



HEALTH LEVEL SEVEN® INTERNATIONAL

The Worldwide Leader in Interoperability Standards

March 2010

What is HL7?

Health Level Seven International (HL7), an ANSI-accredited standards developing organization (SDO), is the global authority on standards for interoperability of health information technology. With members in more than 55 countries, HL7 is deeply involved in worldwide efforts to improve healthcare through information technology and is a founding member of the Joint Initiative Council, an international council on global health informatics standardization that is committed to developing a single standard for a single purpose. HL7 also has an agreement with the International Organization of Standardization (ISO) through which HL7 submits its ANSI-approved or DSTU standards directly to ISO for approval.

Founded in 1987, HL7 is a nonprofit organization comprised of more than 4,000 worldwide members who represent hundreds of healthcare vendors, providers, payors, government agencies, consultants and others. In the U.S. alone, 90 percent of the largest health information system vendors are HL7 members. Volunteers perform HL7's standards development work.

HL7 does not develop software. It creates standards that allow healthcare information to be communicated across and between healthcare enterprises and communities. HL7 standards facilitate the exchange of clinical and administrative data among health information systems. Specifically, HL7 provides a comprehensive framework and related standards for the exchange, integration, sharing, and retrieval of electronic health information that supports clinical practice and the management, delivery, and evaluation of health services.

The most widely used HL7 specifications are messaging standards that enable disparate healthcare applications to exchange key sets of clinical and administrative data. HL7's Version 2.x messaging standard is arguably the most widely implemented standard for healthcare in the world. In 2009, it was published as an ISO standard. In the US, the HL7 Version 2 messaging standard is deployed at most healthcare facilities. The HL7 Version 3 messaging standard is used by U.S. government agencies such as the Food and Drug Administration (FDA) and the Department of Veterans Affairs. Version 3 is also widely used outside the U.S. in countries such as Canada, the United Kingdom, the Netherlands, Germany and Mexico.

The American Recovery and Reinvestment Act in the U.S., has placed an elevated emphasis on healthcare IT. One of the criteria for showing "meaningful use" of an EHR—a requirement for receiving government financial incentives—is to exchange electronic data with other healthcare providers. The HL7 standards already in place and those that HL7 is developing or bringing online in the U.S. will be instrumental in achieving the interoper-

ability that will enable providers to exchange data easily across healthcare communities. In fact, the Office of the National Coordinator for Health Information Technology in the U.S. Department of Health and Human Services selected HL7 Version 2 and the HL7 Clinical Document Architecture (CDA) / Continuity of Care Document (CCD) in its initial set of standards, implementation specifications and certification criteria for EHR technology.

The HL7 Clinical Document Architecture (CDA) is an important step to achieving interoperability. The CDA is an ISO approved standard that provides an exchange model for clinical documents (such as discharge summaries and progress notes) - and brings the healthcare industry closer to the realization of an electronic medical record. There are large-scale CDA implementations in North and South America, Europe and Asia Pacific. The Continuity of Care Document (CCD) is a CDA implementation of the continuity of care record (CCR), created by the American Society for Testing and Materials (ASTM). Disparate information systems can employ the CCD to exchange clinical summaries that contain key data about individual patients, such as diagnoses, medications, and allergies.

Our mission: "HL7 provides standards for interoperability that improve care delivery, optimize workflow, reduce ambiguity, and enhance knowledge transfer among all of our stakeholders, including healthcare providers, government agencies, the vendor community, fellow SDOs and patients. In all of our transactions we exhibit timeliness, scientific rigor and technical expertise without compromising transparency, accountability, practicality, or our willingness to put the needs of our stakeholders first."

HL7 created the EHR System Functional Model, which includes advanced decision support features, to help lay the groundwork for national interoperability of health IT systems. The model also supplies guidance to healthcare providers in preparing for, acquiring, and making the transition to electronic health record systems and was published as an ISO standard in late 2009. The HL7 Personal Health Record System Functional Model, a draft standard, identifies the functions that should be included in a PHR and includes guidelines for data exchange between PHRs and EHRs.

Other SDOs develop standards for particular areas of healthcare, such as laboratory orders and results, electronic prescribing, and medical device connectivity. HL7 is the only SDO that provides the messaging specifications to connect all of the systems in a healthcare enterprise, such as a hospital.

What Does the Name HL7 Mean?

"Level Seven" refers to the highest level of the International Organization for Standardization (ISO) communications model for Open Systems Interconnection (OSI)—the application level. The application level addresses definition of the data to be exchanged, the timing of the interchange, and the communication of certain errors to the application. The seventh level supports such functions as security checks, participant identification, availability checks, exchange mechanism negotiations and, most importantly, data exchange structuring.



HEALTH LEVEL SEVEN® INTERNATIONAL

The Worldwide Leader in Interoperability Standards

March 2010

How is HL7 International Organized?

The organization is managed by a Board of Directors, which is comprised of both elected and appointed positions. The organization is comprised of work groups that are responsible for defining the HL7 standard protocol.

Each work group is chaired by two or more co-chairs. Work groups are organized into four functional groups called steering divisions. Each steering division elects one representative serve on the Technical Steering Committee (TSC). This committee is also comprised of two representatives from the HL7 affiliates, an Architectural review Board representative, and an ad hoc member. The TSC votes on issues related to the standards. These votes are then passed as recommendations to the Board of Directors, who make the final decision.

HL7 Background and General Information

- HL7's impressive suite of standards in the clinical healthcare realm include:

- ARDEN SYNTAX

- FUNCTIONAL SPECIFICATIONS

- Electronic Health Record System (EHR-S) Functional Model
- Personal Health Record System (PHR-S) Functional Model

- VERSION 2 MESSAGES

- VERSION 3 DOCUMENTS

- Structured Product Labeling (SPL)
- Clinical Document Architecture (CDA)
- Continuity of Care Document (CCD)

- VERSION 3 FOUNDATION

- Reference Information Model (RIM)
- Vocabulary Domains & Value Sets

- VERSION 3 MESSAGES

- VERSION 3 SERVICES

- Java
- Web

- VERSION 3 RULES

- GELLO

- VISUAL INTEGRATION

- CCOW

- HL7 Version 2 is well established as the premier messaging standard for the exchange of patient care and clinical information worldwide. HL7 Version 3, the next step in healthcare information exchange, is much more than a simple iteration of messaging standards. Developed using the HL7 Development Framework (HDF) and supported by the HL7 Reference Information Model (RIM), Version 3 brings a new level of interoperability to healthcare information. With the formal binding of standard vocabularies to standard models, and a flexible document architecture employing XML, HL7 Version 3 represents a mechanism for the exchange of "understandable" information able to be reused in multiple application contexts at the highest common level of shared meaning - going beyond messaging to achieve semantic interoperability.

- The Reference Information Model (RIM) is the cornerstone of the HL7 Version 3 development process. An object model created as part of the Version 3 methodology, the RIM is a large pictorial representation of the clinical data (domains) and identifies the life cycle of events that a message or groups of related messages will carry. It is a shared model between all the domains and, as such, is the model from which all domains create their messages. Explicitly representing the connections that exist between the information carried in the fields of HL7 messages, the RIM is essential to our ongoing mission of increasing precision and reducing implementation costs.

- Consistent with its mission, HL7 has developed the HL7 Electronic Health Record-System Functional Model (EHR-S FM) Standard to help lay the groundwork for nationwide interoperability by providing common language parameters that can be used in developing systems that support electronic records. The HL7 EHR-S FM helps guide the provider community in their planning, acquisition, and transition to electronic systems; and it will facilitate a more effective dialogue between providers and vendors. HL7 has also approved the Personal Health Record-System Functional Model for use as a draft standard for trial use (DSTU).

- HL7 convenes three working group meetings per year allowing HL7 members to work on the standards face-to-face. They provide an invaluable educational resource for the healthcare IT community, offering a variety of tutorials as well as HL7 certification testing. HL7 also holds three educational summits per year where a streamlined set of tutorials and HL7 certification testing are also offered. Finally, HL7 provides on-site training by request.

- Numerous HL7 Affiliates have been established around the globe including Argentina, Australia, Austria, Brazil, Canada, Chile, China, Colombia, Croatia, Czech Republic, Denmark, Finland, France, Germany, Greece, Hong Kong, India, Ireland, Italy, Japan, Korea, Mexico, The Netherlands, New Zealand, Romania, Russia, Singapore, Spain, Sweden, Switzerland, Taiwan, Turkey, The United Kingdom, and Uruguay.

Collaboration With Other SDOs

Recognizing that effective standards development requires collaboration with other SDOs and stakeholders, HL7 has established formal agreements with a number of key organizations. These pacts are designed to avoid redundancy in standards creation and work products and to facilitate open communication throughout the industry. Among the organizations with which HL7 has such agreements are the Accredited Standards Committee X12N (ASC X12N), ASTM, the Clinical Data Interchange Standards Consortium (CDISC), Digital Imaging and Communications in Medicine (DICOM), European Committee for Standardization (CEN), GS1, Institute of Electrical and Electronics Engineers (IEEE), Integrating The Healthcare Enterprise (IHE), International Health Terminology Standards Development Organization (IHTSDO), North American Association for Central Cancer Registries (NAACCR), the National Council for Prescriptions Drug Programs (NCPDP) and the Health Story Project. HL7 is also a founding member of the SDO Charter Organization (SCO), a formal collaboration among SDOs which aims to facilitate the creation of industry-wide, interoperable standards that will support meaningful improvements in healthcare outcomes.

HL7 Arden Syntax Overview

The Arden Syntax for Medical Logic Modules (MLMs) is an ANSI-approved American National Standard language for encoding medical knowledge and representing and sharing that knowledge among personnel, information systems and institutions. It is designed for organizations that require or develop systems that automatically assist physicians in decisions and alerts.

