Innovative Solutions for Access Control Management

Ahmad Kamran Malik
COMSATS Institute of Information Technology, Pakistan

Adeel Anjum
COMSATS Institute of Information Technology, Pakistan

Basit Raza
COMSATS Institute of Information Technology, Pakistan

A volume in the Advances in Information Security, Privacy, and Ethics (AISPE) Book Series

Information Science
REFERENCE
An Imprint of IGI Global

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

Library of Congress Cataloging-in-Publication Data

Names: Malik, Ahmad Kamran, 1973- editor. | Anjum, Adeel, 1982- editor. |
Raza, Basit, 1976- editor.
Title: Innovative solutions for access control management / Ahmad Kamran
Malik, Adeel Anjum and Basit Raza, editors.
Description: Hershey, PA : Information Science Reference, [2016] | Includes
bibliographical references and index.
Identifiers: LCCN 2016010969| ISBN 9781522504481 (hardcover) | ISBN
9781522504498 (ebook)
Subjects: LCSH: Database security. | Computer security. | Data protection. |
Computers--Access control.
Classification: LCC QA76.9.D314 I57 2016 | DDC 005.8--dc23 LC record available at https://lccn.loc.gov/2016010969

This book is published in the IGI Global book series Advances in Information Security, Privacy, and Ethics (AISPE) (ISSN: 1948-9730; eISSN: 1948-9749)

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

Chapter 2
Identification and Adaptive Trust Negotiation in Interconnected Systems

Eugene Sanzi
University of Connecticut, USA

Steven A. Demurjian
University of Connecticut, USA

ABSTRACT

Creating an online identity via a username/password does not provide the ability to establish trust with other systems in order to get access to unauthorized information in a time-critical situation. Trust is the ability of two entities to believe one another at some level, so that they can interact in a secure manner, e.g., a physician at one hospital may need to obtain medical data on a patient from another hospital to treat a patient, facilitated if there is a trusted relationship. This chapter explores adaptive trust negotiation that obtains near real-time permission to access a system to which a user has never previously been authorized to, so that the system receiving the request can adjust its security policies depending on the attributes that the requester possesses. To accomplish this, a set of interacting systems (e.g., from different hospitals) can be augmented with identity management and adaptive trust negotiation to create a means where multiple disparate systems can make informed and dynamic security decisions about users relative to their defined security policies.

DOI: 10.4018/978-1-5225-0448-1.ch002

INTRODUCTION

The ability to create, validate, and secure an online identity is a prerequisite for any system that utilizes the user's identity (username/password) to determine access rights to its data and to simultaneously prevent malicious individuals from masquerading as their legitimate service and hijacking user data. Typically, the username/password combination is the system's basis for retrieving and publicly communicating the user's identity and proof of identity to verify that they are the owner of the username and associated identity. In contrast, servers identify themselves by using Public Key Infrastructure (PKI), the domain name system, TLS/SSL, and certificates, which allows users to verify the server's identity. While this model is sufficient for basic client/server interaction involving communication mediums such as email, complex modern time-critical systems necessitate a more robust and responsive system to reach their full potential. For example, the health care domain requires the secure storage and access of information via identity management that has strict requirement on the privacy and security of personal health information (PHI) via the HIPAA standard (HIPAA, 1996). As recently noted (Meyers, 2014), there are increasing attacks on healthcare repositories that contain medical records of patients, including a major attack on a health insurer (Anthem, 2015).

Despite such attacks, there is an increased interested among medical providers (e.g., physicians, clinics, hospitals, imaging centers, testing laboratories, pharmacies, etc.) to share and exchange information (HealthIT.gov, 2014; Kelly, 2013; Mettler & Rohner, 2009) towards quality improvement for timely medical decisions that are able to take advantage of health data stored in multiple locations through the creation and utilization of health information exchange (HIE) (CTDPH, 2013a; JASON, 2014). In such a setting, the username/password combination may be insufficient. For example, a physician using an electronic health record (EHR) to store patient health data and has privileges to see patients at a hospital A which has its own electronic medical record (EMR) for past and current patients. When treating a patient in an emergency situation by a physician at hospital A (authorized to A's EMR) for a cardiac event may need to see an EKG for the patient taken a week ago at hospital B (not authorized to B's EMR) to compare the two EKGs as part of the treatment and assessment process. In such a situation, the physician is attempting to access data at another institution (with a different EMR) to which s/he has no username/password identity. This problem is complicated since patients are often treated by a cadre of medical professionals (specialists, therapists, etc., at multiple locations (State of Connecticut, 2013; The President's Council of Advisors on Science and Technology, 2010)), all with their own systems for storing patient data (Eichelberg, et al., 2005).

There must be a balance struck to allow to access health data in time-critical situations with the granting institution guaranteed that the data is securely shared and not mishandled. This will necessitate the design, development, and deployment of security solutions that can dynamically: vet a user's identity (e.g., physician), ensure the user requires access to certain data in a specific repository (e.g., physician can access patient X's data in hospital A's EMR), and then dynamically generate an identity for the user to actually retrieve the data (e.g., temporary ID for physician to access data in hospital B's EMR). Presently, performing these actions requires the traditional review of the user's credentials where a physician would submit an online application to access hospital B's EMR to which s/he does not have privileges. This is clearly unsuitable in in an emergent care situation. While the creation of one monolithic and universal system to contain login data may be a viable method of allowing credentials to be verified, in reality, the number of systems that would need to be brought together in any kind of realistic scenario is unfeasible. Connecticut has 34 hospitals (with numerous EMRs), 17,294 licensed physicians, and nearly 100,000 other licensed medical professionals (nurses, therapists, dentists, etc. (CTDPH, 2013b)) using hundreds and perhaps thousands of different health information technology systems (HITs). To provide a means for an individual to obtain dynamic access to systems to which s/he has not been previously authorized, there must be the ability to provide some means to demonstrate trust. For example, a physician that has a history of accessing PHI at multiple HIT systems should be trusted for real-time access to information on a needed system that will facilitate prompt patient care.

This chapter explores the concept of trust, the methods currently used to promote trust, and the way that trust can be leveraged, extended, and improved in order to support the aforementioned scenario of health care, with a focus on the reading of information from multiple sources. *Trust* between two interacting entities is defined as the ability of the two entities to believe one another at some level, and is essential when the entities must interact in a secure manner and must be sure of one another's identity. The main objective of this chapter is to be able to dynamically authorize a user (e.g., physician) with a set of credentials (e.g., authorized to hospital A's EMR) to present credentials to a new system (e.g., hospital B's EMR) and is both trusted and accepted. This chapter explores the use of *Adaptive Trust Negotiation* to support this objective with a goal of obtaining near real-time permission to access a system to which the user has never previously been authorized to. This will be achieved by extending *trust negotiation,* a process that two entities undertake with one another with no previous contact in order to exchange information other than their identities that can be utilized to establish mutual trust between them (Ryutov, et al., 2005) to *adaptive trust negotiation* to allow the system receiving the request to adjust its security policies for the requester depending on the attributes that the requester possesses. In the health care example, a physician could submit attributes (akin to a certificate or set of certificates) that demonstrate sustained legitimate usage of a practice EHR and hospital A's EMR, and this in turn can be utilized by the receiving system to grant access to hospital B's EMR. In one adaptive framework based on TrustBuilder and the GAA-API (GAA-API, 2005), an online business can adjust the purchases a customer is allowed to make based on their previous purchase history (a customer's attributes stored in a certificate) and using that in conjunction with a new customer request (which has its own attributes and certificate) to determine if the new request can proceed. Our specific approach in this paper for adaptive trust negotiation leverages and extends work on adaptive trust coupled with the DIRECT project (DIRECT, n.d. d), in order to allow individuals, providers, and organizations to share information with best practices that have trust and privacy considerations. We use the concept of a Health Information Service Provider (HISP) to handle the sending of health data securely from the sender to the receiver, where the sender and receiver each have access to their own HISP which may be provided by their health organization or may be a service that they purchase from a third party. This can be accomplished by a collection of certificates that is gained upon successful access to multiple EHRs/EMRs as a digital wallet, which is one cornerstone of our proposed approach. We are interested in utilizing adaptive trust to complement our work extending National Institute of Standards and Technology (NIST) role-based access control with collaboration (Berhe, et al., 2010) to securely model provider interactions with one another using different HIT systems and our work on role-based access control for XML as applied to a health care setting (De La Rosa Algarin, et al., 2013).

The remainder of this chapter has five sections. In the *Background* section, work relevant for this chapter is presented, including: role-based access control, identity management, public-key cryptography, identity certificates, and the web of trust. In the *Trust in Medical Systems,* the healthcare domain is explored in the context of the concepts from the Background section, allowing the reader to understand the way that sensitive patient data can be securely accessed and shared across multiple HIT systems, for use in the remainder of the paper. Next, we present *An Approach for Adaptive Trust Negotiation* to provide a means for users in time-critical situations to obtain access to read data other systems to which they are not authorized to utilize; this is illustrated by continuing with the health care scenario. Then, the *Future Trends* section identifies new approaches that are evolving in regards to identity control that include the move towards single sign on, the utilization of biometrics, the ever-growing impact of mobile computing in daily life which will necessitate more advanced authentication, and the usage of spatial/temporal information to determine privileges. Finally, the *Conclusion* section draws this chapter to a close.

