HEALTHCARE IT STANDARDS AND THE STANDARDS DEVELOPMENT PROCESS: LESSONS LEARNED FROM HEALTH LEVEL 7

Contributors

Charles Jaffe
Health Level 7

W. Edward Hammond
Duke University

John Quinn
Accenture

Robert H. Dolin
Health Level 7

Index Words

Healthcare Information Standards
Health Level 7

"Health Level Seven (HL7) was organized in 1987 to create standards to support the development of best-of-breed hospital information systems."

Abstract

Standardization of information in healthcare is critical to the ability of a diverse community of caregivers to reliably exchange complex data without ambiguity. Perhaps more importantly, every member of the healthcare team must be able to re-use the data within increasingly diverse applications and environments. In order to improve quality and to constrain the escalating costs of delivering care and preventative medicine, the strategies for developing these standards have attempted to keep pace with the developments in information technology, human biology, healthcare policy, social science, and economics. For more than two decades, Health Level 7 has been at the forefront of those processes for an international community. New technologies, and new applications of existing ones, foster standards development against a complex backdrop of social and political demands.

Healthcare IT Standards: An International Perspective

The history of standards has made the creation of standards difficult. The global community of today did not exist when the healthcare IT community recognized the necessity for health data standards to reduce the costs of interfacing systems and to permit the exchange of data. The health-related community began to be interested in standards during the early 1980s, and several organizations were created during the late 1980s and early 1990s. For clinical systems, the needs were primarily within a hospital and included the domains of patient admission, transfer and discharge; lab test ordering and result reporting; prescriptions and dispensing; materials management; images; reimbursement; and other similar domains. New standards-developing organizations were created, mainly focusing on only one component of the many recognized needs.

Standards developing organizations (SDOs) that were creating standards in communications, banking, manufacturing, and other non-health areas were quite mature by this time. In the United States, ASTM American Society for Testing and Materials created E31 for health data standards, focused initially on standards for the reporting of laboratory test data. Health Level Seven (HL7) was organized in 1987 to create standards to support the development of best-of-breed hospital information systems.

In Europe, the Comité Européen de Normalisation (CEN) formed Technical Committee 251 in the early 1990s to create health data standards. There was little competition between CEN and HL7, simply because there was no apparent common market. In fact, the methodology that was used to create later HL7 standards was influenced by CEN. In 1995, Germany became the first international affiliate of HL7, followed shortly thereafter by the Netherlands. From then on, the competition between HL7 and CEN began. The relationship between members of both organizations was cordial, but both HL7 and CEN were vying to have their standards used in Europe.

"The International Standards Organization (ISO) was founded just after the end of World War II in 1947 and is an international, standard-setting body composed of representatives from these various national standards' organizations."

All countries have a national standards body: the American National Standards Institute (ANSI) in the United States; the British Standards Institution (BSI) in the United Kingdom, the German Institute for Standardization (DIN), the French national organization for standardization, known as Association Française de Normalisation (AFNOR), and so on. The International Standards Organization (ISO) was founded just after the end of World War II in 1947 and is an international, standard-setting body composed of representatives from these various national standards' organizations. ISO Technical Committee 215 was formed in 1998 through efforts of the U.S. and the U.K. to create standards in Health Informatics. Since most countries, by law, require the use of an ISO standard if one exists, ISO became a key player in international standards. In 1991, recognizing that there were not sufficient expert resources available for ISO and CEN to conduct their standardization activities independently, ISO and CEN signed the Vienna Agreement to work together to produce standards. The impact of this agreement, including the joint development of standards and the sharing of standards, was felt independently within each organization.

"After over five years of work, a single standard had still not been defined."

In 2002, HL7, following the lead of IEEE, signed an agreement with ISO, through ANSI, that permitted HL7's work and standards to be brought into ISO upon approval of ISO TC 215. Although both of these agreements resulted in a move toward the consolidation of work, standards continued to be duplicated. In the spirit of harmonization, efforts were put into mapping one similar standard to another across organizations where multiple standards exist. An example of one such activity was the effort to create a single standard for defining data types, by combining similar work from ISO, CEN, and HL7. After over five years of work, a single standard had still not been defined. Similarly, approaches to define a single global Reference Information Model (RIM) failed, because of differences between the CEN standard EN 13606 and the HL7 RIM. This latter issue was further compounded when the HL7 RIM was approved as ISO/HL7 21731:2006-Health Informatics-HL7 version 3-Reference Information Model.

"Over the next few years, HL7 submitted several HL7 standards to the ISO to become joint HL7/ISO standards."

Over the next few years, HL7 submitted several HL7 standards to the ISO to become joint HL7/ISO standards. These included HL7 version 2.5 Messaging Standards; HL7 version 3.0 Clinical Document Architecture, R2; Common Terminology Server, R1; HL7 EHR–Functional Model; and Clinical Genomics-Pedigree. In late 2005, HL7 submitted four regulatory standards to ISO: Structured Product Labeling, Release 1; Individual Case Safety Report; Stability Study; and Annotated Electrocardiogram.

None of these standards were accepted by TC 215 (most nations abstained). In addition, the International Committee on Harmonization (ICH) was very concerned, because they themselves had standards in some of these areas. In 2006, at the International Conference on Harmonization of Technical Requirements for Registration of Pharmaceuticals for Human Use ICH meeting in Yokohama, leaders from ISO TC 215, CEN TC 251, and HL7 made presentations to the group about their respective organizations. As a result, ICH became Class D Liaison with ISO and developed relationships with both CEN and HL7.

"There was a real interest among the leaders of ISO, CEN, and HL7 to work together."

There was a real interest among the leaders of ISO, CEN, and HL7 to work together, largely driven by the limited resources available to produce standards, as well as by the confusion in the marketplace because of multiple standards. At the Global Health Information Technology Standards Summit in Geneva in 2006, presentations by Dr. Yun Sik Kwak, Chair, ISO/TC 215; Kees Molenaar, Chair, CEN TC 251; and Dr. Ed Hammond, Chair-elect, HL7 suggested the three SDOs might be able to work jointly to produce a single global standard for a single business purpose. After additional discussions, a charter was written and was subsequently ratified by all three SDOs. The charter establishes a Joint Initiative Council (JIC) that includes the chairs of the three participating organizations plus two additional representatives from each SDO.

Figure 1 shows the global healthcare landscape. The JIC serves as a collaborative forum for the international standards community. Within the US, the Federal Health Architecture coordinates 55 Federal agencies. Vocabulary standards are depicted in the oval boxes.

"The charter establishes a Joint Initiative Council (JIC) that includes the chairs of the three participating organizations plus two additional representatives from each SDO."