An open standard for representing clinical knowledge, it is intended to be used by individual clinicians, institutions, and vendors to develop clinical rules (rules that directly impact patient care) using a standard format and language. Arden Syntax is also suited for writing business rules that directly interface with clinical data.

The logic for making these decisions or issuing these alerts is encoded into health knowledge bases called MLMs, each of which contains sufficient knowledge to make a single decision. Contradiction alerts, management suggestions, data interpretations, treatment protocols, and diagnosis scores are examples of the health knowledge bases that can be represented using the Arden Syntax.

With an appropriate computer program (known as an event monitor), MLMs run automatically, generating advice where and when it is needed. For example, one MLM warns physicians when a patient develops new or worsening kidney failure.

The American Society for Testing and Materials (ASTM) previously adopted it as document E 1460, under subcommittee E31.15 Health Knowledge Representation. Adopted by ASTM in 1992, this became Arden Syntax Version 1.0.

Beginning in the summer of 1998, sponsorship of this standard was moved to HL7. The Clinical Decision Support and Arden Syntax Work Groups of HL7

oversee maintenance of the standard. Arden Syntax Version 2.0 was formally adopted by HL7 and ANSI in August 1999.

The move of Arden Syntax from ASTM to Health Level Seven originated in 1997. The HL7 Clinical Decision Support Work Group first met in June 1997 and was organized by Carol Broverman. In the Summer of 1999, the Arden Syntax Standard was formally transferred from ASTM to HL7, and Version 2 of the standard was formally endorsed by HL7 and ANSI.

In order to make the overall work group less technology-specific, the Arden Syntax was spun off into its own work group in January 2001, under the sponsorship of the Clinical Decision Support Work Group.

Version 2.1 of the standard was formally endorsed by HL7 and ANSI during 2002. Version 2.1 introduces a structured message as an optional part of the WRITE statement. Using this structure, which is encoded using an XML DTD, authors can specify many different parameters of an output message of a decision support system in a standard fashion. These include time outs and escalation information for alerts; embedded orders; subject; and recommendation.

The Arden Syntax standard is now at Version 2.7 and was approved by ANSI in December 2008. It features an enhanced assignment operator to allow assignment to individual elements in a list and nested attributes in objects. In addition, the object initialization statement has been enhanced to support assigning to named attributes. Corrections /updates to features introduced in previous versions and a reorganization of the evoke slot chapter are also included. It is available in the bookstore area of the HL7 International website: www.HL7.org.

The Arden Syntax for Medical Logic Modules (MLMs) is a language for representing and sharing medical knowledge among personnel, information systems and institutions.

Overview of HL7 CCOW Standard

Aimed at facilitating the integration of applications at the point of use, the Clinical Context Management Specification (CCOW) is an end-user-focused standard that complements HL7's traditional emphasis on data interchange and enterprise workflow.

Using a technique known as context management, the clinical user's experience is one of interacting with a single system, when in fact he or she may be using multiple, independent applications from many different systems, each via its native user interface.

By synchronizing and coordinating applications so that they automatically follow the user's context, the CCOW Standard serves as the basis for ensuring secure and consistent access to patient information from heterogeneous sources. The benefits include applications that are easier to use, increased utilization of electronically available information, and an increase in patient safety. Further, CCOW support for secure context management provides a healthcare standards basis for addressing HIPAA requirements. For example, CCOW enables the deployment of highly secure single sign-on solutions.

IMPACT

CCOW's impact on the healthcare industry is apparent. All of the major HIS vendors are now shipping or planning on shipping both Windows- and Web-based CCOW-compliant applications, while vendors in virtually every segment of the clinical healthcare market have adopted the standard as well. VHA Inc.-a nationwide network of 1,900 leading community-owned healthcare organizations and their affiliated institutions-now requires all of its new business partners to be CCOW-compliant. A growing number of healthcare organizations are also implementing context management solutions to link together diverse multi-vendor, multi-technology IT systems on an enterprise-wide basis.

HOW IT WORKS

The CCOW Work Group became a part of HL7 after starting out as an independent healthcare industry consortium. In a short time, the committee has developed and ratified five versions of the CCOW Standard. This unprecedented pace has been, in part, due to the modular component-based nature of the architecture upon which the standard is based, enabling new specifications to be developed in a complementary and add-on manner.

CCOW's Context Management Architecture (CMA) was founded on the principle that common context can be established across applications by identifying things, such as a patient concepts, or clinical encounters in a manner that different applications can nevertheless recognize.

The core architecture is comprised of three main types of components: applications; a context manager that coordinates and synchronizes applications; and mapping agents that can represent the various synonymous real-world identifiers used to identify clinical patients, users, etc. The architecture defines the roles and responsibilities for each of these components and precisely prescribes the interfaces that enable them to communicate. The architecture does not define or dictate the implementation of any of the components.

The user sets the context using any CCOW-compliant application, for example, to select a patient of interest. The application then tells the context manager that it wants to set the patient context and provides the context manager with an identifier that indicates the context subject, which, in this case, might be the medical record number for the patient of interest.

The context manager then tells the other applications that the context has been changed, and each application obtains the patient's identifier from the context manager. Each application then adjusts its internal state and data display accordingly. This all happens in real-time.

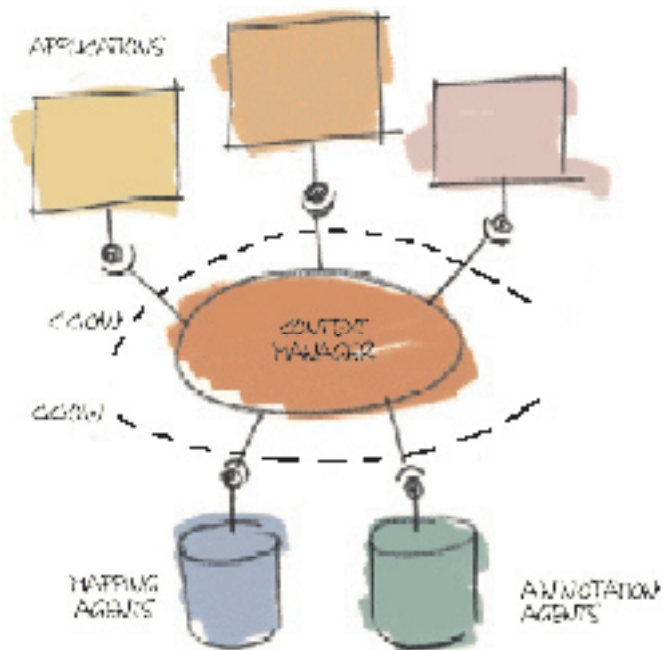
Context links may be *common* or *secure*. Any application may set or get the context

data for a common link. In contrast, only site-designated applications may set and/or get the context data for a secure link. Applications, the context manager and mapping agents use digital signatures to authenticate the messages they send and receive and to ensure the integrity of the data within these messages.

The basic idea is to provide a means for the various CCOW components to trust each other, for example, to enable applications to know that they are communicating with the real context manager (as opposed to a rogue application that is pretending to be a context manager).

One of the more elegant capabilities provided by the architecture is that the use of different context subject identifiers is hidden from applications. An application only needs to know its own identifiers. A mapping agent works in conjunction with the context manager to map the identifiers used by the application that sets the context to identifiers that may be understood by other context-sharing applications. For example, one application may use a hospital-assigned medical record number to identify patients, while another application in the same institution uses clinic-assigned medical record numbers to identify the same patients.

Aimed at facilitating the integration of applications at the point of use, CCOW is an end-user-focused standard that complements HL7's traditional emphasis on data interchange and enterprise workflow.



Detailed How It Works

The user sets the context using any CCOW-compliant application; for example, a patient of interest is selected. The application then tells the CCOW-compliant context manager that it wants to set the patient context, and provides this context manager with an identifier that indicates the context subject, which in this case might be the medical record number for the patient of interest. The context manager then notifies the other applications that the context has been changed, and each application obtains the patient's identifier from the context manager. Each application then adjusts its internal state and data display accordingly.

The CCOW architecture was designed to be easily implemented within all types of healthcare applications and using a variety of technologies. Particular emphasis was given to ensuring that CCOW compliance could be easily retrofitted into existing applications. It is not necessary for an application developer to implement a context manager or mapping agents, as these components are external to the application and can be obtained from other sources.

CCOW V1.0

Approved by ANSI in July 1999, CCOW V1.0 defined: The overall technology-neutral context management architecture (CMA), a core set of data definitions, rules for application user interfaces, and the translation of the CMA to Microsoft's COM/ActiveX technology. The features of 1.0 include:

- General architecture for "linking" applications: Compliant applications "tune" to the same "context subject"—such as a user, patient or encounter.
- Key context management components: The context manager coordinates applications and mapping agents, which map between the various identifiers.
- End-to-end security: Context-based security enforces subject-level access privileges such that only site-designated applications may access the values for a particular subject.
- Patient Link: This is a common subject that enables applications to tune to the same patient.
- User Link: This is a secure subject that enables applications to securely tune to the same user.
- COM Technology Specification: This is a specification of all of the details needed to implement COM-based applications and mapping agents that plug-and-play with a context manager, including all of the necessary COM interfaces.

CCOW V1.1

Approved by ANSI in March 2000, CCOW V1.1 added support for:

- Dependent context subjects: A dependent subject may only be set when one or more other subjects that it depends upon are also set. This ensures that the complete context, which may be comprised of multiple subjects, is always self-consistent.
- Encounter subject: This subject enables applications to tune to the same clinical encounter for a particular patient. The encounter subject is dependent upon the patient subject, so it is not possible for an application to set the encounter subject without also setting the patient subject.
- Custom subjects: A custom subject is one whose data definition has not been published as part of a ratified CCOW specification. The CCOW specification for custom subjects provides a structured way for defining the necessary data definitions and ensures that the data definition for a custom subject will not collide with those defined by other organizations or by CCOW.
- Formal conformance statements: These statements define what an application, context manager, or other type of CCOW-defined component must be capable of doing in order to claim conformance. CCOW does not address the process of validating conformance, but these conformance statements clearly establish what it means to be compliant.