BACKGROUND

This section provides background information of a range of security concepts that are relevant for the chapter. First, we review role-based access control (RBAC) (Ferraiolo, et al., 2001), which provides a means to characterize what a user does by role, and could be useful in proving trust by representing the abilities that a user has been authorized to. Second, we discuss attribute-based access control (ABAC), which shares similarities with the underlying trust negotiation mechanism. Third, identity management (Rouse, 2013) which involves controlling users' rights and restrictions based on their established identity. This could be useful in a trust setting by providing the ability of a user to prove who s/he is. Fourth, to realize identity management, public key cryptography (Rivest, Shamir, & Adleman, 1978) can be utilized, which provides a verifiable way for a trust approach to identify a user by being able to decipher a message with a published key. Fifth, to realize public key cryptography, identity certificates (Housley, et al., 2002) are utilized, to provide the proof of identity, another means to support trust by presenting a certificate that can be verified. Finally, to pull many of the concepts together, the web of trust model (Rouse, 2014) and chain of trust model are described, which provide assurance regarding certificates that are distributed across the Internet.

Role-based access control (RBAC) (Ferraiolo, et al., 2001) is a method for restricting access to secure resources. In RBAC, access is granted based on sets of user roles and permissions. The roles in an RBAC model represent the job that a user is expected to perform along with the implied abilities needed to perform those tasks. These permissions may grant the ability to view patient records, access the e-prescribing system, access auditing logs for review, or manage other users in the system. Users within the system are assigned one or many different roles. Likewise, each role is assigned a set of permissions which represent the ability of a user who was assigned the role to access different parts of the system. When a user attempts to access a resource, the system first checks their role and whether their role has the permission to access that resource. If the user's role does have the required permission to access the requested resource within its set of permissions, then access to the resource is granted to the user. Otherwise the system denies access to the user. RBAC utilizes the concepts of user roles and permissions to organize the security constraints within the system and to simplify user administration. Here, access control is defined as a method for controlling a user's ability to utilize some resource that is only made available to certain users. This restriction of access allows system administrators to protect that resource from improper usage by denying access to unauthorized users. Although RBAC generally simplifies access control since permissions are encapsulated into roles, RBAC may become unwieldy when a system requires multiple similar roles that only differ in the addition or deletion of a small number of permissions. In this case, the number of defined roles may outnumber the number of permissions as the system administrators create roles for many different combinations of permissions. As a result, managing the roles becomes a wieldy administrative task.

Attribute-based access control (ABAC) (Hu, et al., 2014) is an access control method where user access to a resource is determined based on: a set of attributes the user possesses, a set of attributes the resource possesses, the operation to be performed on the object, and a policy. In ABAC, each user has a record of attributes that describe the user such as: organization affiliation, licensing status, or security clearance level. Resources within the system are also assigned a series of attributes. When a user requests access to a protected resource, the access control mechanism uses the policy to read the user attributes, the operation, and the object's attributes and decides whether the user gains access to the resource. Additionally, the policy may be affected by external environmental conditions, such as the time of the request or the user's location. In ABAC, the user's identity may be encoded into his/her set of attributes, or the user's identity may be irrelevant depending on the implemented policy and the resource the user is attempting to access. Rather than the user being authorized explicitly to resources as through RBAC, in ABAC, the user is authorized through the set of attributes created by the credential manager. RBAC can be thought of as a specialized form of ABAC; instead of assigning personalized user attributes as in ABAC and deciding user access with a policy, more generic permissions (create, read, update, delete a particular resource) are assigned to roles, which are then assigned to the users.

Identity management is the process of gaining and maintaining a recognized identity within a computer system. The identity is a means of communicating the specific person who is attempting to access a computer system. In a secure system, establishing a means to communicate and prove the identity of the accessor is imperative since the system grants access based on a security policy that relies on matching a user's identity to the restrictions or allowances on the user's stored security policy. The user's security policy defines the level of access s/he has been granted to the system. Therefore, there are three parts to gaining secure access within a system: the user establishing that they own a particular identity, a valid record of the identity that the system can trust, and the actions that the identity is allowed to perform within the system. Common methods of user identification include: a pre-chosen username, demographic information such as a first name, last name, and place/date of birth, or an identity certificate. Proof of identity is often enforced with many different authentication factors including: user/password, biometric scans, or a digital signature generated with a user's private key. Proper identification and management can mitigate man in the middle attacks (MITM) through the use of public key cryptography. This is especially important for websites that require users to share sensitive information such as PHI and Personally Identifiable Information (PII). Public key cryptography allows both the user and server to identify themselves while also allowing the user and server to establish a secure, encrypted connection for sensitive data transfer.

Public key cryptography (Rivest et al., 1978) utilizes public/private key pairs to form the basis of identification methods between computer systems and is a form of asymmetric cryptography, which means that the key used to encrypt information differs from the key used to decrypt that information. A user that requires the ability to uniquely identify him/herself and open secure channels with other, potentially unknown users generates two encryption keys, a public key and a private key, using an algorithm such as RSA. Any information encrypted by one of the keys can only be decrypted by the other. A required property of the key pairs is that there is no method to determine the structure of a key given only the key's matching pair. The public key can be safely disseminated to the public while the private key is known only to the user that generated the key pair. Public keys are generally made available to the public through a file called a certificate, which binds the public key to an identity and allows others to verify that they are communicating securely with the intended identity. During initial communication, a user that wishes to communicate a request securely to a server or another user first retrieves the public key of the entity s/he wishes to communicate with by retrieving its certificate, verifying the certificate, and extracting the public key from the certificate. After verifying that the public key belongs to the identity of the intended recipient, the request s/he wishes to secure is encrypted using the recipient's public key and the encrypted request is sent to the recipient. Since the request has been encrypted with the recipient's public key, only the holder of the corresponding private key can decrypt it. Note that this exchange of certificates is computationally expensive when the server must maintain secure connections to many users, so in general this initial secure exchange is utilized to allow the user and server to agree on an encryption key to be used with a symmetrical encryption algorithm, which can encrypt and decrypt information much faster. Since the agreement on this encryption key takes place over a secure channel, it can be assumed that only the user and initial recipient are in possession of the symmetrical encryption key, as eavesdroppers on the connection are unable to discern the content of messages passing through.

An *identity certificate* (Housley, et al., 2002) is a file that contains information on the certificate holder's identity, the period of time that the certificate is considered valid, the certificate holder's public key, and the certificate signer. It is the certificate signer's responsibility to vet the holder's identity independently. In order for a certificate to be considered valid by others, it must be digitally signed by a trusted signer who endorses the correctness of the information contained within. The signature is used both to verify the integrity of the certificate as well as to verify the identity of the signer. Identity certificates are the most common type of certificate and are typically encountered by the average person when browsing HTTPS-enabled websites. Identity certificates allow these websites to uniquely identify themselves to the browser and initiate an encrypted connection with the user for added security when users must supply passwords or credit card information. Browsers often have an icon to indicate when a secured connection is present and will warn the user if the server's certificate does not validate properly. In practice, these certificates bind the server's identity to the website's domain and vetting of the certificate holder only confirms that the owner is in possession of the domain. Extended validation certificates have been introduced which offer the user assurance that the owner's identity has been vetted more thoroughly. Generally, holders must prove their legal identity to the certificate signer as well as establish proof that s/he owns the domain in question. Browsers that recognize extended validation certificates often display the owner's name in green next to the URL. A certificate used in a web of trust model may contain multiple signatures, each a signer that endorses that particular certificate.

The *web of trust model* provides an infrastructure to enable the correctness of the certificates to be verified allowing trust to be established with the signer of the certificate in order for the certificate to be considered valid. In the web of trust model, any individual can endorse another's certificate by signing a hash of the certificate with their private key. Each user owns and maintains a certificate store, a collection of the certificates owned by those that the store owner deems trustworthy. During the certificate validation process, if the certificate being verified matches a certificate within the certificate store, the certificate is automatically accepted as valid. If the certificate is not present in the certificate store, then the signatures present in the certificate must be inspected until a signing certificate is found within the certificate store. This process can continue for several levels deep, until a trusted certificate is found or until a limit is reached and the verification process decides that trust cannot be established. A certificate that a user has placed within his/her certificate store is said to be trusted implicitly by the user. In the web of trust model, all users are considered peers and there is no hierarchy for trustworthiness, unlike in the chain of trust model. While this means users have complete control over those that are trusted, it also forces users to manage trustworthiness themselves. In a large network, it may be difficult for new users to decide which other users are trustworthy and it may be difficult to react to changes within the trust network if it becomes fragmented and difficult for new users on the network to gain trust with others. If a user is unable to gain trust they will have issues communicating on the network. This problem is partially mitigated by key signing parties, in which large numbers of users choose to trust and endorse each other's certificates.