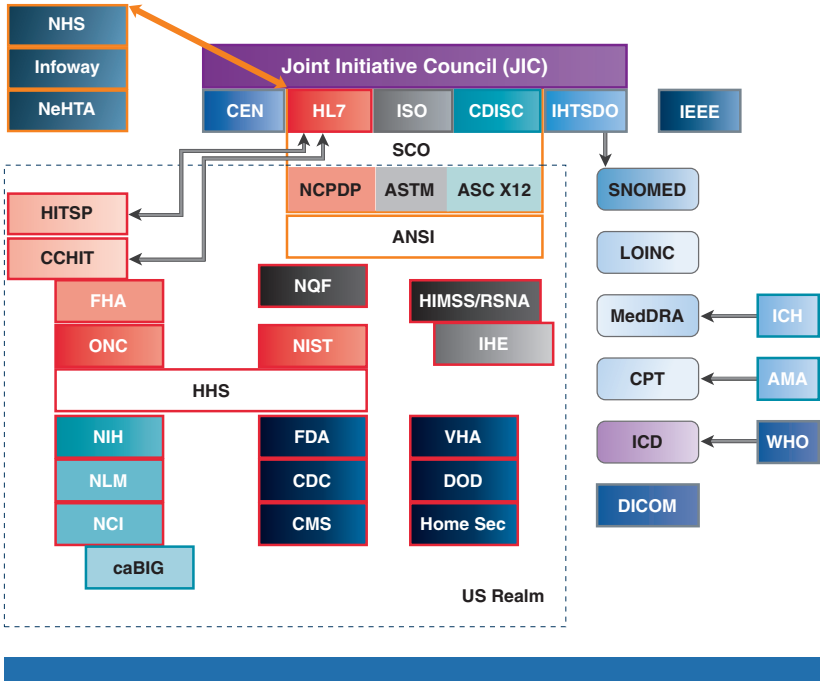

Figure 1: Global Healthcare Standards Landscape
Source: Health Level 7

The word "joint" initiative was chosen specifically to distinguish this activity from harmonization, since harmonization was generally considered to be efforts to map one standard to another competing standard. Any joint project must be agreed to by this group. A Joint Working Group (JWG) was also created, managed by ISO, which includes members from the participating SDOs. The purpose of the JWG is to work through the details of any joint initiative project, identify potential new projects, monitor the work on joint initiative projects, and support, in general, the work of the JIC. Each JIC project would be hosted by one of the participating SDOs, with that SDO providing a project chair or lead for the project. Each of the other SDOs would provide a co-chair to ensure the resulting work met the needs of each participating SDO.

"Harmonization was generally considered to be efforts to map one standard to another competing standard."

The work on any project would be done jointly by volunteer members from each participating SDO, working as a cohesive unit. The three SDOs would be considered equal in the work process. The resulting work product would be balloted simultaneously by each participating SDO, and the comments from each SDO would be aggregated and processed by the joint project team. The resulting standard would be the joint property of each participating SDO and would carry a shared copyright and the logo of each SDO.

"ISO 21090: Healthcare Data Types, was approved in 2009."

The first, now successfully concluded, JIC project was a standard for data types. This standard, ISO 21090: Healthcare Data Types, was approved in 2009, after many years of unsuccessful work to bring disparate groups together. Other projects, listed next, have been accepted as JIC projects:

- Individual Case Safety Report
- Glossary Project
- 13606/HL7 version 3 Implementation Guide
- Identification of Medicinal Products
- Clinical Trial Registration and Result
- BRIDG Data Model

There are a number of other items being proposed as JIC projects. Once the scope and definition of a project is defined, and the participants begin to establish a level of trust in one another and can work together, the projects do move ahead.

"The JIC has developed a policy and procedures for defining projects, for governance, and for bringing other groups into the JIC."

The JIC has developed a policy and procedures for defining projects, for governance, and for bringing other groups into the JIC. Since the creation of the JIC by the three SDOs, two other SDOs have joined the group: the Clinical Data Interchange Standards Consortium (CDISC) became a member in 2008, and the International Health Terminology Standards Development Organization (IHTSDO) became a member in 2009. Other SDOs are in the process of joining the JIC.

There are barriers that still must be overcome for the JIC process to work. The balloting scenario differs among the three SDOs. The length of the balloting period differs as does the publication process and form. Communication among the participating SDOs and their membership is very challenging. Understanding the process has been difficult for all those involved, and managing the process has also been a challenge.

The success of the JIC offers a lot of promise, however. In the United States, an organization similar, at least in purpose to the JIC, has been formed: the SDO Charter Organization (SCO). This organization promises to consolidate efforts within the United States to produce a single standard for a single purpose. Those involved are strongly motivated to work together for a number of reasons. They can share expertise and resources, such as tools and repositories; and working together makes more efficient use of limited funding. Moreover, most SDOs recognize that the global market is expanding, and that there is an urgent need to accelerate the development of standards to avoid both marketplace and regulatory ambiguity. The more confusion, the less likely players will be to use standards in the marketplace.

"Understanding the process has been difficult for all those involved, and managing the process has also been a challenge."

The Role of Architecture in Developing Healthcare Interoperability Standards

In this section, we look at the role of architecture and an architectural framework in developing healthcare interoperability standards, and we present some example environments that use them.

HL7 is a set of standards that supports interoperability of software applications that are designed to support clinical and administrative processes in a healthcare organization. While any healthcare organization can use HL7's interoperability standards, these standards are, for the most part, designed for use by healthcare provider organizations.

"While any healthcare organization can use HL7's interoperability standards, these standards are, for the most part, designed for use by healthcare provider organizations."

Simply put, the architecture of an interoperability standard is a high-level view of the components that make up the standard and the relationships between those components. An architecture framework, as discussed here, is a reference context that we use to view each HL7 standard product and its components. These useful analysis tools are new to HL7.

Over its 22 years, HL7 has both created new and adopted some existing standards. Because of the separate and distinct origins of these standards, a different architecture could have been written for the internal components and their inter-relationships for HL7's major standards products: these products include the HL7 version 2.x; HL7 version 3.0; HL7's EHR standards; CCOW; and Arden Syntax. To begin this remedial effort of creating a common housing to hold the architecture definitions of all HL7's major products, HL7 created an architectural framework called HL7's Services Aware Enterprise Architecture Framework (SAEAF).

"To hold the architecture definitions of all HL7's major products, HL7 created an architectural framework called HL7's Services Aware Enterprise Architecture Framework."

