Claims-based Identity Management Approaches May Offer Support for Health Information Exchanges

{This article originally appeared in substantially the same form in the September 2009 Alert *published by the Computer Security Institute.}*

The need for robust, scalable, yet flexible identity management capabilities for health information exchange presents an opportunity to apply claims-based identity management principles.

Current healthcare modernization initiatives put unprecedented emphasis on health information technology and, in particular, achieving widespread use of electronic health records (EHRs) and exchange of health information stored in EHR systems. One of the key obstacles facing broader health IT adoption is the need to establish appropriate security and privacy protections for personally identifiable health information. Managing identity for the purpose of controlling access to this data is an especially complex challenge, given the large and diverse range of potential participants in health information exchange, ranging from large healthcare organizations, to physician offices and other healthcare providers, to individual citizens seeking to maintain control of the collection, disclosure, and use of their personal health information. The need for robust, scalable, yet flexible identity management capabilities for health information exchange presents an opportunity to apply claims-based identity management principles. This article highlights some of the challenges in implementing widespread health information exchange and ways in which claims-based approaches might be used to address them.

While the term "identity" is often used to connote a single logical concept, in practice identity is really a multi-dimensional and potentially variable set of descriptive information that describes a person, organization, or other discrete entity. When identity is thought of in the physical world it is usually considered to be a singular, absolute thing, but in the online world, it is increasingly common to have more than one identity, each used for a different context. In a claims-based identity model, it is possible to maintain a complete set of descriptive attributes, sub-sets of which can be used to make claims representing specific identities. It is also possible to separate attributes among different assertion issuers, so that the appropriate claim set can be generated and validated when needed to assert a particular identity in a given context. While the potential for a single physical entity to have more than one logical identity comes with its own set of security issues, from the perspective of protecting individual privacy (a central goal of health information technology) this capability makes it a simple matter to limit the personal information provided with a given request to the minimum required by the receiver for authentication and authorization, reducing the risk of unauthorized information disclosure. Associating a given set of claims with the identification and authentication requirements of a specific purpose also gives the service provider the flexibility to declare one or more trusted issuers to which service requesters must go to generate the claim set needed to satisfy the provider's requirements.

A single service provider may have different claim requirements for different types of requests, or even different requirements based on factors such as prior history with the requester. The set of claims required to establish a new relationship between requester and receiver is often different than what is required to fulfill a request once a relationship already exists. For example, to create a new account with a service provider an individual might be required to provide a social security number, date of birth, and address information in addition to his or her name, while in subsequent interactions only a username (authenticated by a password) might be required. In cases involving the most sensitive kinds of transactions, the claims required for enrollment

might preclude a successful online transaction altogether. U.S. government agencies providing transactions evaluated at E-Authentication Level 4[1] are required to conduct initial identify verification in person, adding a physical constraint to the more typical considerations of what claims must be satisfied (and what evidence must be presented to validate the claims). Claims-based identity management constructs can address this diversity of claim requirements by declaring requirements and constraints as metadata applied to the claim set.

Expressing identity as a composition of claims gives information or service providers great flexibility to declare exactly the claim requirements they need, and also specify acceptable sources of the claims themselves or how those claims must be validated.

One of the inherent strengths of a claims-based identity model is the separation of attributes into individual claims, and the ability to group those attributes as needed into claim sets. Expressing identity as a composition of claims gives information or service providers great flexibility to declare exactly the claim requirements they need, and also specify acceptable sources of the claims themselves or how those claims must be validated. Support for varying claim requirements (and for differing methods or technical means of generating those claims) is also enormously helpful for information or service requesters whose requests will be submitted to multiple respondents, each of which might maintain its own claim requirements due to legal, policy, or business reasons. Recognizing identity as a set of claims is essential for health information exchange as currently envisioned, because there are no plans to create unique national identifiers for personal health information, and the de facto national identifier (the social security number) is not considered viable due to concerns over medical identity theft and other threats to privacy. Instead, patient identification is typically based on a combination of demographic factors and personal attributes deemed less sensitive than social security numbers. Not all patient identification methods use the same attributes, but each could describe its basis for patient identification as a claim set.

There are a variety of health information exchange (HIE) programs and initiatives underway across the United States, but with the additional emphasis placed on establishing a national infrastructure in the Health Information Technology for Economic and Clinical Health (HITECH) portion of the American Recovery and Reinvestment Act (P.L. 111-5, Title XIII), the spotlight is focused on the Nationwide Health Information Network (NHIN) initiative currently managed by the Office of the National Coordinator for Health IT in the Department of Health and Human Services. After a series of "trial implementations" during 2008, the NHIN has reached a state of limited production and is poised for broader adoption by both federal agencies and private sector entities. Identity management is a key challenge for the NHIN, often described as "a network of networks" connecting a disparate group of entities participating in HIE through the use of common service specifications, data and technical standards, and legal and policy frameworks. An overarching approach to identity management in the NHIN is needed to ensure appropriate access controls and limits on information disclosure are maintained, and to provide compliance with a complex legal and regulatory system that includes information security and privacy laws at federal and state levels that apply in different ways to different types of participating entities.

