake protocol allows for authentication of the server, so the client is assured of a secure connection with the proper server. OpenVPN uses SSL/TLS technology and has been shown to offer the SSL security properties of authentication, confidentiality, and data integrity^[10]. Client software for the VPN would be set up on all tablet PCs and server software on MET servers. Users accessing the network would do so by logging into the VPN software.

Encrypting all transmissions on the network will secure data from anyone who is not logged onto the network. It is still necessary to implement an access control system to manage permissions and access to patient data for employees who access the software. Role-Based Access Control (RBAC) allows for controlling which users have access to which data on a network^[6]. Users are assigned one or more roles and must authenticate to the system when requesting access to a resource. Each role has permissions assigned to it, dictating which resources are available to that role. We can express read or write access for documents based on the roles of employees using an RBAC system. Policies would need to be created from the employee groupings and rules about which documents they should have access to. Changes in policies would

have to be reflected by changing the permissions for the affected roles.

Role Based Access Control will provide protection from unauthorized access. It would also allow us to limit information disclosure to a strictly 'need to know basis'. Coupled with network encryption this would cover most of our privacy goals. However this may be more work than we need to make our system secure as we will discuss next.

4.2 Encryption Only

It may be unnecessary to employ both encryption and access control as separate technologies in our system. We present here options we considered that combine encryption and access control functionality.

Traditional encryption methods such as public key cryptography (PKC) are cumbersome to apply to access control. In a public key system, each user is assigned a public key and private key. A document is usually encrypted for a single recipient using their specific public key. Only the recipient can decrypt to recover the original document by using their private key. Managing keys in this system has often been a limiting factor in real world applications. Keys need to be created and distributed to all users through the use of digital certificates. Users who leave the system or lose their keys need to have their certificate revoked. Certificates and keys have a finite lifetime and need to be renewed. Old keys however still need to be kept in order to be able to decrypt documents that are encrypted under it for as long as the document's lifetime.

For our system, we need a way to encrypt medical records for access by multiple recipients, who are not necessarily known at encryption time. PKI doesn't offer this flexibility on its own; intended recipients are known and their keys are used in the encryption process. If we wanted to use PKI, we would have to add access control methods to allow for managing multiple recipients. We could then use roles to control access control and public keys to provide encryption when transmitting^[19]. However this adds unnecessary complexity to the system. It would be easier to encrypt all transmissions on the network and use traditional access control methods. Thus PKI does not offer a feasible solution for this case.

4.3 Identity/Role-Based Encryption

Identity-based encryption (IBE) potentially offers more flexibility for our environment than PKI. The main idea behind identity-based encryption is that any string can be used as an encryption key^[17]. For example, a person's unique email address can be used. A document can be encrypted with the recipient's email address. The recipient must identify themselves to a trusted authority to receive the decryption key to recover the original document. This is often how this scheme is described in order to compare it with PKI in sending an encrypted document to a single known recipient. IBE reduces the need for key management, as a user's public key is a well known unique string such as their email address. The private key is obtained by authenticating to a trusted authority (the Private Key Generator, PKG), and so the user doesn't need to keep keys. Managing user accounts becomes easier. When a user leaves the system, they no longer have access

to decrypt because they can't log in to the system to authenticate (e.g., because their password or account has been disabled).

Advantages of this encryption scheme include users not having to manage their own keys and no need for key certificates. From a usability point of view this is ideal. Doctors using our system won't have to worry about details of encryption. When they need to access a document, they will simply identify themselves by logging into the system. However, we will not be encrypting for one identity as mentioned earlier. For our application we need to encrypt for multiple people.

Using a more general approach, rather than encrypting on a unique string such as an email address, we can encrypt for a general grouping such as a role. This allows for multiple people to get access if they belong to that role. But what if we want to control access for multiple roles? For example a doctor and nurse may have access to a patient's record but not administrative personnel. So we want to encrypt for doctors and nurses. Again further generalizing this approach gives us a more flexible solution. Encrypting based on a policy can allow a document to be encrypted for access by multiple roles.