We shall now provide a high-level introduction to HL7's SAEAF. We view SAEAF as it applies to the HL7 version 3.0 standards, including the major components that they share (for example, the HL7 RIM), and the somewhat different purposes that they serve. Our customers, including those working in the UK, Canada, and now the United States, all begin their efforts with provider organizations that use the different architectures of the HL7 version 2.x and the HL7 version 3.x products, that going forward must co-exist and interoperate to achieve country-wide interoperability goals.

HL7 and Architecture: a Brief Introduction

In 1987 the HL7 organization was created with the purpose of creating interoperability standards for the exchange of electronic information between IT systems within and between healthcare providers. Not surprisingly, the technology, workflows, and systems architectures created by HL7 in 1987 were the ones in use by the information technology (IT) industry in 1987. The standard created at that time (HL7 version 2.1) achieved widespread use, starting in 1991 and has evolved through a series of major and minor updates to HL7 2.6 today. No formal architecture was ever planned for any element of the HL7 version 2.

"No formal architecture was ever planned for any element of the HL7 version 2."

About five years later, HL7 volunteers expressed an interest in and began working on a RIM for healthcare. There were no formal expectations at the time for the scope of the effort, the amount of time it would take, or even the expected difficulty that we would encounter. This experience was quite eye-opening, because it forced HL7 developers to document the vastness and complexity of healthcare. Nevertheless, no formal architecture was defined for any HL7 standard that uses the RIM, although the RIM itself did assume a formal internal architecture that addresses its governance, definitions, and the tools related to testing conformance.

"This experience was quite eye-opening, because it forced HL7 developers to document the vastness and complexity of healthcare."

The HL7 RIM was never intended to stand on its own — although as a reference model of healthcare information, it certainly does. Once the RIM was recognized as possible, the organization also turned to the effort of developing a new standard for data interchange that would have these features:

1. Model-driven (i.e., all data items that can be exchanged or used to generate messages, documents, or services have a place and are taken from the RIM where cardinality and relationship(s) to other data are clearly defined.
2. It must be formally connected to a methodology that itself could be instantiated into software tools, giving us the following potential:
    a. Using software tools to manage, validate, assemble, and publish the version 3 Standard's products, including the eXtensible Markup Language (XML) schemas that would be needed to hold the metadata of the information being exchanged.
    b. *Automating* the generation of interchange specifications and creating an environment where two or more individuals working from the same set of requirements could generate identical specifications that could be measured against a similar, software-generated set of conformities.

"Messages support connections between IT systems that are involved in supporting a workflow."

These features evolved into the goals of HL7 version 3, which was first published in 2004. Our goals for HL7 version 3 have also evolved into three distinct delivery mechanisms that can be used for a version 3-based interchange. These delivery mechanisms are discussed next.

An interchange can be a message similar to an HL7 version 2 message. version 3 messages today are produced in XML syntax. Messages support connections between IT systems that are involved in supporting a workflow. HL7 version 2 describes a workflow through a pre-defined and somewhat vague descriptor called a "trigger event." For example, within an institution, a trigger event could be a physician's lab order placed at the patient's bedside to the hospital pathology laboratory. This workflow may be clearly understood by everyone involved in every possible action step required, from the composition of the order to the reporting of the signed final test results to the assigned individual. However, the same workflow within an ambulatory setting can be far more difficult to describe, because the actors in the workflow do not have an integrated understanding of each other's role.

Most importantly, it is not ideal that we lack a formal behavioral model (i.e., a dynamic or workflow model), but we compensate for this by analyzing the existing workflow and then programming our interchange steps to accommodate it. Secondly, we can only create different interchange patterns as we encounter them. In both cases, a formal dynamic model could document the supported workflow; assign a unique identifier to that workflow that could be used by all parties; and serve as a base platform to be modified and re-used when slightly new or different workflows are encountered.

"A formal dynamic model could be modified and re-used when slightly new or different workflows are encountered."

Messages have an expected lifetime bound by the instant in time that they are created and used. A message has no further use once the content of a message is consumed by the target system and optionally used to change the state of its database. Hypothetically, a message might be used as a possible means of assisting in a database recovery action, something that, hopefully, is never needed.

The next delivery mechanism is Clinical Documents, based on the HL7 Clinical Document Architecture (CDA) standard, which is a part of HL7 version 3, and which shares and uses the same underlying artifacts (for example, HL7 RIM, R-MIMS, Data Types, etc.). Clinical Document templates are implementation guides based on the CDA. Clinical Documents contain structured clinical data that also conform to document requirements of being signed (electronically), immutable, and so forth. Documents may indirectly be the product of and/or support a workflow, but they stand on their own and typically have an indefinite life.

"Clinical Documents contain structured clinical data that also conform to document requirements of being signed (electronically), immutable, and so forth."

The third delivery mechanism is a service that is based on the SAEAF. This mechanism (as messages and documents) will be defined by the relevant HL7 Work Group responsible for the interchange's content (for example, Structured Documents for CDA). In HL7, services are currently based on the principles of a Services Oriented Architecture (SOA) that is, itself, based on the antecedent concepts developed in the Referent Model of Open Distributed Processing (RM-ODP).

"This mechanism (as messages and documents) will be defined by the relevant HL7 Work Group responsible for the interchange's content (for example, Structured Documents for CDA)."

HL7's Services-Aware Enterprise Architecture Framework

SAEAF was created by the HL7 Architecture Board (ArB) as a response to a request from the CTO for an HL7 architecture that is driven by HL7's commitment to the following three principles:

1. HL7 produces specifications to enable Computable Semantic Interoperability (CSI) between users of systems implementing those specifications.
2. Instances of CSI between two or more HL7-based systems may cross department, enterprise, and/or national boundaries.
3. An HL7 Enterprise Architecture Specification (EAS) is required, if HL7 is to produce durable specifications that enable CSI in an effective, efficient, and scalable manner.

The ArB recognized early on that three critical components were missing from HL7 and, therefore, they needed to develop SAEAF. These are the missing components:

1. A Behavioral Framework to express interaction semantics.
2. A layered Enterprise Conformance/Compliance Framework (ECCF) to support service integration and run-time assessment of CSI.
3. A Governance Framework to oversee the development and implementation of service (and other HL7 Interoperability Paradigm) specifications.

Computable Semantic Interoperability (CSI)

In order to understand the aforementioned, it is necessary to understand a little of what HL7 sees as the requirements to achieve CSI. These necessary, but not inclusive, requirements include the following:

1. Common static models (for example, the HL7 RIM) across all domains of interest including:
   a. An information model versus a data model.
   b. The semantics of common structures.
   c. Models based on robust data-type specifications.
2. A mutually understood behavioral (or dynamic) model that enables sufficient (as defined by the problem space) understanding of the "use context" of the creation of the data by the producer and its intended use by the consumer.
3. Methodology for binding to concept-based ontologies that support these constraints:
   a. Domain-specific semantics.
   b. Country, regional, or use-domain selection of appropriate ontologies.
   c. Rigorous versioning release cycle management to ensure that individual terminologies are consistently interpreted by both the producer and consumer of the data.

