# Achieving Electronic Health Record Access from the Cloud

Brian Coats and Subrata Acharya

Department of Computer and Information Sciences, Towson University,
Maryland, United States
bscoats@umaryland.edu, sacharya@towson.edu

**Abstract.** There is an impending requirement for healthcare providers to enable widespread access to their electronic health record systems for the patients they serve. Programs such as the Department of Health and Human Services' Meaningful Use are providing monetary incentives to providers for offering this type of access but affording virtually no guidance as to how it could be accomplished. This research proposes a solution to this challenge by creating a flexible, proven framework that sets the stage for ubiquitous patient access to electronic health records, while preserving security and privacy. Using technologies such as OpenID and federated authentication, this research establishes a standardized approach for healthcare providers to follow to bridge their EHR systems to the Cloud and offer the type of pervasive electronic access the connected world demands.

**Keywords:** Healthcare Information Security, Identity Assurance, OpenID, Portable Identity, Identity Management, Federated Authentication.

## 1 Introduction

Healthcare providers are faced with mounting pressure to provide their patients easy and immediate access to their health information. The federal government has insttuted numerous programs and initiatives that account for much of this pressure. While these programs have mandates to provide access, virtually no stipulations have been given for usability or guidance for how this should be accomplished, merely that access provisions must exist. As such providers are left with the daunting task of making the process of electronic patient access simple and straightforward, while ensuring privacy and security. This research addresses the looming requirement of widespread electronic access to electronic health record (EHR) systems by patients.

The healthcare industry, like most industries, entered the digital age with each provider creating its own silos of data stores and corresponding security frameworks to access that data. The traditional model for authentication for all electronic systems, including EHR systems, is credentials used to validate identity are stored within the application being accessed, as depicted in Figure 1. This model involves users - practitioner and patient alike - being issued a credential such as a username and password, within their particular healthcare provider's EHR system. When the user

M. Kurosu (Ed.): Human-Computer Interaction, Part II, HCII 2013, LNCS 8005, pp. 26–35, 2013.

attempts to access the EHR system, they must enter the credential associated with that system to validate their identity. Therefore, if an individual interacts with multiple healthcare providers, thereby needing access to multiple EHR systems, they are required to have provider-specific credentials for each system. The establishment, issuance, and maintenance of digital identities and corresponding credentials creates a usability barrier for patients and similarly an efficiency barrier for healthcare providers. Now, many healthcare providers are finding themselves poorly positioned to enable the types of distributed access that EHR systems are supposed to facilitate. One of the most visible forces driving electronic patient access is the Department of Health and Human Services' (HHS) Meaningful Use programs. These programs authorize incentive payments to healthcare providers that use EHR technology to accomplish specific objectives in care delivery. Amongst the Meaningful Use objectives are requirements to provide patients timely access to their health records[1]. The recently released Stage 2 objectives, that start in 2014, require hospitals to grant patients access to view, download, and transmit their health information online within 36 hours of discharge; Eligible Professionals (EP) must provide this access within 4 business days[2]. These same regulations and programs that are driving EHR adoption provide almost zero guidance on how to address these enormous usability issues and efficiency challenges.

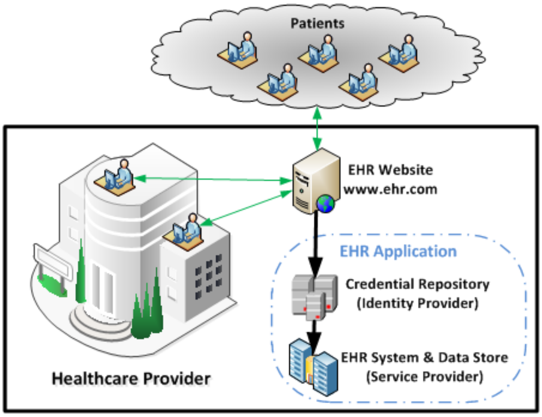

**Fig. 1.** Traditional EHR Access Model

In addition to the regulatory and financial pressures created by HHS, the White House is now creating yet another impetus. In April 2011 the White House released its final draft of the National Strategy for Trusted Identities in Cyberspace (NSTIC). NSTIC is singularly tasked with creating an "Identity Ecosystem" of interoperable technology standards and policies to be used across all sectors to provide increased security and privacy, but most importantly ease of use for individuals[3]. This strategy will force the healthcare industry to structure their identity access approaches to use a distributed model. All federal government agencies, including the HHS, are intimately involved in the development of NSTIC so it is imperative that healthcare providers ensure they are strategically aligned for participation.

