Owner-Controlled Information

Carrie Gates
Faculty of Computer Science
Dalhousie University
Halifax, Nova Scotia
B3H 1W5 CANADA
gates@cs.dal.ca

Jacob Slonim
Faculty of Computer Science
Dalhousie University
Halifax, Nova Scotia
B3H 1W5 CANADA
slonim@cs.dal.ca

ABSTRACT

Information about individuals is currently maintained in many thousands of databases, with much of that information, such as name and address, replicated across multiple databases. However, this proliferation of personal information raises issues of privacy for the individual, as well as maintenance issues in terms of the accuracy of the information. Ideally, each individual would own, maintain and control his personal information, allowing access to those who needed at the time it was needed. Organizations would contact the individual directly to obtain information, therefore being assured of using current and correct information.

While research has been performed on users owning and controlling access to their personal information in an electronic commerce environment, we argue that this concept should be extended to all user information including, for example, medical and financial information. The end goal is not for users to simply maintain copies of this information, but to be the source of this information.

This paper presents the concept of users owning their personal information and introduces some of the issues involved in users being able to control access to this information. The security requirements, including authentication, access control and audit, as well as user interfaces and trust, for this new paradigm are given particular emphasis.

Keywords

security, privacy, architecture

1. INTRODUCTION

Current information systems are based on the premise of storage and control of user information by organizations. That is, whenever a user accesses any system, be it a web-based electronic commerce site, on-line banking, or a hospital visit, that site maintains personal information about the user. This information is meant to provide consistency for the user should they return to that site for further services. Additionally, the site can gain a competitive advantage by storing personal information for a user, and by performing aggregated data mining of all the user information in the database.

There are disadvantages to the scheme where each site maintains separate copies of a user's information. First, there is the issue of privacy for the user, which has been addressed in some countries through legislation. There is also the issue of maintaining correct information. At present, the only solution to this is to put the onus on the user to ensure that all sites have accurate personal information, and to update that site when the information becomes incorrect. However, this presumes that the user will not only take the effort to update all sites, but also knows what sites need to be updated. A third disadvantage is that important information about a user is not always available when needed if, for example, a user visits a different hospital.

Another solution to this proliferation of personal information is to return control of this information to the users who own it. To do this, a user must maintain all of his personal information in his own database, including not only standard information such as name and address, but all of his personal information, such as medical history and financial information. This addresses the issues of privacy, consistency and mobility as mentioned above.

This paper explores the issue of users owning their own personal information, with a particular emphasis on the security requirements of this new paradigm. It begins with a discussion of the basic premise, and tries to define what it means for a user to own their personal information. This is followed by a description of the current system for maintaining user information, along with its advantages and disadvantages. Section 4 presents the new paradigm, followed by a discussion on the security issues raised by this new paradigm in Section 5.

2. BASIC PREMISE

The basic premise behind this paper is that individuals should own and control access to their personal information, where control is provided through technological means.

However, what does it mean to "own" a piece of information, and how can one determine ownership? In this paper, we focus on individuals, and not organizations, and the information that an individual should own and control. One perspective on this is to ask what kind of information can the individual change, or, rather, what kind of information about the individual *can not* be changed by others. Obvious choices here are location information, such as address and phone number. Identifying information such as name and age also fall in this category, as, while an individual can not change his age, neither can anyone else change it for him! Preference information, such as hobbies and sports, favourite books and preferred desktop settings, are also information that only the individual can change.

But what of government-issued identity information? For example, citizenship information such as passports and social security numbers could also be claimed to be owned by the country, as its government can revoke citizenship, and are therefore a property of the individual but not owned by him. It therefore seems reasonable that a user would have a copy of his passport information, but that it would ultimately be "owned" by his country's government. If someone requests passport information from an individual, he could still provide it, and it would be signed by the government of his country to ensure authenticity. Similar to certificates, the requesting person or company would also need to check with the government to ensure that the passport had not been revoked.

Financial information provides an interesting challenge to this model. Certainly an individual owns his money, however other people can affect an individual's bank balance. For example, a bank can automatically deduct service fees, or the user may have automatic debit for particular bills, or direct deposit. However, in all of these cases, the user has made a contractual agreement to allow the credits and debits, and so the final balance still belongs to the user.

Medical information provides, perhaps, the most interesting case to consider. Using this definition, a medical diagnosis is owned by the individual. That is, medical records reflect information about a particular person. While this information is actually determined by another person (such as a physician), it reflects the current state of the individual, and is not a state that is assigned by an outside party. In fact, a strong case can be made for an individual having all of their medical information in their possession, so that regardless of facility or physician visited, a complete medical record is available for that visit.