As a result of these drawbacks, the web of trust model is not recommended in an interconnected hospital environment. It is unlikely that hospitals would want to spend resources maintaining the web of trust on the wider network. Maintaining their own trusted certificate store could also open hospitals to legal liability in the event that a breach occurs as a result of the maintenance of the trust network. The peer to peer nature of the network also means that in order to enter into the network, new medical facilities would need to have their certificates endorsed by other medical facilities, which in the medical field would require extensive vetting of the new facility that established facilities may not have the time or resources for. Instead, a chain of trust model would offer a more organized trust network while offloading the work needed to maintain the trust network to a third party.

The *chain of trust model*, like web of trust, provides an infrastructure for the dissemination and verification of certificates utilizing the X.509 standard for certificate structure. In chain of trust, users are not peers and are unable to endorse each other's certificates to others. Instead, the concept of a Certificate Authority (CA) is used. In the chain of trust model only a CA can generate and sign new certificates. Any certificate conforming to the X.509 standard may only have one signator. During the process of obtaining a new certificate, the user passes their public key along with any information required to vet their identity and ownership of that identity to the CA. The CA vets the information passed by the user and if it is accepted, the CA creates a certificate, signs it with the CA's private key, and delivers the certificate to the user. A root authority is a CA with a self-signed certificate, or a certificate in which the holder information matches the issuer information. These certificates can only be verified against themselves since the certificate signer is also the certificate holder. These certificates must be placed in the user's certificate store, and are only trusted during certificate validation if present in the store. To validate, the CA's own certificate must indicate that the signer has authorized the certificate holder to issue certificates in the basic constraints section. A CA may choose to endorse a certificate holder as a trusted CA by creating a certificate that indicates that the holder of the certificate is authorized to sign certificates on their behalf. During the certificate validation process, the verifier receives the entire "chain" of certificates from the root authority's certificate to any intermediary CA certificates to the user's certificate. The chain is validated by checking the validity of the user's certificate, then walking up through the chain of signatures validating each certificate until the root certificate is found, at which point it must be found in the certificate store or the whole chain is invalid. In the event that a user's private key is stolen, compromising the security of his/her identity certificate, the certificate is added to a revocation list and the user must generate a new public-private key pair and request a new certificate. Any certificate listed on the revocation list is considered invalid.

The chain of trust model is a much better candidate to secure a large-scale medical network. The CA system relieves the hospitals of the responsibility of performing vetting procedures on the entities signing the certificates. The ability to delegate authority to sign certificates enables the creation of an intermediary medical authority, which can be trusted to vet the certificates used to identify the hospitals themselves with a more medical-centric focus. This would allow professionals sharing medical data in the network to know that not only are they accessing the correct domain, they are accessing the correct domain of an entity endorsed by the medical authority as a trustworthy member of the medical subnetwork. Rather than requiring the hospitals to obtain the endorsement of a multitude of peers to prevent a disconnected network, hospitals can simply trust a more limited number medical authorities, just as users on the internet need only trust a limited number of CAs in order to enjoy the benefits of secure web browsing. However, the chain of trust model is dependent on each certificate in the chain being valid. This has the potential to create a single point of failure; if any of the CA's certificates in the chain become invalid, the user's certificate becomes invalid. This contrasts with the web of trust approach where a certificate can be signed more than once and the invalidation of one signature does not automatically make the certificate invalid.

TRUST IN MEDICAL SYSTEMS

Supporting secure identification, user authentication, and authorization among many decentralized, distributed systems is a complex undertaking, particularly in a domain such as health care whose systems were built independently and to different specifications. The hospitals that purchase or commission these systems often customize them to mirror the workflow of the various physicians, nurses, and hospital administrators in order to improve efficiency and reduce medical mistakes. Several major EMR providers such as Epic (Epic, n.d.) advertise their commitment towards customizing their software to fit their clients' workplace culture and workflow. However, it is important to note that there are two significant schisms that often hinder sharing of information across systems in health care. First, hospitals located in a region are often competitors of one another, and free sharing of information may mean that hospitals could lose patients; so while all PHI is available to a patient via HIPAA, its free exchange electronically is not. Second, different EHR/EMR vendors are focused on vendor specific approaches and proprietary formats that inherently limit the ability to share data. The HIT vendor community has been historically hesitant to adopt practices in their development and deployment technologies that are well accepted in other fields through agreements on standards that allow data to easily flow among different systems in a seamless manner. For example, the SQL Standard (ANSI, 1986 and ISO, 1987) allows for the easy porting of databases through ".sql" files. Also, the eXtensible markup language (XML) (W3C, 2015), a de facto way to exchange information, is utilized by many commercial database vendors (MySQL, SQL Server, Oracle, etc.) to provide the ability to export an entire database schema and repository into XML, at which point the database in that format can be moved from one database platform to another in a seamless manner.

Despite these schisms, within the medical domain there is an increasing interest (Mettler & Rohner, 2009) in the secure sharing of patient data among the numerous stakeholders in the health care system to improve the effectiveness of patient treatment and patient satisfaction. Patients are often frustrated when they are required to provide the same information multiple times to different specialists simply because the patients' physicians are unable to access their previously existing medical records. The stakeholders include: physicians treating patients, medical specialists working in a dynamic coalition to treat more complex and specialized patient conditions, medical researchers studying the effectiveness of different medical procedures, insurance companies who use these studies to promote more financially efficient healthcare, the business sector of health care which is interested in improving the effectiveness of treatments and profitability of its hospitals (Deloitte, n.d.), and governmental organizations (NIST, 2014) who wish to use large amounts of medical data in order to observe trends in public health and patient treatment. Ideally, stakeholders would be able to access medical records for any of their clients on a nationwide health network where these stakeholders would share their data for the benefit of patients, and health care as a whole, and the act of sharing data would lead to improvements on current medical procedures.

While these stakeholders are interested in sharing medical data for various reasons, there are complex laws such as HIPAA that dictate the way that medical data must be stored and secured to prevent private patient information from being obtained for nefarious purposes. Patients themselves are also concerned that their medical records could be misused if leaked and accessed by the wrong people. Hospital administrators would like to choose what data is available and to whom, since they have spent resources obtaining this data and the data itself has value. Although in practice a hospital's medical data is generally available to the doctors and nurses, they must have a specific reason to examine a particular patient's records; some records, such as psychiatric history and medications, are even more restricted. Either the records belong to a patient being treated by the physician in question, or the patient's medical data is somehow related to a task the physician is expected to perform. Examining a patient's record even at a physician's own practice simply for curiosity's sake is strictly illegal and all data access must be audited by the system to ensure that no employee of the organization is accessing data illegally. Hospitals are unwilling to leave themselves exposed to lawsuits regarding the improper use or dissemination of patient data. Consequently, if an interconnected health care system is to be created, the hospitals involved in the network must possess some method of verifying a physician's identity (even if the physician is previously completely unknown to the hospital), verifying the physician's need to see medical data, and auditing access.

The nature of this problem requires that every hospital in the medical network be able to establish the truth of a claim to the identity that the requesting physician claims to own. In such a situation, the hospital has two options under the traditional username/password combination form of authentication: the hospital can maintain its own personal collection of identities and shared secrets for each physician authorized to access the medical system in a database, or the hospital can cede this power to a trusted third party which would maintain a repository of physician identities and proofs of identity and offer an API for hospitals to access during an authentication request.

The latter situation is infeasible for a multitude of reasons. First, a single point of failure in the system is created by concentrating user authentication into one mutually shared service. Any medical system must be robust to the failure of any individual component because medical information must be accessible even in the event of failure occurring within the system. Second, it would be very difficult for the idea of a third party authentication service to gain traction within the medical community since hospitals guard their data very closely and are unlikely to allow others to authenticate users in their stead. Finally, a large undertaking for any single authentication service would be required to vet the identities of a large number of users. Any centralized service must have the ability to quickly vet new applicants to the system as well as managing any identity that may be compromised. Given the large number of physicians currently working in the medical community it is unlikely that any single entity would be able to manage this efficiently.

Another problem with the approach is that hospitals would be required to maintain its own records for each physician afforded access to their system. This would mean not only medical personnel that work at the hospital but other medical providers outside of the hospital that would be interested in seeing information on a single patient particularly in a life-critical situation. The hospital or its security administration would have to independently verify each physician that requests access to their system, meaning that physicians who need access to medical data immediately but have no previous history of interaction with the new hospital's system would need to apply for access. This method is clumsy, slow, and completely unworkable in the event that the physician is working in an emergent situation and his/her patient needs medical attention immediately. Furthermore, the medical network as a whole would run inefficiently, as every hospital in the system would need to maintain its own records for every physician who works within the system, forcing hospitals to independently verify the same physicians multiple times and storing more data on them within their systems than necessary. Finally, in the event that a physician's identity is compromised, it would be difficult to manage the repercussions of a stolen identity throughout the entire hospital network. Should an identity be compromised, every hospital that has a record of the physician in their system would need to be alerted of this breach of security and respond accordingly by revoking access until the physician's identity can be secured and reestablished. With each hospital maintaining their own separate records, it becomes difficult to respond to problems such as this.

