Security Models and Requirements for Healthcare Application Clouds

Rui Zhang [1,2] and *Ling Liu* [1]

[1.] College of Computing, Georgia Institute of Technology, Atlanta, GA, USA
[2.] School of Computer and Information Technology, Beijing Jiaotong University, Beijing, China
rui.zhang21@gmail.com, lingliu@cc.gatech.edu

Abstract: With the widespread use of electronic health record (EHR), building a secure EHR sharing environment has attracted a lot of attention in both healthcare industry and academic community. Cloud computing paradigm is one of the popular healthIT infrastructure for facilitating EHR sharing and EHR integration. In this paper we discuss important concepts related to EHR sharing and integration in healthcare clouds and analyze the arising security and privacy issues in access and management of EHRs. We describe an EHR security reference model for managing security issues in healthcare clouds, which highlights three important core components in securing an EHR cloud. We illustrate the development of the EHR security reference model through a use-case scenario and describe the corresponding security countermeasures and state of art security techniques that can be applied as basic security guards.

Keywords: *EHR; cloud computing; healthcare; security; privacy*

I. Introduction

Cloud computing represents a new way, in some cases a more cost effective way, of delivering enterprise IT. As with all major disruptive changes in technology and Internet revolution, cloud computing represents a true democratization of Web computing, and it is not only changing the business models and the way IT infrastructure is being delivered and consumed, but also the underlying architecture of how we develop, deploy, run and deliver applications.

A. Why is cloud computing attractive to healthcare IT?

Many healthcare providers and insurance companies today have adopted some form of electronic medical record systems, though most of them store medical records in centralized databases in the form of electronic records. Typically, a patient may have many healthcare providers, including primary care physicians, specialists, therapists, and other medical practitioners. In addition, a patient may use multiple healthcare insurance companies for different types of insurances, such as medical, dental, vision, and so forth.

Currently, each provider typically has its own database for electronic medical records (EMRs). Sharing information between healthcare practitioners across administrative boundaries is translated to sharing information between EMR systems. The electronic records sharing between different EMR systems are called electronic health records (EHRs). The interoperation and sharing among different EMRs has been extremely slow. Cost and poor usability have been cited as the biggest obstacles to adoption of health IT, especially Electronic Health Records (EHR) systems. Cloud computing provides an attractive IT platform to cut down the cost of EHR systems in terms of both ownership and IT maintenance burdens for many medical practices.

It is widely recognized that cloud computing and open standards are important cornerstones to streamline healthcare whether it is for maintaining health records, monitoring of patients, managing diseases and cares more efficiently and effectively, or collaboration with peers and analysis of data. Many predict that managing healthcare applications with clouds will make revolutionary change in the way we do healthcare today. Enabling the access to healthcare ubiquitous not only will help us improve healthcare as our data will always be accessible from anywhere at any time, but also it helps cutting down the costs drastically. A fundamental step for the success of tapping healthcare into the cloud is the in-depth understanding and the effective enforcement of security and privacy in cloud computing.

B. Security and Privacy Issues in Healthcare

Research on the various security issues surrounding healthcare information systems has been heated over the last few years. ISO/TS 18308 standard gives the definitions of security and privacy issue for EHR [1]. The Working Group 4 of International Medical Informatics Association (IMIA) was set up to investigate the issues of data protection and security within the healthcare environment. Its work to date has mainly concentrated on security in EHR networked systems and common security solutions for communicating patient data [2]. The European AIM/SEISMED (Advanced Informatics in Medicine/Secure Environment for Information Systems in MEDicine) project is initiated to address a wide spectrum of security issues within healthcare and provides practical guidelines for secure healthcare establishment [3,4,5]. US HHS (Health and Human Services) recently published a report about personal health records (PHRs), aiming at developing PHRs and PHR systems to put forward a vision that "would create a personal health record that patients, doctors and other health care providers could securely access through the Internet no matter where a patient is seeking medical care."

In this paper we present an overview of the security and privacy issues in the EHR cloud, including the models and requirements for secure access of EHR data in clouds. We argue that security and privacy protection of cross-institutional electronic patient records is of paramount importance. The following three principles are critical for ensuring privacy of patients and the content authenticity and source verifiability of electronic medical records. *First,* all electronic medical records, be it PHR or EHR or EMR, should be guarded through ownership controlled encryption, enabling secure storage, transmission, and access. *Second,* the creation and maintenance of EHRs should preserve not only content authenticity but also data integrity and customizable patient privacy throughout the EHR integration process. *Third but not the least,* the access and sharing of EHRs should provide end-to-end source verification through signatures and certification process against blind subpoena and unauthorized change in healthcare critical data content and user agreements. We illustrate these security principles through a healthcare scenario that is frequently encountered in many healthcare practices today.

II. Healthcare Cloud: Overview

In this section we first define the concept of Personal Health Record (PHR), Electronic Health Record (EHR), and Electronic Medical Record (EMR). Then we briefly discuss the EMR and EHR systems as well as the security and privacy issues in accessing EMR/EHR systems. We also describe the security issues in different types of cloud service models and cloud deployment models for healthcare applications.

A. *The Definitions of PHR, EMR and EHR*

The terms of EHRs and EMRs are used interchangeably by many in both healthcare industry and the press or health science literature. Strictly speaking, these two terms describe completely different concepts according to HIMSS (Health Information and Management System Society) Analytics [6]. Both EMRs and EHRs are critical to the grand vision of healthcare digitization for improving safety, quality and efficiency of patient care and reducing healthcare delivery costs. EMRs are owned by individual healthcare providers, whereas EHRs are typically composed of some subsets of EMRs. The interoperability of EHRs is a fundamental enabling technology for EMRs to reach its full potential in revolutionizing the healthcare delivery with high quality and affordable cost.

Personal Health Record is typically a health record that is initiated and maintained by an individual. An ideal PHR would provide a complete and accurate summary of the health and medical history of an individual by gathering data from many sources, including EMRs and EHRs, and making this information accessible to those who have the necessary electronic credentials to view the information.