The issues surrounding medical information are less about ownership, and more about access. In particular, some in the medical community do not believe that patients should have access to their medical files, on the premise that doctors are less likely to be truthful or complete in their comments if the patient can later read those comments (e.g. comments about the mental state of the patient may be particularly sensitive, and so statements such as "this patient may have borderline personality disorder" will no longer be made).

However, declaring that an individual owns their personal information does not necessarily imply read access, nor does it necessarily imply write access. The individual should maintain that information, and be able to control access to it. In the case of medical information, it should be possible for the individual to maintain that information in their possession, but also for a physician to encrypt their private notes on the patient such that only other physicians would have the appropriate permissions to view that part of the file.

This raises another important point — there is a distinction between ownership and control [7]. Data ownership specifies to whom the data belongs. However, this person does not necessarily control the data nor how it is used. For example, in some countries, a person's medical information is owned by that person. However, it is controlled by the hospital that person visited. The person does not even maintain a copy of the information he owns.

It is not sufficient for users to own their personal information — they must be able to control access to this information. The difficulty in this is that information, once released, can not currently be controlled by any technological means. Rather, legislative controls exist to provide guidelines on how information can be used and distributed by the parties to whom personal information has been released.

Beyond the issue of control of information, once released, is the issue of social responsibility. Individuals owning and maintaining their personal information implies that the privacy of the individual outweighs the needs of society. That is, in some cases it is in the best interest of society as a whole to have personal information available. One example of this is having medical information available for data mining and analysis, where medical advances might be made from this information. It might also be in the individual's best interest to have this information available to other parties. For example, researchers might recognize that a particular subject has a particular disease, given blinded medical data. Yet if a one-way hash function was used to blind the data, so that it is not possible to determine the subject, then there is no means by which to alert either the subject or his doctor to the situation [25].

Finally, there is the issue of responsibility for the control of information. In this model, rather than relying on various organizations to maintain their information (e.g. banks, hospitals), users are now responsible for maintaining their own information, as well as for controlling access to that information so that only the appropriate individuals can access it. This is not necessarily a responsibility that every person wants, nor is it a responsibility that every person is capable of handling. In particular, if a person is cognitively impaired, they may be incapable of making informed decisions on how their data should be accessed.

3. THE CURRENT STATE — PROS AND CONS

Personal information on various individuals is maintained in many different databases owned by many different corporations and government agencies. It is estimated that information on a particular person is stored in approximately 1000 different databases [3]. For example, information may be kept by government offices at all levels — federal, provincial or state, and municipal — as well as by corporations, such as banks and telephone companies, electronic businesses, such as on-line book stores or clothing stores, non-profit organizations, such as professional affiliations or volunteer organizations, and by other services, such as at hospitals or with family physicians. Currently there is no co-operation between the various entities to share user information.

There are some benefits to such an approach. For example, this model is already well-established and well-accepted. Businesses are able to maintain control over the information they have gathered (as opposed to a centralized repository that all businesses must use, for example). This allows businesses to aggregate their user information and mine it for trends. This can provide benefits as diverse as allowing a business to customize their services to particular target groups, to mining medical data for trends that may indicate disease indicators or progression.

As the information gathered is replicated and distributed, incorrect information obtained by one office is not propagated to other offices. Thus the damaging effects of misinformation is locally contained. Additionally, the distributed, uncoordinated nature makes collusion between, for example, banks, hospitals and insurance companies, difficult to achieve, in addition to being prohibited by legislation.

However, there are a number of disadvantages in relying on every organization to maintain their own database of users. First, given the proliferation of personal information, it is likely that at least some databases contain incomplete or inaccurate information. The onus is on the user to update his information with everyone who may have information on him. Failure to do so can result in potentially serious consequences for the user, such as can arise if a credit agency has incorrect information.

The lack of co-ordination between various offices can also cause difficulties when a user needs to prove his identity. For example, if a person is moving to a different country, proving his identity (and credit history) can be difficult and require extensive paperwork.

Additionally, there are no technological controls on how an office uses someone's personal information, only legislative controls. The result is that a person does not necessarily know how his information is being used. A company can sell a person's information to another company, so that the person no longer even knows who has information on him.

Of greater concern is that information is not necessarily available when and where needed due to its distributed nature and the lack of information sharing between organizations. For example, if a person is admitted to a hospital in a different town, then that hospital likely does not have access to any information on the person's medical history. This can cause complications in treating the individual if, for example, the person has diabetes yet the hospital is unaware of this when selecting treatments.