4.4 Policy Based Encryption

A policy-based encryption scheme offers the greatest flexibility for our security needs. The approach is relatively simple and builds on the idea of encrypting under an arbitrary string. A document is encrypted under a policy which is a combination of rules. Note that in IBE, if the public key can be an arbitrary string, then this string need not be an identity. Rather, the string may be a complete access control policy (or a hash of that policy) for the document that is to be encrypted. A user wishing to have access to a document will authenticate by logging into the system and will obtain a decryption key associated with their role from the PKG. If the user's role satisfies the requirements of the policy, their decryption key will decrypt the document.

The policies can be as simple as combining a few roles or more complex to include other rules such as time constraints. In our implementation for the MET project, a small number of simple policies will be created. Groupings are clearly assigned in the workplace already so policies would be able to be created based on this information. Each document type will have an associated access policy. An example policy could be "doctors or nurses can have read access for lab reports".

Corresponding keys for the decryption process will be created based on a user's role. When a user logs in, a trusted authority (TA) will authenticate them and provide the decryption key associated with the user's role. Then, when the user makes requests for records, if the user's role fits the policy their decryption key will decrypt the document and they will have access.

Keeping usability in mind, it would be favorable to automate the encryption process. Staff should not be asked to enter a policy or choose from a list of policies every time they create a new document. Because of the nature of the hospital setting, we can make policies dependant on the type of document. If a document of a certain type is entered in the system, it will be encrypted based on a corresponding policy. This is possible because of the

finite number of document types in a hospital setting and the access rules for these document types. For example all x-ray lab reports are available to be read by doctors. Therefore these documents can be encrypted under a policy specifically for that type. This will also ensure that we have a finite number of policies to manage.

If a policy is changed or updated, then all documents encrypted under that policy will also have to be updated. This can be done automatically by decrypting and then encrypting under the new policy. It will not affect staff having access since they will receive the updated key when they try to get access. If they have a document open already, then the update won't affect them. If they modify and save the document, it will be encrypted under the new available policy. In our system, the database will have access to decryption keys as well as policies to encrypt under. This will allow for indexing of records which are stored encrypted under the policies. Keys are not private to individual users but rather keys are distributed based on user permissions. Thus, allowing the database to use these keys is not compromising individual users' private keys.

In traditional PKI we have revocation of keys which allows for users to leave the system and to manage compromised keys. In our system, we would have two scenarios that may require key revocation. If a staff member leaves or their role is changed, and if a key is compromised. For the first scenario, if the user's role is changed this will be reflected their account and they will receive the corresponding decryption key from the PKG. If the staff member no longer works at the hospital, their account will be disabled and they won't be able to log in, and thus won't be able to have access to the system. Therefore, key revocation is not needed for these types of changes to accounts. However we do need revocation in the case of a key compromise.

Boneh and Franklin propose adding a time period to keys in order to handle this situation^[5]. Since the policies and keys used in our system are strings, we can also associate them with a time period. Adding a time period to the policy can be done by simply adding another constraint in conjunction with other rules. Decryption keys generated during that time period will contain the proper time to fit the policy. Updating a policy's time period will disable keys which have been acquired and saved from previous decryptions. Users who are using the system properly will be receiving the proper key each time they make a request to have access to a document and so will not need to keep track of updating keys. In the event of key compromise, updating the policy (and updating records in batches) is sufficient to disable the key.

Molva and Bagga propose a policy-based encryption system that looks promising for use in the MET system^[3]. In their encryption scheme a document is encrypted under a policy which is a combination of conditions or rules. Let's consider an easy example of a policy that would be used in our system. An employee with the role of doctor or nurse can read a document. This policy *pol* will specify an action *act* of reading a resource *res* which is a document. In Molva and Bagga's encryption system the document would be encrypted under the following policy

$pol = \langle \text{Doctor, role} \rangle \text{ OR } \langle \text{Nurse, role} \rangle$