Electronic Medical Record is the legal record of what happened to the patient during their encounter at a Care Delivery Organization (CDO) across inpatient and outpatient environments and is owned by the CDO. EMR is created, used and maintained by healthcare practitioners to document, monitor, and manage health care delivery within a CDO. A core EMR system is composed of the clinical data repository (CDR), clinical decision support system (CDSS), controlled medical vocabulary (CMV), computerized provider order entry (CPOE), pharmacy management system, and the electronic medication administration record (eMAR), a functionality in the electronic clinical documentation systems of most vendors.

Electronic Health Record is a subset of EMR record maintained by each CDO and is created and owned by the patient. An EHR typically has patient input and access that spans episodes of care across multiple CDOs within a community, region, or state [6]. Based on ISO/TS 18308 [1] standard, the primary purpose of the EHR is to provide a documented record of care which supports both present and future care received by the patient from the same or other clinicians or care providers. This documentation provides a mean of communication among clinicians contributing to the patient's care.

We can only establish an effective EHR system if the EMRs of various CDOs have evolved to a level that can create and support a robust exchange of information between stakeholders within a community or region. Stakeholders here include patients/consumers, healthcare providers, employers, and/or payers/insurers. Further differentiation between the EMR and EHR is shown in Table 1.

TABLE I. DIFFERENTIATION BETWEEN EMR AND EHR

	EMR	EHR
Definition	The legal record of clinical services for a patient within a CDO.	A subset of EMR from one or more CDOs where the patient received clinical services.
Owner	Owned by the CDO	Owned by patient or stakeholder
Consumer & Usage Scope	EMR systems are supplied by enterprise vendors and installed by hospitals, health systems, clinics, etc.	EHR systems are run by community, state, or regional emergence, or national wide emergence organizations.
Right of patient	Patients can gain access to some EMR information once authorized by the EMR owner.	Patients are provided with interactive access as well as the ability to append information.
Interoperability with other CDOs	Each EMR contains the patient's encounter in a single CDO. It does not contain other CDO encounter data.	Sharing information among multiple CDOs, connected by National Health Information Network (NHIN).

B. *The Relationship of PHR, EMR and EHR*

According to the definition given by HIMSS Analytics [6] and ISO/TS 18308 standards [1], we can easily conclude that medical records of a patient may refer to PHR, EMR and EHR. A part of PHRs can be obtained from the EMR systems of different CDOs and once these EMR data are shared with other CDOs, they become EHR. Due to privacy reason, many patients do not want to place their entire PHRs in EMR/EHR systems. Figure 1 presents the intrinsic relationship of PHR, EMR, and EHR from a patient's point of view. PHR and EMR (or EHR) are partially overlapped. Similarly, EMR and EHR are partially overlapped. The degree of overlap differs from patient to patient due to personalized privacy requirements.

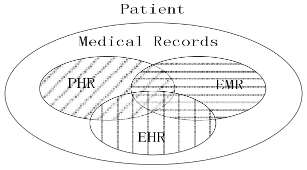

Figure 1. Relationship of PHR, EMR and EHR

C. *Taxonomy of Healthcare Clouds*

We present the taxonomy of healthcare clouds based on the cloud service models and the cloud deployment models.

Based on cloud service models, we can divide healthcare cloud product offerings into three layers:

Applications in the cloud (Software as a Service – SaaS). This layer provides capability for consumers to use the provider's applications running on a cloud infrastructure. For instance, the applications are accessible from various client devices through a thin client interface such as Web browser. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities. In this type of cloud service model, the security and privacy protection is provided as an integral part of the SaaS to the healthcare consumers.

Platforms in the cloud (Platform as a service – PaaS). This layer offers capability for consumers to deploy consumer-created or acquired applications written using programming languages and tools supported by the cloud provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations. In this type of cloud service model, two levels of protection for security and privacy are required. At the lower system level, the cloud provider may provide basic security mechanisms such as end-to-end encryption, authentication, and authorization. At the higher application level, the consumers need to define application dependent access control policies, authenticity requirements, and so forth.

Infrastructure in the cloud (Infrastructure as a Service – IaaS). This type of cloud service model provides the capability for consumers to provision processing, storage, networks, and other fundamental computing resources, in which consumer is able to deploy and run arbitrary software, including operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls). In the Infrastructure cloud model, the healthcare application developers hold full responsibility for protecting patients' security and privacy.

We can also use the cloud deployment models below to give the taxonomy of healthcare clouds.

Private cloud. The cloud infrastructure is operated solely for a healthcare delivery organization (CDO). It may be managed by the organization or a third party and may exist on or off premise. In this type of cloud deployment model, the cloud provider provides the same capability in terms of security and privacy protection as those in the EMR system running by a CDO.

Community cloud. The cloud infrastructure is shared by several CDOs and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It is most likely managed by the third party or the CDOs and may exist on or off premise.

Public cloud. The cloud infrastructure is made available to the general public or a large industry group and is owned by a cloud service provider. In this deployment model, the healthcare application developers and consumers hold full responsibility for protecting patients' security and privacy.

In summary, security and privacy are more than just user privileges and password enforcement. It is a multidimensional business imperative, especially for platforms that are responsible for customer data. Cloud-computing platforms must have detailed, robust policies and procedures in place to guarantee the highest possible levels of physical security, network security, application security, internal systems security, secure data-backup strategy, secure internal policies and procedures, third-party certification. In healthcare cloud, security should be the top priority from day one. We argue that patients' data is protected with comprehensive physical security, data encryption, user authentication, and application security as well as the latest standard-setting security practices and certifications, and secure point-to-point data replication for data backup.

D. Healthcare Security and Privacy Requirements

In healthcare cloud applications, some of the security and privacy requirements are orthogonal to the concrete cloud service model or cloud deployment model used. In this section, we briefly outline these requirements.