CCOW V1.2

ANSI approved in September 2000, CCOW V1.2 added:

- Web Technologies Specification: Web technologies are quite different from the COM/ActiveX technologies defined for 1.0, but the technology mapping defined in 1.2 nevertheless enables interoperability between applications and mapping agents that employ either technology, as mediated by a context manager. Web-based applications and mapping agents send and receive URL-encoded HTTP messages to the context manager. These messages are analogous to the messages COM-based applications, mapping agents, and context managers send and receive, thereby providing the basis for interoperability between technologies.

- Capability to deploy over private and public networks. The security protocols and policies defined in the CMA are used in both scenarios, but Web-based CCOW over public networks adds the use of Secure Socket Layer (SSL) as the communication substrate for all CCOW-based communication.

CCOW 1.3

ANSI approved in June 2001, CCOW V1.3 added the concept of annotations:

- Annotation subjects: Annotation subjects are comprised of data elements that describe something, as opposed to identify something. An annotation subject is always dependent upon an identity subject. For example, the patient subject is used for identifying the patient.

A hypothetical demographics annotation subject would contain the patient's telephone number, address, and so on. The data for an annotation subject may only be set by annotation agent. There is an annotation agent for each annotation subject. After the context is set by an application, each available annotation agent is instructed by the context manager to add the appropriate annotation data to the context. Each site defines the data sources for its annotation agents. This ensures that applications will always see annotation data that comes from the source that is the authentic source for the data.

Two new context subjects were defined as well:

- Observation request subject: This subject enables applications to tune to the same clinical observation request, so that, for example, different applications can display the results of a particular lab test order.
- Certificate subject: This subject which is an annotation subject (see below), enables applications to tune to the same X.509 compatible digital certificate, enabling different applications to nevertheless use the same user security credentials when digitally signing and/or encrypting data on behalf of the user.

CCOW V1.4

CCOW 1.4 was released as an ANSI-approved standard in August 2002. One of the key features is support for multiple context sessions on the same point-of-use device. Each session may be securely "owned" by a different user, although only one context session will be active at a time. Multiple context sessions enable the CCOW standard to be even more flexible when used by caregivers who need to share devices in a kiosk-like manner. Users will be able to quickly and easily access their own sessions, and may "lock" their sessions when not in use, and then return to their sessions at a later time.

Another key feature of CCOW V1.4 is support for action subjects. An action subject enables an application to request, via the context manager, that another application perform a task on behalf of the requesting application. The request is generally issued in response to a user input or gesture. For example, certain clinical applications require that the user be authenticated in order to enter data, even though the user is already signed-on. With an authentication action, a clinical application

can request that the authentication application authenticate the user.

Because this communication is context-based, the two applications do not need to know about each other and yet can nevertheless service the same user.

Other CCOW V1.4 features include: a new context manager interface, ContextFilter, that enables an application to indicate the specific context subjects about which it wants to be notified whenever the context for the subject is set, and a set of data definitions for linking applications based upon DICOM image studies.

CCOW 1.5

CCOW V1.5 was approved as an ANSI standard in August 2004 and reaffirmed in 2009. It is available in the bookstore area of the HL7 website: www.HL7.org.

Approved CCOW Documents

There are three documents based on the CCOW Standard that received ANSI approval in April 2008:

- **HL7 Clinical Context Management (CCOW) Application Protection Package, Release 1**
This document presents the functional security requirements for a CCOW-compliant application.
- **HL7 Clinical Context Management (CCOW) Context Manager Protection Package, Release 1**
This document presents the functional security requirements required in coordinating the component of the CCOW architecture known as the context manager.
- **HL7 Clinical Context Management (CCOW) User Authentication Protection Package, Release 1**
This addresses the need to improve the clinical sign-on process to provide both efficiency and security. It contains functionality specific to programs that are used for authenticating computer system users and which also implement the CCOW Standard for context sharing. More specifically, this protection package describes CCOW-specific functional security requirements required for user authentication by a CCOW-compliant application.

CONCLUSION

The CCOW Standard continues to evolve to address an increasing set of context management capabilities. Leveraging the architectural foundation established at the onset, the standard now embraces multiple technologies, a variety of context subjects, and an increasing array of context management mechanisms. This ongoing evolution will position CCOW as a pivotal resource in the years to come for provider organizations who need an effective way to extend the utility of yesterday's legacy healthcare applications while introducing the breakthrough technologies of tomorrow.

HL7's CDA — Clinical Document Architecture

What Are Clinical Documents?

Clinical documents are the core of a patient's lifetime record. A "History & Physical," "Discharge Summary" or an "X-ray Report" are all examples of clinical documents. Typically, they contain narrative as well as discrete data. While certain structures may apply across document types, like the common SOAP note structure, individual document types vary widely in content.

The HL7 CDA defines clinical documents as having these characteristics:

- Persistence
- Stewardship
- Authentication
- Context
- Wholeness
- Human readability

Why should clinical documents be standardized?

A consistent approach to electronic clinical documents means that the critical information contained in the documents can be used independently of the applications on which they were produced. For example, a discharge summary created by an electronic health record can be rendered on standard browsers and a repository of transcription documents can be indexed with the same metadata as the output of an EHR. Information created today can be migrated to future systems with little or no data conversion. Findings encoded in clinical documents

can be used for third-party decision support and mined at a later date for new applications.

What is CDA?

First published in 2000, the HL7 Clinical Document Architecture (CDA) is a leading standard for the exchange of health-care information and has become a pillar of interoperability for clinical care and public health. CDA Release 2 utilizes a common syntax for all clinical documents. It preserves the integrity and structure of clinical documents. It conveys authenticated content with fidelity and supports discrete data representation that is both extractable and computable.

Clinical Documents are the CORE of a patient's lifetime record. This critical information should be standardized.

CDA Release 2 provides an exchange model for clinical documents (such as discharge summaries and progress notes) - and brings the healthcare industry closer to the realization of an electronic medical record. By leveraging the use of XML, the HL7 Reference Information Model (RIM) and coded vocabularies, the CDA makes documents both machine-readable - so they are easily parsed and processed electronically, and human-readable - so they can be easily retrieved and used by the people who need them. CDA documents can be displayed using XML-aware web browsers or wireless applications such as cell phones.

HL7's CDA — Clinical Document Architecture

Who is using CDA?

There are large scale CDA implementations in North and South America, Europe and Asia Pacific. In the U.S., CDA is being implemented by groups such as NewYork Presbyterian, the Military Health System, the University of Pittsburgh Medical Center, Kaiser Permanente and many others. Groups such as the Healthcare Information Technology Standards Panel (HITSP) and Integrating the Healthcare Enterprise (IHE) are also utilizing CDA in their work. The CDA implementation guide, the Continuity of Care Document (CCD) was recently also selected by the U.S. Office of the National Coordinator for Health Information Technology as part of its initial set of standards, implementation specifications and certification criteria for EHR technology.

CDA Implementation Guides

Several implementation guides based on the CDA have been published or are now available as Draft Standard for Trial Use (DSTU) status. They are as follows:

CDA IG Public Health Case Reporting, Release 1 – Common data elements found in multiple states' reportable condition forms were compiled and standardized in a project initiated in 2007 by the CDC National Center for Public Health Informatics (NCPHI) and Council of State and Territorial Epidemiologists' (CSTE) Case Report Standardization Workgroup (CRSWg) and leveraged in this project by NCPHI. This IG will allow healthcare facilities/providers to communicate these data elements to the state and local public health departments in CDA

format, an interoperable, industry-standard format.

CDA IG Imaging Integration; Basic Imaging Reports in CDA and DICOM, Release 1 – The implementation guide for this informative document was developed by DICOM, with support from the HL7 Imaging Integration Work Group and the Health Story Project. It describes constraints on the CDA header and body elements for Diagnostic Imaging Reports, which contain a consulting specialist's interpretation of image data. It is intended to convey the interpretation to the referring (ordering) physician and become part of the patient's medical record.

CDA IG Care Record Summary, Release 1 – The purpose of this document is to describe constraints on the CDA Header and Body elements for Care Record Summary documents. A Care Record Summary document contains patient's relevant health history for some time period. It is intended for communication between healthcare providers.

CDA IG for Reference Profile for EHR Interoperability, Release 1 – This guide describes characteristics of interoperable EHR records. An EHR record is a persistent artifact which may be independent of the EHR or other system from which it originated. This profile shows how HL7's CDA, Release 2 fulfills requirements of the Common EHR Record Unit, as specified in the HL7 EHR Interoperability Model DSTU. It is the result of an ongoing collaboration between the HL7 EHR, Structured Documents, and Security Work Groups.

HL7's CDA — Clinical Document Architecture

CDA IG for Healthcare Associated Infection (HAI) Reports, Releases 2-4

– This implementation guide was developed in conjunction with the Structured Documents Work Group and the Division of Healthcare Quality Promotion, National Center for Preparedness, Detection, and Control of Infectious Diseases (NCPDCID), Centers for Disease Control and Prevention (CDC). The purpose of this implementation guide is to specify a standard for electronic submission of Healthcare Associated Infection (HAI) reports to the National Healthcare Safety Network (NHSN) of the CDC. It defines the overall approach and method of electronic submission and develops a set of appendices defining specific HAI report types. As reports are modified and new report types are defined, additional appendices will be developed and published by CDC and HL7.

CDA IG for Consultation Notes, Release 1– The implementation guide for this HL7 Draft Standard for Trial Use (DSTU) was developed in conjunction with the Health Story (formerly CDA4CDT) project, which has an Associate Charter Agreement with HL7. The guide reuses templates developed for the HL7 Continuity of Care Document (CCD) and for the History and Physical DSTU and is suitable for consultation notes.

