

Health Information Standards Manifesto

Keywords: EHR, messages, software engineering, health

Author: Thomas Beale

Revision:2.5

Pages: 46

Copyright © 2001

Thomas Beale

email: thomas@deephought.com.au **web:** www.deephought.com.au

Introduction

Health information systems today suffer from a number of significant problems. Challenges which need to be met by the systems of tomorrow include:

- Support for a life-long health record
- True interoperability among all parties and systems used in patient care
- Intelligent decision support
- Domain size and rate of change
- Systems obsolescence
- Multi-contact healthcare system and mobile patients
- Multiple medical cultures, including developing world, asian
- Support for domain experts to have direct control over the information design and change management of their systems

Current work in health standards, notably by HL7, CEN, ISO and the OMG attempts to address some of these problems, as does implementation-based work including a number of publicly-funded efforts in the UK, Australia and the US, university work in the UK, US, Holland and Germany, and numerous open-source initiatives.

This document discusses the challenges for health information systems of the (immediate) future, and offers some suggestions for how the work of both standards bodies and implementation efforts might be brought together in the form of global standards for health information systems, particularly EHR (electronic health record; note that in this document, the term EHR is used to mean all variants, e.g. CPR, EMR and so on).

A key aspect of the discussion here is that there is a need for an independent set of criteria for judging standards and other work efforts in the health informatics arena. Currently, standards in health informatics tend to be judged in terms of themselves, against particular local requirements, or against each other. Whilst all of these approaches throw up useful facts, they are not in themselves sufficient. What is needed is an objective, independent set of criteria by which any standard or proposal can be judged.

It is the purpose of this paper to establish a general basis for thinking about the health information environment and users, and to describe such criteria, in order to make better progress with standards for health information systems. The intention is not to describe requirements for health records or health information systems *per se*; such requirements are to be found in the standards themselves, national health strategies, and systems development projects.

The sections are organised so that some background is provided first, before proceeding to a description of critical attributes of “good” health information systems.

The Health Information Environment

Overview

As a first step we need a picture of the things health interoperability standards seek to standardise. These are not information models, message structures, or use cases on paper, but the information and interactions of real systems. By extension, such standards say something about the nature of the systems themselves, since no system is independent of its information or messages.

The criteria for judging standards which we want to develop must necessarily relate directly to the *IT costs* and *quality of service* of these systems, if they are to materially improve things for the users in the domain.

FIGURE 1 is an attempt to illustrate a notional health information environment. It can also be understood more broadly as a “landscape” of issue areas or focal points in the area of electronic health information. Most of these areas have one or more standards currently available. One of the big challenges in an integrated health information environment is to make these work together. (A primary inspiration for the separation of interests shown in this document has been Corbamed [19.], now known as OMG HDTF - Health Domain Taskforce; the HDTF standards specify interfaces for most of areas on the diagram).

The diagram should be read from the inside out, starting at the level of a “minimally functional” EHR environment, in which some basic level of patient health information is available, along with terminology, reference data (e.g. drug data), patient identification and clinical models (an addition to the health information environment which this paper proposes as being a necessary element in effective HIS environments of tomorrow).

The next level out, “fully functional”, contains other services which would be expected in a fuller environment, such as at a hospital, including decision support, guidelines and protocols, and mobile computing. Note in particular, that at this level, the notion of the EHR has been extended to include events, workflow, multimedia and genetic information.

The “provider” level includes further services typical within provider organisations seen as economic entities, and also as cooperative institutions in a larger network of public or private health information facilities.

Security and access control services appear across all levels of the provider, indicating that they provide a level of support appropriate to the other services available at each level.

The same general model of a provider organisation applies to both secondary care institutions and community-based general practice and allied health, although clearly not all services will be needed in the community or primary care contexts. However, it should be remembered that, depending on where the computing infrastructure lies, there is no reason why the most sophisticated decision support or care pathway planning would not be visible by GPs; all this would require is a system in which primary and secondary healthcare institutions in each community share computing infrastructure at some level. There are no technical impediments to such a scenario, only cultural and political ones.

Details of the services are discussed under the following headings.