4. A formally-defined process for specifying structures that contain both data and defined actions to be exchanged between machines, i.e., a data exchange (such as a message or a service) or an electronic document or services specification.

Looking at the first three requirements of CSI just outlined, it is useful to understand them as the dimensions of variability for any data element exchanged between two HIT systems. As Figure 2 illustrates, three properties must be defined to exchange data elements: the data, terminology, and process or behavior. Taken together, these three dimensions provide a context for the information to be shared. We need to know data type (e.g., integer, text, image), terminology (i.e., the specific ontology being used), and process or behavior (e.g., a diagnostic process).

These three requirements, in addition to the fourth requirement of actually moving the data, are necessary for CSI. Figure 2 depicts this.

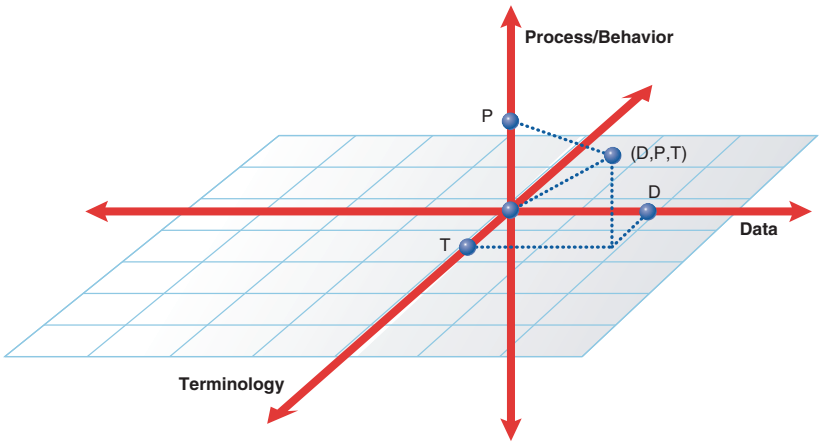

Figure 2: Information Context for Data Elements
Source: Health Level 7

Summary and Next Steps

HL7 has begun the effort of defining an architectural framework called SAEAF that is both accommodating of services-based interoperability in an SOA and capable of providing a framework for holding, comparing, and analyzing the similarities and differences among HL7's products. Going forward, SAEAF is also the basis for EAS that will be used in the future to apply HL7 to user SOA environments.

In our next section, we move on to describe the HL7 CDA, and we examine its role in healthcare in the United States today.

"HL7 is also developing standards for the representation of clinical documents (such as discharge summaries and progress notes)."

The Templated CDA Strategy: Scalable and Incremental Interoperability

Many people know of HL7 as an organization that creates healthcare messaging standards. HL7 is also developing standards for the representation of clinical documents (such as discharge summaries and progress notes). These document standards make up the HL7 clinical document architecture (CDA). The HL7 CDA, Release 1, became an ANSI-approved HL7 standard in November 2000. CDA Release 2 became an ANSI-approved HL7 standard in May 2005, and it is now in widespread use across the globe.

What follows is an introduction and overview of the CDA specification, along with a description of how and why CDA has emerged as a cornerstone component of the United States healthcare interoperability strategy.

"Approximately 1.2 billion clinical documents are produced in the United States each year. Dictated and transcribed documents make up around 60 percent of all clinical notes."

The need for a clinical document standard stems from the desire to unlock the considerable clinical content currently stored in free-text clinical notes, and to enable comparison of that content in documents created on widely different information systems. Approximately 1.2 billion clinical documents are produced in the United States each year. Dictated and transcribed documents make up around 60 percent of all clinical notes. These documents contain the majority of physician-attested information and are used as the primary source of information for reimbursement and proof of service.

The challenge, addressed by CDA and the *templated CDA* strategy, is to continue to meet the needs of front-line clinicians today, who are heavily dependent on largely narrative documents, while at the same time providing a migration pathway for greater and greater discrete data. In other words, the goal of CDA is not simply to provide yet another format for the exchange of clinical documents. If that were all we wanted, we could use PDF, MS Word*, or any other file format. Instead, our goal is to get at the clinical content within those documents and make it computable, accessible to a computer, so it can be used for such things as decision support and quality reporting. We want to do this in a way that fits in to real-world clinical workflows so that we can tackle the problem incrementally, without the need for massive process redesign.

"Our goal is to get at the clinical content within those documents and make it computable, accessible to a computer, so it can be used for such things as decision support and quality reporting."

From a technical perspective, the HL7 CDA is a document markup standard that specifies the structure and semantics of a clinical document. A CDA document is a defined and complete information object that can exist outside of a message, and it can include text, images, sounds, and other multimedia content. Just as you can create a document in MS Word, in PDF, etc., you can create a clinical document in CDA format. CDA documents are encoded in XML and derive their machine *processable* meaning from the HL7 RIM.

CDA is based on a principle of incremental interoperability, whereby an implementer can begin with a simple CDA, and then add structured data elements over time. CDA R2 consists of a single CDA XML schema, and the *architecture* arises from the ability to apply one or more *templates* that serve to constrain the richness and flexibility of CDA. Professional society recommendations, national clinical practice guidelines, and standardized data sets can be expressed as CDA templates.

"All CDA documents must use the same narrative markup so that you can receive a CDA document from anyone in the world."

From a practical perspective, incremental interoperability means that one can easily create a minimally-conformant CDA document — not really much different than creating an HTML document. CDA documents have a header that identifies and classifies the document and provides information on authentication, the encounter, the patient, etc. There are a handful of required fields and a number of optional fields. The body of the CDA can be purely narrative, by using markup very similar to XHTML. All CDA documents must use the same narrative markup so that you can receive a CDA document from anyone in the world, and, following a defined algorithm, render the document such that the receiving clinician correctly views content that was attested to by the originating clinician.

In addition to its narrative markup, CDA provides XML markup for formally representing the clinical statements within the narrative. A complete encoding of all clinical utterances can be hard if not impossible, and there is no model of healthcare that can fully and formally represent everything a clinician might say. While the HL7 RIM is a richly expressive model that can represent much of clinical narrative, the *templated CDA* strategy shields developers from the need to learn all the RIM nuances. Developers only need to understand those templates that have been recommended by the Healthcare Information Technology Standards Panel (HITSP) or put on the Certification Commission for Healthcare IT (CCHIT) roadmap. Developers simply map their internal data stores to the prioritized templates.