---

[1] Federal agencies are required to review new and existing electronic transactions to ensure that authentication processes provide the appropriate level of assurance, using a Level 1 to Level 4 categorization more fully described in the National Institute of Standards and Technology's Special Publication 800-63, *Electronic Authentication Guideline*.

If participating entities can document as a claim set the basis by which they grant users access to NHIN-connected systems able to generate requests, other participating entities will have a much easier time assessing the relative rigor of that basis and its impact on the trust relationship among the entities.

User authentication and authorization across the NHIN as currently implemented can best be described as a work in progress. Beginning with the trial implementations, a public key infrastructure was established for the NHIN, providing one technical means for entities receiving information requests to validate those requests before responding. The NHIN Authorization Framework[2] specifies the use of Security Assertion Markup Language (SAML) assertions to bind claims about requesting subject (i.e., user) to the message containing the request. The Authorization Framework specification lists the required attributes (i.e., claims) that must be included as part of the SAML assertions to enable the receiving entity to validate the request and make a determination to response. These include User Name, User Organization, User Role, and Purpose for Use; Subject Authorization details; and Authorization Method, Digital Signature, Key Information, and Issuer (of the digital certificate used to sign the assertions). These attributes collectively can be thought of as a claim set, with the claims validated by the use of a digital certificate, with the evidence supporting the assertion provided by the issuer of the certificate (see Sidebar, "Whom do you trust?"). The attributes in this claim set are all required, even though not all requesters may need all of the claims in order to make their determination of whether to respond to the request. For example, if a hospital with an EHR system receives a request for a patient record from another hospital, the receiving entity may be more concerned with the intended purpose for the requested data (e.g., treatment) and with the role of the subject making the request (e.g., medical doctor) than with the user name of the subject, since the receiver would not be expected to have users internal to other entities represented in their own user directories.

The current state of health information exchange emphasizes one-to-one or multi-party information sharing agreements executed among organizations. In the NHIN, for example, entities interested in participating enroll through a governance process, execute a data use and reciprocal support agreement (the DURSA[3] – a multi-party legal document establishing participant obligations, expectations, and rights), and implement gateway software or other technical means of connecting securely to other participants over the Internet. Once enrolled, entities receive a digital certificate issued by a central certificate authority, and the entity's public and private keys are used to provide authentication, authorization, and non-repudiation for NHIN transactions. Through the DURSA, NHIN participants agree to ensure that all users making requests via the entity's connection to the NHIN are authorized to do so. A responding entity is required to accept as sufficient the user access policies of requesting entities, even if they result in greater levels of access than the responding entity's own access policies would allow. The resulting disparities in individual user authentication and authorization requirements is another area in which claims-based identity management constructs can help. If participating entities can document as a claim set the basis by which they grant users access to NHIN-connected systems able to generate requests, other participating entities will have a much easier time assessing the relative rigor of that basis and its impact on the trust relationship among the entities. (see Sidebar: "User-level authentication and role-based access")

[2] NHIN Cooperative Technical and Security Working Group, "NHIN Trial Implementations Authorization Framework Service Interface Specification v1.9.1," January 2009

[3] NHIN Cooperative DURSA Workgroup, "Data Use and Reciprocal Support Agreement" (DRAFT), January 23, 2009.

Using a claims-based approach cannot resolve the issue that imbalances do in fact exist in identity verification and authentication policies, but it would allow each participant to make a consistent decision whether the level of trust in other entities is sufficient to warrant participation in the NHIN.

With participation in health information exchanges typically managed at the entity – rather than individual user – level, each organization has broad latitude in terms of what internal mechanisms it uses for identity verification, user authentication, and assignment of permissions such as access rights.

With participation in health information exchanges typically managed at the entity – rather than individual user – level, each organization has broad latitude in terms of what internal mechanisms it uses for identity verification, user authentication, and assignment of permissions such as access rights. Similarly, aside from legal constraints imposed by HIPAA and other regulations, entities have broad latitude to determine the basis upon which they agree to respond to requests for information or services from other entities. To the extent that HIE participants require non-repudiation (for legal, regulatory, policy, or other business reasons), the use of digital certificates issued at the entity level is likely insufficient, especially as the universe of potential HIE participants grows to include individuals on their own behalf (as opposed to representing a participating entity). Leaving aside

Whom do you trust?