One other possible way to address this issue is via RBAC, which is widely used in the medical field due to its simplicity and the fact that its model closely mirrors the manner in which information is controlled within the hospital environment. In the medical field, those requesting access to hospital resources are: practicing physicians, surgeons, nurses, insurance companies, medical researchers, or administrators. Each of these people need to be able to access some part of the hospital's secure computer systems in order to properly execute their tasks, but none of them requires complete, unbridled access to the entire system. By restricting their access to that of only the part of the system they need to properly execute their tasks, their ability to compromise the hospital by inappropriately exposing or editing sensitive medical data is reduced. For instance, in a hospital setting, there may be a separate role for the physicians, a role for the nurses, and a role for members of the hospital administration staff. Permissions represent the ability to access a resource within the system. Resources that need to be protected from this type of damage may include: patient records (both regular and psychiatric), billing systems, e-prescribing systems, or de-identified data repositories. These resources must be controlled so that unauthorized users are unable to compromise the privacy of patients in the system and to ensure that the hospital is in compliance with HIPAA regulations regarding proper protection and dissemination of patient data. Improper access to patient records or billing is a breach of patient privacy that could result in legal action against the hospital, while an unauthorized user accessing e-prescribing systems could illegally prescribe themselves drugs. Owners of de-identified data repositories are interested in protecting their research data since it is gathered at great expense and is necessary to conduct research on the effectiveness of medical treatments or leveraged to find newer, more effective treatments. The proper restriction of access to these systems protects the hospital from damage.

Another secure and efficient way to address the problem of identity verification is to leverage public key cryptography and associated infrastructure with a certificate system whereby each physician obtains an identity certificate which is digitally signed by a hospital that has allowed the physician access to its medical data. Under this scheme, the hospital is responsible for verifying the identity of the physicians on its payroll. Creating and signing a personal identity certificate for the physician is proof that the hospital has vetted their physician's identity and proof that the hospital trusts the physician to access sensitive medical data contained within their system. During a request for data access, the physician can present their identity certificate, use the TLS/SSL protocol to identify him/herself, and establish a secure, encrypted connection to the server. During validation of the identity certificate, the server receiving a request for medical data can inspect the physician's identity certificate and prove its validity by inspecting the physician's hospital's digital signature. If it is valid and the server trusts the physician's hospital, the server can be certain that the presented identity is in fact the physician's actual identity.

The hospitals can establish trust with each other through mutually shared medical authorities, certificate authorities that provide proof of identity in the form of an identity certificate for hospitals involved in data transfer. This allows hospitals to identify themselves securely to hospitals that they have no previous access history with and eliminates the n^2 problem posed by independent verification by reducing the number of background checks that must be performed on the physicians to only the physicians that the hospital itself employs. The data that must be stored in the hospital's systems is reduced to a certificate store containing the root certificates of a small number of medical authorities. Hospitals establish trust with each other if the certificate of the medical authority that signed another hospital's certificate is found within their certificate store during the certificate validation process.

A real world scenario that will be built on for the remainder of the chapter describing the necessity of such a system is as follows: A physician working at Hospital A working in the ER of a major hospital receives a patient in critical condition. The patient has no medical history at Hospital A, but the patient's identity has been discovered using some form of personal identification. Using a Master Patient Index (MPI), the physician is able to locate the patient's records at Hospital B. However, the physician has never been in contact with Hospital B before, and both the physician and Hospital A are completely unknown to Hospital B. The physician now faces the challenge of identifying him/herself to Hospital B in able to gain access to the patient's data in a time critical situation. Under the current system, it would be impossible for Hospital B to properly examine and validate the physician's credentials in a short enough timeframe to be useful to the physician and patient. In this situation, the physician would be forced to treat the patient without access to the patient's medical record.

AN APPROACH TO ADAPTIVE TRUST NEGOTIATION

In this section, we present an approach to *Adaptive Trust Negotiation* that can be utilized as a means for users in time-critical situations to obtain access to read data in other systems to which they have not been not previously authorized to access. Our approach focuses around three concepts and leveraging ongoing work on the DIRECT project. The first concept, *identity management*, is utilized to allow users to construct a unique identity for themselves as well as relevant information pertaining to that identity for the purpose of gaining access to new systems. The second concept, *trusted identities*, can be employed to uniquely identify a user by including user relevant attributes. These attributes are matched to a security policy constructed by the facility the user wishes to obtain access to and the trusted identity offers assurance that the information provided is trustworthy. The third concept, *attribute certificates,* provides the ability to build a profile of a user in terms of their usage of information, which is an important part of our proposed approach to adaptive trust negotiation. To bring these concepts together, we explain and apply existing ongoing work on DIRECT (DIRECT, n.d. c) which allows individuals, providers, and organizations to share information in a trusted manner. The section concludes with an examination of potential security threats against our approach and strategies for mitigation. Throughout the discussion, we continue with the health care scenario from the previous section.

In terms of *identity management* and *trusted identities*, there are wide range of stakeholders that require access to various kinds of health care information, including: patients, medical providers (e.g., physicians, nurses, therapists, etc.), various laboratory and testing facilities, medical researchers, and support personnel throughout the health care system. Across this wide spectrum of stakeholders are individuals that create and access all types of information stored in EHRs and ancillary systems such as imaging, laboratories, etc., and include: structured text, free text, images/scans, test results, prescriptions, PHI, PII, etc. Three complementary requirements dictate the need for identity management and trusted identities: the need to increase the availability of the data in emergent situations in real-time; the ability to provide that data from multiple sources via HIE while insuring that the security and privacy of patient information is protected at all costs; and, the ability to dynamically authenticate to use a system never previously authorized. As personalized medicine (McCarty & Wilke, 2010) increases, there will be need for all types of genetic information to be available and accessible. The work on adaptive trust presented in the chapter seeks to overcome a significant barrier to integrated patient care data access: when any of these stakeholders seek to access information from some source they are often bound by the authentication credentials utilized to access their own specific systems and are not easily able to be authenticated to access information from systems that they have not been previously authorized to use. Our approach to *adaptive trust negotiation* is to focus on X.509 certificates and their ability to be extended via *attribute certificates*. This allows for adaptive certificates to be dynamically generated as needed for authentication in order to provide a level of trust as users attempt to utilize information from sources to which they have yet to be authorized to in a dynamic/real-time manner.

Our approach to adaptive trust negotiation focuses on the idea that a user will acquire multiple X.509 certificates over time based on their activity being authorized to utilize different systems. Each certificate allows access to a specific system, e.g., a physician would acquire certificates from the EMR at his/her medical practice, the EHR at the hospital he/she has privileges, the EMR at the free clinic that he/she volunteers at, etc. Each of these X.509 certificates can be augmented with multiple *Attribute Certificates* (AC) (Farrell & Housley, 2002) that are secure documents containing attributes associated to the holder by the issuer structured using X.509 and signed by an *Attribute Authority* (AA). The advantage of multiple certificates (one per work setting/EMR/EHR) is to minimize the impact for failure; with a single certificate with multiple attribute certificates (one for each work setting), failure may compromise multiple settings, while with multiple certificates (one per work setting), failure of one should have no impact on the others. In a realistic scenario like health care, each work setting (EHR/EMR/HIT system) can have their own security infrastructure and algorithms to generate public-private keys; the concept of multiple certificates each with multiple ACs attached is akin to a wallet with multiple cards issued from different sources (Mavridis, et al., 2001). Our work also considers Oblivious Attribute Certificates (Li & Li, 2006), where a certificate holder can select which attributes to use and the way to use them, and the user obtains a service if and only if the attribute values satisfy the policy of the service provider, yet the service provider learns nothing about these attribute values and the certificate holder learns nothing of the requirements specified in the service provider's security policy. This too could be an approach used by the certifier in consultation with a third party to allow the emergent authentication. Related efforts in adaptive certification include: a framework for secure e-Health authentication using a multiple factor approach where physicians would provide multiple pieces (e.g., ACs) of information in emergent situations (akin to our multiple certificate approach) (Boonyarattaphan, Bai, & Chung, Sept. 2009); a framework for adaptive trust negotiation that establishes trust based on attributes other than identity (Ryutov, et al., 2005); and, a dynamic adaptive authentication policy that could be utilized as a model for a certificate-based identity representation (Ventuneac, Coffey, & Salomie, 2003).

Our approach to adaptive trust negotiation is similar to the concept of ABAC discussed in the *Background* section. The attributes a user gains throughout his/her career are recorded in the user's attribute certificates, which are read by a policy that decides whether the user is allowed to gain access to the requested health records. ABAC can establish a method for an unknown user to gain secure access to an unknown system since ABAC does not depend on identity, as long as the system can verify the accuracy of the credentials and the credentials follow a pre-established standard. However, our approach in this chapter allows for the user to generate new credentials dynamically as a natural step in the authentication and authorization process. This creates more options for security management by incorporating the ability to tailor the service provider's actions in regards to data dissemination to the request. Instead of simply deciding whether the user is allowed access to the resource in question as in a traditional access control scheme, we utilize adaptive trust negotiation to allow a policy that can also determine a level of access to a requested resource and allow additional actions to be undertaken, such as audit notifications, depending on the level of trust established during the negotiation phase.