Recall the definition and relationship of PHR, EMR and EHR, a patient may have several EMRs stored in different CDOs, in addition to his PHR and EHRs. From the viewpoint of a patient, there may exist a number of EHRs about the patient. Some of EHRs are obtained from various EMR systems after the patient visits practitioners in hospital or other CDOs, and some are hold by the patient himself or patient's family members, such as historical health information. There are three important security and privacy challenges: First, we need to address the question of how to manage and control the access of the EMR data in the EHR system as accessing EMR data are typically controlled through authorization models. Second, a patient may not want to divulge some of his sensitive health information in his EHRs to some family members or some healthcare providers who will offer healthcare for him due to varying concerns. Thus, we need to address the requirement of privacy preserving access to EHRs. Finally, we also need to address the authenticity of EHR data with respect to both content authentication and source verifiability.

From the viewpoint of clinician or practitioner, she or he may have patients in different EMR systems, depending on the type of healthcare and treatment procedures performed. Therefore, one of the most important functionality for practitioners is to provide secure mechanisms to obtain patients' information from multiple EMR/EHR repositories accurately, securely and fast. The next important challenge is patient-consent enabled access control. When a practitioner wants a patient to provide his historical medical records that are stored in other CDOs, he needs to obtain both the patient's consent and the authorization from the respective CDOs, which involves multi-stakeholders. The third security and privacy challenge is the data integrity guarantee. For example, the amount of data update in EHR after the practitioner has offered healthcare to a patient should be carefully controlled and managed.

III. Healthcare Cloud: Basic Security Concepts

In this section we outline the basic concepts in EHR security and privacy. One unique concept in healthcare clouds is "patient-centric" view, which is a term used mostly in community healthcare systems. Community healthcare system offers an open platform for patient to collect, store, use, and share health information in a controlled manner with ubiquitous accessibility. It also offers secure storage and management of patients' EHRs for multiple applications (e. g. disease treatment, lab research, insurance, and other social-networking applications). Most of the community healthcare cloud service models, such as Microsoft HealthVault and Google Health, adopt a centralized architecture with patient-centric views. By patient-centric, it means that the information stored in the community EHR system is imported by patients and only can be made available to a variety of applications under the control of patients (see Figure 2).

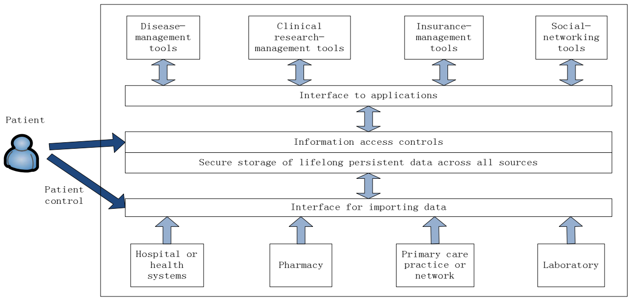
Disease-management tools
Clinical research-management tools
Insurance-management tools
Social-networking tools
Interface to applications
Patient
Information access controls
Secure storage of lifelong persistent data across all sources
Patient control
Interface for importing data
Hospital or health systems
Pharmacy
Primary care practice or network
Laboratory

Figure 2. Patient-Centric and Initial (Centralized) Healthcare Cloud

The common security issues shared by healthcare cloud applications are ownership of information, authenticity, authentication, non-repudiation, patient consent and authorization, integrity and confidentiality of data.

1) Ownership of information: In general, the owner is defined as the creator of the information. Establishing the ownership of the information is necessary for protection against unauthorized access or misuse of patient's medical information. The "owner" can refer to the person responsible for the information or the organization creating and storing the information. The term of "owner" may refer to "creator", "author" and "manager" of the information.

"Creator" indicates the person generating the data. In healthcare system, practitioner or laboratory staff is the creator of medical data about a patient. **"Author"** means the person or entity responsible for the content of the information. In healthcare system, author is the creator of the information, be it the clinician or the CDO which the creator belongs to. **"Manager"** is for the person or entity responsible for management, provision and protection of information. In patient-controlled healthcare system, manager is the patient self. While in decentralized healthcare system, manager may refer to a trusted third party, who is authorized by the patient or healthcare providers.

Ownership of information can be protected through a combination of encryption and watermarking techniques.

2) Authenticity and Authentication: Authenticity in general refers to the truthfulness of origins, attributions, commitments, and intentions. Authentication is the act of establishing or confirming claims made by or about the subject are true and authentic. The authentication of information can pose special problems, especially man-in-the-middle (MITM) attacks, and is often implemented with authenticating identity. Most cryptographic protocols include some form of endpoint authentication specifically to prevent MITM attacks. For instance, Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols that provide security for communications over networks such as the Internet. TLS and SSL encrypt the segments of network connections at the Transport Layer end-to-end. Several versions of the protocols are in widespread use in applications like web browsing, electronic mail, Internet faxing, instant messaging and voice-over-IP (VoIP). One can use SSL or TLS to authenticate the server using a mutually trusted certification authority. In a healthcare system, both healthcare information offered by providers and identities of consumers should be verified at the entry of every access.

3) Non-repudiation: Non-repudiation implies one's intention to fulfill its obligations to a contract. It also implies that one party of a transaction cannot deny having received a transaction nor can the other party deny having sent a transaction. Electronic commerce uses technology such as digital signatures and encryption to establish authenticity and non-repudiation.

4) Patient consent and authorization: Patient can allow or deny sharing their information with other healthcare practitioners or CDOs. To implement patient consent in a healthcare system, patient may grant rights to users on the basis of a role or attributes held by the respective user.

5) Integrity and confidentiality of data: Integrity means preserving the accuracy and consistency of data. In the health care system, it refers to the fact that data has not been tampered by unauthorized use. Confidentiality is defined by the International Organization for Standardization (ISO) in ISO-17799 as "ensuring that information is accessible only to those authorized to have access". Confidentiality is one of the design goals for many crypto systems and made possible in practice by the techniques of modern cryptography. Confidentiality can be achieved by access control and encryption techniques in EHR systems.

