U.S. Department of Health and Human Services
Office of the National Coordinator for Health Information Technology

Privacy and Security Tiger Team
Health Information Technology Policy Committee

Consumer Choice Technology Hearing
June 29, 2010
Washington, D.C.

HIPAAT
Written Public Testimony
Kelly Callahan

1) Describe how the technology implements the patient's consent and the granular choices given to the patient.

HIPAAT (hip • ăt) provides interoperable, web-based consent management and auditing software to support protected/personal health information (PHI) privacy. Our set of software tools has been designed to meet the consent requirements of any jurisdiction.

The software is informed by industry standards such as Health Level 7 (HL7) Community Based Collaborative Care (CBCC), OASIS eXtensible Access Control Markup Language (XACML) and Cross-Enterprise Security and Privacy Authorization (XSPA) to:

- Allow authorized users at care delivery organizations and HIEs to create granular consent directives on behalf of patients: e.g. "do not disclose my PHI to pharmacists for treatment purposes." This is done via our Service Oriented Architecture-based (SOA-based) consent engine, *Privacy eSuite,* which is accessed through a clinical desktop, provider portal or standard web browser.

- Enable patients to confirm or refuse participation in health information exchange (HIE), and to create, edit, store and withdraw health information privacy policies: e.g. "do not disclose my immunization information" or "do not disclose my PHI for research purposes." This is done using simple Web-based consent forms through our patient consent interface, *myConsentMinder,* which would be accessed through a personal health record or patient portal.

- Enable authorized administrators to create and record organizational privacy policies, e.g. "allow all providers to access all patients' PHI" (an example default policy of an opt-out model) and jurisdictional policies, e.g. "restrict disclosure of mental health records."

- Enforce policies network-wide in real time by allowing or denying access to PHI across disparate applications, through privacy-based access control mechanisms.

How it works

1. Patient consent directives are recorded in the consent engine and access control rules are created.

   HIPAAT's *Consent Management Service (CMS)*, also known as the policy administration point, is *Privacy eSuite's* SOA-based web service that enables patient/consumer, organizational and jurisdictional privacy policies to be created and administered, and converted into access rules. Policies may be created directly in *Privacy eSuite's* graphical user interface (GUI) by healthcare providers/privacy officers, or indirectly through *myConsentMinder* by the patients themselves.

   Policies are expressed as PDF (human-readable) and XACML (machine-readable) documents. They are stored in an IHE XDS.b repository as an HL7 CDA R2 document.

2. A healthcare provider requests access to PHI, and the request is evaluated against the patient's policies.

   HIPAAT's *Consent Validation Service (CVS),* also known as the policy decision point, is *Privacy eSuite's* web service that automatically adjudicates requests to access PHI against PHI access rules, and provides a response of "Allow," "Deny," or "Allow through override." This access control is defined according to OASIS XACML/XSPA.

3. The policy decision point's response is then enforced by the policy enforcement point.

   An application, e.g. a clinical application or HIPAAT's *Privacy Manager*, enforces consumer consent preferences by allowing or denying access to PHI in accordance with the decision received from the CVS.

4. All PHI-related actions are logged by a centralized audit service.

   HIPAAT's *Universal Audit Repository (UAR)* is a centralized, standards-based repository of audit events that logs all access and attempted access to PHI. The UAR follows Integrating the Healthcare Enterprise's (IHE's) Audit Trail and Node Authentication (ATNA) Profile.

Granular choices given to the patient:

Our consent management system includes the tools to allow individuals to create and modify privacy policies to direct who shall have access to their electronic PHI, for what purposes and under what circumstances, within a given organizational and jurisdictional policy framework.

Patients/consumers may create very simple **or** granular consent directives – either directly, or with the assistance of a healthcare provider or privacy officer – based on any or all of the following:

- Consent type – collection, use or disclosure of PHI
- Purpose of use – treatment, payment, healthcare operations, healthcare oversight, etc.
- Who may or may not access the PHI – all healthcare providers, specific individuals (e.g. Dr. Smith), roles (e.g. radiologist) in accordance with ASTM/SNOMED structured roles, department (e.g. ER), facility (e.g. General Hospital), pharmacy, payer, jurisdiction (e.g. state), etc.
- What PHI may be accessed – all PHI, PHI of a specific date or date range, PHI related to a specific medical condition, specific PHI types (e.g. prescription history), HL7 attributes for confidentiality codes, category codes and permissions
- 'Emergency Override or 'break the glass' – the ability for healthcare providers to override PHI restrictions, as permitted by the patient and legislation.