To place this into a real context, consider the example given in Figure 1, where on the left side, a particular user (e.g., a physician) has been granted multiple certificates in his or her multiple roles (e.g., clinician to treat patients, on-call physician, hospital service, medical researcher, etc.) across different systems, including a EHRs at an ambulatory practice (Asylum Hill Family Medicine), a EHR at a hospital (St. Francis Hospital and Medical Center), and a de-identified research i2b2 data repository that includes both clinical and genomic data (UConn Health Center;). Holding these three certificates, the physician is able to present some or all of these certificates to the verifier to allow authentication to utilize a system that s/he does not yet have access (right side of Figure 1). This process would be adaptive in allowing the certifier to utilize a third party to verify the existing credentials and use them to dynamically generate new credentials (certificates) on the fly. For example, in a emergent situation where a physician needs to access data at another hospital (Hartford Hospital), a physician could present his/her X.509 certificates for his/her practice EMR (Asylum Hill Family Medicine) and hospital (St. Francis Hospital and Medical Center) and based on the information in these certificates, a process can be defined which is initiated by the certifier that will allow a third party to analyze the presented X.509 certificates (along with their attribute certificates) and dynamically decide to generate a new X.509 identity certificate with attribute certificates or only a new attribute certificate based on the presented X.509 certificate associated with Hartford Hospital to satisfy the verifier and allow the physician access. In the right side of Figure 1, when the physician requests access to data from Hartford Hospital by sending a subset of his/her certificates, the hospital verifies the certificates through a recognized certificate authority (CA) and then generates a new X.509 certificate with respective AC or only AC to the physician granting access to the requested data. These certificates are then passed back to the physician and are added to the physician's digital wallet. The physician now has permission to access Hartford Hospital's data, and additionally can present these new certificates in subsequent accesses to other hospitals to aid in the verification process

Figure 1. Example Scenario with ACs

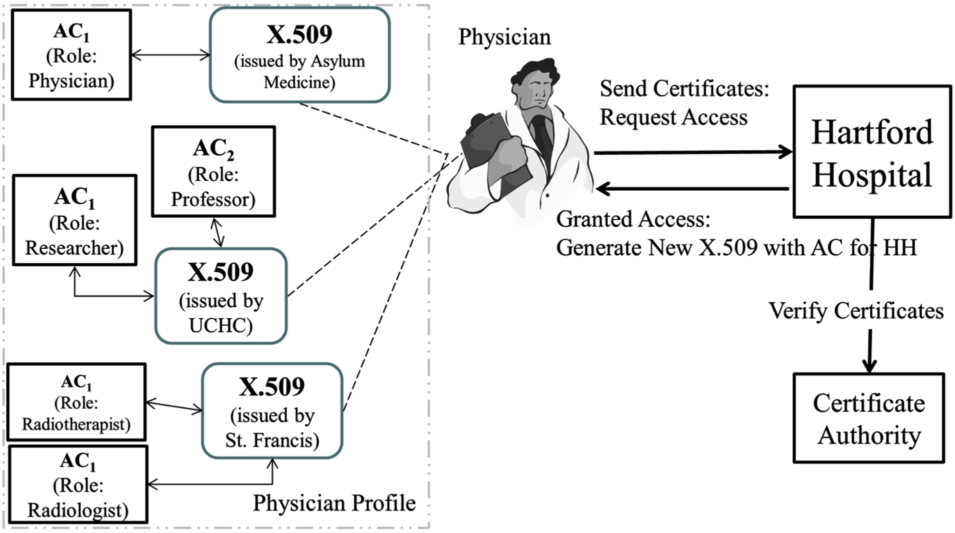

Traditionally, adaptive authentication research focuses on biometric approaches that act as second or third steps in user authentication after a more traditional approach such as usernames and passwords. For example, research in (Boonyarattaphan, Bai, & Chung, Sept. 2009) provides a framework for secure e-Health authentication using a multiple factor approach where physicians would provide multiple pieces of information in emergent situations (akin to our multiple certificate approach). In the work of (Ryutov, et al., 2005), a framework for adaptive trust negotiation is presented which facilitates "… authentication by establishing trust based on attributes other than identity"; again this mirrors our multiple certificate approach and the use of a third party by the certifier. Both of these efforts can be applied to our approach in terms of the identity X.509's and the degree of robustness that they are able to support in the adaptive authentication process. Lastly, work in (Ventuneac, Coffey, & Salomie, 2003) presents a dynamic adaptive authentication policy that could be utilized as a model for a certificate-based identity representation such as our proposed multiple certificates. Their work utilize a two-step authentication method (first username and passwords, then the X.509 certificate against an LDAP registry), but it can be modified by removing the first step and adding credential checking after the LDAP lookup to allow the certifier to verify with a third party and generate a new certificate to allow authentication to the new system in an emergent situation.

Our approach to achieve this adaptive and dynamic certificate generation and verification process is to leverage existing ongoing work on DIRECT which is to allow individuals, providers and organizations to share information (DIRECT, n.d. a) and provides a set of best practices (DIRECT, n.d. b) that have trust and privacy considerations that are very consistent to the privacy emphasis of our work. Health Information Service Provider (HISP) has been used by the DIRECT project, both to describe a function (the management of security and transport for directed exchange) and an organizational model (an organization that performs HISP functions on behalf of the sending or receiving organization or individual). HISP is similar to the internet IP domain providers who are responsible for managing network connection between the user and the internet services. A HISP is a separate business organization from the sending and receiving organization required to have Business Associate Agreements (BAAs) with HIPAA Covered Entities as shown in the Figure 2. All HISPs must have contractually binding legal agreements with the sender or receiver of directed exchange of PII, including all terms and conditions required in a BAA and including other protections as noted in this best practice. The Privacy and Security Tiger Team (HealthIT.gov, n.d.) has published recommendations on transparency on use, retention, data handling, and other activities of HISPs. HISPs must include all data collection, use, retention, and disclosure policies (including rights reserved but not exercised) in BAAs or other service agreements. In a simplistic

Figure 2. Overview of HISP and HIE

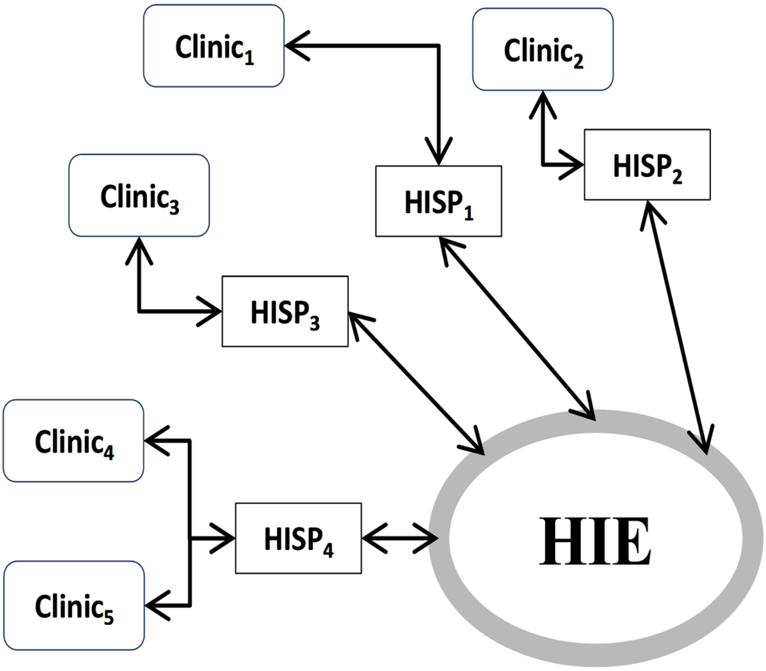

situation, the sender and receiver take sole responsibility for encryption/decryption activities. The interpretation of the encrypted package, by any parties other than the sender or receiver is exposed to data that are "rendered unusable, unreadable, or indecipherable to unauthorized individual through the use of the NIST-recognized AES128 or AES256 encryption algorithms".

For our approach to adaptive trust, we can transmit *identification information* encrypted using X.509 certificates across HISP to various hospitals for authenticating sender's credentials. X.509 certificates can be augmented with *Attribute Certificates* (AC) that are secure document containing attributes associated to the holder by the issuer and are invalid without a proper authentication from the issuer. A collection of AC certificates in the course of time will build a *profile* for the physician over time that could be passed to the certifier which would consult a third party to dynamically and adaptively provide a new certificate for the system that the physician has yet to be authenticated (or authorized to previously use) as was shown in Figure 1. In this scenario, the physician holds three X.509 certificates: one issued by Asylum Medicine with an AC with a Physician role, one from St. Francis hospital with an AC with a Radiologist role, and one from UCHC with AC with a Researcher role). In the example, we used role as an AC for the three different certificates. In a health care domain, there may be ACs for: HIPPA accessibility, de-identified data only for research, patient clinical data for emergent care, genomic access, number of trusted logons, etc., or defined with broader categories such as ACs for the data level (HIPPA, FERBA, DE-IDs), an AC for the situation (Urgent care, Primary Care, Inpatient Care), the type of data (patient, genomic), etc. A certificate for a given system would denote the types of access that a person has had (we can define lots of types) and the combination of all of these from multiple certificates would allow the certifier, through a third party, to generate a certificate for the system to be authenticated against that the user needs. Such an approach to processing attribute certificates as shown in the right side of Figure 1, would be overlaid to be an integral part of the HISP/HIE architecture as given in Figure 2.