[Note that not all possible links have been drawn in the diagram - in particular, services like security and access control have only been connected to a couple of other entities, when in

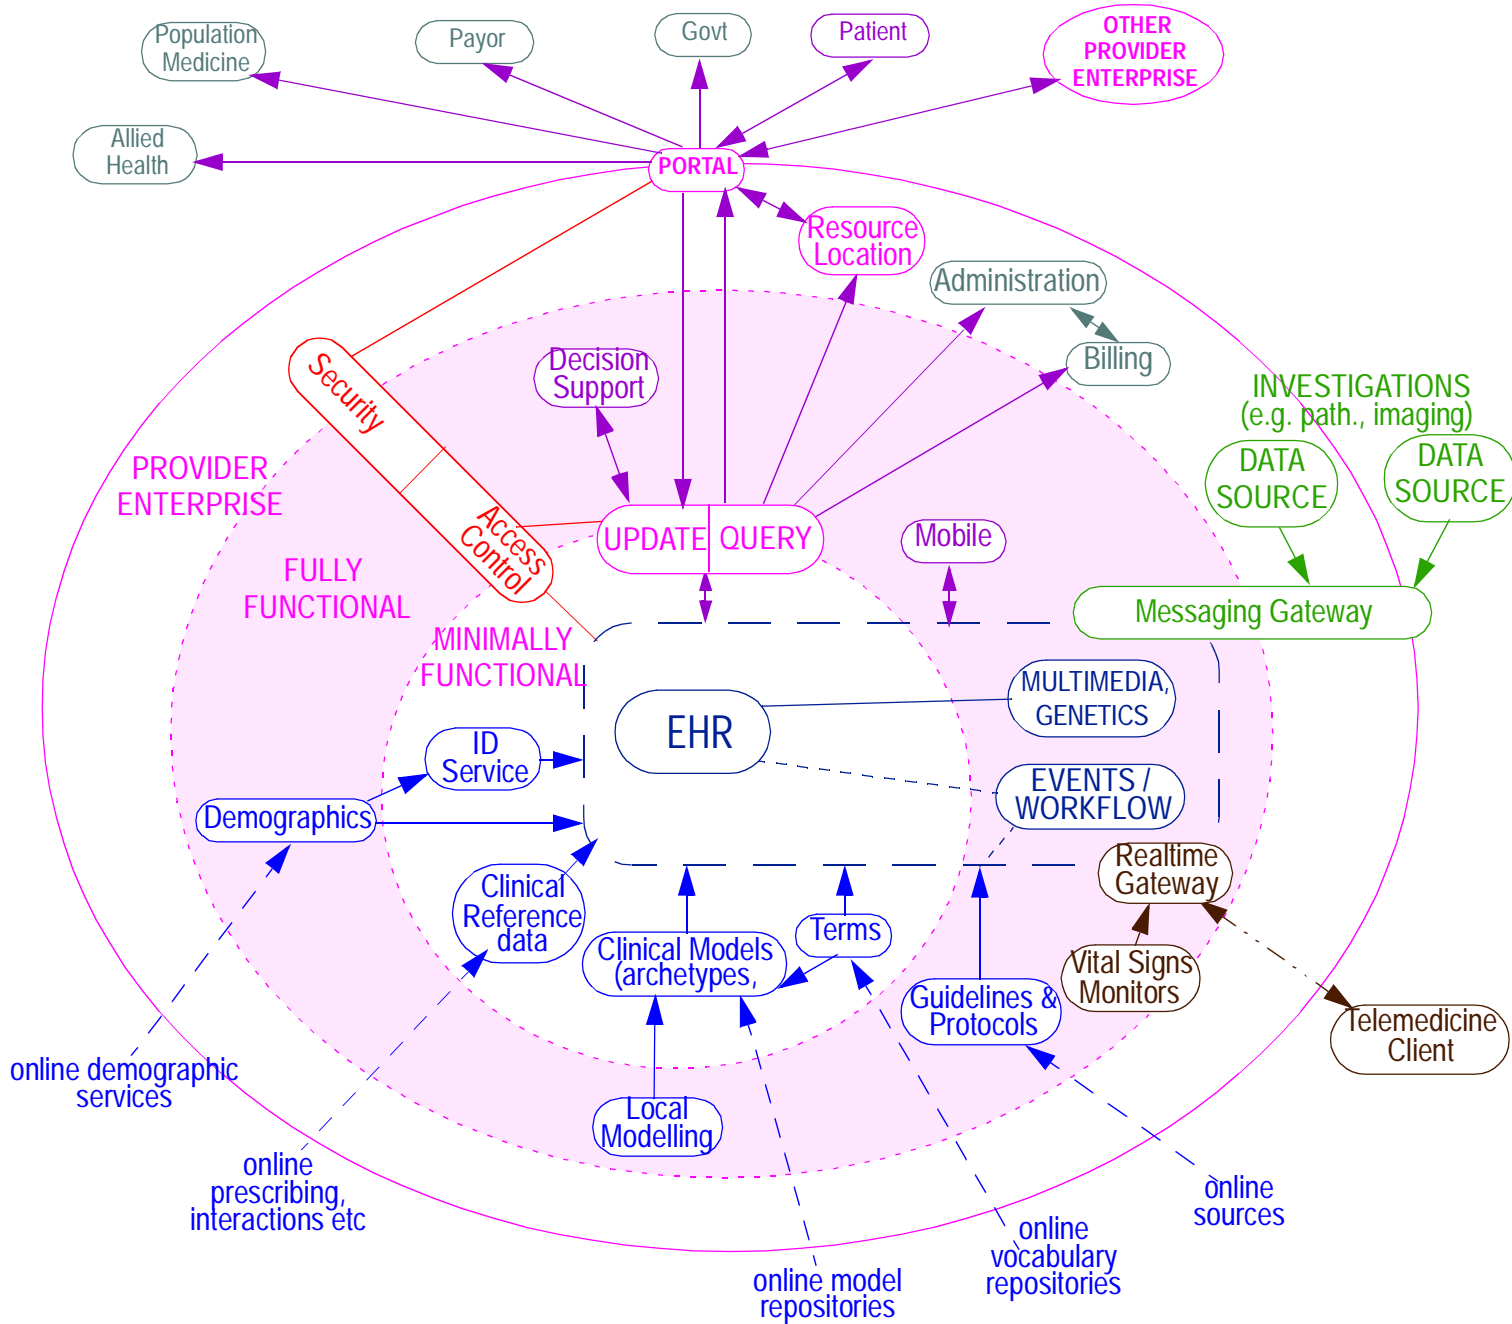


FIGURE 1 A Health Information Environment

fact they tend to be ubiquitous. This kind of diagram is thus designed as a high-level basis for discussion, not an accurate software engineering diagram.]

Infrastructure Services

Security & Access Control

Security services enable demographic entities to be *authenticated* and *authorised* at all levels in the health computing environment. The access control service defines the mapping between user stereotypes and available functions in the environment. Both content and services are profiled in security definitions which are informed by *user consent* specifications.

A Minimally Functional EHR Environment

The EHR

In the ideal world of FIGURE 1, the electronic health record (EHR) is a repository of patient-centred records containing contributions by carers in the course of a lifelong care process. Carers may be from within or outside the provider institution. In today's real world, the EHR is of course a much messier affair, with non-standardised fragments and partial records to be found in most provider institutions. The difficulties for clinical care and administration (and not infrequently, significant costs) encountered routinely today only serve to underline the importance of the goal of a patient-centred, standardised health record for the future.

Party Identification Service

A minimal identification service is needed in any health information environment, containing sufficient demographic information on patients, health care professionals, and health care institutions, to *disambiguate* parties, including if extracts from the EHR are sent to other sites. This is distinct from any more sophisticated demographics service in the environment.