6) Availability and utility: For any EHR system to serve its purpose, the information must be available when it is needed. This means that the computing systems used to store and process the EHR data, the security controls used to protect it, and the communication channels used to access it must be functioning correctly. High availability systems aim to remain available at all times, preventing service disruptions due to power outages, hardware failures, and system upgrades. Ensuring availability also involves preventing denial-of-service attacks, and preserving uitility of EHR data. Utility here refers to the ability to preserve the usability of EHR data after excerising and enforcing security and privacy protection and HIPPA compliance.

Audit and *archiving* are two optional security metrics to measure and ensure the safety of a healthcare system. Audit means recording user activities of the healthcare system in chronological order, such as maintaining a log of every access to and modification of data. Auditing capability enables prior states of the information to be faithfully reconstructed. Archiving means moving healthcare information to off-line storage in a way that ensures the possibility of restoring them to on-line storage whenever it is needed without the loss of information [7].

IV. EHR Reference Security Models

In this section we present an EHR security reference model for healthcare clouds and our vision in terms of important security and privacy challenges and countermeasures.

A. EHR Usage Scenario

We use an EHR usage scenario in this subsection to introduce the key components of the healthcare cloud security reference model.

A patient, named Alice, is recently diagnosed a gastric cancer. Surgical removal of the stomach (gastrectomy) is the only curative treatment. For many patients, chemotherapy and radiation therapy given after surgery improve the chance of a cure. Alice entered a cancer-treatment center at her chosen hospital. Alice has a general practitioner whom she regularly visits. Upon entering the hospital, Alice also has an attending doctor from the hospital. Alice's health condition has caused some complications that her attending doctor would like to seek for expert opinion and consultation for Alice's treatment from different CDOs, including Alice's specific general practitioner because he is fully informed about Alice's medical history. Note that the invited practitioners are specialized in different subjects, and some of them are specialists, and others are general practitioners. In such a group consultation, every participant needs to obtain the medical records they request based on the HIPAA minimal disclosure principle. Furthermore, the consultation result, such as the diagnosis and treatment suggestions, should be signed and certified by this group of specialists and practitioners. The medical certificate with their signatures is sent to Alice. If Alice would like to share this medical information with her loved ones and her family physician, she can put the new medical certificate into her PHR database.

In this scenario, a trusted third party is needed to serve as the group manager, who is trusted by all group members, and responsible for choosing the practitioners who may attend the consultation to compose the group, and revoking the group after completion of the diagnosis consultation for the patient. In addition, the providers of EHRs for the patient in this scenario are manifold, given Alice may have other health problems such as diabetic, heart disease, etc. Finally, Alice may have some historical health information in her PHR and some EHRs, to which the group of specialists and practitioners do not have the access.

Now we describe the security and privacy issues involved in this usage scenario from patient's view and practitioner's view respectively. For practitioners, the security and privacy issues can be characterized in two folds: (i) How to securely obtain the EMRs of patient Alice, which is relevant to her gastric cancer treatment, with the compliance of HIPAA minimal disclosure. This concerns the problem of secure EHR collection and integration. (ii) How to certify the authenticity of EHRs obtained from different CDOs or information from Alice's PHR upon authorization by Alice. This relates to the problem of secure storage and management of EHR. Similarly, for patient, Alice needs to be ensured that the diagnosis from the group of doctors can be trusted with a true medical certificate from the group of practitioner. This is the problem of secure EHR usage models. Figure 3 summarizes this example scenario and the related security requirements and techniques. Furthermore, this use-case scenario also involves privacy issues from both practitioner and patient. Alice may prefer to disclose the minimal amount of her sensitive medical information and her family health history. The practitioners in the consultation group wish to inform the specific patient Alice the complication diagnosis and the treatment recommendation as a group based medical decision instead of bounding it to a specific individual practitioner in the group. The patient only needs to verify that the medical certificate is authentic and generated by the specific group of doctors and specialists.

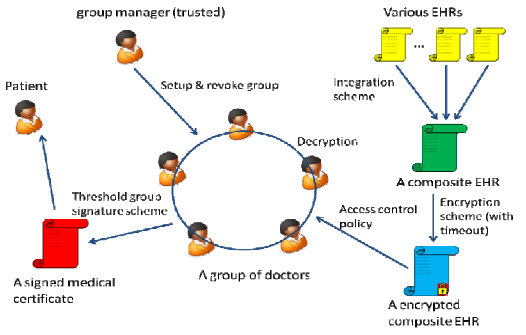

Figure 3. An EHR usage scenario and related security requirements

B. The Proposed EHR Security Reference Model

Based on the use-case scenario shown in Figure 3, we present a basic EHR security reference model for both healthcare professionals and patients in the context of patient care delivery. This security reference model consists of three core components and secure interactions among them to address and implement security and privacy requirements for patient care delivery in an EHR cloud.

(1) EHR secure collection and integration

Based on the definition of EMR, EHR, and PHR given in Section 2, the first core component of the basic security reference model is secure collection and integration of EHR data from multiple EHR repositories, created and maintained independently by care delivery organizations (CDOs). This component is the first mandatory step to take for any CDO to securely share its EHRs with other CDOs (and healthcare providers). The key functional requirement of this component is the EHR integrator. It is responsible for two important tasks: (i) It needs to verify various EHRs provided by different CDOs in terms of authenticity, confidentiality, integrity, and ensuring non-repudiation as well as HIPAA compliance; and (ii) it combines and integrates the successfully verified EHR data into a new composite EHR with a security certificate signed by the integrator. Recall our use-case scenario, Alice's attending doctor wants to organize a group of medical specialists to hold a consultation about Alice's treatment. EHR integration is required to combine Alice's EHRs from different CDOs Alice has visited in the past into one composite EHR before consultation. Semantic interoperability is the key functional issue to be considered. The storage format of EHRs must enable the ability to share data between various EHR systems and easily and efficiently combine EHRs from multiple repositories into a composite EHR. The format of the composite EHR should be conveniently search and access by practitioners and should be verifiable in terms of security, privacy, and HIPPA compliances.