In the HISP/HIE environment, extended with attribute certificate to support adaptive trust notification, each entity (EHR/EMR/HIT system) that holds private medical data that it wishes to make available for use by others is responsible for maintaining an authority through which it can issue new certificates and defining a security policy. In the event that a request is received and the user's certificates have been properly validated, the entity's authorization system is now in possession of the subset of the user's profile that s/he has chosen to present during the request for access. However, the security policy may authorize different levels of access depending on the credentials presented, allowing it to be adaptive in its analysis. For example, a physician at a local practice recently acquired in a business deal by a larger hospital may be granted full access to the larger hospital's EMR since their patients are now the hospital's patients and vice versa. In contrast, a physician attempting to access the EMR of a different hospital entirely may only be allowed to access a data warehouse specifically provisioned with a subset of patient records that the hospital has authorized for outside use, and the access may trigger an additional level of auditing on the data retrieved by the physician. The security policy offers a level of automated flexibility, allowing medical data to be shared securely when needed while also protecting private data from unwanted access.

In addition to the data recorded in the physician's ACs, the trend towards the utilization of mobile devices in medical applications demands their consideration during the adaptive trust negotiation process. Figure 3 presents an Hybrid Extensible Architecture (HEA) for adaptive trust for use in HIE where a total of three HISPs have been included above the various HIT systems (Hospital A, Hospital B, Local Pharmacy, Medical School, i2b2 repository, and Physician Portal). HEA is primarily comprised of remote data servers that have been connected by health providers in a HIE network through a set of common health standards. The HISPs at the top of Figure 3 are establishing trust amongst the HIT systems each contain a HISP focused

Figure 3. Overview of a secure HIE solution

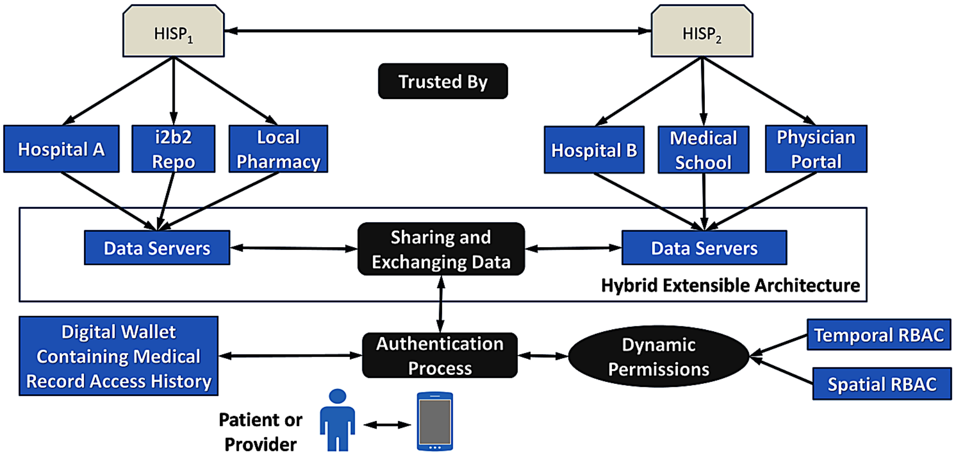

certificate authority (CA) to oversee certificate interactions for the connected HIT systems in support of the adaptive trust negotiation process. This allows a user to access information he/she is not previously authorized to via the attribute certificates as previously described and shown in Figures 1 and 2. In addition, the CA of each HISP also interacts with a more global certificate authority (upper center of Figure 3) in order to coordinate certificate exchanges and verifications across HISPs.

Figure 3 also explores many other potential issues related to the secure access of information particular in regards to the different stakeholders that are likely attempting to access via mobile devices. For example, a patient, accessing a secure patient portal on his/her tablet mobile device, or a health provider/physician accessing an EMR through their mobile device, may require the initiation of a versatile and changeable authentication process dependent on various factors. Specifically, the authentication process is likely to involve the addition of a dynamically changing set of permissions that change based on the user's physical location as well as the time period during which authentication occurs. In a medical office or practice, a physician's set of permissions may change when his/her shift ends or during a time period during which he/she is on call. At the practice during the day, the physician is focused on his/her patients. At night, the physician may get calls from any of the patients in the practice, or may receive calls from other practices that are being covered for care and may require access to EHRs/EMRs at other practices or hospitals. This is another situation where adaptive trust negotiation with attribute certificates may have a significant role. Likewise, permissions may change when physicians travel between different portions of the hospital during their shift (e.g., a physician moving from a research wing to pediatrics). A successful authentication allows access to a Hybrid Extensible Architecture with an underling HISP infrastructure from DIRECT thereby supporting adaptive trust negotiation. Trust between the health providers in the HIE is provided through a set of common medical authorities, in this case the HISPs, that undertake the responsibility of vetting health organizations and enforcing strict security regulations.

In order to be implemented as an authorization method, the trust negotiation process must be resilient to attack, both external and internal. An external attack in the scenario presented is defined as an attempt to obtain sensitive PHI without participating in the trust negotiation process or otherwise disrupting the trust negotiation process in a manner that makes it impossible for a user and healthcare organization to complete the trust negotiation process. An internal attack is an attempt to obtain sensitive PHI by manipulating the trust negotiation process. Examples of an external attack are man-in-the-middle (MITM), stealing the private key and certificates of a digital wallet owner, or overloading the healthcare organization's connection with a denial-of-service (DOS) attack. An internal attack could involve manipulating CAs or AAs to sign false information, or making repeated trust negotiation attempts to ascertain the security policy of the healthcare organization. A MITM attack involves an outside observer intercepting communication between the user and the healthcare organization. The attacker may attempt to steal either the user's digital wallet credentials, private key, or the PHI as data is transferred. A MITM attack may be mitigated by incorporating SSL, which uses PKI to both verify that the identity of the server is correct and provide a means of encryption so that an eavesdropper is unable to obtain usable data. The DIRECT project's HISP, mentioned previously in the section, is an entity that assists in the secure data transfer process. Private keys can be protected with passphrases and encoded with a hardware security module (HSM), which is hardware that stores private keys and is resistant to external attack. Should private keys be compromised, the affected certificates can be revoked in a certificate revocation list, as in the X.509 standard. The work in (Ryutov, Zhou, Neuman, Leithead, & Seamons, 2005) provides a methodology for mitigating DOS attacks on a trust negotiation system by creating a scoring system that rates how likely a request is merely a DOS attempt. A DOS attack on a trust negotiation system may arise from opening multiple requests for trust negotiation or disclosing a complex credential disclosure policy.

Since the medical authorities are responsible for ensuring that all of the service providers that participate in trust negotiation are legitimate medical enterprises, responsibility for unauthorized credential signing resides with the medical authorities and the organization. Therefore, organizations authorized to maintain credential signing services must have a large enough bureaucracy of trustworthy individuals to distinguish individuals requesting data from individuals that maintain the credential signing systems. This is to prevent a conflict of interest where a user that requires access to protected data simply generates new, phony credentials without having legitimate access to the protected data for the purpose of illegitimately passing the trust negotiation process. In smaller organizations where this is not possible, such as in a small private medical practice, the users must affiliate themselves with a larger medical organization or a HISP. Additionally, a trust level may be established between the organizations that have signed the user's credentials and the organization from which data has been requested beyond the baseline medical authority endorsement. An internal attack that attempts to discover the security policy of the trust negotiation server may be mitigated by simply denying the user from making multiple attempts at access in the case that access is denied in the initial attempts at access. Note that the work presented in (Li & Li, 2006) describes a method for disclosing credentials such that the user is unable to discover any information regarding the security policy of the provider.

FUTURE TRENDS

As private data continues to be entrusted to interconnected systems, even as attacks against those systems become more common and sophisticated, the measures utilized to protect those systems must also become more sophisticated. As such, there are emerging technologies that are gaining traction and show promise in reducing the complexity of managing the multitude of credentials required for access to each system. There are two advantages to the emerging technologies: to ease the difficulty users have in juggling multiple identities; and to strengthen weak points in the systems against attack. Furthermore, there is a demand for more fine-grained access control to secure these systems by determining system permissions based on user location and time of access. Our objective in this section is to explore emerging technologies in: single sign on, biometrics, the increasing impact of mobile computing, and spatial/temporal information for dynamic privileges.

Single sign-on (SSO) (United States of America Patent No. 5,684,950, 1996) reduces the necessity for users to remember multiple passwords for multiple services by aggregating all of the services a user has been authorized to use into a single log in service. When the user successfully logs in to the single sign-on service, he/she then gains secure access to any compatible service connected to the single sign-on account. Services that may be connected to a single sign-on account include: email, social networking accounts, corporate virtual private networks (VPNs), and e-commerce websites. Single sign-on increases security by allowing the authentication mechanism to use authentication tokens or passwords more complex than a human user would be capable of entering or remembering. Since the single sign-on service is responsible for authenticating the user to all of his/her services, it reduces the number of users utilizing the same low entropy passwords amongst multiple security domains to prevent forgotten passwords. SSO techniques are frequently employed to authenticate users to distributed systems, as in order to fully utilize the system, the user must also authenticate to each physical or virtual server contained within. Well known works on SSO systems for distributed systems include Kerberos (Neuman & Ts'o, 1994) and Shibboleth (Shibboleth Consortium, 2016) and there is ongoing research into adapting SSO for mobile devices (Yu, Wang, & Mu, 2012).