"CDA provides XML markup for formally representing the clinical statements within the narrative."

Assume that next year, HITSP defines the way in which we should communicate medical conditions and allergies. Assume that CCHIT takes these HITSP patterns and creates corresponding certification requirements. For those using CDA, there is no need to change to a new XML schema, and there is no need to change the approach to communicating narrative notes. CDA templates for medical conditions and allergies are created. An application maps their internal data stores to these templates, and it includes corresponding structured markup for medical conditions and allergies into CDA, whilst making no change to the narrative. On the recipient side, there will continue to be those who only render the document, whereas there may also be those who can parse out the formally encoded medical conditions and allergies for use in decision support, disease management, personalized medicine, and many other critical healthcare delivery requirements. Figure 3 shows the CDA constrained by the data elements of the Continuity of Care Record (CCR, in red) named the Continuity of Care Document (CCD). In this example, the CDA template includes additional data elements not found in the CCR (such as Chief Complaint and Discharge Diagnosis).

"For those using CDA, there is no need to change to a new XML schema, and there is no need to change the approach to communicating narrative notes."

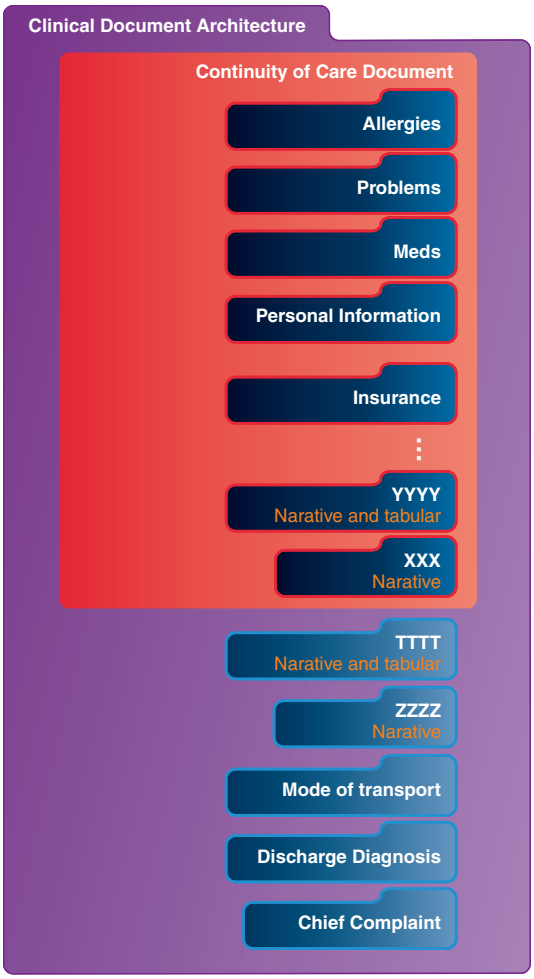

Figure 3: CDA constrained by Additional Data Elements
Source: Health Level 7

Next year, HITSP and CCHIT will prioritize new templates; the following year more templates will be produced, and so it goes — scalability arises from the fact that one XML schema does it all; incrementalism arises from the fact that only well-described changes need to be introduced over time, on a defined roadmap.

CDA is attractive for these reasons:

- *Scope*. There is a single XML schema for all CDA documents.
- *Implementation experience.* CDA has been a normative standard since 2000, and it has been balloted through HL7's consensus process. CDA is also widely implemented.
- *Gentle on-ramp to information exchange.* CDA is straight-forward to implement, and it provides a mechanism for incremental semantic interoperability. One can begin with a simple CDA narrative document implementation and add discrete data elements over time, based on national priorities.
- *Improved patient care.* CDA provides a mechanism for inserting evidence-based medicine directly into the process of care (via templates), thereby enabling the application of evidence-based medicine.
- *Lower costs.* CDA's top-down strategy lets you implement CDA once, and reuse it many times for new scenarios.

These attributes of CDA have lead to its adoption as a core healthcare information technology component in many countries, including the United States, in particular by HITSP. Interoperability on a use-case by use-case basis can lead to a disjointed set of standards that do not support re-use or internal consistency. What is needed is an overarching approach to interoperability, one that is both widely encompassing and scalable, as new use cases are developed.

A strategy being exploited by HITSP, as well as the Integrating the Healthcare Enterprise (IHE) and other international HIT initiatives, is to base a growing number of Interoperability Specifications (ISs) on the CDA, or, more precisely, on *templated CDA.* From its inception, CDA has supported the ability to represent professional society recommendations, national clinical practice guidelines, and standardized data sets as "templates" or constraints on the generic CDA XML. Perhaps the best known example of a templated CDA specification is the ASTM/HL7 Continuity of Care Document (CCD) specification (see Figure 4), where the standardized data set defined by the ASTM Continuity of Care Record (CCR) is used to further constrain CDA, specifically for summary documents.

"What is needed is an overarching approach to interoperability, one that is both widely encompassing and scalable, as new use cases are developed."

Subsequent to its adoption of CCD as the basis for HITSP/C32 "Summary Documents using HL7 CCD," HITSP recognized that a top-down strategy, whereby one learns CDA once and then re-uses it in other ISs, leads to greater cross use-case consistency. HITSP has since created a library of CDA templates (HITSP/C83 "CDA Modules Component") that are used within a growing number of CDA-based specifications (for example, HITSP/C28 "Emergency Care Summary," HITSP/C32 "Summary Documents Using HL7 CCD," HITSP/C105 "Patient Level Quality Data Document Using HL7 Quality Reporting Document Architecture (QRDA)," HITSP/C48 "Encounter Document constructs," HITSP/C84 "Consult and History & Physical Note Document," and HITSP/C78 "Immunization Document").

Additional templated CDA specifications being developed within IHE and HL7 include the QRDA, the Public Health Case Report Document, Operative Reports, Personal Health Monitoring Reports, and Minimum Data Set, version 3. What is true across the spectrum of specifications is that all the specifications conform to the underlying CDA XML, and templates are re-used to the extent possible. CDA, coupled with a library of re-usable templates, forms the basis for a growing number of HITSP ISs, and CDA represents a national interoperability strategy that is both widely encompassing and scalable as new use-cases are developed.

Considerable discussion is taking place now in the United States around the notion of *meaningful use*, and about ensuring that our interoperability strategy supports *meaningful use*. At the heart of meaningful use is simply a requirement for *data re-use* as in re-using clinical trial data in the construction of decision support rules, re-using clinician-captured data for quality reporting, public health reporting, etc. Imagine, for instance, that there were separate models and XML schemas for immunization data, medication administration, pharmacy dispensing, lab, and clinical summaries, and that the onus was on the implementer to reconcile the differences in the data in order to support data re-use. CDA and templated CDA address the concept of meaningful re-use, by maximizing data re-use.