4. PARADIGM SHIFT

An alternative to distributed information amongst various offices is a paradigm shift to one centralized location for all user information. In this new system, a user would own all of his personal information. It is anticipated that users will be able to keep their personal information with them at all times through small, specialized devices. Infrastructure changes will be required to systems in order to communicate with individual devices for personal information, rather than contacting a centralized database.

The advantages of such a system are centered on the control and access a user has to his own personal information. In this system, a user can ensure that his personal information is correct, needing to update any changes in only one location. Additionally, a user can control access to his personal information, including being able to specify who can access his information, what portions of his information can be accessed, and the conditions under which it can be accessed. By having this central repository, it becomes easy to provide complete information to whoever may need it (e.g. when needing to prove identity to gain access to new services).

Perhaps of greatest benefit in this new paradigm, however, is that personal information will be able to be travel with a user, allowing access to that information whenever it is needed. One example of where this is of particular value comes from a user who may be traveling, and need medical attention. By having all his information on-hand, he can provide his medical history to the attending physicians to ensure that proper care is received.

However, there are also disadvantages with this method. While it provides a user with the capability to ensure that his information is current and provide them with control over access to that information, this is a double-edged sword, requiring the user to understand that responsibility and be able to provide appropriate access. Additionally, with this centralized keeping of information, it will be easier for adversaries to perform identity theft, as gaining illegal access to some information will likely result in illegal access to all a user's information.

A less obvious disadvantage to this paradigm shift is that businesses will no longer be able to perform data mining, unless they have been able to store the information they need. This will impact on the competitive advantages now enjoyed by businesses who perform this type of analysis. It will also impact on some areas of research, such as medical research, if patient information is not available in a central location for data mining.

4.1 It's Already Started....

There are several examples that such a paradigm shift is already occurring. For example, Microsoft Passport [14] is a single sign-on facility for people accessing the Internet. A user's Passport would contain some minimum amount of personal information, and the user can control what information is released to other companies. The passport can also contain e-wallet information, such as billing and shipping information. Then, when the user contacts a site that accepts MS Passports, the user can have his Passport provide his information, rather than needing to type in the same information for each site visited.

Extending the Passport and .NET concepts even further, Microsoft had also been working on a product called Hailstorm [15], which would incorporate Passport information, as well as calendar and scheduling information, default application settings, and personal preferences, among others. The goal was to use XML to allow users access to all of their information — calendars, phone lists, address books, documents — from any device, with Passport being used as the authentication mechanism. In a statement released to the press, Bob Muglia, MS Vice President of the .NET Services Group, commented "Hailstorm turns the industry debate over online privacy on its head.... It starts with the fundamental assumption that the user owns and controls their personal information ...." [15]

In response to Microsoft's release of Passport, the Liberty Alliance was formed, currently consisting of more than 160 member organizations [12]. The goal of the Liberty Alliance is to create a federated identity infrastructure, where links between an individual's identity information within various organizations are kept, rather than maintaining all of a user's identity on a central server or with one organization. As a result, the network identity for an individual has a broad definition, as "...the sum of their financial, medical, and personal data...." [13]

The Persona Project extends the work by Microsoft and the Liberty Alliance [29]. This project, based out of Oregon State University, goes beyond having a web page that allows single sign-on, to having a "persona", or consumer-centered identity model, that is distributed across multiple systems so that it can be accessed via desktops, personal digital assistants (PDAs), cell phones, and even from cybercafés. According to Toth and Subramanium, the persona is "an active software agent that encapsulates private and personal data and performs a range of authentication and personalization services on behalf of its owner." [29] The persona holds a user's personal information, including identity, passwords, preferences and e-wallet information. The basic premise is that a user will authenticate himself to his persona, who will then act on behalf of the user to supply on-line information such as billing information or personal schedules. Access to this information is moderated by the access control rules employed by the user (e.g. so that only a limited number of companies can access credit card information, for example).

Trusted Computing Platform technology [20] can also be used as a basis for providing increased privacy for users in electronic commerce settings. Pearson has described how users can employ a self-profiling approach, storing profiles (e.g. of their on-line shopping habits) on their home system, and even using different profiles for interactions with different sites [19]. Users can choose to provide their profiles to various web sites as they interact with that site, and can even provide the profile anonymously should they choose to not release their identity. By employing a trusted computing platform, the site receiving the profile can be assured of the integrity of the profile through the use of a public key infrastructure in an agent-based system.