SSO systems are further subdivided into pseudo-SSO systems and true SSOs (Pashalidis & Mitchell, 2003). In a pseudo-SSO, the SSO service merely stores separate user credentials (username/password, certificates) for each service and utilizes these stored credentials to log the user into these services when the user successfully authenticates to the pseudo-SSO. In this case, the pseudo-SSO is merely making requests for authentication on behalf of the user. In a true SSO, an Authentication Service Provider (ASP) handles authentication for the user and the generation of secure authentication assertions. These authentication assertions are utilized by the user to confirm successful authentication to the ASP to the services s/he wishes to connect to. This method requires more formal integration between the ASP and the services, as there is a required level of trust between the ASP and the services that accept the ASP's authentication assertions, and the services must have a mechanism to read the chosen assertion format of the ASP. A true SSO offers a level of privacy to the user by generating multiple authentication assertion tokens for the user, which also allows the user to create separate roles for each identity on the service network. The same assertion may be used to interface with multiple services, or the user may choose to use multiple assertions in conjunction with the same service.

Biometrics (biometrics, n.d.) are also becoming increasingly utilized for multifactor authentication. Biometric authorization involves the utilization of unique biological signature as a method to both identify and verify the user during the authentication process. A biometric system may utilize fingerprint scans, handprint scans, facial scans, retinal scans, gait recognition, or voice recognition for identification. Biometrics possess the advantages that the user generally always has the means to access a system secured with their biometric data and that biometric security is capable of performing identity retrieval and identity verification with only the biometric data. However, the nature of biometric data means it cannot easily be changed in the event that a user's biometric data is compromised. Systems secured with biometric data may put users at risk if attackers are sufficiently motivated to target the users themselves. There is also a danger that biometric access may become impossible if the user's biometric data changes. For instance, losing a hand in an accident would make it impossible to use that hand as a password in the future.

There exist many attempts to unify biometrics systems with mobile devices to allow for secure authorization techniques while partially alleviating the need for a diverse set of user passwords. In (Hwang, Cho, & Park, 2009; Chang, Tsai, & Lin, 2012), the authors utilize Keystroke Dynamics-based Authentication (KDA) as a form of multi-factor authentication. While the user enters passwords, the device takes note of several types of data while the password is being entered including such as password entry speed or accelerometer data, and compares this data to a previously stored digital fingerprint. Should the fingerprints match, the user has been authenticated with both the password and the biometric data obtained from password entry, all while under the impression that only a password was needed. The inclusion of a built-in microphone makes voice biometrics (Authentify, n.d.) a natural security mechanism for mobile devices such as smartphones and tablets. Voice biometrics utilizes a recording of the user's voice converted into a digital fingerprint. When access to a device is desired, the device listens to the user's voice and attempts to match it to the previously stored fingerprint. If the device detects a match, use of the device is granted to the user. In (Baloul, Cherrier, & Rosenberger, 2012), the authors realize speaker recognition through a challenge-based method that can detect and defeat a replay attack.

The ever-growing impact of mobile computing in daily life poses new challenges with regards to authentication, user security, and data sharing. Attacks on mobile devices have been increasing in number (Ruggiero & Foote, 2011) and the nature of mobile devices requires a diverse set of technologies to authenticate users, communicate securely, and enable the translation of health data to mobile-friendly formats. Devices may include: smartphones, tablets, or mobile health sensors. Mobile devices may also be lost or stolen while the rightful user is still considered logged in on the device, and any mobile authentication system would need to properly detect untrustworthy access and reject subsequent requests for data. Allowing for the modification of a patient's medical record from the mobile device further complicates the situation to control PHI, since there must be care to never delete information but to create a longitudinal record of medical treatment (Wiedemann, 2010). Historically, the paper-based medical chart contains an entire history of all visits, test results, etc., and this feature must be maintained in an electronic representation in an EMR/EHR. As a result, the most significant danger that the use of medical data on mobile devices face is deletions that can cause significant damage to the patient's health caused by a malicious change to his/her medical record, which may occur before the change is detected and fixed. Unification of mobile security technologies allows devices to remain secure while recognizing the diverse set of mobile technologies that can be utilized to improve health. The ever-increasing utilization of mobile devices in everyday life for a wide range of activities will require authentication solutions that would augment what has been proposed in this chapter on adaptive trust negotiation via attribute certificates.

The final future trend that is discussed is Spatio-temporal access control (Ray & Toahchoodee, 2007) which refers to limiting the user's permissions based on their location and the time of access. For instance, in the medical domain, physician or nurse permissions to access an EMR varies based on the current time. At different times in the workday, a nurse may be expected to administer medications to patients or meet with recently admitted patients to perform an initial assessment of his/her health. As the nurse's assignments throughout the day change, his/her permissions within the system must change to reflect his/her current task so that necessary data may be readily accessed. The location of the user may also affect user permissions, as explored in (Bertino, Catania, & Damiani, 2005). The location may be the actual physical location of the user, or a more abstract type of location like the user's administrative network location. In a hospital, the user's physical location may grant him/her access to a different set of permissions if s/he is within the hospital building than when they may be accessing the hospital's systems from home. A physician accessing the hospital network from home may only be allowed to access their own financial and compensation data from home, while accessing the hospital network from within the hospital building may also grant access to the EMR. The administrative location, or the position of the user within the administrative network, may also be used to control access as in (Mavridis et al., 2001).

CONCLUSION

This chapter has introduced and presented *adaptive trust negotiation* as a means for a user to access a system to which he/she has never been authorized, in a near-real-time period, allowing the system receiving the request to adjust its security policies for the requester depending on the attributes as represented by certificates that the requester possesses. The concept of trust allows two entities to believe something about one another; we believe such a relationship is critical in a health care setting to allow, for example, a physician in an emergency room treating a patient to have almost instantaneous access to a patient's medical data from an external location in order to ensure that all relevant information is available for successful treatment. The chapter provided information on role-based access control, identity management, public-key cryptography, identity certificates, and the web of trust in the *Background* section. Using this as a basis, the healthcare domain was reviewed in terms of stakeholders, health information technology systems, and the way that information is securely shared, in the *Trust in Medical Systems* section. Both of these sections set the stage for the *An Approach for Adaptive Trust Negotiation* section that introduces the concept of certificate attributes and a certificate authority, describes a time critical scenario for a user to present credentials to a certificate authority to obtain access, and leverages existing work on the DIRECT framework for the supporting infrastructure. In the process, a Hybrid Extensible Architecture (HEA) was presented as a means to bring together all of the work in this chapter into a realistic context. To complete the chapter, the *Future Trends* section discussed identity management issues including: single sign on, biometrics, the increasing impact of mobile computing, and spatial/temporal information for dynamic privileges.

REFERENCES

Anthem. (2015). *Anthem Facts*. Retrieved from https://www.anthemfacts.com/

Authentify. (n.d.) *Voice Biometric Authentication*. Retrieved from http://authentify.com/solutions/authentication-concepts/voice-biometric-authentication/

Baloul, M., Cherrier, E., & Rosenberger, C. (2012). Challenge-based speaker recognition for mobile authentication. In *Proceedings of the International Conference of the Biometrics Special Interest Group (BIOSIG)* (pp. 1-7). Darmstadt: IEEE.

Berhe, S., Demurjian, S., Saripalle, R., Agresta, T., Liu, J., Cusano, A., . . . Gedarovich, J. (2010). Secure, Obligated and Coordinated Collaboration in Health Care for the Patient-Centered Medical Home. *AMIA Annual Symposium Proceedings*, (pp. 36-40). Academic Press.

Bertino, E., Catania, B., & Damiani, M. (2005). GEO-RBAC: a spatially aware RBAC. In *Proceedings of the tenth ACM symposium on Access control models and technologies* (pp. 29-37). Stockholm, Sweden: ACM.

biometrics. (n.d.). *Dictionary.com Unabridged*. Retrieved from http://dictionary.reference.com/browse/biometrics

Boonyarattaphan, A., Bai, Y., & Chung, S. (2009). A Security Framework for e-Health Service Authentication and e-Health Data Transmission. In *Proc. of 9th Intl. IEEE Communications and Information Technology* (pp. 1213-1218). Icheon: ISCIT.