CDA IG for Operative Notes, Release 1– The implementation guide for this HL7 Draft Standard for Trial Use (DSTU) was developed in conjunction with the Health Story project, which has an Associate Charter Agreement with HL7. The guide reuses templates developed for

the HL7 Continuity of Care Document (CCD) and is suitable for any type of operative report.

CDA IG for Quality Reporting Document Architecture (QRDA), Release 1– This HL7 DSTU was supported by the Child Health Corporation of America (CHCA) with participation from the American College of Physicians, American Health Information Management Association (AHIMA), Alliance for Pediatric Quality, Iowa Foundation for Medical Care, The Collaboration of Performance Measure Integration with EHR Systems (“The Collaborative”), HITSP, Integrating the Healthcare Enterprise (IHE) and others. The guide covers patient-centric quality data reporting and lays out a framework for aggregate, population-based quality reports.

CDA IG CDA Framework for Questionnaire Assessments, Release 1– The purpose of this IG is to specify a standard for electronic submission for CDA Questionnaire Assessments that will allow healthcare facilities to communicate reports in an interoperable, industry-standard format.

CDA IG Care Record Summary, Release 2; Discharge Summary, Release 1– The implementation guide for this HL7 Draft Standard for Trial Use (DSTU) was developed in conjunction with the Health Story project. The guide describes constraints on the CDA header and body elements for Discharge Summary documents.

HL7's CDA — Clinical Document Architecture

CDA IG for History & Physical Notes, Release 1 – The implementation guide for this Draft Standard for Trial Use was developed in conjunction with the Health Story Project with support from groups such as Kaiser Permanente, Mayo Clinic, Military Health System and others. This guide defines additional constraints on the CDA Header and Body used in a History and Physical document in the U.S. realm, and provides examples of conforming fragments in the body of the document and an example of a conforming XML instance.

CDA IG for Personal Healthcare Monitoring Reports, Release 1 – This implementation guide was co-developed by Continua Health Alliance, which has a Liaison Agreement with HL7. The guide conforms with the HL7 CCD and describes how to use CCD templates for communicating home health data to an electronic health record.

There are five more CDA Release 2 Implementation Guides currently in the balloting process. They include Genetic Testing Reports, Neonatal Care Report, Consent Directives, Procedure Note and Unstructured Documents.

The HL7/ASTM Continuity of Care Document (CCD)

What is the Continuity of Care Document?

The CCD is a joint effort of HL7 and ASTM to foster interoperability of clinical data to allow physicians to send electronic medical information to other providers without loss of meaning, which will ultimately improve patient care. It passed balloting in February 2007 and was endorsed by the Healthcare Information Technology Standards Panel (HITSP) as the harmonized format for the exchange of clinical information including patient demographics, medications and allergies. It was recognized by the U.S. Department of Health and Human Services as part of HITSP's first set of interoperability standards in January 2008. It was also selected by the U.S. Office of the National Coordinator for Health Information Technology as part of its initial set of standards, implementation specifications and certification criteria for EHR technology.

The CCD is a CDA implementation of ASTM's Continuity of Care Record (CCR). It is intended as an alternate implementation to the one specified in ASTM ADJE2369 for those institutions or organizations committed to implementation of the HL7 Clinical Document Architecture.

The CCD represents a complete implementation of CCR, combining the best of HL7 technologies with the richness of CCR's clinical data representation, and does not disrupt the existing data flows in payer, provider, or pharmacy organizations.

The CCD is an XML-based standard that specifies the structure and encoding of a patient summary clinical document. It provides a "snapshot in time," constraining a summary of the pertinent clinical, demographic, and administrative data for a specific patient. From its

inception, CDA has supported the ability to represent professional society recommendations, national clinical practice guidelines, standardized data sets, etc.

The HITSP Summary Documents Using HL7 Continuity of Care Document (CCD) Component (CS32) describes the document content summarizing a consumer's medical status for the purpose of information exchange. The content may include administrative (e.g., registration, demographics, insurance, etc.) and clinical (problem list, medication list, allergies, test results, etc.) information. This Component defines content in order to promote interoperability between participating systems such as Personal Health Record Systems (PHRs), Electronic Health Record Systems (EHRs), Practice Management Applications and others.

Integrating the Healthcare Enterprise has the CCD for Patient Care Coordination in seven of its profiles. IHE's XDS Medical Summary for referral and discharge is also being built upon HL7's CCD.

Implementation Guides Using CCD

There are currently four CDA, Release 2 implementation guides that make use of CCD templates. These include: CDA for History and Physical Notes, Release 1; CDA for Consultation Notes, Release 1; CDA for Operative Notes, Release 1; CDA for Diagnostic and Imaging Reports, Release 1; CDA for Quality Reporting Document Architecture (QRDA), Release 1; and Personal Healthcare Monitoring Reports. CDA for Procedure Notes is currently under development.

HL7's Version 2.x Messaging

The Standard in Data Exchange for Healthcare

HL7's Version 2.x messaging standard is the workhorse of electronic data exchange in the clinical domain and arguably the most widely implemented standard for healthcare in the world. There have been seven releases of the Version 2.x Standard to date. HL7 Version 2 was also recently selected by the U.S Office of the National Coordinator for Health Information Technology as part of its initial set of standards, implementation specifications and certification criteria for EHR technology. Version 2.5 was also published as an international standard by ISO in June 2009.

The HL7 Version 2.x Standard covers messages that exchange information in the general areas of:

- Patient Demographics
- Patient Charges and Accounting
- Patient Insurance and Guarantor
- Clinical Observations
- Encounters including Registration, Admission, Discharge and Transfer
- Orders for Clinical Service (Tests, Procedures, Pharmacy, Dietary and Supplies)
- Observation Reporting, including Test Results
- The synchronization of Master Files between systems
- Medical Records Document Management
- Scheduling of Patient Appointments and Resources
- Patient Referrals – Specifically messages for primary care referral
- Patient Care and problem-oriented records.

Version 2.6 represents a major revision to Versions 2.5 and 2.5.1, refining and updating existing messages and adding new messages and domains all based upon proposals submitted and accepted by the HL7 membership. Modifications from Version 2.5.1 include:

- The addition of a new segment, UAC – User Authentication Credential, to ALL messages

- The replacement of the TS – Timestamp data type with the DTM – Date/Time data type
- The replacement of the CE – Coded Element data type with either the CNE – Coded with No Exceptions data type or the CWE – Coded with Exceptions data type
- The deprecation of the CNN, NDL, LA1 and LA2 data types
- The inclusion of “external” tables referencing a set of coded values defined and published by another standards organization assigned an HL7 number but without designation as an HL7 table (as was previously the practice)
- The revision of examples in all chapters to support HIPAA compliance

HL7's Version 2.x messaging standard is the workhorse of electronic data exchange in the clinical domain and arguably the most widely implemented standard for healthcare in the world.

The inclusion of a new chapter supporting electronic messaging transactions of claims and reimbursement data (which is produced for implementations of HL7 outside of the United States; in the United States, HIPAA law mandates an already in-use set of implementation guides of X12 messages for these purposes)

- The inclusion of a new chapter supporting electronic messaging transactions of supply chain management data within healthcare facilities

Meanwhile, Version 2.7 is currently being balloted at the normative level. Due to its widespread use, Version 2 will, no doubt, continue to play an integral part in healthcare messaging, even with the HL7 Version 3 Normative Edition.

HL7 is committed to supporting and extending Version 2 in parallel with Version 3, providing continuity for current installations.

HL7's Version 3

Normative Edition —

The Foundation of Healthcare Interoperability

The Health Level Seven Version 3 (V3) Normative Edition—a suite of specifications based on HL7's Reference Information Model (RIM)—provides a single source that allows implementers of V3 specifications to work with the full set of messages, data types, and terminologies needed to build a complete implementation.

The 2009 edition represents the fourth publication of a complete suite of V3 specifications, each of which has received formal approval as either a Normative Standard or a Draft Standard for Trial Use. It includes standards for communications to document and manage the care and treatment of patients in a wide variety of healthcare settings. As such, it is a foundational part of the technologies needed to meet the global challenge of integrating healthcare information, in areas such as patient care and public health.

Throughout the course of V3 development, HL7 has focused on a few salient features that are its hallmarks. In brief, these are:

- A focus on semantic interoperability by specifying that information be presented in a complete clinical context that assures that the sending and receiving systems share the meaning (semantics) of the information being exchanged;
- It's designed for universal application so that the standards can have the broadest possible global impact and yet be adapted to meet local and regional requirements;
- Model-based specifications that provide consistent representation of data laterally across the various HL7 domains of inter-

est and longitudinally over time as new requirements arise and new fields of clinical endeavor are addressed;

- Technology-neutral standards that allow HL7 and the implementers of HL7 standards to take advantage, at any point in time, of the latest and most effective implementation technologies available; and
- It's founded on a development methodology and metamodel that assures consistent development and the ability to store and manipulate the specifications in robust data repositories rather than as word-processing documents.

The Version 3 project represents a new approach to clinical information exchange. It is built from the ground up around a single object model, the RIM, and a rigorous UML-based methodology that ties model to messages and finally to the message's expression in XML syntax.

Version 3, HL7's exciting, new next-generation standard, takes a completely new approach to messaging—a new precision to healthcare standards.

The V3 specification is built around subject domains, for each of which it provides storyboard descriptions, trigger events, interaction designs, domain object models derived from the RIM, hierarchical message descriptors (HMDs) and a prose description of each element. Implementation of these domains further depends upon a non-normative V3 Guide and normative specifications for: data types; the XML implementable technical specifications (ITS) or message wire format; message and control “wrappers”; and transport protocols.

Members logged in to the HL7 website may download V3 standards at: <http://www.hl7.org/implement/standards/index.cfm>

HL7's Reference Information Model (RIM)

The Reference Information Model (RIM) is the cornerstone of the HL7 Version 3 development process. It is the combined consensus view of information from the perspective of the HL7 working group and the global HL7 affiliates. The RIM is the ultimate source from which all HL7 Version 3 protocol specification standards draw their information-related content.