This research proposes a solution to this challenge with a framework for healthcare providers to allow access of their EHR systems by their patients from the Cloud. The framework involves creating identity assurance profiles that follow National Institute of Standards and Technology (NIST) e-Authentication specifications. By conforming to the NIST standards, the profiles will introduce varying degrees of trust and assurance for the different identities the providers manage or interact with. Healthcare providers can then establish trust relationships based on these profiles with external authentication systems or Identity Providers. This would enable patient access to their EHR system using the patients' familiar Cloud credentials. These trust arrangements work within the HIPAA compliance guidelines to meet the Meaningful Use objectives while preparing providers to become engaged in cross-industry initiatives such as NSTIC. Specifically, the key contributions of this research to the healthcare information technology industry are:

- A comprehensive framework for healthcare providers to follow to enable external authentication systems to be used for patient access;
- A set of identity assurance profiles for Identity Providers to follow to ensure their practices conform to industry standards and meet HIPAA guidelines;
- Enhanced patient access for a national healthcare provider that assisted in the qualification for Meaningful Use Stage 1.

The remainder of the paper is as follows: Section 2 describes the federated access model and its components; Section 3 details the criteria external Identity Providers must meet in order to participate; Section 4 describes how the Cloud is specifically incorporated as an authentication source for a healthcare provider; Section 5 explains how this research is already being applied and benefiting a national healthcare organization; finally Section 6 summarizes the goals of this research and its importance to the landscape of information security in healthcare.

## 2 Federated Access Model

When it comes to electronic access to applications, there are 3 core questions to be addressed: 1) who does the digital identity belong to, 2) how does the individual prove their identity, and 3) what should the user be allowed to access or carry out in the application? Within the digital identity space, these questions are known as identity management (IdM), authentication, and authorization respectively. The IdM aspect of access consists of the systems that establish and track who an individual is and allows other systems to relate a digital identity to an physical human. The majority of individuals have any number of identifiers that make up their digital identity and it is the IdM system that correlates that information. Authentication and authorization are often confused and mistakenly used interchangeably, but it is important to understand their distinct purposes. The authentication step is how users assert their identity to an application whereas authorization deals with what that user can do within that application such as read, write, or modify data. It is important to understand this differentiation as this research is specifically aimed at the authentication portion of the access equation. The authentication event itself can be broken down into 3 key pieces: the user, known as the Subject, with possession of some set of credentials; an authentication system that can validate credentials, known as the Identity Provider (IdP); and the application to which identity is attempting to gain entry, known as the Service Provider (SP). As Figure 1 shows, traditional systems have the credential repository or IdP built into the application itself. This model creates a dependency that in order to access the application the corresponding, internal credential must be used. A primary objective of this research is to eliminate this dependency. While authorization decisions must inherently be made with the application itself, the authentication decisions can almost always be externalized. The premise of this research contends that all EHR systems should allow the authentication process to be externalized from the rest of the EHR application. Fundamentally, EHR applications need to be able to use other identity stores to validate credentials, beyond those stored in the local EHR database. Luckily, this basic functionality is supported by all the major commercial EHR offerings in some fashion and the real effort lies in getting these EHR systems to play by the same basic rules.