(2) EHR secure storage and access management

The EHR storage and management component is comprised of two main entities: the secure storage server and the access control engine. The former stores the encrypted composite EHR data and allows only authorized access. The access control engine manages a collection of role-based or attribute-based access control policies and HIPPA compliance policies, and enforces the access control policies to prevent the data from unauthorized access. Only authorized practitioners can obtain the access to the authorized portions of the encrypted EHR data through identity and authorization based decryption mechanisms. For example, an authorized EHR requestor of Alice can be granted the permission to access the authorized portion of Alice's composite EHR using her private key.

(3) EHR secure usage model

The secure usage model is the third component of our EHR security reference model, which provides source verifiable content access for consumers of EHR data, including both patients and healthcare practitioners. Thus the two basic functional building blocks in this component are signature and verification. Figure 4 shows a sketch of the system components. We illustrate the key requirements for signature and verification in the EHR secure usage model using the scenario in Figure 3.

Signature. Once the practitioners who participate in the consultation of Alice reach the medical conclusion regarding the next step treatment for Alice, they sign the medical certificate of the corresponding EMR with appropriate signature algorithm. When the portion of the EMR is used to create an EHR record for Alice regarding this consultation, the certificate is sent to Alice alone with the corresponding EHR.

Verification. Consequently, the patient Alice can verify the authenticity of the consultation result made available in the form of EHR to her through the use of this medical certificate and practitioners' digital signature. Note that by respecting for the privacy of the practitioners, the patient, say Alice, needs not to know the group of practitioners who signed the medical certification of the consultation results. In the case of dispute later on, the signature can be "opened" to reveal the identities of the practitioners who sign the consultation result.

In addition to the three core components discussed above, we also need to ensure that the interactions among these components and the communications between any two parties are secure. Concretely, the information transferred between any two parties in one component or two components should be encrypted, as shown in the dashed lines in Figure 4. This is to prevent attackers to sniff and intercept unauthorized information in public networks. Techniques used to protect information transferred in public network are well developed and deployed, such as Secure Socket Layer (SSL), Transport Layer Security (TLS), Internet Protocol Security (IPSec) and so on. The details of these techniques are out of scope in this paper. Note that information is transmitted from storage server to practitioner in the form of non-encrypted transmission as the information has already been encoded by the storage server.

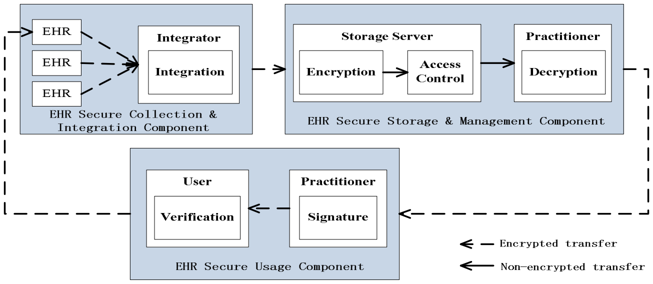

Figure 4. System components

V. EHR SECURE INTEGRATION MODEL

Secure EHR integration is one of the most important building blocks in healthcare cloud development and an essential step towards effective implementation of the secure storage and access management and secure usage component. The following two security elements should be addressed in secure EHR integration: First, we need to verify the authenticity and integrity of each EHR provided by a legitimate CDO through validating the signature signed by the owner of the EHR record. For example, the EHR integrator can verify the authenticity of an EHR with the public key of the owner who signed the record. On the other hand, the EHR integrator should not integrate any EHRs that have no legitimate signatures or the owner of EHRs is untrustworthy, such as EHRs created, or modified, and signed by the patient self. We assume that only the EHRs that have been signed by legitimate healthcare professionals are used and integrated to ensure that all EHRs integrated into a composite EHR are authentic and genuine. Second, we need to define the structure and format of composite EHR such that not only EHRs of different formats from different CDOs can be easily and correctly integrated into a composite EHR but also the data encryption and access control of individual EHRs can be incorporated without any unexpected compromises. For instance, the changes of EHRs due to integration should be minimal. The EHR integration model should support dynamically and expediently adding EHR into the composite EHR. Furthermore, the secure integration model should enable the selective sharing of a composite EHR such that fine-grained authorization can be applied. We below discuss two types of composite EHR model, each is designed in the close synchronization with either attribute-based or role-based authorization and access control hierarchy.

A. *Attribute-based Composite EHR Model*

An attribute-based EHR model adopts a hierarchical structure to define a composite EHR record, aiming at enabling fine-grained selective sharing of a composite EHR. An example attribute-based hierarchical structure for modeling composite EHR was proposed in [8]. The authors propose to logically divide an EHR document into subcomponents such that fine-grained authorization can be applied by labeling each subcomponent with properties of privacy sensitivity, intended purpose, and object type as well as the authorization policies to determine whether a specific sub-object is allowed to be exchanged or not and under what condition and with whom.

The advantage of the attribute based composite EHR is its ability to offer fine-grained authorization and access control. However, such fine-grained authorization and access control also brings a number of disadvantages. First, we need to develop the corresponding enforcement of fine-grained authorization and access control. The complexity of such enforcement can be high when the number of nodes in the attribute-based hierarchy of a composite EHR is large. When the system has large number of users and each user has different access rights, it can be difficult and cumbersome to define authorization rules for each user at very fine granularity, while maintaining the desired access efficiency and availability of the EHR system. For example, when a new practitioner needs to be added into the consultation group of a patient, say Alice, we need to define the authorization on each and every node in the ConsultationNote EHR according to the patient's consent documented, the new practitioner's interest in terms of the set of attributes he wishes to access, and the compliance to the minimal disclosure of HIPPA. Furthermore, the EHR composition process is complex. For each EHR record node to be integrated into the current composite EHR, it needs to traverse each node of the attribute-based EHR hierarchy to find a correct insertion point and check the consistency of the newly inserted node with the existing node at the insertion point in terms of the fine-grained authorization. The time complexity of this integration process grows with the total number of nodes in the attribute-based EHR hierarchy as well as the total number of users for each composite EHR hierarchy.