The RIM is a static model of health and health-care information as viewed within the scope of HL7 standards development activities. It is an object model and graphically represents the clinical data (domains) and identifies the life cycle of events that a message or groups of related messages will carry. As a shared model between all the domains and the model from which all domains create their messages, the RIM is essential to HL7's ongoing mission of increasing precision of data. The RIM became an ANSI-approved standard in late 2003 and was published as an International Organization for Standardization (ISO) standard in September 2006.

The RIM provides a static view of the information needs of HL7 Version 3 standards. It includes class and state-machine diagrams and is accompanied by use case models, interaction models, data type models, terminology models, and other types of models to provide a complete view of the requirements and design of HL7 standards. The classes, attributes, state-machines, and relationships in the RIM are used to derive domain-specific information models

that are then transformed through a series of constraining refinement processes to eventually yield a static model of the information content of an HL7 standard.

The HL7 Version 3 standard development process defines the rules governing the derivation of domain information models from the RIM and the refinement of those models into HL7 standard specifications. The rules require that all information structures in derived models be

traceable back to the RIM and that their semantic and related business rules not conflict with those specified in the RIM. Therefore, the RIM is the ultimate source for all information content in HL7 Version 3 standards.

The RIM is used by HL7 affiliates to extend HL7 Version 3 standards to meet local needs. Through a process known as localization, Version 3 standard specifications are extended using the RIM as the source for new information content. This new information is derived from the RIM and refined in the same manner used to create the original specification.

Explicitly representing the connections that exist between the information carried in the fields of HL7 messages, the RIM is essential to HL7's ongoing mission of increasing precision and reducing implementation costs.

** Please note that the Reference Information Model is updated on a quarterly basis, with minimal changes.*

The RIM is the cornerstone of the HL7 Version 3 development process representing the combined consensus view of information from the perspective of the HL7 working group and global HL7 affiliates.

HL7's Common Terminology Services (CTS) Standard

The ANSI-approved HL7 Common Terminology Services (HL7 CTS) standard defines an Application Programming Interface (API) that can be used by HL7 Version 3 software when accessing terminological content. Before proceeding, we need to first state some things that the CTS specification is not designed to do:

- The current version of CTS is not intended to be a complete terminology service. The scope of CTS is restricted to the functionality needed to design, implement and deploy an HL7 Version 3 compliant software package. In much the same spirit as the XML/SGML relationship, the HL7 CTS is meant to represent a proper subset of functionality that may be provided by more sophisticated APIs such as that represented by the OMG TQS specification.
- CTS is not intended to be a general purpose query language. It is intended to specify only the specific services needed in the HL7 implementation.
- CTS does not specify how the service is to be implemented. It is intentionally silent when it comes to service advertising and discovery, establishing and maintaining connections, and the delivery and routing of messages. It is presumed that a CTS implementation will use the underlying architecture that is most appropriate for the given implementation circumstances.

The Health Level Seven (HL7) Version 3 standards are based on the Reference Information Model (RIM), which is flexible and general in structure. Representation of information within this model is dependent on the availability of terminological resources which can be used to populate the properties of the model with appropriate semantic content. Whenever possible, the HL7 Version 3 Standard ref-

erences existing terminological resources instead of attempting to create a new resource within the standard itself.

These external terminological resources can vary considerably in both content and structure. The HL7 standard needs to be able to identify a minimum set of characteristics that any terminological resource must possess if it is to be used in a HL7 messaging environment. One approach to this task would be to specify a common data structure which all terminological resources would have to fit. This approach, however, is not without drawbacks. First, a common data structure would have to represent a 'least common denominator', which could mask more advanced content and functional characteristics that might be particular to a specific terminology. Another drawback is that this approach puts much of the responsibility for maintaining and updating the content on the HL7 standards

body rather than the individual terminology developers.

The Common Terminology Services (CTS) specification was developed as an alternative to a common data structure. Instead of specifying what an external terminology must look like, HL7

has chosen to identify the common functional characteristics that an external terminology must be able to provide. As an example, an HL7 compliant terminology service will need to be able to determine whether a given concept code is valid within the particular resource. Instead of describing a table keyed by the resource identifier and concept code, the CTS specification describes an Application Programming Interface (API) call that takes a resource identifier and concept code as input and returns a true/false value. Each terminology developer is free to implement this API call in whatever way is most appropriate for them.

This document describes a set of API calls that represent the core functionality that will be needed by

This document describes a set of API calls that represent the core functionality that will be needed by basic HL7 Version 3 applications.

HL7's Common Terminology Services (CTS) Standard

CTS Release 2 is now available as an HL7 Draft Standard for Trial Use. This document is the Service Functional Model Draft Standard For Trial Use (DSTU) for the Common Terminology Services 2 specification, which is specified under the Service Development Framework process under the auspices of the Healthcare Services Specification Project (HSSP). One key point to note is that the Service Functional Model (SFM) provides a Service Interface specification, NOT the specification of a Service implementation. This is a critical distinction in terms of Service Oriented Architecture. There could be many different ways of implementing all or part of the functionality to support the behavior described in this specification.

The purpose of an HL7 SFM is to identify and document the functional requirements of services deemed important to healthcare. Accordingly, the CTS R2 service provides a critical component within the larger context of service specifications in that it defines both the expected behaviors of a terminology service and a standardized method of accessing terminology content. This consistent approach to terminology interaction will benefit other business context services by providing a level of terminology interoperability that currently only exists in a limited form.

The goal of the CTS R2 specification stack is to provide a standardized interface for the usage and management of terminologies. Terminologies provide the atomic building blocks of shared semantics, concepts. In a shared semantics environment, CTS R2 provides a modular, common and universally deployable set of behaviors which can be used to deal with a set of terminologies chosen by the users of the service in their deployment environment. The service will contribute to interoperability by supporting an easy access to the foundational elements of shared semantics. It will also foster the authoring of high-quality terminologies via its authoring profile. This goal is realized via the expansion of the original functionality outlined in HL7's CTS Specification. CTS R2 defines the functional requirements of a set of service interfaces to allow the representation, access, and maintenance of terminology content either locally, or across a federa-

tion of terminology service nodes.

The CTS R2 specification strives to expand on the original functionality outlined in HL7's Common Terminology Service specification, specifically looking to:

- Establish the minimal common structural model for terminology behavior independent from any specific terminology implementation or interchange model, and how it is related to meta-data (information about data) and data (the information itself)
- Integrate into CTS R2 the functional coverage outlined in the existing CTS specification.
- Specify both an information and functional model that addresses the relationships and use of terminology, e.g. how value sets are built and queried, and how terminological information is validated.
- Specify the interactions between terminology providers and consumers – how terminology users can submit unambiguous requests for corrections and extensions and how revisions to content are identified, distributed and integrated into running systems.
- Specify how mapping between compatible terminologies and data models is defined, exchanged and revised.
- Specify how logic-based terminologies can be queried about subsumption and inferred relationships
- Engage broad community participation to describe the dimensions of use and purpose for vocabularies and value sets. This aim will attempt to harmonize these efforts

This version of the document includes several changes:

- Updated conceptual model of terminology
- More generic Detailed Functional Models
- Parameters for the Detailed Functional Model have been made consistent

You can download a copy of the DSTU at the following URL: <http://www.hl7.org/dstucumments/>

Claims Attachments...

Not Just for HIPAA Anymore

In response to the Federal mandate under HIPAA, HL7 and ASC X12 collaborated for a number of years to develop standards for claims attachments. After seeing the model for use with claims attachments, the X12 Work Group addressing Pre-certification and Pre-authorization decided that these standards would also benefit their transaction (X12N 278). As a result we now have ability to do attachments for the X12 278 also. *The HL7 standards referenced above are used to convey clinical and other supporting information. The model moves information from provider to payer, either unsolicited (attachment is sent with claim, or 278 for example) or solicited (request is made by payer for the attachment – in our solution it's made by an X12N 277 transaction).*

Even though this initiative started out only as development of claims attachments, many different attachment needs have been expressed to the HL7 Attachments Work Group in the recent past. As time permits, these requests will be addressed.

So far we've seen two very successful pilots of the claims attachments standards (that we know of), both of which moved into production.

The HL7 attachment standards are based on the Clinical Document Architecture (CDA) and have been proposed by the Department of Health and Human Services (HHS) as the standard for claims attachments under HIPAA (Health Insurance

Portability and Accountability Act of 1996). In the HHS proposal, six attachment types developed by HL7 have been put forward for adoption: Clinical Reports; Rehabilitation Services; Laboratory Results; Medications; Ambulance Services; and Emergency Department (ED), but we expect that the ED will be rolled into the Clinical Reports.

HL7's XML Encoding Syntax Standard for Version 2

HL7 Version 2 XML Encoding Syntax was approved by the American National Standards Institute (ANSI) as an American National Standard in June 2003.

This specification, informally known as HL7 V2.xml, defines the Extensible Markup Language (XML) encoding rules for traditional HL7 Version 2 message content. It primarily addresses the expression of HL7 Version 2 messages in XML as an alternative to the traditional "vertical bar" encoding, and describes the underlying rules and principles. XML schema definitions are provided for all Version 2 message types, including the corresponding data type descriptions necessary for this specification. Due to their greater expressiveness, schemas are the preferred way to describe a set of constraints on message instances. Document Type Definitions (DTDs) are also provided as an informative appendix.

THE BENEFITS OF XML

"The XML capability of HL7 V2.xml makes messages web-enabled, provides more widespread acceptance, and allows up-front validation that reduces the chance of error," said Kai U. Heitmann, M.D. of the Institute for Medical Statistics, Informatics and Epidemiology at the University of Cologne, Germany. Heitmann is a past Affiliate Director on the HL7 Board

of Directors; Past Co-Chair of the HL7 XML Work Group; and editor of the Version 2 XML Encoding Syntax standard.