Basic identification is often provided by an LDAP service, or may be an integrated part of a more sophisticated demographic service, such as provided by PIDS (the OMG party identification service).

Terminology

Controlled vocabularies are a key element in structured health data, enabling proper automatic processing such as decision support to be performed. Terms from multiple vocabularies are typically made available by a single service in the HIS environment, and are used ubiquitously in the EHR, by decision support and by clinical models.

The OMG HDTF TQS (terminology query service) is probably the best known public specification for interfacing with terminology.

Reference Data

Basic data in health care includes prescribing data and drug/drug, drug/disease, drug/allergy (etc) interactions. This data is essential for prescribing in GP or hospital situations, and decision support regarding suitability of medications or therapies. Actual prescriptions for drugs or therapies are recorded in the EHR.

Most computer systems supporting prescription, including general practice, offer at least some level of reference data, namely drug descriptions and at least first order interactions (multi-drug and condition interactions are not widely documented).

Clinical Models

Formal models of clinical concepts are the next revolution in health information systems. Research by GEHR Australia and University College London show that health information systems based on formal domain models known as “archetypes” enables significant improvements in both the clinical quality of systems, as well as their economics. The HL7 v3 standard is also leaning in a similar direction, with discrete models of domain concepts as the basis for message definitions.

Such models may be authored in the enterprise, or come from online libraries. The GEHR and SynEx system both include archetype repositories.