In contrast to the web-based services described above, the United States Department of Defense (DoD) has employed smart cards, dubbed Common Access Cards (CACs). The DoD CAC is being used initially for security, allowing access to buildings and computer networks and software, and for authentication for on-line transactions [27]. The CAC contains the user's private key, allowing encrypted communications and digital signatures. The DoD is moving towards the use of biometrics with the CAC, rather than passwords, as well as using the cards for more detailed information on soldiers. One example given was the ability to record medical information on the CAC so that medications and treatments of an injured soldier could be tracked through the medical system without requiring a network.

4.2 A Natural Extension

There are, however, issues with the approaches taken above, particularly those that are web-based. In particular, Microsoft has received considerable negative publicity. For example, security vulnerabilities have been discovered in both Passport and Hotmail that allows an adversary to obtain a user's credit card information from his e-wallet [28]. Microsoft was also nominated for the 2001 Austria Big Brother Awards [11], in part for Passport due to its potential for privacy violations, and has had articles written commenting on the privacy violations of Passport and .NET (such as by Diffie and Landau [6]). Finally, Microsoft has also faced legal problems for Passport and Hailstorm (see the Electronic Privacy Information Center (EPIC) [8] for information). As a result, while Microsoft still provides Passport, the Hailstorm project has been abandoned.

Much of the concern with Microsoft's approach was centered around one large corporation having access to all of a user's information. This is being addressed by the Liberty Alliance's use of a federated system, while the United States Department of Defense (DoD) addressed it through having each user carry a smart card, rather than rely on a central system.

However, the paradigm shift presented here goes beyond that of Passport, the Liberty Alliance and the Persona Project, and most closely resembles that of the DoD's Common Access Cards. Rather than provide information to only on-line services through the use of, essentially, a single sign-on facility, the authors suggest that *all* of a user's personal information, including his complete medical history, financial records, dental history, etc., should be owned, maintained, and controlled by the data owner.

The current state of industry and research indicates that some form of centralized management of personal information is inevitable. However, users reject the notion of a central repository of all of their personal information under the control of a single corporation. While the federated approach mitigates this concern, it is limited to only those services that are on-line, which largely limits users to an electronic commerce forum. Yet it is important for users to be able to access off-line information, such as medical history. Therefore the approach taken by the DoD for the Common Access Cards seems to be the most promising.

However, the DoD approach was not designed to provide control to users over their personal information, but rather for convenience and mobility of important information. As a result, the user has no control over who can access his information, when, how much and under what circumstances.

A merging of these two approaches is required, where a user has access to his information at all times and that it is not kept in a central location, yet he can still control who else can access that information. In order to provide both of these abilities, some form of mobile device that contains a user's information is required, where the access controls put in place control access to this local information. This personal information should be held in separate databases, grouped based on context.

In order to ultimately protect a user's privacy, the authors argue that the only copy of a user's information should be on the user's device, rather than maintained by multiple entities, such as various hospitals, banks, shops, etc.

5. SECURITY ISSUES

This paradigm shift from the current distributed manner of dealing with personal information to a user-centered, user-controlled system raises several security issues in all areas of security, including authentication, access control and audit. Additionally, the issues relating to user interfaces and trust have been included as security concerns. Without a properly designed interface, users may not understand how to deploy the security mechanisms, thus leaving their systems vulnerable. And without trust in the system, regardless of the comprehensiveness of the security mechanisms, the system will not be used.

5.1 Authentication

If a user is to control his own information, it is very important to ensure without any doubt that the identity of the person accessing or controlling the information is the owner of that information. While there are many methods for authentication, such as passwords, tokens and biometrics, each method has its disadvantages. Passwords can be forgotten, and tokens can be lost. Biometrics provide a nice alternative to passwords and tokens, and have been shown to adhere to the universal access paradigm [10]. However, multiple modalities should be available to users based on a balance of personal preferences, flexibility, device capabilities and security requirements.

It is equally important, and perhaps more difficult, to ensure the identity and authenticity of external users who need access to an individual's information. For example, if someone is dealing with a bank to perform a financial transaction, there needs to be some form of authentication to ensure that the access to the user's financial information is a legitimate access by the user's bank. Not only will institutions need to be authenticated, but also other individuals. For example, some form of identifying that an individual is a medical doctor before allowing access to write to a user's medical database is required. One possible solution to this form of authentication is through the use of attribute-based certificates (such as simple public-key infrastructure (SPKI) [22][9]) or role-based certificates (such as proposed by Park and Sandhu [18]).