B. Role-based Composite EHR Model

Instead of attribute-based composite EHR model, in this paper we advocate a patient-centered role-based composite EHR model. Each patient has a patient *PID*, and we use a token, denoted by $Token_i = H(PID_i)$, as the pseudonym of a patient to collect all the EHR documents of this patient, while preserving the privacy of both the patient and the integrator. *H* is a one way hash function, which ensures that it is hard to compute *PID* from token. Each token is mapped to a logic path where the tree structure of token stored.

We construct an EHR tree for each patient using Token$_i$ as the root of EHR tree. For the same patient, say Alice, the same token is assumed. We can set the initial role based hierarchical structure of an EHR in terms of hierarchical template, say five levels of hierarchy, as shown in Figure 6. The root of the tree is at Level 0, Level 1 is the role nodes of practitioners, and their children nodes are labeled with unique identity of practitioners within each corresponding CDO, such as hospital1 and hospital2 in Figure 5 (Level 2). The nodes in Level 3 are medical diagnosis nodes and other correlative inspection nodes. Finally, the leaf nodes of the tree represent EMR records, such as prescriptions and diagnosis and so on. For easy retrieval, we want to sort all child nodes by alphabetical ordering of their tokens or node IDs from left to right except the diagnosis nodes in Level 3, where we place diagnosis nodes as the leftmost child node of their parent nodes, since they are more important than other sibling nodes. Obviously in this structure, all record nodes are nested according to a role node, so that they can be expediently retrieved by different roles of practitioners.

We combine EHRs held by various practitioners into one composite EHR tree. Each branch of the root represents one role of practitioners, and each role node has one or more subtree(s) corresponding to one (or more) practitioner(s). In other words, each practitioner only holds the records under the subtree of his identity. When combining two EHR trees, we cluster the Level 2 nodes that have the same role node together (e.g., Physician, Surgeon, Dentist, etc.) and set them as the child nodes of the specific role node (see Figure 5). Actually the sequence of subtrees is the lexicographic order of sources. Videlicet the sequence of the CDOs has precedence over the sequence of doctors' identity. This merging policy can be easily expanded to integration of multiple EHR trees, just add the new EHR tree into the composite EHR tree one by one.

When a new practitioner is added into the group, the EHR held by the practitioner may need to be merged into the composite EHR tree in the same way as is done by the above merging policy. The time complexity only depends on the number of role nodes. Note that when a practitioner left from the group, the EHR which originally belongs to him should not be deleted from the composite EHR.

Role-based composite EHR model has higher integration efficiency and at the same time it conforms nicely to the role-based authorization by assigning access rights to various practitioners with different roles.

VI. Encryption and Access Management

We stress that all medical information should be stored securely in a private medical record so that patient's information can be tracked from one doctor to another. To ensure secure storage and access management, we need to address the following fundamental security issues.

Selection of encryption scheme. The encryption scheme should be efficient, easy to use by both patients and healthcare professionals, and should be easily extensible to include new EHR records. Furthermore, the number of keys hold by each party should be minimized. Although various encryption algorithms have been developed and deployed relatively well, the proper selection of suitable encryption algorithms to enforce secure storage remains a difficult problem.

Establishment of privacy preserving index. The index of EHR should not leak any sensitive information of a patient. Furthermore, the index should be effective and efficient to speed up the search of various types of EHR records and easy to expand when new records are added.

Access Control. Access control is of particular importance when the database storing the composite EHR is using a database-as-a-service (DAS) paradigm, where an organization's database is stored at an external third-party service provider. The access control policy is typically based on the privilege and right of each practitioner authorized by patient or a trusted third party. We argue that access control policies should be consistent with the structure of the stored EHR record and the usage of the encryption scheme. A number of solutions have been proposed to address the security and access control concerns. Role-Based Access Control (RBAC) [9] and Attribute-Based Access Control (ABAC) [10,11] are the most popular models for EHR. There are pros and cons of RBAC and ABAC when they are used alone in medical system as shown in Table 2. To satisfy requirements of fine-grained access control yet security and privacy preserving, we suggest adopting technologies conjunction of other security techniques and access control method.

Cryptographic Access Control (CAC) [12] is a new distributed access control paradigm and it defines an implicit access control mechanism, which relies on cryptography to provide confidentiality and integrity of data managed by the system. CAC schemes are typically modeled in the form of a partially ordered set of security classes that each represents a group of users requesting access to a portion of the data on the system. Hierarchical cryptographic access control scheme is more general and capable of providing security in different contexts without requiring extensive changes to the fundamental architecture. There is little research to date on access of EHR with CAC method. One possible reason is that it is hard to address the problem of key management and distribution for different user. For instance, how to distribute keys to every healthcare professional so that each of them can decrypt the data authorized by the patient, while cannot decrypt any other data containing sensitive information or no patient's consent is granted.

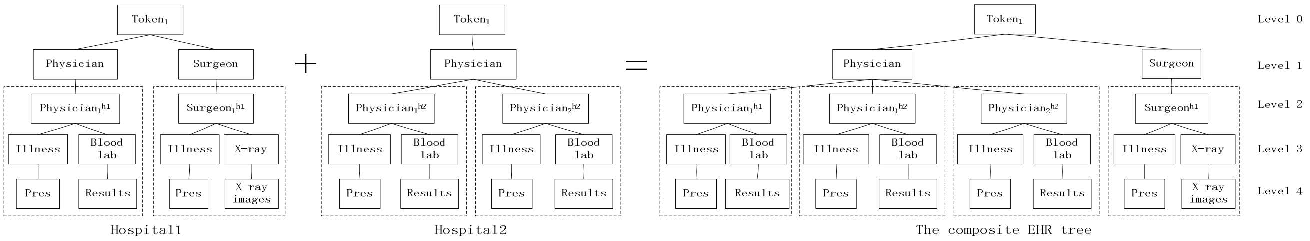

Figure 5. Integration of RBS EHRs

TABLE II. COMPARES RBAC AND ABAC IN COMPOSITE EHR

	RBAC	ABAC
Fine-grained	Poor	Better
Anonymity	Poor	Better
Efficiency of search	Better	Poor
Scalability of policy	Better	poor

Key Management. Key management is another important issue closely related to encryption and access control. Here are a number of desired features for key management. First, the number of keys both held by patient and doctors should not be large. The keys should be easy to store and consume low space complexity. Second, the update of keys should be convenient and efficient in terms of time complexity. Third, all the keys should not contain any private information of any parties. All the keys should be traced and revoked when they are expired or when a user leaves from the group.