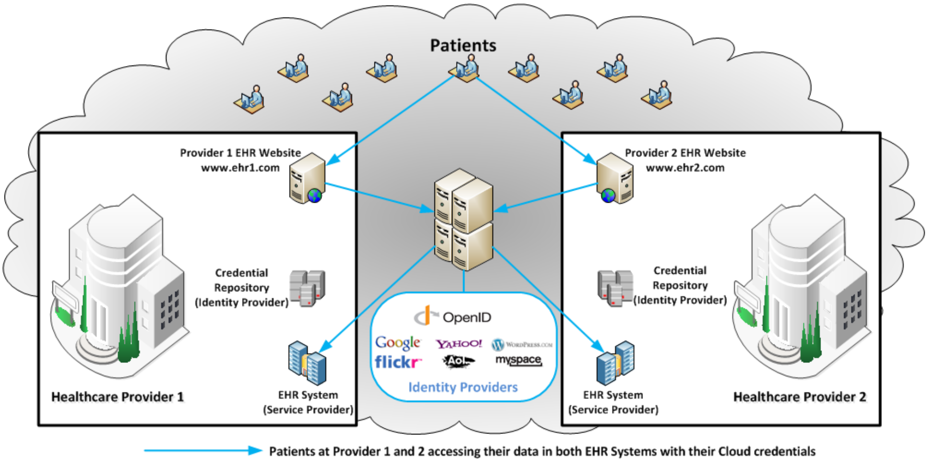

**Fig. 2.** Federated EHR Access Model using the Cloud

By leveraging the ability to separate the authentication process from the EHR application, this research proposes a framework by which multiple authentication systems can be used by a single EHR system, as shown in Figure 2. As long as the authentication is performed by a trusted Identity Provider, the EHR system can be assured the user has adequately verified their digital identity. The true value of this model is realized in healthcare providers effectively outsourcing the entire patient credentialing effort to other entities that have already made significant investments in that arena. The central function of an IdP is to be an authoritative source for establishing and maintaining identities and corresponding credentials. As such, potential IdPs could be commercial vendors like Verizon, Comcast, or AT&T that have existing business relationships with individuals. These companies already have processes in place for validating the identity of their customers and providing them credentials. Likewise, an IdP could be an entity such as Google, Yahoo!, Microsoft, or MySpace, that may have a different type of business relationship but nonetheless tracks relevant identity information and credentials.

It is essential to recognize that while the federated access model moves the authentication of users to an external system, the healthcare provider must have an internal identity management system to map the external IdP's identifiers to EHR users. It is extremely unlikely that the healthcare provider and the external IdPs all have the same key identifier. For many of the free IdPs, this identifier is an email address, whereas most EHR systems will likely use something entirely different such as a Social Security number, Patient code, or some other such identifier. Once these different identification schemes are reconciled, a healthcare provider's IdM system can translate an external identifier into the internal identifier. Beyond the identifiers of a digital identity, it is imperative to analyze the security requirements and practices of each Identity Providers for establishing their identities and credentials. Depending on how these practices are carried out, IdPs should effectively be extended a proportionate amount of trust for their digital identities. It is upon this concept of varying trust or levels of assurance (LOA) that this research builds a foundation for regulating external credentials for EHR access.

When examining trust for identities, there are 2 basic qualities that dictate assurance: 1) the degree of confidence in the vetting process for establishing the identity and matching credential, and 2) the degree of confidence that the user of the credential is the owner of the credential. The more confidence achieved for each of these aspects, the higher the level of assurance connected systems can have in the external credential. Depending on the requirements of the system being accessed, a minimum LOA can be required of the credentials allowed to be used. In order to align the healthcare industry with national standards, this research proposes identity assurance profiles that map directly to the NIST e-Authentication specifications and their corresponding levels of assurance.

## 3 Identity Providers Profiles

In 2003, NIST was mandated by the Office of Management and Budget (OMB) to establish technical standards to support 4 key levels of assurance. As a result, NIST published the Electronic Authentication Guideline[4] which still serves as the regulatory standard for all electronic authentication of federal agencies. The four levels are:

- Level 1: Little or no confidence in the asserted identity's validity.
- Level 2: Some confidence in the asserted identity's validity.
- Level 3: High confidence in the asserted identity's validity.
- Level 4: Very high confidence in the asserted identity's validity.

In response to the NIST standards, the Centers for Medicare & Medicaid Services (CMS) issued detailed requirements for e-authentication and levels of assurance when accessing electronic protected health information (ePHI) covered by the Health

Insurance Portability and Accountability Act (HIPAA)[5]. CMS requires the equivalent of NIST LOA 2 identity assurance for accessing your own health information and LOA 3 for accessing someone else's. Therefore potential IdPs for EHR systems must be able to achieve the appropriate identity assurance equivalence depending on the activity. This research provides detailed identity assurance profiles for IdPs to follow to achieve NIST Levels 1-3, enabling them to reliably guarantee the LOA their digital identities assert. For the purposes of EHR access, the LOA 1 profile is out scope, although it does have a number of other useful applications. While many of the criteria are the same for all the profiles, key difference is the higher the LOA of the identity being asserted, the higher the standard for how the identity was established, how the credentials issued, how the user asserts their identity, and the general integrity of the business practices of the IdP. A summary of the criteria for each profile is provided in Table 1.