In the future, it is likely (or at least hopeful!) that both terminology, clinical models, and clinical protocols and guidelines are developed and standardised within a collaborative framework. Compatibility between the designs of terms, archetypes, and process-oriented guidelines will be essential for correct functioning of sophisticated health information environments.

Query/Update Service

This is the primary interface to the EHR used by most applications and users. The use of terminology and archetypes enables intelligent queries to be formulated.

The Fully Functional Environment

Workflow Management

In modern health systems, each episode of care can be delivered by multiple professionals. For example, a patient goes to a GP, is referred to a diabetologist who orders various tests, interprets the tests, makes a diagnosis, formulates a care plan, and updates the GP on the patient’s situation. Such activities can take days or weeks, and constitute a workflow, i.e. a network of actions by professionals or providers, linked in time. In order for them to progress smoothly, someone has to manage the process. Today, this is typically an *ad hoc* process managed by medical secretaries, various letters of referral, and the patient him/herself.

Workflow systems can support such processes, by executing agreed models of care, and reminding the relevant parties of events to be performed.

Major workflow events cause updates to the EHR, while the numerous minor events which occur in most workflows will usually only be recorded in their primary database, since they are not of interest in the patient care process.

Event Management

Events and orders constitute a large proportion of the actions which occur in a health system. The EHR includes a record of the major actions, such as prescriptions, requests for pathology, and surgery. The many intermediate steps which occur in some of these actions (e.g. the more complex pathology tests, surgery) are not of interest in the EHR, but are essential for order management. Many orders take place in the context of well-known workflows.

Multimedia & Genetics

Test results in imaging (e.g. x-ray, ultrasound, MRI, nuclear scans), electrophysiology (e.g. ECG, EEG), and genetics data constitute the most detailed, and generally most bulky information in the health environment. Such data are logically integrated with the EHR, but usually served from specialist computers with special storage devices designed to satisfy the retrieval

patterns associated with large data; specific items or sections, usually of diagnostic significance, may be added to the EHR.

Demographics / Party Identification

Most providers record more demographic information than needed by the simple ID service in a basic environment. Such information is usually for practice management, billing and administration, but may also be clinically relevant, e.g. attributes such as ethnicity. Most national and/or regional health bodies have a “minimum data set” which includes a large number of demographic details for patients and health professionals.

PIDS (the OMG party identification service) is possibly the most appropriate service for demographic information and identification in health environments today.

Guidelines, Protocols and Care Pathways

In the era of evidence-based medicine, guidelines and protocols are the decision-making and planning tools of the clinician. Their use needs to be recorded in the health record, so that other users can understand the reasons for decisions, and so that medico-legal investigations can be effective.

Guidelines and protocols may in the future be managed in a “care pathway” framework, i.e. a time-based care planning paradigm for patients. If this is the case, the models describing this area may end up resembling the kinds of models used in project management systems, and will form the framework for workflow events.

Work in the guideline area includes the Arden Syntax [16.], the GLIF language [22.], Asbru [17.], and ProForma [26.].

Decision Support

Decision-support systems require disciplined patient data in order to function. They rely heavily on coded terms, and are likely to rely equally heavily in the future on formal clinical models in order to navigate data intelligently.

In secondary care contexts of the future, decision support may be the primary reader of EHR data.

Clinical Modelling

Clinical models may be locally authored in more advanced provider institutions, typically those attached to universities or medical colleges. Such organisations are likely to be a source for clinical models for other providers to use.

Mobile Computing

Hand-held computing devices are becoming more common in hospitals and allied health. The EHR needs to be available for viewing and potentially update via relatively low-bandwidth interfaces, which may also need to support asynchronous transfers.

The Provider Enterprise

Administration & Billing

In the provider enterprise as a business, administration and billing are the main non-clinical users of core health data, using it to manage beds, operating rooms and other resources, and to communicate relevant financial documents between patients, other providers and health insurers or the government.

Resource Location

This service exists at the enterprise level in order to respond to distributed queries in the open environment. It enables requestors to know what information is available inside the provider environment.

Portal

The interface visible at the enterprise level has to be quite varied, since it needs to satisfy the needs not just of clinical users, but of population queries, government, insurance organisations, and patients themselves.

While patient-accessible health records do not exist yet in production, a number of prototypes have been built indicating likely directions. Patients would be able to control *access* to their records, via an “eConsent” mechanism; they would be able to enter their own data (particularly useful for self-monitoring chronic patients such as diabetics), and they would be able to perform basic queries and summaries of their health record. Guardians and parents having legal consent would be able to see the records of their wards.