The value of the RIM, and the rationale for its use as the underlying formalism for encoding clinical statements in a CDA document include the following:

- *Consensus.* A consensus process is used in the development of the RIM that encompasses many years, many vendors, many countries, and many use-cases.
- *Expressivity.* This allows for the formal representation of many types of clinical statements.
- *Data re-use.* All version 3 specifications are derived from a common model.
- *Concrete.* While not perfect, RIM is here today for us to use.
- *Governance and maintenance.* There is a defined consensus process for revisions.

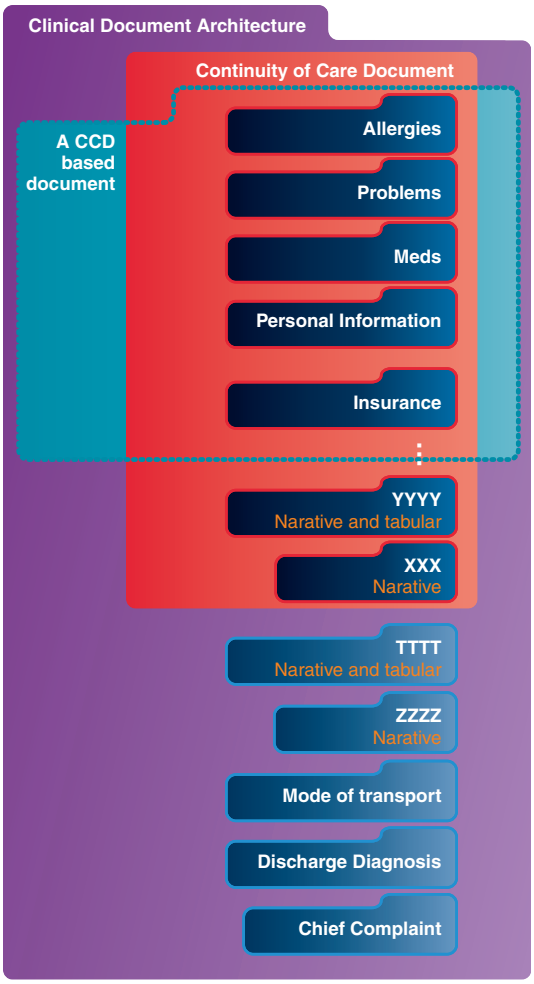

Figure 4: *Templated CDA* Specification
Source: Health Level 7

"CDA and templated CDA address the concept of meaningful re-use, by maximizing data re-use."

"It is anticipated that several ballot cycles will be required to work through the list of new requirements, and that CDA Release 3 will become an ANSI-approved HL7 standard sometime in 2011."

CDA Release 1 became an ANSI approved HL7 standard in November 2000. CDA Release 2 became an ANSI-approved HL7 standard in May 2005. Balloting on CDA Release 3 is slated to begin towards the end of 2009. Given that HL7 has three ballot cycles per year, and given the widespread adoption of CDA Release 2, it is anticipated that several ballot cycles will be required to work through the list of new requirements, and that CDA Release 3 will become an ANSI-approved HL7 standard sometime in 2011. CDA Release 3 requirements are cataloged as formal proposals and as suggested enhancements.

The main feature enhancement expected for CDA Release 3 is the inclusion of much more of the HL7 RIM, thereby enabling a wider range of clinician narrative to be expressed.

Standards and Standards Development Organizations: A New Collaboration

"Standards development is a political process."

Standards development is a political process. As described previously, the stakeholders are diverse but not every interest is equally represented. In many countries, the dominant voices in this process are the government agencies that determine the business requirements and fund the development initiatives. The other end of the spectrum is represented by the vendor community that has always made a substantive contribution to the definition of the standards' artifacts. In the past, the caregiver community has been significantly disenfranchised.

That is not to say that credentialed professionals have not been leaders of standards development organizations and government agencies. Even when clinicians lead these organizations, however, decision-making processes are often blunted by layers of administrative bureaucracy and regulatory overhead. Within the healthcare IT community, the chief medical information officer has often provided critical guidance and commanded substantive influence.

There is a growing body of critical data and published studies that describe the successful deployment of large-scale IT solutions within multi-disciplinary, healthcare systems. In fact, the application of eHealth processes has been shown to be transformative in various large, healthcare provider organizations. Unfortunately, recent exposés in the public media have recounted instances of dire unintended consequences of new or updated healthcare IT solutions. More often than not, the problem with these installations has not been technical. The apparent proximal cause has been the significant failure by system designers and software developers to understand the workflow and daily business requirements of the clinical end-users.

"The application of eHealth processes has been shown to be transformative in various large, healthcare provider organizations."

At the technical end of the implementation spectrum lies the need to ensure seamless, unambiguous interchange of data. The inability to achieve this is rarely the failure of the technical standards (specifications) themselves. Often there is a lack of formal agreement on vocabulary, the definition of individual data elements, and the application of the required terminology within the patient care continuum. This is exacerbated when requests are made or requirements defined across clinical boundaries. In simple terms, a physical therapist may not apply the same meaning to a term as the orthopedic surgeon who first used it. In the in-patient setting, this is repeated daily in a failure to achieve unambiguous communication between physicians, nurses, pharmacists, and laboratory staff.

"In simple terms, a physical therapist may not apply the same meaning to a term as the orthopedic surgeon who first used it."

Attempts to rectify these ambiguities have often failed to overcome parochial and economic hurdles. In a practical sense, there is an ongoing tension among primary caregivers and the specialist community. The gatekeeper model of the managed healthcare system failed to achieve either fiscal or clinical outcome metrics because of the dissension that was exacerbated when practice variability was confronted with ambiguous, often counter-productive payment schema. In those managed care systems with closed practice systems and a single information system (e.g., Kaiser Permanente), the results are uniformly significantly more successful.

Outside of the United States, stringent government regulation has often led to improved outcomes following implementation of enhanced healthcare information systems. Perhaps this is more related to a single payer system and a more homogeneous patient population than to the technology itself. At times, regulatory oversight has been an enormous obstacle to data interchange. For example, the global regulated research community (chiefly pharmaceutical and biotech industry) has embraced a structured vocabulary called the Medical Dictionary for Regulatory Activities (MedDRA). This terminology is principally required for data encoding in the reporting of adverse events in clinical trials and in post-approval pharmacovigilance. The system is largely incompatible with SNOMED CT (Systematized Nomenclature of Medicine — Clinical Terms), which is in widespread use for patient care worldwide. Both clinical research and patient care suffer because of the artificial barriers to effective information exchange.