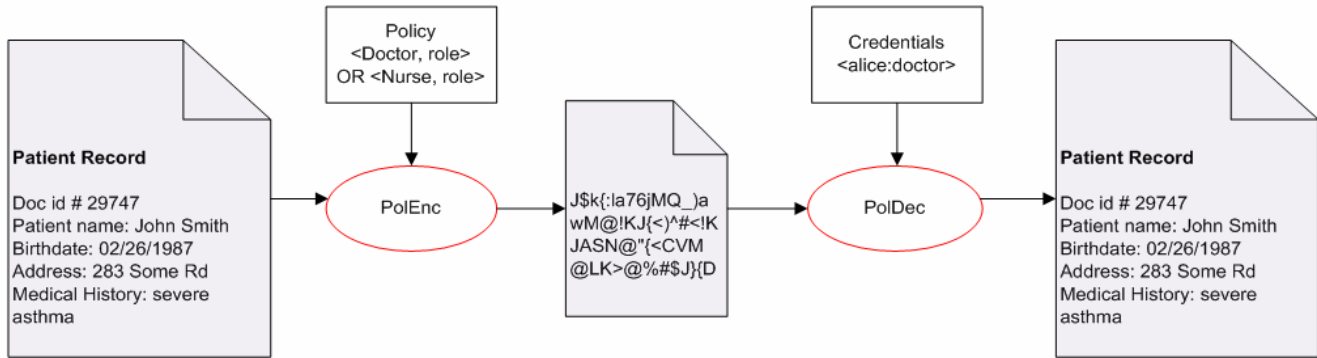


Figure 1. Policy-Based Encryption

The document res will be encrypted using policy pol as follows

$$c = \text{PolEnc}(res, pol)$$

A user can decrypt c by providing their credentials $cred = (alice:doctor)$ that satisfy the policy pol

$$res = \text{PolDec}(c, pol, cred)$$

See Figure 1 for a diagram demonstrating the encryption and decryption process. A brief description of the way their encryption works is that each set of conjunctions is assigned a mask. Each disjunction is assigned a random key value. Then, each key value is encrypted by each mask. A user who satisfies any of the sets of conjunctions will be able to retrieve each key value, and decrypt the entire document.

Their scheme works well for a policy that specifies one action for a specific resource. In our case, we would need for the policy to be applied to more than one action. For example if it is a doctor asking for access to a patient's record then they can read and write, but if it's a nurse then they can only read. Molva and Bagga's scheme allows for policies that are made of combinations of rules, but don't include information about the action allowed on the resource for those who fit the policy.

Therefore, we add expressive functionality to their system by allowing actions to be specified for each conjunction. Since the idea is to be able to encrypt under any arbitrary string, it seems fairly reasonable to say that we can add this information to the policy to encrypt under. A policy could therefore be represented as

$$pol2 = \langle \text{Doctor, role} \rangle \text{ AND } \langle \text{Write, action} \rangle$$

Using default values for actions, all users would be able to perform these actions. However, a user would only be allowed to perform the action if they also satisfy the other rules in the conjunction. In the example of $pol2$, upon presenting a credential for the role of a doctor, a user could decrypt and have write access. The action rule will signify what permissions the user has for that particular document. This will result in treating keys for each credential in a different way, allowing default key values for the actions.

The efficiency of the Molva and Bagga encryption system depends on the number of conditions that are combined to form the policy. The complexity increases with the number of computations required to encrypt for each rule. For our

application, we plan to have few rules combined in a very simple manner. We will also have limited number of credentials for users, as their roles will be the credentials and we have a limited number of actions possible. They also state ways of maximizing the efficiency of encrypting by pre-computing and caching some values. Decryption is more efficient than encryption in their scheme and would cut down on the computation time in the overall system.

By extending their scheme to allow for more expressive policies, we can adapt this system for use in our MET project. Using simple policies that specify actions allowed on the document encrypted under that policy we strive to keep the encryption process efficient. We also have minimal number of credentials for users based on their roles in the hospital. Finally, we can automate the encryption process by specifying policies for each document type, thus allowing staff members to enter documents in the system without worrying about encryption details.