Data Sources - Investigations

Raw data comes to the EHR from numerous sources including pathology laboratories, electrophysiology clinics, but may also come from sources in the home such as wearable monitors. The common aspect of these sources is that the computing systems are not generally based on sophisticated or standard information models, they are often dedicated to just a narrow kind of information. For this reason, data is likely to be formatted using a messaging standard such as HL7 or EDIFACT, which guarantees uniformity for the information receiver, usually an EHR system.

Outside the Provider Enterprise

Parties outside the provider enterprise include the patient, the government, insurance organisations, other providers and research users, including statisticians, epidemiologists and educators. All of these interface with the services provided inside the enterprise via a secure gateway, typically implemented as a web portal interface, and for the most part, read-only (the exception will increasingly be the patient).

A patient’s health record information may reside in more than one provider environment, and is logically assembled via distributed web clients which make use of the resource location facility to determine where data exists for a given patient.

Online Resources

There are numerous online resources which are beginning to appear, which are or will be used operationally by provider EHR environments. These currently include terminologies, clinical guidelines, drug and interaction data, and will in the future include “public” demographics, and libraries of domain models such as archetypes. In the future it will be essential for provider organisations to be able to make use of such resources, in order to stay up to date with the domain in a cost-effective way.

Domain terminologies and archetypes particularly require sharing, since they are too involved for each institution to develop single-handedly.

The Clinical View

The view of things shown in FIGURE 1 is a systems one, and corresponds to a useful thematic separation of standards. However, from the point of view of clinicians, healthcare is about the management of patient problems and issues in time, and being able to share information in a community care environment.

The key aspects of the clinical viewpoint (albeit one not necessarily promoted by all health professionals) are:

- Patient-centered health records
- Community based care, where patient records are shared by many providers, including GPs, allied health professionals, social workers, the local hospital, emergency services, pathology laboratories and pharmacies.
- A care pathway approach to managing patient health. Care pathways treat problems and issues as threads in a time-based view of patient health events, applying a project management mentality to coordinating resources and tasks.

By way of illustration, FIGURE 2 shows how a care pathway view of the patient might appear on a clinician screen. While this might seem more related to interface design than underlying standards, it is essential that the clinician view of information and process be understood by standards developers, so that standards and systems properly support the integration of all information to a patient-centred EHR.