Cross-industry efficiencies are another important benefit of XML-based standards.

"Implementing HL7 V2.xml is also less expensive because its tag-oriented format makes available to users a rich suite of free, off-the-shelf XML tools," Heitmann said.

The traditional HL7 Version 2 Standard, the workhorse of clinical information exchange, is the most implemented standard for healthcare information in the world. It has been used for the exchange and management of clinical and administrative data since March of 1987. It should be noted that though HL7 Version 3 has been released, HL7 Version 2 will still be necessary to continue support for the large base of existing systems that employ it—hence the need for HL7 V2.xml.

"The XML capability of HL7 V2.xml makes messages web-enabled, provides more widespread acceptance, and allows up-front validation that reduces the chance of error."

HL7's Structured Product Labeling (SPL) Standard

The Structured Product Labeling (SPL) specification is a document markup standard that specifies the structure and semantics of the content of authorized published information that accompanies any medicine licensed by a medicines licensing authority. These documents are known as “product label,” “package insert,” “prescribing information,” “product information,” “medicines information,” and under many other names. The precise definition and content of product labeling usually varies depending on the country. For example, in the U.S., all written, printed, or graphic matter accompanying a medicinal product is called “labeling.” For human prescription drugs, the “content of labeling” includes all text tables and figures in the labeling described in 21CFR 201.57. Implementers of this standard should refer to regulations, definitions and guidances applicable for the realm in which the standard will be used.

An SPL document is created by an organization that is required by law to submit product information document because it is responsible for the creation or marketing of a product, or any other person or organization compelled by other motives to submit information about products, whether originally created or not. This includes original manufacturers, repackagers, relabelers, and public agencies or private information publishers that submit product information documents. Recipients of product label documents are any person or organization, including the public at large, or an agent of the public (such as a regulatory authority). The need to create SPL documents is typically governed by legal statutes which set points such as the completion of a new drug application (NDA), the change of product information or annual reports as requiring submission of an SPL document.

This specification includes a detailed description of an information model for structured product labeling documents as well as the XML representation of that model. The information model is based on the HL7 Reference Information Model (RIM) and uses the HL7 Version 3 Data Types.

SPL is based on the HL7 Clinical Document Architecture (CDA), which specifies the structure and semantics of “clinical documents” for the purpose of exchange (see 3.1.1 Relationship of the SPL Specification to CDA). The SPL Schema is defined as an XML entity. An SPL document references the SPL Schema.

This specification includes a detailed description of an information model for structured product labeling documents as well as the XML representation of that model.

THE PURPOSE OF THE SPL

The major purpose of the SPL specification is to facilitate the review, editing, storage, dissemination of, and access to, product labeling document content. It is intended to:

- Facilitate provision of the content of product labeling both electronically and in a human readable format. SPL documents can be exchanged across systems without the need for additional transformation steps.
- Improve dissemination of product labeling (both new product labeling and product labeling updates) to users of product labeling. The ability to provide the most up-to-date product labeling in a timely manner is considered to be critical to improving risk management of regulated products.

HL7's Structured Product Labeling (SPL) Standard

- Facilitate more efficient evaluation of labeling changes by allowing more effective use of computer technology to compare different versions of labeling on a section by section basis.
- Promote more coordinated data collection throughout the regulatory agency and improve processing, storage and archiving capabilities. Reduce or eliminate redundancies in data collection.
- Improve access to information and enhance the ability to query and report on the content of labeling, allowing better support for specific analyses such as sub-population assessments of differences in products based on gender, race, age, and geographic location.
- Improve interoperability of the regulatory agency's systems with other clinical information systems.
- Use standards to improve integration of clinical data.
- Enhance patient safety by helping to provide prescribers and consumers with improved access to information needed to make better risk management decisions in a format that will enhance integration with other technical and clinical applications.
- Support retention of legacy product labeling in databases.

Even though the SPL specification was designed analogous to the HL7 Clinical Document Architecture, there are fundamental differences between the two specifications, for example:

The major purpose of the SPL specification is to facilitate the review, editing, storage, dissemination of, and access to product labeling document content.

- CDA documents involve a Patient - SPL documents do not.
- CDA documents involve one or more Providers - SPL documents do not.
- CDA documents involve an Encounter - SPL documents do not.
- The potential for authentication is subtly different for product labeling documents than for CDA documents. While a product labeling document may be authenticated, and may even have a requirement for legal authentication in some realms, this authentication occurs on the officially approved version of the document rather than on each

minor revision of the document in the process of finalizing it.

Although SPL does not give priority to delivery of patient care in the same way as CDA documents, which are directly associated with patient encounters, the goal of providing timely information about medical products ultimately serves patient care.

The most recent version, Release 4, received ANSI approval in March 2009 and is available for purchase at the HL7 bookstore at www.HL7.org. Release 4 extends the SPL to cover medical devices and to enable the use of SPL in biologics and veterinary product labeling.

HL7's Electronic Health Record System Functional Model (EHR-S FM)

In February 2007, the HL7 Electronic Record Health System Functional Model (EHR-S FM) became the healthcare industry's first ANSI-approved standard that specifies the functional requirements for an electronic health record system (EHR-S). In November 2009, the standard was published by the International Organization for Standardization (ISO), and became the first international standard to specify functional requirements for an EHR system. This version of the standard is referenced by HL7 as EHR-S FM Release 1.1 and by ISO as ISO 10781. The model will enable vendors worldwide to develop, and clinicians to use, EHRs based on one set of functional requirements.

The standard has received broad industry input from more than a thousand clinicians, as well as from EHR vendors, payers, researchers, and others across the industry. The EHR-S FM outlines important features and functions that should be contained in an EHR system. The standard's Functional Model contains approximately 1,000 conformance criteria across 160 functions, including medication history, problem lists, orders, clinical decision support, and those supporting privacy and security. The functions are described from a user perspective and enable consistent expression of EHR system functionality, while the conformance criteria serves as a reference for purchasers of EHR systems and vendors developing EHR software.

The EHR-S Functional Model is versatile, adaptable, and applicable across the continuum of care, supporting key advances in EHR systems. HL7 encourages healthcare stakeholders to participate in the development of profiles to support specific uses across the continuum of care. To date, a number of profiles

have been developed. Many of these profiles have been submitted to the Certification Commission for Health Information Technology (CCHIT) for review. In fact, a number of the criteria from these profiles have been adopted by CCHIT.

In April 2007, the Emergency Care Functional Profile, which can be used to develop, refine and evaluate EHR systems for emergency departments, was registered. The Child Health Functional Profile, which provides critical electronic health record system functions to care for children in the United States, was developed in August 2007, and became an ANSI-approved standard in December 2008.

The Behavioral Health Functional Profile can be used by treatment provider organizations to select or build their own EHR systems; EHR software developers to guide their future product development efforts; certification organizations to certify EHR software; and healthcare payers as part of their criteria for pay-for-performance and other incentives. It also became an ANSI-approved standard in December 2008.

The Long-Term Care Functional Profile reflects the unique mandates and practices of the long-term care setting. This end product is an invaluable tool as LTC providers and IT vendors work to advance technology that enhances: patient safety; quality of care; efficiency; and continuity of care. This profile was approved as an informative standard in September 2008.

The Long-Term Care Functional Profile reflects the unique mandates and practices of the long-term care setting. This end product is an invaluable tool as LTC providers and IT vendors work to advance technology that enhances: patient safety; quality of care; efficiency; and continuity of care. This profile was approved as an informative standard in September 2008.

The EHR-S FM standard will facilitate key advances in electronic health record systems across the continuum of care to enhance quality, safety and efficiency of patient care.

HL7's Electronic Health Record System Functional Model (EHR-S FM)

The Electronic Health Record (EHR)/Clinical Research (CR) Profile provides high-level requirements necessary for using electronic health record data for regulated clinical research, as well as a roadmap for integrating the data collection for both patient care and clinical research ("collect once, re-purpose many times."). This functional profile encourages EHR vendors to incorporate functions into their products so that electronic health records can be a direct data source for clinical studies as appropriate. This profile was approved as an ANSI standard in July 2009.

The Records Management and Evidentiary Support Functional Profile provides functions in an EHR system that can help an organization maintain a legal record for business and disclosure purposes, help reduce a provider's administrative burden, and reduce costs and inefficiencies caused by redundant paper and electronic record keeping.

The development of the Vital Reporting Functional Profile is now underway. The first profile to focus on secondary data use, this profile outlines the functions for collecting data at the point of care that can support the reporting of births and deaths.

A number of countries are already using or plan to use HL7's EHR-S Functional Model, including The Netherlands, Ireland, Japan, Korea and Thailand. For example, in Korea, the Center for Interoperable EHR developed a national EHR functionality standard to be implemented in public hospitals last year based on the EHR-S FM. In addition, the prototype EHR-S developed in 2009 for national university hospitals incorporated the EHR-S FM. Japan also expects that the EHR-S FM can be leveraged as a framework for EHR systems in the country. The General Practice Information Technology (GPIT) Group in Ireland has used the EHR-S FM to certify family practice software.

Work is currently underway on the second release of the EHR-S Functional Model. Recently approved as a project in the Joint Initiative Council, an international SDO collaborative effort, the EHR-S FM, Release 2 will be a truly global standard in scope. The EHR-S FM, Release 2 is expected to open for ballot in late spring or early summer of 2010. Those interested in both the planning and development efforts are welcome to participate. Contact the HL7 office for more information

For those thinking about creating a profile, the *How to Guide for Creating Functional Profiles* is available on the Work Group's Functional Profile webpage (www.HL7.org/ehr). In addition, the HL7 EHR Work Group is available to provide further guidance. Any functional profile that conforms to the EHR System Functional Model standard can be registered with HL7. This registration involves self-attestation of conformance by those submitting the functional profile for registration via a questionnaire that is completed at submission time. Registration can facilitate the adoption of the profile by making it publicly available for use. All registered profiles are available to the public through a searchable registry at:

<http://www.nist.gov/profileregistry>.