4.5 Network Encryption & Access Control vs. Policy-Based Encryption

4.5.1 Access Control with Network Encryption

Access control and network encryption using a VPN plus a role-based access control mechanism would allow us to store the records in readable format. This has the advantage that for storing over long term, records won't be made obsolete from out of date encryption methods. However, this allows anyone who gets physical access to the servers storing the records access to the records in readable format. If they were encrypted then the attacker would have no advantage by getting the encrypted format records. Encrypting the database separately requires an added step, and each document would require decryption and re-encryption under SSL for transmission on the network. The policy-based encryption would be a simpler solution for providing complete encryption and access control at the same time.

4.5.2 Policy-Based Encryption

The policy-based encryption method has the advantage that encryption and access control are in one package. Not only would it provide for encryption during transmission but also for storage. Records will be transmitted to the devices encrypted, and be decrypted once on the PC. This ensures all sensitive information being transmitted is secure. The encryption scheme also allows us to implement access control as only those who satisfy the policy can decrypt and have access to the record. There is no need for users to manage their own keys nor is there need to distribute

keys. There will be no difference for the user experience between the two choices. In either case, the user will log in and will have access to documents for which they have permission.

A general disadvantage of any encryption scheme is the fact that the records are stored encrypted. The decryption keys must be kept for the lifetime of the record. The decryption algorithm must also be available to recover the original documents. In order to ensure these are both available, it may be necessary to keep a backup of the keys and algorithm. Or it would be possible to store the records in readable format somewhere physically secure and not connected to a network. We would need to store backup copies of the records in plaintext regardless of which system we use, due to possibilities of hardware failure.

Overall, the policy-based encryption method provides the most compliance with PIPEDA regulations with the least amount of complexity. We get a high level of security for the sensitive data. The data is protected by encryption while stored and transmitted protecting it from malicious people who may be eavesdropping. The records are also protected from unauthorized access as only users of the system with the proper account permissions set up are given a decryption key. We have limited the access on a 'need to know' basis by forming policies constraining access to documents based on roles. The users of the system will not have to know about the security technologies used and their only security task will be to log in as they already do on other hospital computer systems.

5. AUTHENTICATION

The security of the encryption system relies on strong authentication. The system is only as secure as its ability to prevent unauthorized access. Currently in the hospital, the standard username and password method is used without any restrictions on password choice. We wanted to explore authentication options that would be both secure and usable so as not to change the user experience too much.

The username and password combination is the most popular authentication method. There are many reasons why this is not a secure method^[1]. Users will choose very easy to remember passwords that are also easy to break. If we imposed restrictions on the password choice to force users to have stronger passwords, they would do what they could to make logging in easier, by writing them down for instance. We have already had feedback from doctors who would not be happy having password restrictions being put in place and encouraged us to look at other options. Passwords are subject to social engineering attacks which could be easy to manipulate in an emergency setting. If a malicious person claims there's an emergency and asks for a staff member's password, they may be more likely to divulge information.

A two-factor authentication would be desirable to provide improved security. Our motivation for finding a proper two-factor authentication mechanism lies in working with what we have. One factor will be the username and password system without restrictions on the password. The second factor could be a biometric fingerprint or RFID tag. The tablet PC is equipped with both a fingerprint reader and RFID reader. We discuss both and what we chose for the second factor in our authentication scheme.

5.1 Fingerprint Biometrics

One of the most common biometrics in use for authentication is the fingerprint. Using the fingerprint readers has an advantage that users don't need to 'remember' to bring a token with them to work. However some usability problems may make this less desirable to use. In order to provide your fingerprint for verification, a user must place their finger in a certain position. Factors such as placement, heat, cold, and perspiration can all affect how accurate the system is^[1]. We want our users to be able to log in every time because the setting is the emergency room of a real hospital.

Another factor with fingerprint biometrics is that not everyone can give the fingerprint to enroll. Fingerprints damaged by injuries could be a problem. We want to make sure everyone can use the system, including current employees and future employees. In the healthcare setting many employees may be wearing hygienic gloves which would not allow them to use the fingerprint reader without first removing them. This is a real usability problem in our setting. When discussing the fingerprint option with doctors who will be using the system, they seemed reluctant to use fingerprints in the authentication stage and wanted something easier.