"Outside of the United States, stringent government regulation has often led to improved outcomes following implementation of enhanced healthcare information systems."

These problems transcend the functional requirements for interoperability. For our purposes, it is best to rely on the definition established by IEEE. Interoperability is the ability of two or more systems or components to exchange information. Semantic interoperability is the ability to use the information that has been exchanged. Faxing an EKG tracing between two professional offices provides a high degree of interoperability, if both clinicians agree on the parameters to use for interpretation. Like most data that are not encoded, the EKG cannot be re-used in any meaningful way. Parenthetically, HL7 has developed a standard for encoding an annotated EKG.

"Interoperability is the ability of two or more systems or components to exchange information. Semantic interoperability is the ability to use the information that has been exchanged."

"These practice guidelines are predicated on more data being derived from clinical research and have been referred to as evidence-based medicine."

Reducing the ambiguity of the data is increasingly important as healthcare managers and researchers attempt to measure quality. For our purposes, the enhanced clinical quality and patient-care outcomes are realized when practice guidelines are adhered to by individual caregivers and system-wide care requirements. Usually these guidelines are established by central authorities (such as the US National Institutes of Health) and by professional societies (for example, the American College of Cardiology). Globally, quality estimates (outcomes) are provided by the World Health Organization.

In the United States, these guidelines have been historically written by the Agency for Healthcare Research and Quality (AHRQ), an agency of the Department of Health and Human Services. More recently, non-profit organizations, such as the National Quality Forum, have been created to take a more proactive role in writing these guidelines. These groups establish priorities for quality evaluation as well as establishing the parameters for measuring success. These practice guidelines are predicated on more data being derived from clinical research and have been referred to as evidence-based medicine. These guidelines are critical to improving patient outcomes and reducing costs. Unfortunately, the clinical guideline for one disease may be contraindicated when complying with the guideline for a co-morbid disease. For example, the use of a relatively common anti-inflammatory drug for arthritis, such as ibuprofen, may be contraindicated in the management of hypertension. These two ailments are often concurrent in the Medicare-aged population, but practice management algorithms would indicate a violation of quality patient care. The development of treatment algorithms that weigh the relative severity (importance) of two co-morbid diseases is the foundation for ongoing research.

"The development of treatment algorithms that weigh the relative severity (importance) of two co-morbid diseases is the foundation for ongoing research."

Other instances of conflicts of standard practice guidelines are less easily resolved. These conflicts may represent the differences in the interpretation of medical evidence between two medical societies or the same specialty in different regions or countries. It is easy to imagine that the recommendations offered by the American Academy of Orthopedic Surgeons would not coincide with those of the American Chiropractic Association. HL7 is working closely with the National Quality Forum and the Agency for Healthcare Research and Quality to standardize these metrics.

"Decision support systems are highly complex, predicated on the integration of a vast amount of patient care and research data, and reliant upon the harmonization of technical standards."

Delivering evidence-based medicine at the point of care (inpatient, ambulatory, emergent, or chronic home care) is enabled by a technology that is broadly classified as decision support. Decision support systems are highly complex, predicated on the integration of a vast amount of patient care and research data, and reliant upon the harmonization of technical standards. Within HL7, the Decision Support Working Group has advanced the technical functionality of these systems. Practical implementation is more complex. Individual system vendors have implemented decision support by using a vast array of technologies, alerts, graphical interface representations, and workflow modification. At the most simple level, this technology may be implemented

in the form of a clinical alert, for example, when two drugs with a potential for adverse interaction are prescribed concurrently. Another example of decision support might be a recommendation for an annual PAP smear at the time of an outpatient visit. The most sophisticated systems guide complex disease management and require data from lab testing and medication administration, as in guidance for insulin dosing. HL7 is working closely with several decision support initiatives, such as the Clinical Decision Support Consortium.

"HL7 is working closely with several decision support initiatives."

The most innovative program does not directly involve the development of standards. It is an initiative in social engineering. Nearly four years ago, HL7 undertook a program for including specific clinical input into the standards development process. This Work Group recognized the contributions of a wide range of members of the caregiver community, including physicians, nurses, pharmacists, and even medical librarians. While important progress was made towards the identification of critical development pathways, the Work Group sorely lacked a focus on patient care, public health, and program management. In 2008, HL7 recognized the immediate need to align with the caregiver communities. With funding from the AHRQ, an experiment was first undertaken. The experiment went by the name of Bridging the Chasm.

"In 2008, HL7 recognized the immediate need to align with the caregiver communities."

At first blush, the chasm to be bridged was between the healthcare community and the organization of healthcare IT professionals. It was most evident to the clinicians that everyone was speaking a different language, filled with techno-speak, jargon, and acronyms, with which they were embarrassingly unfamiliar. Moreover, concepts, such as knowledge representation, so commonplace in IT, were lost on the clinicians. To re-purpose an old adage, "the clinicians were mad and they weren't going to take it anymore." Of course, there was another, much older, and often more acrimonious divide to bridge. That was the divisions erected over hundreds of years between specialists and primary care physicians, between physicians and nurses (you can substitute almost any other caregiver function), and between many of the ancillary but critical roles of pharmacist, dietician, physical therapist, pathologist, and many others.

In April 2009, a conference was convened in Washington to begin the bridge building. Funded by AHRQ and led by HL7, over 100 professional societies met with one objective in mind: begin to define the terminology, workflow processes, business requirements, and the like that everyone had in common. Of course, standards development was to be a collateral outcome, but only after the difficult task of bridge building was underway. This was not about fence mending, since it was hoped that the fences would be torn down eventually. For two days there was no jargon. Not an acronym was uttered. The results were unprecedented, because this heterogeneous body moved forward with unanimity of purpose.

"Standards development was to be a collateral outcome, but only after the difficult task of bridge building was underway."

"In the office setting, in the home, and in the hospital, the patients will be the beneficiaries."

With the creation of a new professional society, the Clinical Information Interchange Collaborative, development of common terminology seems possible. Integration of workflow into eHealth systems seems achievable. Articulating a business case for the harmonization of practice parameters and clinical guidelines has begun. In the office setting, in the home, and in the hospital, the patients will be the beneficiaries.

Standards development for healthcare will always be about politics. The problems to be solved have and will always have highly technical solutions. Interoperability is not a goal, but rather the means to improving quality and reducing costs. Creating innovative solutions to complicated issues of unified vocabulary and seamless data integration can be realized in an environment of collaboration.