Authentication is particularly important to prevent identity theft. One solution to this issue is to ensure that all of a user's information is encrypted at all times. However, if a criminal gains access to the underlying system, they might also find a copy of the key. It is possible to avoid this issue by encoding the key so that the user does not need to store it locally, such as through the use of tokens. This is not an ideal solution, as tokens can be stolen. Another possible solution is to utilize a user's biometric information to create a key (see work by Fabian Monrose at Bell Labs and Michael Reiter at Carnegie Mellon University, such as [16] for current research in this area). However, this is still not a complete solution, as it requires the user to authorize all accesses to his information, which may not be possible (e.g. if the user is unconscious in a hospital, yet to decrypt his medical record requires the user's voice). There is also the requirement for a third party to be able to decrypt a user's information, such as in the case where a user dies and his estate goes to probate.

At first, it would appear that research is needed to create a cryptosystem such that the user can encrypt and decrypt as required, yet his encryption scheme can also be decrypted by a second, generic key owned by a trusted third party. However, such a system has been developed in the past and rejected by the public. The Clipper Chip (for voice communications) and Capstone (for data communications) both used a secret cryptosystem known as SkipJack, which was developed by the United States National Security Agency (NSA) [17]. However, this system had numerous flaws [21], and after an initial flurry of discussion, seems to have disappeared circa 1995.

An alternative, perhaps, is a cryptographic system where the owner can specify who else can decrypt his data and under what circumstances. This could take the form of partial key escrow that obey the secret sharing property (that any k pieces of the key can reconstruct the key, but that no t pieces provide information about the key, where $t < k$) [2]. Thus other users can maintain portions of the complete key, with appropriate access control used to determine if those users can use their keys on the data.

5.2 Access Control

In addition to the issues faced in authentication, access control in such a system also raises serious research issues. A single user will have multiple contexts (e.g. medical, financial), each needing its own database. Access to a user's databases should be based on the context in which the access is being granted. For example, if someone has the authority to access a user's medical database, they should not at the same time be able to access a user's financial database. Thus the access control system needs to be aware of the context in which it is being accessed, and potentially restrict access to certain databases based on that context regardless of the authority of the accessing user.

The owner of the data will need to be able to grant access to a large number of external users, based on several criteria. One obvious criteria is the role that external user plays, such as a doctor or bank manager. Different roles should have access to different parts of the owner's database, so that, for example, a medical doctor can not view a user's bank balance. Role-based access control [23] can be employed in this situation. However, roles will need to be defined based on the owner's perspective in terms of who they will contact. This is opposed to the current system where roles are defined within the context of an organization. In the case of a data owner, appropriate roles may be both doctors and bank managers. However, they are unlikely to require roles internal to an organization, such as financial officer or systems analyst. This is not to say that a data owner might not encounter these roles in other contexts, such as in his own place of employment. However, should the owner encounter such roles at his employing organization, he will not be dealing with his own personal information, and so is likely using a different system with its own access control requirements.

Criteria other than roles may also be required, such as the location of the owner and the external user. In this manner, access can be restricted so that an owner's information is only accessed at appropriate places (e.g. a medical doctor can only access a patient's information if the patient is in a hospital at the time). Access can be further restricted based on when the access request occurs. Thus transactions, such as financial transactions, can be restricted to only occur during regular business hours. A final criteria for granting access that needs to be available is proximity of the owner to the external user. This will allow the owner to control the circumstances under which an external user might access his personal information. One example of this might be that an emergency room physician can only access a patient's medical information not only if both of them are in an emergency room, but that the physician is in the same emergency room as the patient.

Attention needs to also be paid to the access given to the owner of the data. In current access control systems, the security officer has the responsibility for generating the appropriate rules [24], and therefore has the ability to assign to herself complete access. However, in a user-owned information system, the user plays the role of the security officer for his own data. But that does not imply that a user should be able to write or delete any part of his data (e.g. a user should not be able to modify his financial records!), nor does it even imply that a user should necessarily be able to read his own data (e.g. the current state of medical data does not allow read access to patients). The access control system needs to be designed so that the user does not have super-user like powers to access and modify his own information. This also implies larger design issues, where a user should not have access to read or modify the underlying database systems, operating systems or applications.

A final issue in terms of access control is the need for the owner of the data to be able to delegate partial or complete access. For example, if an owner is elderly, it may be desired to delegate access from the owner to the owner's children. Along the same vein, some form of information over-ride is required so that access to personal information, such as financial records, can be obtained in the event the owner dies unexpectedly.