2) How far along is the technology in terms of implementation? What steps or technological advances need to be made in order to implement the system in health information exchange?

Our software technology has undergone significant changes to bring it to where it is today for implementation in health information exchange. Here is a brief history:

In 2007, our Version 1.0 was offered as a bundled consent management and auditing solution. This version underwent successful clinical testing at a major university hospital in Toronto later that year.

In 2008, we unbundled our solution into SOA components for Version 2.0. It was commercialized in February of that year with our consent engine **implemented** as part of the IBM led Nationwide Health Information Network (NHIN) hosted solution for North Carolina Health Information and Communications Alliance (NCHICA). As such**,** it was included in the NHIN II Forum 5 demonstrations in Washington in December 2008.

In July 2009, we introduced our Version 3.0 product suite designed to follow HITSP-recommended interoperability standards. Further advancements to support the "Accounting of Disclosures" requirement are now in the testing phase.

Since our v3.0 release, the technology has been incorporated in the designs for implementation as the centralized Consent Directive Management Service (CDMS) in a multi-state Beacon Initiative, a 24-hospital HIE and a provincial multi-domain initiative.

3) What are the advantages to your approach to obtaining patient consent?

We believe there are several advantages to our approach:

- our set of standards-based software tools has been designed to meet the consent requirements of any jurisdiction and to adapt to additional standards as they are finalized

- our solution is an SOA-based, 3rd party managed service layer that acts as an extension of existing EHR technologies – no 'rip and replace' required

- equally suitable as a centralized service to an HIE, RHIO or state, or to small and medium-sized provider environments

- supports PHI requests from any sized EMR/EHR, on a standards-based request/response basis which does not require a generational change.

In addition, there are two specific advantages to our approach that we'd like to highlight: leveraging SOA and supporting the Accounting of Disclosures requirement.

Underlying Issues:

All data processing nodes in a health information network need to be privacy-aware. However, the application nodes (e.g. clinical workstations) at the edges of the network – which provide PHI to users of various roles – need to enforce access control of the PHI, as this is where all the factors affecting the privacy decisions are known. These factors include **when** the PHI is requested, **what** PHI is being accessed, **who** is requesting the PHI and **why** the PHI is required.

It is possible for an application node at a point of service – or a network application server – to obtain the privacy and consent policies of the consumer in question and then adjudicate whether to allow access to the PHI in question. However, this is not necessarily achievable in the case of every application or clinical device, as they come in a range of capabilities from a diversity of vendors.

Using SOA:

What SOA does is allow these weighty decisions to be offloaded to a specialized 3rd party web service that is optimized for this purpose. Instead of dealing with consent only locally, the following occurs: when a user requests access to PHI, the clinical application sends the known attributes of the PHI, the requester and the intended purpose of use for the PHI to the trusted Consent Validation Service (an SOA-based Web service). It then simply needs to enforce the CVS' answer of: Allow access, Deny access, or Allow through override.

SOA

Our software leverages SOA to manage patient consent. This offers many advantages:

- Reduced costs – using SOA, one CVS system can support a large network of existing clinical applications, systems and technologies
- Proactive health information privacy management, allowing prevention of inappropriate access to PHI – versus the reactive approach of relying solely on audit trails to determine who accessed PHI inappropriately
- minimal overhead and integration: the standards-based network interface, using XACML with HL7 vocabulary, is simple to implement
- provides consistent privacy-based access control capabilities to PHI-related applications in and across organizations, HIEs and jurisdictions
- PHI privacy may be managed both locally and centrally
- Non-disruptive – virtually no impact on workflow
- the 'heavy lifting' of validating patient-centric PHI access permissions is moved away from diverse applications (e.g. clinical applications) to specialized third-party, web-based privacy services
- accommodates granular directives – consumer, organizational and jurisdictional; consumers may restrict access to specific, sensitive portions of their record
- consent preferences and organizational and jurisdictional policies are rigorously managed in one or more central servers, with policy changes available network-wide in real time
- provides real-time auditing: the CVS generates an audit trail for all access and attempted access to PHI.