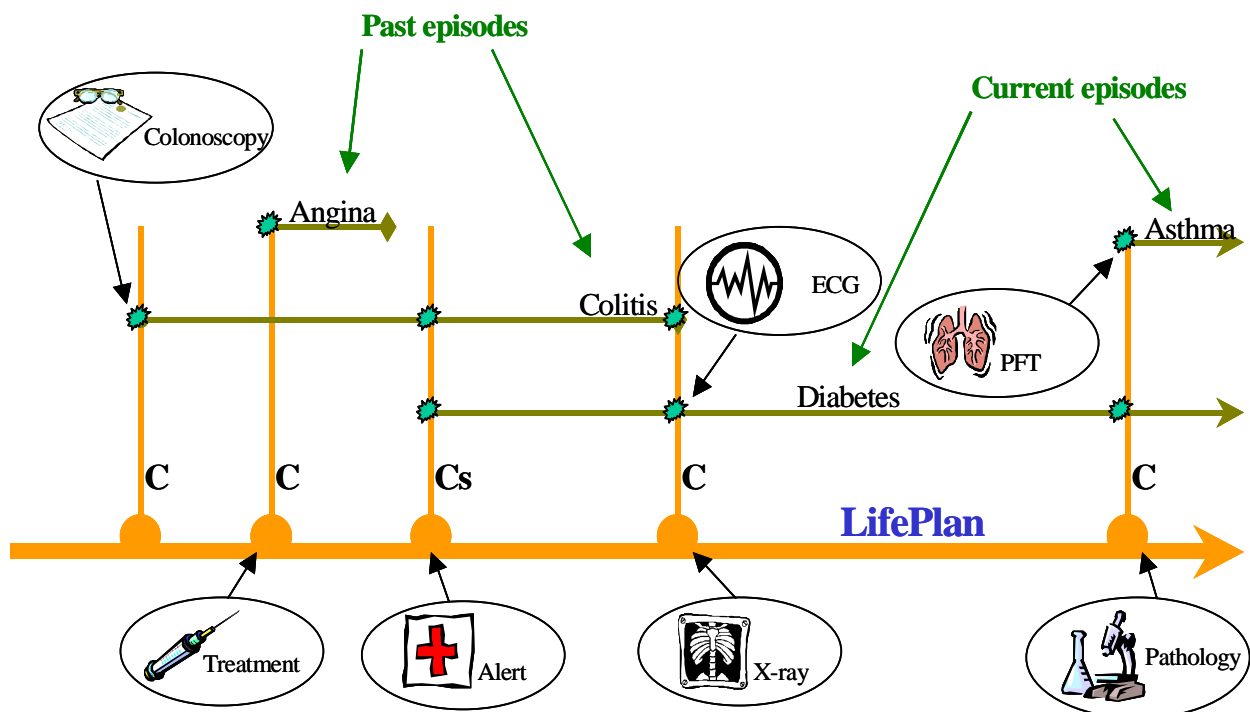


FIGURE 2 Care Pathway
(with permission from P Ameline)

A number of requirements are implied by the care pathway approach, for example:

- Support for multiple, asynchronous users of shared health record information.
- An integrated access control model for multiple healthcare professionals

- Clear patient consent and privacy mechanisms
- Support for state-based care processes
- Support for workflow event management in the EHR, within the framework of clinical guideline and care pathway models.

The EHR

Definition

This sometimes confusing term has been a stumbling block for some people in understanding what EHR systems are about. The definition used by the Australian Electronic Health Records Taskforce (see [11.]) is:

An electronic longitudinal collection of personal health information usually based on the individual, entered or accepted by health care providers, which can be distributed over a number of sites or aggregated at a particular source. The information is organised primarily to support continuing, efficient, and quality health care. The record is under the control of the consumer and is stored and transmitted securely.

This could be summarised more succinctly as follows:

EHR: a longitudinal record of information concerning a subject of care, for use by carers.

A few things to note about this definition:

- In health, the subject of care is normally a person (the “patient”), but might also be a family or group (in some cultures, such as the Maori culture in NZ, the family is the logical unit of health care) or an animal (veterinary health). Outside of health, it might be something entirely different, such as a record of interactions with a company, or a history of care events in a managed forest.
- Patient-centred records are a more desirable paradigm than organisation-centered records, since the reality today is a multi-contact health care system, where patient data is created at multiple sites.

What Should Be In It?

The Core

Basic data in any EHR include:

- “Persistent” information - data which remains valid or nearly so for the life of the record, including:
 - Patient identification
 - Past medical history
 - Problem list
 - Current medications and therapies
 - Therapeutic precautions (including allergies, drug intolerances etc)
 - Family history
 - Social history
 - Vaccination record
- Test results.
- Contributions by healthcare professionals, such as histories, examinations, diagnoses, orders, care plans and summaries.

A type of data of primary importance in the EHR, but not well understood in naive analyses is “context”. Each item of knowledge in the EHR is *contextualised* around the time, place, clinician and patient particularities during care. This means that every item that is added to the

record includes context-related attributes, i.e. the who/what/when/where/why of each context. Models of context have been described in [4.] and [14.]. A good model of context is essential to support not only clinical care, but medico-legal and research investigations. A model based on these and other sources is described in [6.] and illustrated in FIGURE 3.