5.2 RFID Reader

The RFID reader offers an alternative to the fingerprint reader. Doctors carry employee badges which can be equipped with a barcode. To sign in a doctor swipes their card (or, if it is a proximity reader, simply has their badge somewhere on their person) and provides their login password. Everyone can be given a badge and it will always be accepted. They don't have to have a high entropy password which may be hard to remember. They may be less likely to simply tell someone their password to let them have access, since they would also have to provide their card. However this option is still available if the doctor wishes to delegate his tasks to another employee for a period of time. This provides for a flexible authentication system compared to using a fingerprint which cannot be delegated. We decided to go ahead with using the doctor's badges and the RFID reader as the second factor in our authentication system.

Unlike the fingerprint method, a doctor can forget their employee badge. If they borrow someone else's pass or get a temporary one, this will not work with the reader since it will be a different pass. We could manage to have temporary badges available for the day which can be associated with their account. One way around this is to use a common backup method of asking the user a series of pre-answered questions. If the user answers correctly and provides their usual login information along with it, then the user can log in for the day. Proper tracking of these events is important for being able to notice and document malicious behavior.

6. OTHER PRIVACY CONCERNS

There are other security and privacy issues that we wanted to address for our MET project. The specific security needs arise because of the environment and physical location of the project. Some security measures we will include in our system are audit logs, integrity mechanisms, and automatic purging of sensitive files.

6.1 Audit Logs

An important privacy requirement and a feature our system must ultimately include is an audit log. Audit logs allow for tracking user's activity on the system. If by any chance someone is misusing the system, then a record of that activity must be available. For example if a doctor is accessing large amounts of patient records that clearly aren't all their patients, this would be worth investigating. It could be that a malicious user has gotten access to that account and is stealing patient information. It would be a good addition to our system to add logging functionality of all access and changes to patient records. Specifically we want to track when a user logs on and off, which records they request, and which records they make changes to.

6.2 Integrity Mechanisms

Another privacy aspect is not only protecting the data from unauthorized access but also from unauthorized modification. Ensuring proper access control for read/write permissions is essential. However we also want to include some sort of integrity mechanism to be able to check that no changes have been done maliciously. For example if someone changes a record stored in the database we need to be able to check this.

A Message Authentication Code (MAC) applied to a record would provide us with a way of checking for any changes made^[8]. We could do this on an individual record basis, creating a MAC of the entire record, storing the MAC value elsewhere on the system, and comparing the stored MAC with a freshly-computed MAC for every record read event.

6.3 Purging Files

Consider the scenario where a doctor has their tablet PC which they bring home everyday. At the end of their shift, they may have seen a number of patients and have their data saved on the PC. The doctor goes home and someone in their household uses the PC for other purposes. It could be connected to the Internet at this point and anyone with access to the PC could have access to the sensitive records.

We need a system of preventing records from leaving the hospital which is facilitated by the use of the tablet PCs. A doctor may not notice or remember that he still has files on his PC when leaving the hospital. It's much more obvious with paper-based records, where a doctor has to actually carry them out or knowingly put them in a bag. If it becomes a habit of taking the tablet PC home, then the doctor may not remember to erase patient data each time. This means we need to know when a tablet PC is being taken out of the hospital so that we can erase those files that haven't been deleted.

One simple thing to do is to purge all files on shutdown. However sometimes PCs are not shut down properly especially in the case of a tablet PC which may lose battery power. Therefore it would be better to implement on system startup. Each time the laptop is booted, the files are purged. Since doctors wouldn't be turning it off and on all day, they shouldn't lose any records until they are done their work. We don't need to implement this on hibernate states, so that if a doctor leaves his PC for a while with no activity then he'll still have the information he needs on it.