For more information on the EHR-S FM and the HL7 Electronic Health Records Work Group, please visit: www.HL7.org/ehr.

HL7's Personal Health Record System Functional Model Draft Standard for Trial Use (PHR-S FM DSTU)

The Personal Health Record System Functional Model (PHR-S FM) Draft Standard for Trial Use (DSTU) is the industry's first draft of a technical standard to specify functionality for PHR systems. It defines the set of functions that may be present in PHR systems, and offers guidelines that facilitate health information among different PHR systems and between PHR and EHR systems.

"The PHR-S FM serves as a general model that can be customized to the specific PHR models already in existence, such as stand-alone, web-based, provider-based, payer-based and employer-based systems," said Donald T. Mon, PhD, vice-president of practice leadership, American Health Information Management Association (AHIMA), HL7 Board member and co-chair of the HL7 EHR Work Group. "While the PHR-S FM is general in scope and was developed with an eye towards what is achievable today, it contains the flexibility necessary for product innovation and sets a vision for future PHR systems."

The HL7 EHR Work Group formed a PHR work group in 2005 in response to the growing aware-

ness that personal health records are a valuable tool consumers can use to help them make informed healthcare decisions. While an abundance of PHR systems exist in today's market, the industry currently lacks a functional standard to which these systems can conform. The creation of a PHR standard is essential because it outlines guidelines for systems to follow, facilitating the

The PHR-S FM is the industry's first standard to specify functionality for PHR systems and offers guidelines that facilitate health information exchange among different PHR systems and between PHR and EHR systems.

exchange of health information among different PHR systems as well as between PHR and EHR systems.

HL7's PHR-S FM has benefited from the input of a broad range of stakeholders. The model was developed by a work group consisting of consumers, providers, health plans, vendors and health information management and information technology professionals. It is critical that a PHR system standard include criteria that are universal across a variety of PHR system models, yet at the same time, be easily adaptable to encourage product innovation.

As a DSTU, the PHR-S FM allows the industry worldwide to work with a stable standard for up to two years while it is being refined into an

HL7's Personal Health Record System Functional Model Draft Standard for Trial Use (PHR-S FM DSTU)

American National Standards Institute-accredited version. During the DSTU period, consumers can begin requesting standards-based functionality when they select PHR systems for their use, vendors can begin incorporating the model's requirements into their products and organizations that certify PHR systems can begin using the model's conformance criteria for certification development and testing purposes. Groups such as the Certification Commission for Healthcare Information Technology (CCHIT) and the Centers for Medicare and Medicaid Services have already begun using components of the PHR-S FM.

The PHR-S FM can be applied to specific PHR models (stand-alone, web-based, provider-based, payer-based, or employer-based models). At the same time, the Functional Model is flexible enough to encourage product innovation. The PHR-S FM was developed with broad stakeholder input, resulting in a well-balanced and versatile functional model that can be applied across the continuum of care. Because the model can be adapted to a variety of care settings, a number of profiles are already under development as subsets of the Functional Model.

PHR-S FM Provides Guidance to Health Authorities & Consumers

Based on the PHR-S FM, the Health Authority-Based PHR System Profile represents an effort to derive the capabilities that are relevant for personal health record systems provided by health authorities. It provides a list of capabilities a health authority such as a county or state public health or behavioral health agency, should consider when selecting or developing a PHR-S. The profile also educates consumers regarding what functions they might consider accessing and using if their health authority offers a PHR, and what functions a health authority should request if it is considering selecting a PHR.

Payer-Linked Profile to Support Health Benefits Plans

The Payer-Linked Profile is aimed at developing an HL7 Informative Functional Profile for personal health record (PHR) systems that are used between payers and their members. The profile provides essential general functions and specific conformance criteria that are important to include in any payer-linked system through which a member might access, store and communicate their health-care information. The model is meant to support all types of health benefits plans including medical, dental, vision, and pharmacy.

HL7 Clinical Genomics Standards

Since its formation, the HL7 Clinical Genomics Work Group (CGWG) has worked to develop HL7 Version 3 standards (and recently Version 2 implementation guides) to enable the exchange of inter-related clinical and personalized genomic data between interested parties. In many cases, the exchange of genomic data occurs between disparate organizations (healthcare providers, genetic labs, research facilities, etc.). Therefore, acceptable standards are crucial for the usefulness of genomic data in healthcare practice and research. We envision that the use of genomic data in healthcare practice and research will become ubiquitous. Today, there are already several examples of the use of genomic data in healthcare and a few of them are presented in detail in our HL7 specifications. The HL7 Version 2 implementation guide is particularly focused in this area, and these LOINC coded terminologies are portable to Version 3 models. The use of genomics data in clinical trials research has also seen a wider adoption especially in the following areas: increased understanding of molecular pathways, cohort identification, biomarker discovery, drug efficacy, drug toxicity, drug metabolism and companion diagnostics.

Within the Version 3 Clinical Genomics Domain, the CGWG first developed the Genotype Topic, which includes the core models of representing genomic data associated with phenotypic data such as clinical observations. The Genotype Topic was approved as a Draft Standard for Trial Use

(DSTU) in 2005. It consisted of various types of genomic data relating to a specific DNA locus, including sequencing, expression and proteomic data. Common bioinformatics markups representing raw data received from genomic facilities are utilized within the Genetic Variation model in a seamless way to the user. This enables the encapsulation of raw data such as full sequencing along with bubbling-up the most clinically significant data to be associated with phenotypic data by decision support applications. Examining and constraining bioinformatics markups is an ongoing process involving collaboration with the bioinformatics communities. The DSTU Topic was deprecated recently while much of its modeling work has been propagated to the HL7 Clinical Genomics Domain Information Model.

As members of the team tested the standard with different types of data, they were able to come up with new ways to optimize the model for the area of genetic variations. This resulted in the creation of an informative Version 2 implementation guide for the clinical environment as well as the creation of normative Genetic Variation Topic and a Genetic Variation CMET constrained for conveying results of genetic variations testing. This CMET passed normative ballot in January 2010.

The Genetic Variation model specifies the structure and semantics for the transmission of information created during single or multiple

gene testing and analysis of a single subject with chromosomal based DNA. The model is further constrained to genetic variation analyses based upon sequence variation and derived from a set of scientific laboratory methods (such as SNP probes, sequencing and genotype chip arrays) that focus on small scale genetic changes, usually in the coding region(s) of one or a small number of genes.

The CGWG also developed the Pedigree Topic, which was approved by ANSI in July 2007. The Pedigree Topic includes the Family History Model, which describes a patient's pedigree with genomic data. This model currently utilizes the Genotype DSTU models (e.g., GeneticLocus) to carry the genomic data for the patient's relatives. Based on the recent approval of the Genetic Variation CMET, the CG team will be modifying the Pedigree model to incorporate this new CMET. The requirement for an elaborated Clinical Genomic Family History Model was identified while working on the Breast Cancer Storyboard. The CGWG recognizes it as key to the utilization of these standards as well as for personalized healthcare in general.

The HL7 Version 3 Pedigree standard is also in the process of becoming an ISO standard. The U.S. Healthcare Information Technology Standards Panel (HITSP) selected the Pedigree standard as the standard for communication between EHR systems and decision support applications.

The domain of family health history is a convergence point of EHR, PHR and Genomics in a way that enables clinical decision support (CDS) applications to run effectively, in particular when it comes to prevention and early detection of hereditary disease. A breakthrough in EHR-PHR

communication of family history data has been recently achieved: the U.S. Surgeon General's web tool for family history, My Family Health Portrait, has adopted the HL7 Version 3 Pedigree specifications, and can communicate with professional tools compliant with the HL7 Pedigree, such as HughesRiskApps developed at Massachusetts General Hospital. Information exchange between those systems (and others) based on the HL7 Pedigree was highlighted during a special HHS meeting that can be viewed here <http://videocast.nih.gov/launch.asp?14803> (Videocast of the Surgeon General Next-Generation Family Health History Tool, November 25, 2008).

The Clinical Genomics team has been working on the development of Gene Expression data exchange standard. The first step in this process was the development of a Clinical Genomics Domain Analysis Model which contained the foundation UML classes for various -OMICS technologies. The Gene Expression area of the domain analysis model was further developed to serve as data requirements for a Gene Expression CMET.

For more information on the HL7 Clinical Genomics standards, please contact Amnon Shabo at shabo@il.ibm.com.

HL7 Service-Oriented Architecture: PRACTICAL GUIDE FOR SOA IN HEALTHCARE

The *Practical Guide for SOA in Healthcare* is an informative document that was produced jointly by the HL7 SOA Work Group in collaboration with the OMG Healthcare Domain Task Force under the auspices of the Healthcare Services Specification Project (HSSP). The Practical Guide was created to help assist organizations that are either considering or are doing projects involving SOA to make effective decisions and to understand how SOA might fit into their organizational landscape. Many people find that it difficult to determine “how to get started” and lack an approach for doing so— a void this document hopes to address.

Using a fictitious organization—SampleHealth—as a backdrop, the Practical Guide discusses a multi-step approach on how to plan, design, implement, deploy, and support a SOA environment. Consciously limited to 50 pages, this document provides a summary overview for understanding the key issues involved with such a program and one approach that has been used successfully in many organizations. Within the document, services such as those specified by HL7 appear as specific examples, demonstrating how these industry standards might be used to promote interoperability and viability within an organizational context.

Why produce an informative reference instead of a standard?

As participants in standardization work, we are cognizant that industry standards ascribe ways to achieve a specific objective and are prescriptive about how to do so. As a group, we felt that it would be overstepping to define how organizations should implement SOA, opting instead to provide guidance that would be useful with the expectation that needs change and this advice would be tailored to meet specific situational requirements.