5.3 Audit

There are also issues in terms of the auditing of such a system. It is obvious that an audit trail is needed that can be trusted and used in the event of a dispute. For example, should either the user or a bank dispute the information in someone's financial record, a non-repudiated audit trail will be required to determine all accesses to the user's financial information, and the form of those accesses (e.g. reads, writes). In a manner similar to today's court system, a trusted third party will be required to view the audit trails. One security issue this raises is the need to control access to the underlying operating and database systems so that users can not modify or delete the audit records. One approach to this is to design the system so that it uses a tamper-resistant hardware, such as the Trusted Computing Platform (TCP) with a protected storage area [20]. The access can be designed so that audit logs are maintained in an encrypted form in the protected storage area where it can not be deleted.

Not only will an audit trail need to be available and trusted, but some mechanism will need to be in place to alert an institution that a user may have tampered with his records. Continuing with the financial example, a bank will require more than trust that the user did not modify his bank balance! There are two approaches here. The first is that, when a user updates his balance with a bank, the bank inserts a digitally signed copy of the new balance into the user's financial database. In addition, the bank also signs a hash of the relevant tables in the database, so that the user is not able to withdraw money, and then delete the row indicating the withdrawal, returning to his previous (higher) balance. When a transaction is about to be made, the bank can confirm both the user's balance and if there has been any tampering by confirming that the signed hashes of the tables match the hashes of the current tables. This does not address the case where the database has been completely deleted, although the use of tamper-resistant hardware to store the databases and signed hashes could be used. A second approach, that does not exclude using the first approach concurrently, is that the bank could keep a copy of the balance, although this would need to be done in such a manner so as to maintain the user's privacy. One approach to this is the use of a pseudonym, such as described by Chaum [4], that links a balance to a user without identifying the user.

In addition to being able to detect the reads and writes to a user's data, some form of intrusion detection (and preferably intrusion prevention!) is required. Should someone with criminal intent gain unauthorized access to a user's system, and then copy all of a user's information, that person is then in a position to perform perfect identity theft. Thus there are open research questions in terms of determining if someone has bypassed the standard authentication protocols and gained access to the underlying databases of personal information.

Another issue, related peripherally to audit, is that of copies of released personal information. That is, a user may legitimately release some of his personal information to a third party (such as during a financial transaction). How can the user ensure that the third party has not made a copy of this information? And if the user has authorized the copying of some of his personal information, how can the user ensure that this information is not kept for longer than authorized, and not used for any unauthorized activity? The issue of expiring data is an open research question, and is particularly being addressed in the entertainment arena through items such as limited-use MP3s and digital rights management (DRM) research.

However, some places have a legitimate need to keep some personal information on various individuals. For example, medical institutions should be able to keep medical information so that they can perform medical research through data mining and data analysis. Thus the system should be configured to allow specific information to be copied by other institutions, where the information available, and the institutions who can view it, should be configurable by the owner of the data.

Another issue that falls under audit is that of backup and recovery procedures. It would be disastrous for a user to have all of his personal information on one mobile device, and then have that device's hard drive crash! Yet relying on users to back up their own systems is unreliable at best. Some method for easy backup is required, along with a reminder system to ensure that the data owner does perform the backup. Beyond this, the backup needs to be protected so that, if stolen, it can not be used to replicate someone else's identity. One possible solution here is to encrypt the entire backup with the owner's public key. Additionally, backups should be performed after every transaction so that, should the backup be required, there is no information loss. Another possible solution to this is to have the user's data distributed across servers such that the mobile device is only an access point (this assumes that the user's information is encrypted so that it can not be viewed by the owners or administrators of the servers). If the servers follow regular backups and employ hardware solutions such as RAID, then the user should not need to worry about performing backups or losing data.

5.4 User Interface

While not traditionally listed as one of the key areas in computer security, the user interface is an extremely important aspect of any model that allows the end user to control his own personal information. The user interface must be designed in a manner that not only allows users with various levels of skill to configure the access control system, but that also allows the user to understand the consequences of his configuration decisions, alerting the user to any conflicts in configuration.

In addition to configuration, it must also be clear to a user when others are attempting to access his information and for what purpose. The system should be designed to allow users to approve any data transfer, rather than having all transfers occur automatically, as users want to be involved in any data transfer process [5]. This system should also allow the user to specify the granularity of the information to be provided. For example, rather than assuming that a user would want to transfer all of his contact information, it should be possible for the user to submit only their mailing address, but not their phone number. This flexibility is required because users view different methods of contact differently. For example, users who do not object to providing their email address may object to providing their phone number [5].

The user interface needs to be considered when the system is being designed, before it has been implemented. Otherwise the user interface portion will consist of trying to retrofit an interface to a design, rather than knowing the limitations at design time. Given the difficulty of designing a security interface that can be understood by users (see, for example, Whitten and Tygar [30] for experiences with e-mail encryption), this area will require much research.