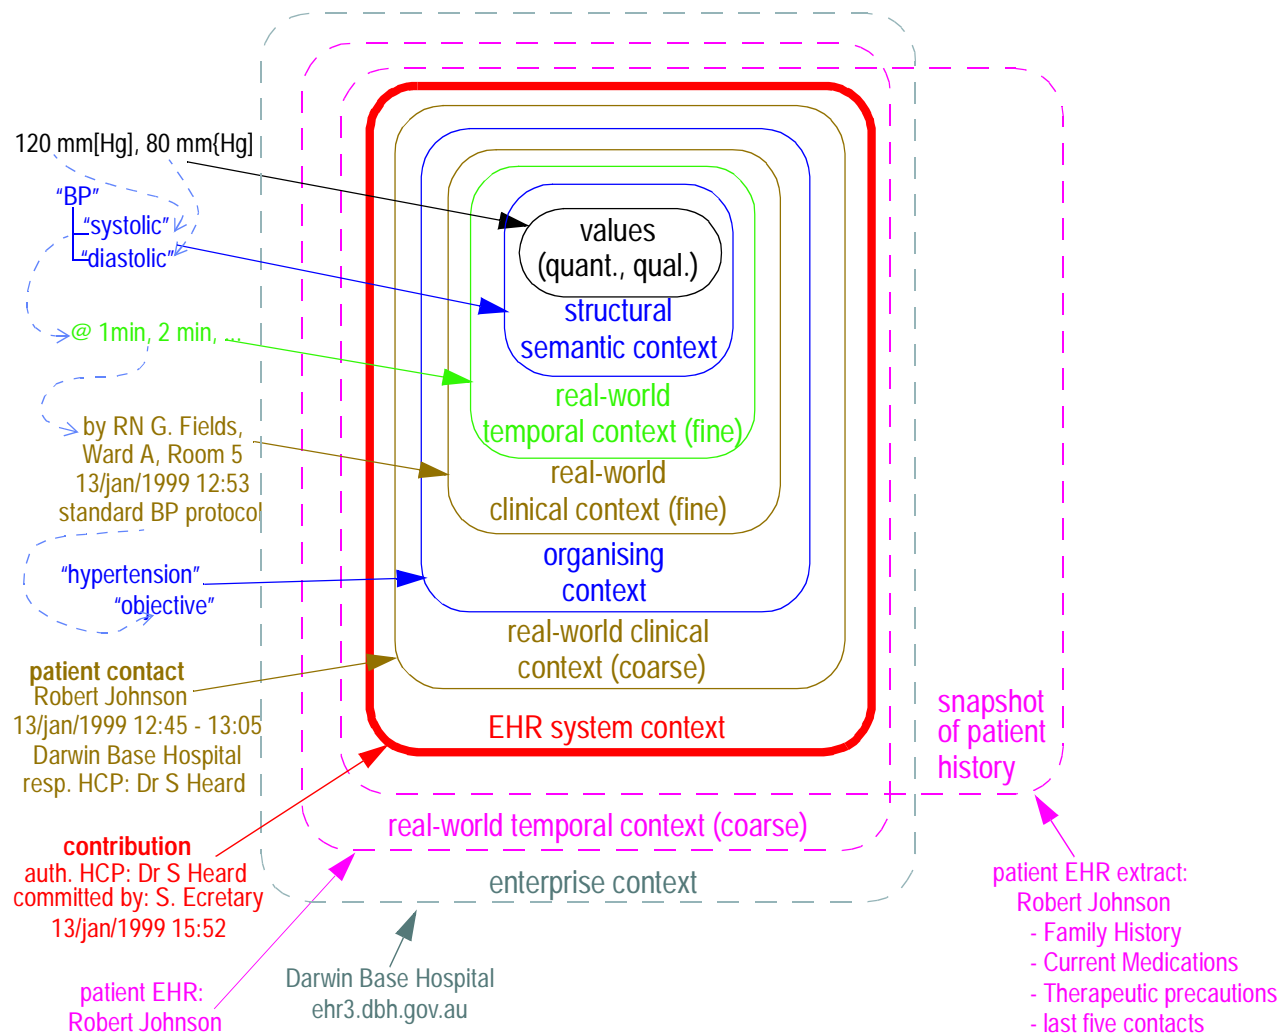


FIGURE 3 Context Architecture of the EHR

Complex Data

Examples of information which need to exist in some form in the record, but not necessarily in their full detail include:

- Orders and pathology. Most doctors want to see a request for a test, and the results included in the EHR. However, they don't typically care about the resources required by the laboratory for the test, or how the test samples were transported to the laboratory. Over time such data would become voluminous in the record, and serve no useful clinical purpose.
- Vital signs data. In general, vital signs monitors generate too much data over a hospital stay to consider including the whole lot in a patient's record. However, some data may be of interest. Healthcare Professionals often express the wish that "diagnostic sequences" - that is, significant (usually abnormal) sample sequences from which a diagnosis or differential diagnosis can be made - be included in the record.

- Imaging, electrophysiology and other large multimedia items. With cheap and massive storage capabilities, storing at least some multimedia items is no longer a problem, but the storage of large numbers of such items can still be problematic both in terms of sheer storage required, and reasonable latency in servicing requests. In an integrated middleware environment, the question of “inclusion” of this sort becomes academic, since physical representation is a separate issue from the integrated logical model.
- Demographic Information: EHRs have historically included some patient data, but with the advent of open distributed systems in which services such as the Person Identification Service (PIDS; [19.]) are available, the need to include demographic information is reduced to that required by sharing of information between systems. In reality, complete sharing of all demographic data is unlikely for political and social reasons, creating the problem of what to do when a piece of the health record is sent to another computation environment.
- Guidelines. Clinical guidelines will be managed by specialised systems, but certain guideline data will need to appear in the EHR. In some cases, it will be by reference, sometimes by inclusion. The question is to what extent EHR models need to represent the detail of guideline concepts, or whether they can just be included opaquely.