For our motivating use-case scenario, we consider two key management schemes: hierarchical key management, which is corresponding to the hierarchical cryptographic access control [13]; and time-bound key management [14], which sets time-based termination condition on healthcare professionals' access of EHR after the healthcare is delivered.

VII. SIGNATURE AND VERIFICATION

Digital signature is a very useful tool for providing authenticity, integrity and non-repudiation while it has seldom been considered to provide user privacy by its own. In our scenario and many other applications such as e-voting, e-auction, we need to protect a signer's identity from being known by eavesdroppers or other parties in a system. In this section we focus our discussion on three important signature techniques for healthcare applications in cloud environments: (i) Anonymous signature, represented by group signature techniques and ring signature techniques; (ii) Threshold signature, and (iii) digital credential.

A. *Anonymous Signature*

The basic idea of anonymous signature is that a signature scheme itself can guarantee the anonymity of the signer. Although there are a good number of anonymous signature schemes, group signature and ring signature are the two most basic and important anonymous signature schemes.

1) Group Signature: A Group signature scheme is a method for allowing a member of a group to anonymously sign a message on behalf of the group. Using a group signature scheme, the members of a group can sign a message with their respective secret keys. The resulting signature can be verified by everyone who knows the common public key, but the signature does not reveal any information about the signer except that she is a member of the group. Essential to a group signature scheme is a group manager, who is in charge of adding group members and has the ability to reveal the original signer in the event of disputes. Comparing our scenario, we also need a person who is responsible creating and revoking the group and dynamically adding and deleting practitioners to make consultation.

The concept of group signature was first introduced in 1991 [15]. The state of the art in group signature research is represented by the ACJT2000 [16], BBS04 [17], BS04 [18]. ACJT2000 introduced a provably secure and efficient group signature, which is coalition-resistant under the strong RSA and the decisional Diffie-Hellman assumption. BBS04 published a novel group signature scheme based on bilinear maps. BS04 is a short group signature scheme that supports Verifier-Local Revocation (VLR). In this model, revocation messages are only sent to signature verifiers as opposed to both signers and verifiers.

2) Ring Signature: In cryptography, a ring signature [19] is another typical anonymous signature algorithm, which is a type of digital signature that can be performed by any member of a group of users that each have keys. The name "ring signature" comes from the ring-like structure of the signature algorithm. A message signed with a ring signature is endorsed by someone in a particular group of people. One of the security properties of a ring signature is that it should be difficult to determine which of the group members' keys was used to produce the signature. Ring signatures are similar to group signatures but differ in two key ways: first, there is no way to revoke the anonymity of an individual signature, and second, any group of users can be used as a group without additional setup.

Suppose that a group of entities each have public/private key pairs, (PK_1, SK_1), (PK_2, SK_2),…, (PK_n, SK_n). Party i can compute a ring signature σ on a message m, on input (m, SK_i, PK_1,…, PK_n). Anyone can check the validity of a ring signature given σ, m, and the public keys involved, PK_1,…, PK_n. If a ring signature is properly computed, it should pass the check. On the other hand, it should be hard for anyone to create a valid ring signature on any message for any group without knowing any of the secret keys for that group.

TABLE III. COMPARISON OF GROUP SIGNATURE AND RING SIGNATURE

	Group Signature	**Ring Signature**
Involvement of trusted third party	Has group manager	No group manager
Organization	Setup by group manager	Self-organization by the signer, no setup, no negotiation with any other members
Adding and deleting members	Dynamically add and delete members by group manager	Statically, members are selected only once by the signer
Revocability	Group revoked by group manager	No revoke
The choice of signature algorithm and key	Decided by group manager	Chosen by the signer
Traceability of signature	Identity of signer can be revealed by group manager	Identity of signer cannot be revealed, but signatures can provide linkability
Multiple signers	Allow multiple signers	Does not allow multiple signers

In summary, group signature and ring signature as the two main types of anonymous signature algorithms have their own advantages. We summarize respective characteristics in Table 3. In our medical consultation example, there may need more than one doctors to sign the diagnosis, in case that it needs at least one (or more) surgeon, one (or more) physician and one anesthetist sign the medical certificate and implement the operation for Alice. Thus, we take threshold group signature into account.

B. *Threshold Signature*

Threshold signature [20] is another signature technique which can be used in our scenario for signing the medical certificate. Suppose the signer is not only one practitioner but a subset of the group of practitioners. Therefore, we need not bother to consider who has the capacity to represent the entire group members to sign the medical certificate. Besides, the medical certificate signed by a number of practitioners is more convincing and reasonable in this scenario.

In threshold cryptography, in order to decrypt an encrypted message a number of parties exceeding a threshold are required to cooperate in the decryption protocol. The message is encrypted using a public key and the corresponding private key is shared among the participating parties. Let n be the number of parties. Such a system is called (t, n)-threshold, if at least t of these parties can efficiently decrypt the ciphertext, while less than t have no useful information. Similarly it is possible to define (t, n)-threshold signature scheme, where at least t parties are required for creating a signature. Threshold versions of encryption schemes can be built for many public encryption schemes. The natural goal of such schemes is to be as secure as the original scheme. Correspondingly, threshold cryptography can be developed to the signature schemes.