**Table 1.** Criteria for Identity Provider LOA Profiles

| Category | Criteria | LOA 1 | LOA 2 | LOA 3 |
|---|---|---|---|---|
| A. Organizational Requirements | 1. Certification | ♦ | ♦ | ♦ |
| | 2. Legal Status | ♦ | ♦ | ♦ |
| | 3. Liability Provisions | ♦ | ♦ | ♦ |
| | 4. Policies and Practices | ♦ | ♦ | ♦ |
| B. Infrastructure Guidelines | 1. Software Security | | ♦ | ♦ |
| | 2. Physical Security | | ♦ | ♦ |
| | 3. Network Security | | ♦ | ♦ |
| C. Identity Creation and Proofing | 1. Identity Establishment | | ♦ | ♦ |
| | 2. Identity Proofing | | ♦ | ♦ |
| | Existing Relationship | | ♦ | ♦ |
| | In-Person Proofing | | ♦ | ♦ |
| | Remote Proofing | | ♦ | ♦ |
| | 3. Record Retention | | ♦ | ♦ |
| D. Identity Management Practices | 1. LOA Classification per Identity | ♦ | ♦ | ♦ |
| | 2. Consistent Data Definitions | ♦ | ♦ | ♦ |
| | 3. Informed Consent | ♦ | ♦ | ♦ |
| E. Credential Management | 1. Subject Interactions | | ♦ | ♦ |
| | 2. Revocation | | ♦ | ♦ |
| | 3. Reissuance | | ♦ | ♦ |
| | 4. Record Retention | | ♦ | ♦ |
| F. Authentication Guidelines | 1. Unique Identifier | ♦ | ♦ | ♦ |
| | 2. Minimum Entropy of Authentication Secret | 14 bits | 20 bits | 64 bits |
| | 3. Protection of Authentication Secrets | ♦ | ♦ | ♦ |
| | 4. Assertion Security | ♦ | ♦ | ♦ |
| | 5. Multi-Factor Authentication | | | ♦ |
| G. Risk Mitigation | 1. Acceptable Use Policies | ♦ | ♦ | ♦ |
| | 2. Business Continuity | ♦ | ♦ | ♦ |
| | 3. Attack Resistant | ♦ | ♦ | ♦ |
| | 4. Single Sign-on (SSO) | ♦ | ♦ | ♦ |
| | 5. Credential Sharing Resistant | ♦ | ♦ | ♦ |

The Organizational Requirements category details the basic guidelines for each IdP to obtain certification for each level of assurance. IdPs must demonstrate they are a legitimate entity and qualify to be recognized as an authoritative source of identity for other organizations. IdPs must also establish they can provide appropriate levels of liability for their actions. Finally, IdPs must guarantee they possess documented policies and procedures and their actual practices are consistent with those documents.

The Infrastructure Guidelines section provides guidelines the Identity Provider's IT environment must follow. All software used for: transactions of identities, credentials, and assertions; the authentication process; credential issuance and maintenance; and identity data storage must be kept up to date and patched to ensure appropriate security. Similarly, IdPs must have adequate physical and network security at the locations where their identity data is stored.

The Identity Creation and Proofing category covers how identities are created, vetted, and proofed. While LOA 1 provides no true confidence that the identity being asserted matches an actual person, LOA 2 and 3 must verify the identity data collected is based on public records or government-issued IDs. Following identity registration, the identity must be proofed to ensure the information collected represents an actual person, that the information can uniquely distinguish a single individual within the IdP's system, and that the person requesting the registration matches the identity being registered. There are 3 basic methods that can be used to perform the identity proofing: the person is already known through an existing relationship; the person can be proofed in-person; or the person can be proofed remotely using additional verification checks against established accounts at financial institutions or utility companies. This category also includes requirements for record retention.

The Identity Management category describes how each Identity Provider defines, asserts, and releases identity information. IdPs must assign their digital identity to a specific LOA and address the possibility of accidental LOA elevation. This section lays out a standard set of data definitions for identity data for all IdPs to utilize to ensure interoperability with EHR systems. Finally, IdPs must incorporate informed consent capabilities into their transactions such that users are presented the specific data being released about them and have the ability to consent or deny its release.