Chang, T., Tsai, C., & Lin, J. (2012, May). A graphical-based password keystroke dynamic authentication system for touch screen handheld mobile devices. *Journal of Systems and Software*, *85*(5), 1157–1165. doi:10.1016/j.jss.2011.12.044

CTDPH, Connecticut Department of Public Health. (2013a). *Health Information Technology and Exchange*. Retrieved from http://www.ct.gov/dph/cwp/view.asp?a=3936&q=463000

CTDPH, Connecticut Department of Public Health. (2013b). *Practicioner Licensing and Investigtations Statistics.* Connecticut Department of Public Health. Retrieved from http://www.ct.gov/dph/lib/dph/practitioner_licensing_and_investigations/plis/statistics/2013.pdf

Dare, T. S., Ek, E., & Luckenbaugh, G. (1996). *United States of America Patent No. 5,684,950.* Retrieved from http://www.google.com/patents/US5684950

De La Rosa Algarin, A., Demurjian, S., Ziminski, T., Rivera Sánchez, Y., & Kuykendall, R. (2013). Securing XML with Role-Based Access Control: Case Study in Health Care. In A. Ruiz-Martínez, F. Pereñíguez-García, & R. Marín-López (Eds.), *Architectures and Protocols for Secure Information Technology* (pp. 334–365). IGI Global.

Eichelberg, M., Aden, T., Riesmeier, J., Dogac, A., & Laleci, G. (2005, December). A survey and analysis of Electronic Healthcare Record standards. *ACM Computing Surveys*, *37*(4), 277–315.

Deloitte. (n.d.). *ConvergeHEALTH*. Retrieved from www.converge-health.com

DIRECT. (n.d.a). *Best Practicies for HISPs*. Retrieved from http://wiki.directproject.org/Best+Practices+for+HISPs

DIRECT. (n.d.b). *Best Practicies for HISP-HISP Agreements*. Retrieved from http://wiki.directproject.org/Best+Practices+for+HISP-HISP+Agreements

DIRECT. (n.d.c). *The Direct Project Home*. Retrieved from http://wiki.directproject.org/home

DIRECT. (n.d.d). *The Direct Project Overview*. Retrieved from http://directproject.org/content.php?key=overview

Epic. (n.d.). *Epic*. Retrieved from http://www.epic.com/

Farrell, S., & Housley, R. (2002, April). *An Internet Attribute Certificate Profile for Authorization*. Retrieved from The Internet Engineering Task Force (IETF®): https://www.ietf.org/rfc/rfc3281.txt

Ferraiolo, D. F., Sandhu, R., Gavrila, S., Kuhn, D. R., & Chandramou, R. (2001). Proposed NIST standard for role-based access control. *ACM Transactions on Information and System Security*, *4*(3), 224–274. doi:10.1145/501978.501980

GAA-API. (2005). *Generic Authorization and Access-control API (GAA-API)*. Retrieved from http://gost.isi.edu/info/gaaapi/

HealthIT.gov. (2014, May 12). *What is HIE?* Retrieved from HealthIT: http://www.healthit.gov/providers-professionals/health-information-exchange/what-hie

HealthIT.gov. (n.d.). *Privacy & Security Tiger Team*. Retrieved from https://www.healthit.gov/FACAS/health-it-policy-committee/hitpc-workgroups/privacy-security-tiger-team

HIPAA. (1996). *Health Insurance Portability and Accountability Act (HIPAA)*. Retrieved from http:// www.hhs.gov/ocr/privacy/

Housley, R., Polk, W., Ford, W., & Solo, D. (2002, April). *Internet X.509 Public Key Infrastructure Certificate and CRL Profile*. Retrieved from The Internet Engineering Task Force: http://www.ietf.org/rfc/rfc3280.txt

Hu, V. C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., & Scarfone, K. (2014). *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*. NIST Special Publication. doi:10.6028/NIST.SP.800-162

Hwang, S., Cho, S., & Park, S. (2009). Keystroke dynamics-based authentication for mobile devices. *Computers & Security*, *28*(1-2), 85–93. doi:10.1016/j.cose.2008.10.002

JASON. (2014). *A Robust Health Data Infrastructure*. Agency for Healthcare Research Quality. Retrieved from http://healthit.gov/sites/default/files/ptp13-700hhs_white.pdf

Kelly, V. (2013, Dec.). *Global Healthcare Stakeholders Want Standards-Based Interoperability and Communications, According to IEEE*. Retrieved from IEEE Standards Association: https://standards.ieee.org/news/2013/ieeesa_mhealth-summit.html

Li, J., & Li, N. (2006, October-December). OACerts: Oblivious Attribute Certificates. *IEEE Transactions on Dependable and Secure Computing*, *3*(4), 340–352. doi:10.1109/TDSC.2006.54

Mavridis, I., Georgiadis, C., Pangalos, G., & Khair, M. (2001, January-March). Access Control based on Attribute Certificates for Medical Intranet Applications. *Journal of Medical Internet Research*, *3*(1), e9. doi:10.2196/jmir.3.1.e9

McCarty, C., & Wilke, R. (2010). Biobanking and pharmacogenomics. *Pharmacogenomics*, *11*(5), 637–641. doi:10.2217/pgs.10.13

Mettler, T., & Rohner, P. (2009). Increasing the Networkability of Health Service Providers: The Case of Switzerland. *Sprouts: Working Papers on Information Systems, 9*(1).

Meyers, J. (2014, September 6). *Hackers threaten health care industry's patient records*. Retrieved from Boston Globe: http://www.bostonglobe.com/news/nation/2014/09/05/health-care-industry-ill-prepared-for-vicious-cyberthreats/ZdvDGaipJi7VSN0TogezkL/story.html

Neuman, B., & Ts'o, T. (1994, September). Kerberos: An Authentication. *IEEE Communications Magazine*, *32*(9), 33–38. doi:10.1109/35.312841

NIST, National Institute of Standards and Technology. (2014). *Nationwide Health Information Network*. Retrieved from http://www.nist.gov/healthcare/testing/nhin.cfm

Pashalidis, A., & Mitchell, C. (2003). A Taxonomy of Single Sign-On Systems. *8th Australasian Conference, ACISP. 2727*, (pp. 249-264). Wollongong, Australia: Springer-Verlag Berlin Heidelberg.

Ray, I., & Toahchoodee, M. (2007). A spatio-temporal role-based access control model. In *Proceedings of the 21st annual IFIP WG 11.3 working conference on Data and applications security* (pp. 211-226). Redondo Beach, CA: Springer-Verlag.

Rivest, R. L., Shamir, A., & Adleman, L. (1978). A Method for Obtaining Digital Signatures and Public-key Cryptosystems. *Communications of the ACM*, *21*(2), 120–126. doi:10.1145/359340.359342

Rouse, M. (2013). *What is identity management (ID management)?* Retrieved from http://searchsecurity.techtarget.com/definition/identity-management-ID-management

Rouse, M. (2014). *What is PKI (public key infrastructure)?* Retrieved from http://searchsecurity.techtarget.com/definition/PKI

Ruggiero, P., & Foote, J. (2011). *Cyber Threats to Mobile Phones.* US-CERT, United States Computer Emergency Readiness Team. Retrieved from https://www.us-cert.gov/sites/default/files/publications/cyber_threats-to_mobile_phones.pdf

Ryutov, T., Zhou, L., Neuman, C., Leithead, T., & Seamons, K. (2005). Adaptive Trust Negotiation and Access Control. In *Proceedings of the tenth ACM symposium on Access control models and technologies* (pp. 139-146). New York: ACM.

Shibboleth Consortium. (2016). *Shibboleth.* Retrieved from https://shibboleth.net/

State of Connecticut. (2013). *Connecticut Healthcare Innovation Plan.* State of Connecticut. Retrieved from http://www.healthreform.ct.gov/ohri/lib/ohri/sim/plan_documents/ct_ship_2013_12262013_v82.pdf

The President's Council of Advisors on Science and Technology. (2010). *Report to the President Realizing the Full Potential of Health Information Technology to Improve Healthcare for Americans: The Path Forward.* Washington, DC: Executive Office of the President.

Ventuneac, M., Coffey, T., & Salomie, I. (2003). A Policy-Based Security Framework For Web-Enabled Applications. In *Proceedings of the 1st international symposium on Information and communication technologies* (pp. 487-492). Dublin, Ireland: Trinity College Dublin.

W3C. (2015). *Extensible Markup Language (XML).* Retrieved from https://www.w3.org/XML/

Wiedemann, L. (2010). Deleting Errors in the EHR. *Journal of American Health Information Management Association*, 53–54.

Yu, J., Wang, G., & Mu, Y. (2012). Provably Secure Single Sign-on Scheme in Distributed Systems and Networks. In *Proceedings of the 2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications* (pp. 271-278). Washington, DC: IEEE.

KEY TERMS AND DEFINITIONS

Adaptive Trust Negotiation: Trust negotiation in which the request receiver adjusts its security policies based on the user's credentials.

Attribute Certificate: A structured tamper-resistant file that is associated with an identity through an identity certificate and is used to list data in a key-value pair format.

Certificate Authority (CA): An entity endorsed by another authority that vets user identities and signs identity certificates.

Digital Wallet: A collection of credentials a user earns through being granted access to secure systems.

Electronic Medical Record (EMR): A data representation of a collection of patient medical histories in digital format.

Health Information Exchange (HIE): The sharing of health data between stakeholders over a secure medical network, or the computer system that facilitates data sharing.

Identity Certificate: A structured tamper-resistant file that is used to identify an individual and provide assurance for secure connections.

Root Authority: An entity that signs user certificates with a self-signed certificate, users must add the certificate to their certificate store to establish trust.

Trust Negotiation: The process two entities without prior contact undertake to establish trust based on credentials other than identity.