Health information exchange participants include business entities such as health plans, providers, payers (insurance companies), hospitals, labs, and state, local, and regional health clearinghouses, as well as individuals, either directly or through personal health record systems. This diverse participant environment makes a strong use case for federated identity management, and technical efforts to date tend to rely on Security Assertion Markup Language (SAML) to provide this capability. Claims-based identity principles are easily translated into SAML applications, as SAML is intended to communicate one or more security assertions (claims) bound to specific requests for information or functions provided by a service. Because claims contained in a SAML message may come from different sources, entities receiving requests supported by SAML assertions have to consider the source of the claim (and how much trust they have in the source) before making a decision to grant access or fulfill the request. The SAML specification provides for two different ways (termed "confirmation methods") for requesters to back up their claims: holder-of-key and sender-vouches. With holder-of-key, claims are signed with a private key corresponding to a digital certificate issued by an entity the receiving party trusts. The receiver does not directly trust the requester, but does trust the issuer of the requester's credentials. This model allows receiving entities to require requesters to provide claims issued by a particular (trusted) authority, and also delegates the responsibility of initial identity verification to such a trusted authority. In the sender-vouches model, an attesting entity vouches for the verification of the subject (i.e., the user making the request). The receiver in a sender-vouches exchange also does not trust the requester, but does have an existing trust relationship with the attesting entity. A simple example might be a pre-established business partner agreement between two parties, so that subsequent requests from employees of one of the parties are authorized by the employer. In this case the status of the requester as an employee is one claim that might be required by the receiving entity in order to grant access. A given participant might choose to support both of these methods, but in general, when the necessary trust between information exchange participants is at the organization level, sender-vouchers may be a more appropriate choice. In contrast, where strong individual-level authentication is required, holder-of-key may be more suitable. Also, where a central certificate authority is in used, holder-of-key obviates the need for individual exchange partners to trust each other directly, by instead placing trust with the issuer.

for the moment the major undertaking of rolling out and managing a public key infrastructure for tens or hundreds of millions of users, digital certificates alone provide little to facilitate the submission of different claims for different purposes or in different contexts. Given the diversity of organizations likely to be holding personal health information about a particular individual, claims of affiliation may be equal or greater in importance to conventional identity claims. Depending on the context, an individual might need to provide claims regarding personal identity, employment, membership (in a health plan), eligibility, and medical history, evidence for each of which might need to be provided by a different entity. Regardless of the underlying methods of asserting identity, claims-based identity management provides the structural foundation and flexibility to support an evolving set of authentication and authorization needs.

Regardless of the underlying methods of asserting identity, claims-based identity management provides the structural foundation and flexibility to support an evolving set of authentication and authorization needs.

Stephen Gantz, CISSP-ISSAP, CEH, CGEIT, CIPP/G, is founder and Principal Architect of SecurityArchitecture.com. He can be reached through the "Contact Us" link on the SecurityArchitecture.com website or by email at sgantz@securityarchitecture.com.

User-level authentication and role-based access

One approach to providing identification and authentication information across multiple independent entities is federated identity management, often using SAML as a supporting technical standard. Consistent with the benefits of a claims-based framework, SAML enables multiple different privilege-granting schemes to be used by different entities, yet express the resulting assertions in a common token format that can be evaluated and processed the same way by receivers of requests. What SAML does not provide is semantic normalization across authorization contexts, so additional effort is required with SAML-based federated identity management to accurately associate assertions stemming from different policy or privilege frameworks. An OASIS Technical Committee began work last year on a cross-enterprise security and privacy authorization (XSPA) profile, which among other objectives is intended to support access control standards declared for use in the NHIN, and to allow entities to exchange information about privacy policies, consent directives, and other factors influencing decisions about responding to requests for information. The Technical Committee is also producing health-specific profiles for key standards, including SAML, XACML, and WS-Trust. Within the health care domain, an alternative approach exists using the Cross-Enterprise User Assertion (XUA) profile developed by Integrating the Healthcare Enterprise (IHE), a standards development organization deeply involved in health information interoperability whose technical standards have been incorporated into federal health IT initiatives such as the NHIN. The XUA profile includes the provision of user identity for each transaction, with the intent of allowing (and enforcing) some level of role-based access control for information requested from providers. IHE also developed a standard profile for Enterprise User Authentication (EUA) that health enterprises can use to establish authoritative user names for internal authentication and authorization purposes such as single sign-on. The Cross-Enterprise User Assertion profile depends in part on each participating organization using a consistent internal mechanism for identification and authentication (at the technical level, EUA is an implementation of Kerberos that also uses the HL7 Clinical Context Object Workgroup specification of a user subject). The establishment of a clearly defined context for user authorizations facilitates consistent role-based access decisions, although multiple instances of this model may need to be defined for different healthcare contexts. This approach also standardizes to some extent the set of attributes (or claim set) that will be used to establish an assertion in a specific health information exchange context.