Further Reading

Blobel, Bernd. "Making hospital IT interoperable." HITE 2008:1:2: Summer 2008:12-15. Available at ***www.hospitaliteurope.com***

Clinical Data Interchange Standards Consortium (CDISC)
This website contains standards and policy documents cited in this article.
http://cdisc.org/

Hammond, William Edward; Jaffe, Charles; Kush, Rebecca Daniels, "Healthcare Standards Development: The Value of Nurturing Collaboration" Journal of AHIMA, July 2009 80/7, pages 44-40.

Health Level Seven, Inc.
www.hl7.org
On the website can be found more detailed supporting material cited in this article:
CDA Tutorial, 2008.
HL7 version 2.5 and HL7 2.6 Standards.

Dolin RH, Alschuler L, Boyer S, Beebe C, Behlen FM, Biron PV, Shabo A. "HL7 Clinical Document Architecture, Release 2." *J Am Med Inform Assoc.* 2006;13:30–39. Available at ***www.jamia.org.***

HITSP/C32: "Summary Documents Using HL7 Continuity of Care Document (CCD); Component." July 8, 2009. Version 2.5.

HITSP/C37: Lab Report Document Component. July 8, 2009. Version 2.3.

HITSP/C75: Healthcare Associated Infection (HAI) Report Component. July 8, 2009. Version 1.1.

HITSP/C83: CDA Content Modules Component. July 8, 2009. Version 1.1.

HITSP/C105: Patient Level Quality Data Document Using HL7 Quality Reporting Document Architecture (QRDA). June 30, 2009. Version 0.0.1.

European Committee for Standardization (CEN) Technical Committee 251 (Healthcare)
This site contains standards and policy information cites in this article.
http://www.cen.eu/cenorm/sectors/sectors/isss/index.asp

ISO Technical Committee 215 (Healthcare Informatics)
This workspace contains documents and standards cited in this article:
http://isotc.iso.org/livelink/livelink?func=ll&objId=529137&objAction=browse

Acronyms

Agency for Healthcare Research & Quality (AHRQ)
American National Standards Institute (ANSI)
Association Française de Normalisation (AFNOR)
British Standards Institution (BSI)
Certification Commission for Healthcare IT (CCHIT)
Clinical Data Interchange Standards Consortium (CDISC)
Clinical Document Architecture (CDA)
Comité Européen de Normalisation (CEN)
Computable Semantic Interoperability (CSI)
Continuity of Care Document (CCD)
Continuity of Care Record (CCR)
Enterprise Architecture Specification (EAS)
Enterprise Conformance/Compliance Framework (ECCF)
Extensible Markup Language (XML)
German Institute for Standardization (DIN)
Health Level Seven (HL7)
Healthcare Information Technology Standards Panel (HITSP)
Integrating the Healthcare Enterprise (IHE)
International Health Terminology Standards Development Organization (IHTSDO)
International Standards Organization (ISO)
Interoperability Specifications (ISs)
Joint Initiative Council (JIC)
Quality Reporting Document Architecture (QRDA)
Reference Information Model (RIM)
Referent Model of Open Distributed Processing (RM-ODP)
Services Aware Enterprise Architecture Framework (SAEAF)
Standards Developing Organization (SDO)
SDO Charter Organization (SCO)
Services Oriented Architecture (SOA)

Author Biographies

Charles Jaffe, MD, PhD, FACP. Since 2007, Dr. Jaffe has served as the Chief Executive Officer of HL7. He began his career at Intel in 2005 as the Senior Global Strategist for the Digital Health Group. He has also served as Vice President of Life Sciences at SAIC and as the Global Director of Medical Informatics at AstraZeneca Pharmaceuticals. He completed his medical training at Johns Hopkins and Duke Universities, and he was a post-doctoral fellow at the National Institutes of Health and at Georgetown University. Formerly, he was President of InforMed, a consulting firm for research informatics. Over his career, he has been the principal investigator for more than 200 clinical trials, and he has served in various leadership roles in the American Medical Informatics Association. He has been a board member on leading organizations for information technology standards, and he served as the Chair of a national institutional review board. Currently, he holds an appointment in the Department of Medicine at the University of California at San Diego. He has been the contributing editor for several journals and he has published on a range of subjects, including clinical management, informatics deployment, and healthcare policy.

W. Ed Hammond, PhD, is Professor-emeritus, Department of Community and Family Medicine and Professor-emeritus, Department of Biomedical Engineering, Duke University and Adjunct Professor in the Fuqua School of Business at Duke University. He has served as President of the American Medical Informatics Association (AMIA), President of the American College of Medical Informatics, and as Chair of the Computer-based Patient Record Institute. He is currently serving his third term as the Chair of HL7. He was Chair of the Data Standards Work Group of the Connecting for Health Public-Private Consortium. Dr. Hammond was a member of the IOM Committee on Patient Safety Data Standards. He was awarded the Paul Ellwood Lifetime Achievement Award in 2003 and the ACMI Morris F. Collen Award of Excellence in November 2003.

John Quinn is the Chief Technology Officer of HL7 and has served in this full-time position since September 2007. HL7 is one of several ANSI -accredited SDOs operating in the healthcare arena.

Quinn is also a Senior Executive, Chief Technology Officer and Thought Leader in Accenture's Health Life Science's Provider Practice focusing on healthcare information system's technologies, architectures, and data and messaging standards; and he serves in his role as HL7's CTO in part through a significant contribution by Accenture. He has over 32 years experience in Healthcare IT and 35 years experience in computer industries, respectively. He has participated on the HL7 Board of Directors since its inception 21 years ago and served as the Chair of its Technical Steering Committee from 1989–2007.

Robert H. Dolin, MD, MS, serves as the Chair-elect on the HL7 Board of Directors and has been involved with the organization since 1996. Dr. Dolin is an internationally renowned and innovative physician expert in healthcare information technology standards development with more than 20 years of clinical experience. After receiving his medical degree, he served as chief resident at UCLA Department of Medicine in 1989, where he developed Harbor-UCLA first outpatient electronic medical record system. Dr. Dolin is a Fellow in the American College of Physicians, and he was elected into Fellowship in the American College of Medical Informatics in recognition of his work on standards development. Dr. Dolin currently owns Semantically Yours, a healthcare standards consulting firm. Previously, he spent 18 years as a Hospitalist at Kaiser Permanente, Department of Internal Medicine. At Kaiser, he was the physician lead for the enterprise terminology services team, responsible for deploying SNOMED in their national electronic health record. Dr. Dolin also co-chairs the HL7 Structured Documents Work Group, and is co-editor of the HL7 CDA and the CCD specifications. He has also served on the SNOMED CT Content Committee, and currently co-chairs the US HITSP Foundations Committee. Dr. Dolin's work has also been published in both technical and clinical peer reviewed journals.

Copyright