Support for the Accounting of Disclosures requirement

Another advantage of our approach to obtaining patient consent is our support for the accounting of disclosures requirement.

Cost-effective technology is already in place to support the accounting of disclosures requirement to log disclosures made for treatment, payment and healthcare operations.

IHE's ATNA Profile accommodates all of the related fields currently discussed in the Interim Final Rule. This includes reporting disclosure date, time, user ID and Patient ID with the description of disclosure (reason for disclosure) recorded in an optional field.

A comprehensive SOA-based consent management system, when complemented by a web-based manual disclosure tool (such as one supplied by HIPAAT), can feed the audit repository with information required for a complete accounting of disclosures. For example, with HIPAAT's 3rd party consent management solution, every time a clinician changes from one patient to another, or from one study to another, a validation that this interaction is permitted is performed by the CVS. The CVS then provides the system (e.g. EMR or HIE) with a response as to the appropriateness. This validation check is the ideal filtered trigger point to generate audit events in the ATNA audit repository which can then be used to capture disclosures.

There are various ways to record the details of disclosures in an ATNA-compliant repository. EHRs that are ATNA-compliant may already do this automatically. For those that are not, or to supplement automation, the online manual disclosure tool can be used.

HIPAAT's UAR offers a simple, patient-centric Accounting of Disclosures report which leverages audit logs generated by *Privacy eSuite's* CVS to supply information such as the reason for disclosure, the method of disclosure and whether or not the disclosure was authorized.

4) Is the technology scalable so that small and medium-sized providers could implement it?

Our consent management solution is scalable and standards-based. It was architected based on SOA to act as a 3rd party service to physician practices, clinics, hospitals, electronic medical records/electronic health records and HIEs. As a centralized solution it need not be part of any EHR but available to serve all EMRs/EHRs large and small. (We make interface software available to small office implementations at low cost.) As such, consumer preferences can be implemented and enforced consistently and cost-effectively at all levels of health information exchange.

5) ***Is the consent technology being developed interoperable with other systems? (i.e. can the patient's preferences be passed to other HIEs?)***

Yes, our software is interoperable and vendor-agnostic. It accommodates any consent model that allows consumer choice, from full opt-out, opt out with exceptions, full opt-in, opt in with restrictions, etc. Patient/consumer preferences may be stored locally and/or centrally. Our software facilitates cross-organizational, cross-HIE and cross-jurisdictional health information exchange. Any combination of human-readable and machine-readable patient policy documents, as well as standards-based request/response interoperability, are available to HIEs as required according to jurisdictional policies.

6) ***If the consent is not currently interoperable, what are the barriers that stand in the way of this?***

Not applicable to our solution set.

7) ***What resources are necessary to implement the consent system in its current form? What further resources would be necessary to offer further granular consent choices?***

Our consent management service can be implemented for a similar resource allocation as that of an enterprise-class master patient/provider index (eMPI) with a similar size and scope of service area. However, without knowing the environment, further investigation must be given to determine if any additional integration issues require resolution.

As to granular choice, our consent engine and related tools can provide for as much granularity that is natively available/capable in any HDO, IDN, HIE or statewide exchange.

To confirm, the current version can accommodate access control of detailed granularity limited only by the level of identification of the document or domain under consideration.

8) How many users does the system serve currently, if applicable, and how many will it serve when it is fully operational?

The seventh largest jurisdiction in North America is the Province of Ontario in Canada. This is a jurisdiction where privacy legislation provides for a patient to restrict a clinician from sharing (disclosing) personal/protected health information (PHI) with another clinician for healthcare purposes in a PHI access granular environment. In this case, the recipient clinician is informed that some part of the medical record has been “locked” by the patient.

Our COTS system has been designed, built and tested to scale for the volume of network traffic of such a jurisdiction or health information exchange to accommodate consent validation of each PHI access request serving a population of 12-13 million. This essentially captures 89% of all the jurisdictions in North America

Mr. Kelly Martin Callahan
President &COO
HIPAAT International, Inc.