The requirement for security and the user interface to be considered at design time is underscored by studies that have shown how unwieldy interfaces and unrealistic security practices have resulted in users developing insecure practices to deal with the security requirements. One example of this is in the case of companies with multiple passwords and strict policies regarding how often passwords need to be changed [26, 1]. Users circumvented security by writing down their passwords, choosing insecure passwords, and sharing their passwords. If users are to maintain their own personal information, security will need to be designed so that users do not feel the need to circumvent the system, and that they are motivated to keep their information secure.

Finally, the system should be designed with various levels of user in mind, including varying levels of security expertise, computer literacy and cognitive ability. Rather than designing the interface with the assumption that the user will understand the impact of his configuration changes, the system should be designed with some defaults that allow users to not need to understand the underlying concepts of security, such as access control and encryption.

5.5 Trust

Finally, there is the issue of trust. If users own and maintain their personal information, then organizations such as hospitals and banks must be able to trust that the information provided by a user is correct and accurate. This may be the most difficult hurdle in moving to a user-owned information paradigm. Banks, for example, would need to be assured that users have not bypassed the security mechanisms installed on their local devices to, for example, adjust their bank balance. It is much easier for a bank to ensure the security of a central system that it owns, than to trust the security of highly-distributed, mobile devices over which it has no control.

In order to gain trust, a number of security mechanisms will need to be in place: tamper-resistant hardware, operating systems, databases, applications and audit trails, in addition to strong authentication, encryption services, secure protocols, and reliable software.

A very promising approach to providing the level of trustworthiness required in this system is that of Trusted Computing Platforms (TCP) [20]. In this system, tamper-resistant hardware, similar to that found in a smart card, is provided as part of the computing platform. This hardware can be used to maintain the private keys of the system's user, as well as providing protected storage for signed audit logs. In [19], Pearson provides a very good example of how TCP can be used for user self-profiling, allowing a user to provide information to web sites in a manner that allows the web site to trust the information, without requiring that the user reveal his identity. This same approach can be used and extended to provide a trusted platform for owner maintained information.

Individuals will also need to trust the system. One of the reasons that Microsoft Passport has not been particularly successful is that users do not trust providing all of their personal information to one central location that is under the control of a single corporation. While the approach presented in this paper will not suffer from this particular issue, users will still need to gain trust in both the hardware and software if the system is to be used. For example, how can a user trust that some error in the software won't cause his bank balance to be reduced? And, if something were to happen, what recourse does a user have?

One final concern relating to trust is that of network pervasiveness and reliability. This model assumes that a network connection is always available. However, if a network connection is not available, and the user can not complete the desired transactions, then he will lose trust in the system. Conversely, if the user's system is not available to a business entity when required (for example, if a bank needs to confirm a balance in order to cash a cheque), then the participating business entities will also lose trust, and not be willing to use or support this model.

6. CONCLUDING REMARKS

This paper argues that users should own their personal information, and should have control over both access and distribution of this information. A paradigm shift is therefore needed from the prevailing system where information is distributed across multiple organizations, to a system where users maintain their information in one location, such as on a small, specialized device. This ensures that information is correct and up-to-date and that it is available when needed.

The security implications caused by this new paradigm have been presented in this paper, and affect all areas of security, including authentication, access control and audit, as well as user interfaces and trust.

7. ACKNOWLEDGMENTS

The authors would like to gratefully acknowledge the support of the IBM Canada Centre for Advanced Studies, the National Sciences and Engineering Research Council of Canada (NSERC), and the Canada Foundation for Innovation (CFI). We would also like to thank all of the participants at NSPW 2003 for their feedback and comments.

8. REFERENCES

[1] Anne Adams and Martina Angela Sasse. Users are not the enemy. *Communications of the ACM*, 42(12):40 – 46, 1999.

[2] Mihir Bellare and Shafi Goldwasser. Verifiable partial key escrow. In *Proceedings of the 4th ACM Conference on Computer and Communications Security*, pages 78 – 91, Zurich, Switzerland, 1997.

[3] Stefan A. Brands. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*. MIT Press, Cambridge, Massachusetts, 2000.

[4] David Chaum. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM*, 28(10):1030 – 1044, 1985.

[5] Lorrie Faith Cranor, Joseph Reagle, and Mark S. Ackerman. Beyond concern: Understanding net users' attitudes about online privacy. Technical Report TR 99.4.3, AT&T Labs, April 1999.