The Credential Management section covers how credentials are to be used in transactions. IdPs must ensure users reassert their identity for each transaction in some reliable fashion. IdPs will also guarantee that credentials will be revoked immediately if they are no longer valid for any reason. Additionally, if credentials are ever reissued, users must provide information from prior transactions like pre-registered questions and responses before the identity is reinstated. Lastly, IdPs are required to maintain a record of all credential management activities including issuance, revocation, expiration, and reissuance for a period of at least 180 days. This amount of documentation is vital for IdPs to satisfactorily establish non-repudiation for the user's transactions.

The Authentication Guidelines category covers the requirements of the authentication process for the different levels of assurance. This includes IdPs ensuring all issued credentials are universally unique to a single individual. The authentication secret portion of the credential - often a password - must meet a minimum entropy or resistance to guessing, depending on the LOA. Entropy can be impacted by a variety of methods such as the length of the password, complexity requirements of characters included in the password, and expiration period and reuse of the password. For LOA 2 and 3, the CMS-approved LOA for ePHI access, a minimum entropy for the authentication secret is 20 bits and 64 bits respectively. This means LOA 2 secrets must have no fewer than a 1 in 1,048,576 or $2^{20}$ chance of being guessed and LOA 3 resistance is $2^{64}$. Additionally, IdPs that assert LOA 3 identities need to employ multi-factor authentication when validating the user. All levels require that IdPs use industry-standard encryption algorithms to provide ample protection to their identity data both while at rest and during all transmissions.

Risk Mitigation is described in the last section of the profile. Every IdP must possess acceptable use policies and record their users' periodic agreement to said policies. IdPs must also make efforts to minimize the chance of system failures to ensure normal business continuity. Further, if a failure does occur, the IdP must make certain the failure wouldn't compromise the security of their system or allow an inaccurate identity assertion to be sent to an EHR system. Additionally, IdPs are required to show their authentication systems are resistant to various attacks including replay and eavesdropping. For IdPs that utilize any type of single sign-on (SSO) technologies, industry-standard techniques and encryption must be employed to guarantee the integrity of the identity assertions. Lastly, each IdP is responsible for enacting safeguards to resist credential sharing, either accidental or intentional.

## 4 Integrating the Cloud

The identity assurance profiles provide all participating entities a known set of technical and functional rules in which to operate. Before an actual implementation can begin, it is vital for the participants to standardize on a specific technology to facilitate the actual sharing and exchanging of identity information. There are actually quite a number of organizations and foundations currently working in the identity space related to portable digital identities and a handful of mature standards have emerged. The prominent standards that have emerged are: 1) Security Assertion Markup Language (SAML), 2) OAuth, 3) WS-Trust, and 4) OpenID. While all these technologies can potentially offer a similar solution, this research proposes that OpenID is the most suitable identity standard currently available. As many organizations decide to adopt one standard or the other, significant work is being done in parallel to erect bridges between the technologies to expand the possibilities of interoperability even farther. Therefore it is arguable that the specific standard decision is as critical as the commitment to adopt a standard and then move quickly and surely to make the necessary organizational and technical choices.

There is a clear advantage to choosing a standard that has wide adoption already as it lowers the barriers for entry. OpenID consists of the most common Identity Providers in the Cloud including Google, Yahoo!, Flickr, MySpace, and AOL. Its corporate members include such companies as Microsoft, PayPal, Symantec, and Verizon, which combine to form an organization with momentous market share in the digital identity space. Over a billion OpenID enabled accounts exist already and are in use by more than 50,000 websites today[6]. The federal government legitimized

OpenID as a key integration technology by certifying an OpenID profile for LOA 1 and expansion to include LOA 2 and LOA 3 is currently underway. This only further signifies the wide adoption of this technology by the public and private sectors. OpenID is a seasoned, established standard and has been incorporated into this proposed framework to be the underpinning for identity and credential creation.