If a tablet PC leaves the hospital while still turned on, we still need to purge the files. In this case we can't rely simply on the

files being purged when the PC is turned off or rebooted. The connection to the wireless network can be used as an event trigger to purge all files. When the tablet PC is taken out of range and no longer has a connection with the network, the patient data files will be erased. A doctor who is in the wireless network area doing his work will not lose any data. It would only affect tablet PCs that are taken far enough away from the hospital department with the network connection available.

7. SYSTEM ARCHITECTURE

The tablet PCs being used in the trial are the Motion Computing C5 models^[13]. The PCs are equipped with a fingerprint reader and RFID scanner. The devices will ensure that doctors are able to travel patient to patient while having access to the information they need on the wireless network.

The MET software and privacy functionality will be implemented in Java using the JADE development environment. The system is agent-based with multiple agents each assigned a specific task in the software. A central temporary storage area, called a blackboard, will be used to store patient records and medical information that agents are using. Adding an encryption agent seemed reasonable but has not become a practical solution. All agents will need access to encryption services for pulling and pushing information to and from the databases, the blackboard, and the tablet PC. Multiple encryption agents would be needed for the servers, the tablet PC, and the databases. Integrating the security functions as a layer, or set of services, available to all agents in the system proved a better solution. The MET architecture is reflected in figure 2 below. The security and encryption services layer will provide encryption, decryption, authentication, and account management services. In addition, a monitoring agent on the tablet PC will be necessary to track when to purge the files in the event of a disconnection from the network or system reboot.

8. RELATED WORK

8.1 Encryption and Access Control

Role-based access control systems using public key cryptography have been proposed. Wilkinson, Hearn, and Wiseman describe an access control system that uses encryption to control access to documents^[19]. Documents are encrypted under a group's public key, and members of that group can decrypt with the group public key. They also describe how symmetric keys could be used as the group keys, however they conclude that asymmetric cryptography will provide better protection from key compromise at proxies. While the scheme does provide security measures that we are looking for, it still suffers from the key management issues of public key cryptography schemes. It also doesn't seem straightforward how to encrypt a document for multiple groups or multiple roles.

Kapadia, Tsang and Smith propose an attribute-based encryption system that allows for role based access control^[9]. Their system relies on hidden credentials and policies. In our case, the policies will most likely be public since there are a finite set of policies created based on well-known rules. It would be better to have a system that doesn't rely on having secret policies. Other attribute-based encryption schemes aren't as efficient as the Molva and Bagga scheme discussed next^[7].




Figure 2. MET Architecture

Most policy-based encryption schemes that have been proposed are based on the Boneh-Franklin ID-based encryption scheme using bilinear pairings over elliptic curves^[5]. Smart proposed a scheme extending this for encrypting on multiple identities^[18]. Molva and Bagga further extend their work to propose a policy-based cryptography scheme including an encryption scheme and signature scheme^[3]. They propose using a policy as a public key to encrypt a document. A user obtains their decryption key based on their credentials and can decrypt if these credentials satisfy the policy rules. As mentioned earlier their scheme is a good basis for us to extend our work on.

8.2 Examples of Privacy Technologies Used in Health Care Environments

Voltage is great example of a company using the technologies discussed here for security and privacy solutions in health care environments. They offer an identity-based encryption for email messaging that is currently being used in hospitals in the United States^[20]. Similarly, Secure Computing offers policy-based cryptography products. They implemented a token based authentication system with audit logs to ensure HIPAA compliance for a system in a hospital^[16].

Mont, Bramhall, and Harisson from the Hewlett Packard Lab in Bristol, UK have developed a messaging service using identity-based cryptography for a hospital^[12]. Their scheme uses the fact that any string can be used to encrypt on including a role. When a user wants to send a message they choose a role to encrypt it under, and recipients of the message can decrypt if they belong to that role. Anyone who doesn't belong to the role cannot see the message. This provides a secure email system for use within the hospital.