[6] Whitfield Diffie and Susan Landau. Commentary: The threat of Microsoft's .NET. *kingpublishing.com*, 2001. Last visited: 17 July 2003.

[7] John Dobson. Private communication between John Dobson and Steven J. Greenwald at NSPW 1996 as referenced by Greenwald, 2003.

[8] Electronic Privacy Information Center. Microsoft passport investigation docket. `http://www.epic.org/privacy/consumer/microsoft/passport.html`, 2003. Last visited: 17 July 2003.

[9] Carl M. Ellison, Bill Frantz, Butler Lampson, Ron Rivest, Brian M. Thomas, and Tatu Ylonen. SPKI certificate theory, 1998. Internet draft, work in progress.

[10] M.C. Fairhurst, R.M. Guest, F. Deravi, and J. George. Using biometrics as an enabling technology in balancing universality and selectivity for management of information access. In *Universal Access. Theoretical Perspectives, Practice, and Experience: 7th ERCIM International Workshop on User Interfaces for All*, pages 249 – 259, Paris, France, 2002. Springer-Verlag Heidelberg. Lecture Notes in Computer Science 2615. October 24-25, 2002.

[11] John Lettice. Big Brother Award nomination for WPA, Passport pains MS. *The Register*, 2001. 1 April 2003. Last visited: 17 July 2003.

[12] Liberty Alliance. Liberty Alliance project. `http://www.projectliberty.org/`, 2003. Last visited: 17 July 2003.

[13] Paul Madsen. The Liberty Alliance. *webservices.xml.com*, 2003. 1 April 2003. Last visited: 17 July 2003.

[14] Microsoft Corporation. Microsoft .NET passport: One easy way to sign in online. `http://www.passport.net/`. Last visited: 17 July 2003.

[15] Microsoft Corporation. Microsoft announces "Hailstorm," a new set of xml web services designed to give users greater control. `http://www.microsoft.com/presspass/features/2001/mar01/03-19hailstorm.a%sp`, 2001. Last visited: 17 July 2003.

[16] Fabian Monrose, Michael K. Reiter, Qi Li, and Susanne Wetzel. Cryptographic key generation from voice. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, pages 202 – 213, 2001.

[17] National Institute for Standards and Technology. Escrowed encryption standard (EES), 1994.

[18] Joon S. Park, Ravi Sandhu, and Gail-Joon Ahn. Role-based access control on the web. *ACM Transactions on Information and System Security*, 4(1):37 – 71, 2001.

[19] Siani Pearson. Trusted agents that enhance user privacy by self-profiling. Technical Report HPL-2002-196, HP Laboratories, 2002.

[20] Siani Pearson. Trusted computing platforms, the next security solution. Technical Report HPL-2002-221, HP Laboratories, 2002.

[21] Ronald L. Rivest, Martin E. Hellman, John C. Anderson, and John W. Lyons. Responses to NIST's proposal. *Communications of the ACM*, 35(7):41 – 54, 1992.

[22] Ronald L. Rivest and Butler Lampson. SDSI—a simple distributed security infrastructure. `http://theory.lcs.mit.edu/~rivest/sdsi10.ps`, 1996. Presented at CRYPTO '96. Last visited: 27 May 2002.

[23] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinnstein, and Charles E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38 – 47, 1996.

[24] Ravi S. Sandhu and Pierangela Samarati. Access control: principles and practice. *IEEE Communications*, 32(9):40 – 48, 1994.

[25] Martina Angela Sasse. Private communication at NSPW 2003 with Carrie Gates. 20 August 2003.

[26] Martina Angela Sasse, Sacha Brostoff, and Dirk Weirich. Transforming the 'weakest link' — a human/computer interaction approach to usable and effective security. *BT Technology Journal*, 19(3):122 – 131, 2001.

[27] Secure Electronic Transactions – Devices, U.S. Army. Common access card (CAC). `https://setdweb.setd.army.mil/cac/whatiscac.htm`. Last visited: 17 July 2003.

[28] Marc Slemko. Microsoft passport to trouble. `http://alive.znep.com/~marcs/passport/`, 2001. Last visited: 17 July 2003.

[29] Kal Toth and Mahesh Subramanium. The persona concept: a consumer-centered identity model. `http://eecs.oregonstate.edu/~ktoth/Other/TrustBus03-Persona-Toth&Subram%aniumV1.pdf`. Last visited: 15 July 2003.

[30] Alma Whitten and J.D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proceedings of the 8th USENIX Security Symposium*, pages 169 – 184, Washington, D.C., 1999. Usenix. August 23-26, 1999.