While the OpenID standard accommodates the authentication event, it is critical providers have a mechanism by which a user's OpenID identity can be correlated to the organization's record of that identity. This mapping process can happen any number of ways, either with end-user involvement or not. A simplistic approach offered by this research is a user-driven registration process. This process involves the patient going to a registration site, hosted by the healthcare provider, and entering key pieces of personally identifiable information to establish their identity with the provider. Next, the patient would choose an OpenID IdP and enter the corresponding credentials. After the credentials are validated by the OpenID IdP the correlation is complete and the user then can use their Cloud credentials to access the healthcare provider's systems. This straightforward approach is used extensively within the Cloud today by many merchants and web resources, presenting options such as 'Register with Google'. Healthcare providers can easily emulate this process by letting their patients attach a Cloud credential to their identity in the provider's EHR.

## 5 Pilot Implementations

To demonstrate the feasibility of this research, a partnership was formed with a large national hospital to host a series of pilots and proof-of-concept activities. The partner hospital has over 800 licensed beds and more than 300,000 patient admissions every year. As such, this hospital wanted to leverage industry standards and technologies, without taking on additional overhead, to solve their patient access issues. Using this research's framework, the hospital was able to integrate with Cloud IdPs as well as act as an IdP itself for some federal government resources. By acting as an IdP, the hospital was able to provide their practitioners access to 11 National Institutes of Health (NIH) resources using their hospital credentials, most notably PubMed, the Clinical Translational Sciences Award (CTSA) Management System, the Flow Cytometry Experiment and Reagent Management System (FERMS), and the Address Lookup Tool (ALT) for National Children's Study. The hospital also implemented an OpenID pilot project for a variety of scheduling applications including the radiology and diagnostic testing. These integrations allowed their patients to schedule, modify, and view appointments using their Cloud credentials and then viewing results following their visits. These pilots have noticeably improved the usability of these systems for patients, while reducing the associated support costs of the hospital. Reducing user support overhead is a key benefit of the Cloud access model. With only a very modest time investment, healthcare providers can effectively outsource patient account support to the Cloud and get out of the business of supporting an internal system that issues, maintains, and revokes identities and credentials for all their patients. The projects and pilots at the partner hospital have demonstrated this research offers practical solutions to the dilemma of providing widespread patient access to a healthcare provider's resources. The hospital is evaluating how this model can be extended to other and e-Prescription applications based on the success thus far.

## 6 Conclusion

There is significant work being done in the digital identity space across all industries. It is not surprising all this work is moving in the same relative direction resulting with all industries and technologies converging to form a larger interoperable community. Ubiquitous access is rapidly becoming both a reality and expectation of our connected society. The Meaningful Use programs are just a piece of this larger evolution and force healthcare providers to enable patients greater and easier access to their health information. Concurrently, NSTIC is building a solid foundation for cross-industry collaboration of user access to a multitude of electronic resources within both the private and public sectors. With healthcare providers contemplating when not if, it is imperative that they adopt scalable and interoperable solutions to not only satisfy the immediate needs but be poised for the future. This research combines and builds on many of the lessons learned by other industries to provide a practical solution to a potentially overwhelming issue. Even with the early success being realized by the adoption of this initial research, there is much work left to complete to further broaden its application. The next stage of this research involves examining how the different federating technologies and standards can work together. The next permutation of the proposed access model is to become technology-agnostic to only expand the horizon of possible integrations even further. This framework, with OpenID at the core, will bridge the gap from the healthcare industry to the commercial/social identity space and ensure interoperability with all other industries into the future.

## References

1. United States. Department of HHS. CMS. CMS EHR Meaningful Use Overview, `http://www.webcitation.org/6ElKQGZLj` (last accessed June 2012)
2. United States. Department of HHS. CMS. Stage 2 Overview Tipsheet, `http://www.webcitation.org/6ElLmvSlB` (retrieved December 2012)
3. United States. Department of Commerce. NIST. About NSTIC, `http://www.nist.gov/nstic/about-nstic.html` (last accessed November 2012)
4. United States. Department of Commerce. NIST. Electronic Authentication Guide (rev 1), `http://csrc.nist.gov/publications/nistpubs/800-63-1/SP-800-63-1.pdf` (retrieved December 2011)
5. United States. Department of HHS. CMS. CMS System Security and e-Authentication Assurance Levels by Information Type, `http://www.webcitation.org/6ElMqJfeW` (retrieved November 2012)
6. OpenID Foundation. What is OpenID? `http://openid.net/get-an-openid/what-is-openid/` (retrieved November 2012)