9. CONCLUSION & FUTURE WORK

We presented alternative security and privacy technologies considered for use in our system architecture for an e-hospital environment. The motivation for the inclusion of high security technologies comes from the requirements by privacy legislations PIPEDA and PHIPA. Our system therefore includes encryption for data confidentiality, integrity mechanisms, authentication, authorization, and audit logs. An additional security measure put in place for our project also involves automatic deletion of sensitive data when tablet PCs are taken out of the hospital area.

Our main contribution is the use of a policy-based encryption scheme in providing data encryption and access control. We propose extending Molva and Bagga's work to suit a health care environment for access control with a patient's records database. Policy-based cryptography looks promising for use in different settings. The uses of this type of system could span many environments including corporate settings and email systems. The usefulness in creating keys based on roles and the simplicity of key management give policy-based encryption many advantages over current encryption schemes.

10. REFERENCES

- [1] A Adams, M Sasse, "Users are not the enemy", In Communications of the ACM, pp 40-46, 1999
- [2] Array Networks Inc. SSL VPN vs IPsec VPN, Jan. 2003. white paper.
- [3] W Bagga and R Molva, "Policy-Based Cryptography and Applications", In Lecture Notes in Computer Science, pp. 72-87, Springer Berlin / Heidelberg, 2005.
- [4] A Bittau, M Handley, J Lackey, "The Final Nail in WEP's Coffin", The 2006 IEEE Symposium on Security and Privacy SP, pp. 386-400, 2006.

- [5] D Boneh and M Franklin, "Identity-Based Encryption from the Weil Pairing", In Proceedings of CRYPTO 2001, pages 213-229, Springer-Verlag, 2001.
- [6] D Ferraiolo, J Cugini, R Kuhn, "Role-based access control (RBAC): Features and motivations", In Proceedings of the 11th Annual Conference on Computer Security Applications, pp 241-248, 1995.
- [7] J Holt, R Bradshaw, KE Seamons, and H Orman, "Hidden credentials", In Proceedings of the 2003 ACM Workshop on Privacy in the Electronic Society, ACM Press, 2003.
- [8] RR Jueneman, SM Matyas, CH Meyer, "Message Authentication", IEEE Communications Magazine, pp 29-40, 1985.
- [9] A Kapadia, P Tsang, SW Smith. "Attribute-Based Publishing with Hidden Credentials and Hidden Policies." In 14th Annual Network and Distributed System Security Symposium (NDSS '07), pp. 179-192, 2007.
- [10] S Khanvilkar, A Khokhar, "Virtual Private Networks: An Overview with Performance Evaluation", IEEE Communications Magazine, pp 146-154, October 2004.
- [11] G Lassmann, "Some results on robustness, security and usability of biometric systems", In Proceedings of the 2002 IEEE International Conference on Multimedia and Expo ICME '02, pp 577-579, 2002
- [12] MC Mont, P Bramhall, CR Dalton, and K Harrison, "A Flexible Role-based Secure Messaging Service: Exploiting IBE Technology in a Health Care Trial", Hewlett-Packard Laboratories, technical report HPL-2003-21, 2003.
- [13] Motion Computing, "Motion C5", 2007, <http://www.motioncomputing.com/>
- [14] Personal Health Information Protection Act (PHIPA 2004) http://www.health.gov.on.ca/english/public/legislation/bill_31/personal_info.html
- [15] Personal Information Protection and Electronic Documents Act (PIPEDA 2000) http://www.privcom.gc.ca/legislation/02_06_01_01_e.asp
- [16] Secure Computing, 2007, <http://www.securecomputing.com/>
- [17] Shamir, "Identity-based cryptosystems and signature schemes", In Proceedings of CRYPTO 84 on Advances in cryptology, pp. 47-53. Springer-Verlag New York, Inc., 1985.
- [18] N Smart. "Access control using pairing based cryptography", In Proceedings CT-RSA 2003, pp 111-121. Springer-Verlag LNCS 2612, April 2003.
- [19] T Wilkinson, D Hearn, and S Wiseman, "Trustworthy access control with untrustworthy web servers", In Proceedings of the 15th Annual Computer Security Applications Conference, pp 12. IEEE Computer Society, 1999.
- [20] Voltage, <http://www.voltage.com>