How does this relate to the OMG Technical Specification and to the Healthcare Services Specification Project?

The Practical Guide was developed in concert with several of the HSSP specifications and, in fact, cites many of these standards as examples throughout the document. Note that HSSP standards are both HL7 standards and often OMG standards as well. One of the principal reasons for authoring the document was in part to address the numerous queries received about these standards and how they were intended to be used. By design, HSSP standards tend to be fairly generic and abstract, maximizing their flexibility for use in multiple situations. Unfortunately, that very design benefit makes

HL7 Service-Oriented Architecture: PRACTICAL GUIDE FOR SOA IN HEALTHCARE

them harder to understand and less self-evident on how they would be used and useful. The Practical Guide attempts to resolve these ambiguities by placing the abstract specifications into a real-world context.

How does this relate to the HL7 Services-Aware Enterprise Architecture Framework?

In 2008, HL7 embarked upon developing its Services-Aware Enterprise Architecture Framework (SAEAF). SAEAF will be aligning the broad range of HL7 specifications – including services, messages, and document standards – to a consistent approach across the whole of HL7's offerings. We anticipate that some elements of the SAEAF will influence the steps outlined in the Practical Guide. Those portions of the SAEAF that are appropriate to the target audience of this document will be incorporated into future revisions of the document.

HL7 Service-Oriented Architecture: IDENTITY CROSS-REFERENCE SERVICE FUNCTIONAL MODEL (IXS SFM)

The *Identity Cross-Reference Service Functional Model (IXS SFM)* – formerly known as the Entity Identification Service (EIS) – was adopted as a normative HL7 specification in late 2009. It defines the functions, responsibilities, inputs, outputs, and expected behavior of a system component for managing identities, such as would be used in a Master Patient Index (MPI). Not limited to use for Patients, the EIS SFM can be equally applied to manage identities for staff, providers, facilities, or any other “entities” needing identity management.

Why produce an industry standard Identification Service?

Quite simply, the Identity Cross-Reference Service defines the collective set of behaviors that one would expect system components such as an MPI to perform. This functionality is required of most healthcare organizations and supported by many vendors and products. The challenge is that each one does things a bit differently, and making them work together becomes exponentially complex. Without clearly defined expectations of what falls within the responsibility of the Identification Service, variants abound and interoperability suffers.

Don't industry standard services (such as the Entity Identification Service) limit vendor competition?

Not necessarily. While industry standards specify how a consumer interacts with the service, these specifications have expressly left how these functions are supported out-of-scope. In other words, there is no predetermined matching algorithm, system design platform, or approach that is advocated in the standard. Vendors are able to compete based upon quality-of-service and the benefits of their specific implementation. Further, the specification includes mandatory and supplemental requirements (“nice-to-haves”), which can further stratify marketplace offerings.

How does this relate to the OMG Technical Specification and to the Healthcare Services Specification Project?

For IXS, HL7 has partnered with OMG to produce technical specifications supporting Service Functional Models within OMG's technology adoption process. To the healthcare industry, this provides value via OMG's rigorous architectural approach, review process, and distributed systems expertise. For OMG, this relationship brings deep healthcare experience via HL7's extensive international participation and vertical industry expertise.

HL7 Service-Oriented Architecture: IDENTITY CROSS-REFERENCE SERVICE FUNCTIONAL MODEL (IXS SFM)

The Healthcare Services Specification Project (HSSP) is a moniker for this collaboration. The technical specification supporting IXS continues to bear its former name – the Entity Identification Service – and is freely available from <http://www.omg.org/spec/EIS>.

How does this relate to the HL7 Services-Aware Enterprise Architecture Framework?

In 2008, HL7 embarked upon developing its Services-Aware Enterprise Architecture Framework (SAEAF). SAEAF will be aligning the broad range of HL7 specifications – including services, messages, and document standards – to a consistent approach across the whole of HL7's offerings. SAEAF also specifies usage constraints for deployed HL7 standards. Several HSSP services predate SAEAF. Over time, we expect all HL7 Service Functional Models to come into alignment with SAEAF.

HL7 Service-Oriented Architecture: RETRIEVE, LOCATE, UPDATE SERVICE FUNCTIONAL MODEL (RLUS SFM)

The *Retrieve, Locate, Update Service Functional Model (RLUS SFM)* Draft Standard for Trial Use (DSTU) is a draft functional standard defining the functions, responsibilities, inputs, outputs, and expected behavior of a system component capable of querying information and returning data and metadata between systems. Although fairly abstract in nature, RLUS is actually very simple to understand – it provides a generic query and retrieval mechanism that can be used for a multitude of information content via a standard access mechanism, promoting consistency within a heterogeneous environment.

For example, an organization may have a benefits/enrollment system, an electronic health record system, and a personal health record system all within the organization. Integrating information from among these systems can be complex. RLUS could be used to further this integration by building an RLUS-compatible interface into each of the above systems, and making a distributed RLUS call to retrieve pertinent information for a specific patient. The RLUS specification standardizes how disparate information types are managed and aggregated into a single result. Further, since RLUS provides a mechanism for new, richly structured information content to be supported, integration based upon RLUS allows for new systems to come online and

integrate into the organization's infrastructure. RLUS specifies a collection of behaviors needed to manage this inquiry, query, and retrieval of content. The "location" function allows for the return of candidate information, indicating the availability of matching records without actually returning their instances (*for example, does the system have any positive zzzz lab results for patient X*). Other interface behaviors specify how the targeted information can be retrieved or updated.

What is the benefit to using RLUS?

Integration of new software packages in a heterogeneous environment creates an exponential integration problem. Each new system-to-system interface creates tremendous integration burden in terms of creating messages, mapping data fields, and potentially transforming data for use. Interface engines have a role to play here, but do not directly support complex queries among participating systems. RLUS allows for these complexities to reside within the systems that hold the data, shielding requestors from unnecessary detail. RLUS simplifies the ask-answer pattern, providing both rigor and clarity while supporting rigorous information content.

HL7 Service-Oriented Architecture: RETRIEVE, LOCATE, UPDATE SERVICE FUNCTIONAL MODEL (RLUS SFM)

How does this relate to the OMG Technical Specification and to the Healthcare Services Specification Project?

For EIS, HL7 has partnered with OMG to produce technical specifications supporting Service Functional Models within OMG's technology adoption process. To the healthcare industry, this provides value via OMG's rigorous architectural approach, review process, and distributed systems expertise. For OMG, this relationship brings deep healthcare experience via HL7's extensive international participation and vertical industry expertise. The Healthcare Services Specification Project (HSSP) is a moniker for the collaboration between groups collaborating on service standards, of which HL7 and OMG are participants. The RLUS technical specification is freely available from <http://www.omg.org/spec/RLUS>.

How does this relate to the HL7 Services-Aware Enterprise Architecture Framework?

In 2008, HL7 embarked upon developing its Services-Aware Enterprise Architecture Framework (SAEAF). SAEAF will be aligning the broad range of HL7 specifications – including services, messages, and document standards – to a consistent approach across the whole of HL7's offerings. SAEAF also specifies usage constraints

for deployed HL7 standards. Several HSSP services predate SAEAF. Over time, we expect all HL7 Service Functional Models to come into alignment with SAEAF.

HL7 Service-Oriented Architecture: DECISION SUPPORT SERVICE FUNCTIONAL MODEL (EIS SFM)

The *Decision Support Service Functional Model (DSS SFM)* Draft Standard for Trial Use (DSTU) is a draft functional standard defining the functions, responsibilities, inputs, outputs, and expected behavior of a system component for evaluating patient data to reach patient-specific conclusions. A DSS, for example, can evaluate a patient's health summary as encoded in a Continuity of Care Document (CCD) and provide structured recommendations regarding the patient's health maintenance and chronic disease management needs.

Why produce an industry standard Decision Support Service?

Quite simply, the Decision Support Service defines the collective set of behaviors that one would expect a clinical decision support engine to perform. This functionality is required of most healthcare organizations and allows for the data collected in electronic health records and other clinical information systems to provide enhanced value for patients, clinicians, healthcare providers, and payors. The challenge is that the lack of a standard makes the use of decision support services more costly and difficult. Without clearly defined expectations of how Decision Support Services should interface with health information systems, variants abound and interoperability suffers.

Don't industry standard services (such as the

Decision Support Service) limit vendor competition?

Not necessarily. While industry standards specify how a consumer interacts with the service, these specifications have expressly left how these functions are supported out-of-scope. In other words, there is no predetermined knowledge modeling formalism, system design platform, or approach that is advocated in the standard. Vendors are able to compete based upon quality-of-service and the benefits of their specific implementation. Further, the specification includes mandatory and supplemental requirements ("nice-to-haves"), which can further stratify marketplace offerings.

How does this relate to the OMG Technical Specification and to the Healthcare Services Specification Project (HSSP)?

For DSS, HL7 has partnered with the OMG to produce technical specifications supporting Service Functional Models within OMG's technology adoption process. To the healthcare industry, this provides value via OMG's rigorous architectural approach, review process, and distributed systems expertise. For OMG, this relationship brings deep healthcare experience via HL7's extensive international participation and vertical industry expertise. The [Clinical] Decision Support Service technical specification

HL7 Service-Oriented Architecture: DECISION SUPPORT SERVICE FUNCTIONAL MODEL (EIS SFM)

is freely available from <http://www.omg.org/spec/CDSS>. HSSP is a moniker for the collaboration activity between the groups.

How does this relate to the HL7 Services-Aware Enterprise Architecture Framework?

In 2008, HL7 embarked upon developing its Services-Aware Enterprise Architecture Framework (SAEAF). SAEAF will be aligning the broad range of HL7 specifications – including services, messages, and document standards – to a consistent approach across the whole of HL7's offerings. SAEAF also specifies usage constraints for deployed HL7 standards. Several HSSP services predate SAEAF. Over time, we expect all HL7 Service Functional Models to come into alignment with SAEAF.