C. *Digital Credential*

Medical certificate used our use-case scenario is in the form of digital credential. Digital credentials are the digital equivalent of paper based credentials, such as passport, credit cards, health-insurance cards. Credentials are issued by organizations that ascertain the authenticity of the information and can be provided to verifying entities on demand. A credential is a proof of qualification, competence, or clearance that is attached to a person. Similarly digital credentials prove something about their owner. The digital medical credential in our running scenario is used for several purposes: (1) The foremost one is a proof of medical results containing diagnosis, prescription, and so on for a specific patient from the group of practitioners. (2) When a practitioner logins the EHR system to access the data about a patient, he needs to show his authorization credential issued by the group manager. (3) When a practitioner is added into the consultation group of doctors, he must display his identity credential to the group manager, and (4) when practitioners in the group discuss with each other, they also should exchange the identity credentials to guarantee that their chat is not eavesdropped by others on the public network.

In summary, for the privacy of practitioner, we advocate the use of anonymous digital credentials [21, 22]. The main idea behind anonymous digital credentials is that users are given cryptographic tokens which allow them to prove statements about themselves and their relationships with public and private organizations anonymously. Such credentials, while still making an assertion about some property, status, or right of their owner, do not reveal the owner's identity.

VIII. Discussion and Conclusion

We have taken a methodical approach to investigating security models and security requirements for healthcare application clouds. Meanwhile, we have discussed important concepts related to EHR sharing and integration in healthcare clouds and analyzed the arising security and privacy issues in access and management of EHRs. Then we present an EHR security reference model for managing security issues in healthcare clouds, which highlights three important core components in securing an EHR cloud. Finally, we illustrate the development of the proposed EHR security reference model through a use-case scenario and describe the corresponding security countermeasures and possible security techniques.

Acknowledgement: This work is partially supported by grants from NSF CISE NetSE program, NSF CISE CyberTrust program, and an IBM faculty award, an IBM SUR grant, and a grant from Intel Research Council. The first author performed this research as a visiting PhD student at the Distributed Data Intensive Systems Lab (DiSL) in the School of Computer Science, Georgia Institute of Technology, supported by China Research Council and Beijing Jiaotong University.

References

[1] ANSI, ISO/TS 18308 Health Informatics-Requirements for an Electronic Health Record Architecture, ISO 2003.

[2] R. Bakker, B. Barber, R. Tervo-Pellikka, A.Treacher, (eds.), Communicating Health Information in an Insecure World, in: Proceedings of the Helsinki Working Conference. 43:1, 1995. 2.

[3] B. Barber, D. Garwood, P. Skerman, In: Security in Hospital Information Systems, Security and data protection programme presented at the IMIA WH10 Working conference, Durham. 1994.

[4] S. M. Furnell, P.W. Sanders, Security management in the health-care environment, in: R.A. Greenes, H.E. Peterson, D.J. Protti, (eds.), MEDINFO '95, Proceedings of the eighth World Congress on Medical Informatics. Canada. p. 675–678.

[5] A. Patel, I. Kantzavelou, Implementing network security guidelines in health-care information systems. In: MEDINFO '95. Proceedings of the eighth World Congress on Medical Informatics. Vancouver Trade and Convention Centre, Canada. p. 671–674.

[6] D. Garets and M. Davis, A HIMSS Analytics White Paper. Electronic Medical Records vs. Electronic Health Records: Yes, There Is a Difference. January 26, 2006.
http:// www.himssanalytics.org/docs/wp_emr_ehr.pdf

[7] H. Linden, D. Kalra, A. Hasman, J. Talmon. Inter-organization future proof HER systems-A review of the security and privacy related issues. International Journal of Medical Informatics, 78(2009), 141-160.

[8] J. Jin, G.-J. Ahn, H. Hu, M. J. Covington, and X. Zhang. Patient-centric Authorization Framework for Sharing Electronic Health Records. Symposium on Access Control Models and Technologies, Proceedings of the 14th ACM symposium on Access control models and technologies, 2009, 125-134.

[9] Science Applications International Corporation (SAIC). Role-Based Access Control (RBAC) Role Engineering Process Version 3.0. 11 May 2004.

[10] A. Mohan, D. M. Blough, An Attribute-Based Authorization Policy Framework with Dynamic Conflict Resolution, Proceedings of the 9th Symposium on Identity and Trust on the Internet, 2010.

[11] M. Hagner. Security infrastructure and national patent summary. In Tromso Telemedicine and eHealth Conference, 2007.

[12] A. Harrington, C. Jensen. Cryptographic access control in a distributed file system. Proceedings of the eighth ACM symposium on Access control models and technologies, 2003, 158-165.

[13] S. G. Akl and P. D. Taylor. Cryptographic solution to a problem of access control in a hierarchy. ACM Transactions on Computer Systems, 1(3):239–248, August 1983.

[14] W.-G. Tzeng, A Time-Bound Cryptographic Key Assignment Scheme for Access Control in a Hierarchy, IEEE Trans. On Knowl. and Data Eng., 14(1), 182–188, 2002.

[15] D. Chaum and E. van Heyst. Group signatures. Advances in Cryptology - EUROCRYPT '91, volume 547 of Lecture Notes in Computer Science. 1991, pp. 257–265.

[16] G. Ateniese, J. Camenisch, M. Joye, *et al.* A practical and provably secure coalition-resistant group signature scheme. Advances in Cryptology-CRYPTO, LNCS 1880. Heidelberg: Springer-Verlag, 2000.

[17] D. Boneh , X. Boyen , H. Shacham. Short Group Signatures. In proceedings of CRYPTO '04, LNCS series, 41-55.

[18] D. Boneh and H. Shacham. Group signatures with verifier-local revocation. Conference on Computer and Communications Security, Proceedings of the 11th ACM conference on Computer and communications security, Washington DC, USA, 168 – 177.

[19] R. Rivest, A. Shamir, and Y. Tauman. How to leak a secret, ASIACRYPT 2001. Volume 2248, 552-565.

[20] Y. Desmedt. Threshold cryptography, European Transactions on Telecommunications, 5(4), 1994.

[21] D. Chaum. Security without identification: transaction systems to make big brother obsolete. Communications of the ACM, October 1985, 28 (10): 1030–1044.

[22] J. Camenisch, E. V. Herreweghen. Design and implementation of the idemix anonymous credential system. Proceedings of the 9th ACM conference on Computer and communications security, 2002, 21-30.