How is the EHR Used?

Human Users

EHRs in operational systems are *sinks* of clinical-level knowledge, from sources including:

- clinicians (recorded observations, summaries, care plans)
- allied health workers (social workers, nurses)
- incoming raw data (test results)
- patients.

At the technical level, health records may exist in a federated environment in which they can accept data from multiple “feeder” systems (see below).

Data in the EHR should express knowledge at a clinical level of abstraction, i.e. as used by carers and patients to devise and execute time-based care processes. In other words, not every event which concerns the subject of care is of interest to EHR users, who mainly want to see plans and results.

Healthcare professionals need to be able to use the data in the EHR to support views and high-level concepts such as:

- *Trends*. Most patients with chronic conditions exhibit identifiable trends in their data, which are useful in making diagnoses and managing drug or other therapeutic plans.
- *Derived views*, such as timeline views and problem threads (e.g. as illustrated in FIGURE 2), previous versions in time and so on.

Automatic Processing and Decision Support

Automatic processing systems, such as decision support, statistical and epidemiological systems also have requirements, which can be mainly summarised by the following needs:

- To be able to access EHR data distributed over multiple systems
- To be able to process EHR data in a disciplined way, making reliable assumptions about the structure and semantics of the data, preferably at a domain concept level.

- To be able to access a de-identified version of the data.

The EHR needs to be structured to take account of the relative importance of these various data, and their typical access patterns. In the future, the EHR may include contributions by patients, which will be a particularly useful form of data gathering for patients with chronic conditions.

Proposed Architectures

A general structure which has been developed by various efforts, include GEHR Australia, CEN ENV 13606, the SynEx project, HL7v3 Clinical Document Architecture and others is as follows:

- EHR or EHR_EXTRACT or message consisting of:
 - Containers, each containing content, defined as follows:
 - Headings under which are found ...
 - Entries, represented as...
 - Hierarchies, containing...
 - Data Values, which come in various flavours

The general model can thus be described as “containers and content”. In this analysis, the “container” can be equated with:

- Documents (including in the web sense of a unit of information transmission)
- The unit of committal to a record systems
- The unit of security setting
- The unit of transmission

The container is called a “transaction” in GEHR, and a “component” in CEN 13606.

The fact that numerous initiatives have the same general model means that at least a basic level of standards convergence is possible, and also that the model can be considered widely acceptable.

The Federated EHR

In a distributed systems conception of the EHR, such as that implied by the OMG HDTF standards, the problem of what is and what is not included in the EHR might seem to go away, since everything exists somewhere in the environment. Technically this may appear to be the case within an enclosed computing environment, since EHR information can easily be integrated on the screen through middleware, without having to worry about (for example) whether large images have to be fitted into databases not designed for them.

However, there are two real world factors which usually prevent such convenient access always being the case. The first is that many patients have EHR information at multiple locations, at least at a local GP clinic, as well as hospital(s), alternative practitioners and so on. It is not generally the case that these environments have a tightly coupled computing infrastructure, nor will they have a common security system, or common information models. This means that considerably more work is needed to integrate the patient’s “virtual EHR” than in tightly-coupled environments, potentially necessitating the actual movement and integration of data from one site to another.

The second case is that extracts from the EHR may be required by other parties, which may be other providers, payors (insurers), health departments and so on. Such extracts normally have

to be considered “min-EHRs”, i.e. a composition of all data items found in the EHR, since the requestor is typically not part of the EHR computing environment.

Given that these scenarios are likely to be very common, and that there is no realistic likelihood of totally integrated health computing environments on a national level, EHRs or snapshots thereof will commonly move en masse. Consequently, some consideration has to be given to the question of what is in or out of the EHR.

Who Owns It?

The issue of who owns the health record has been vexed in the past, mainly due to confusion about the difference between ownership of the technical infrastructure on which the records reside (typically a provider institution), who is responsible for managing the data, and who is the legal owner of the content.

Consensus appears to exist around the world on the right of the consumer as the legal “owner” of health record content, which effectively means having the ability to control access rights to the record. It appears that more complex questions to do with the “rights” of government, research or commercial institutions may eventually be solvable with the definition of an overarching e-consent model in which even the rights to de-identified data for research or epidemiological purposes may be controlled by the consumer.

The Message / Record Dichotomy

In recent years, there has been a certain amount of debate about the relative merits of a philosophy of health information based on messages versus one based on health records. Not infrequently, such debates have become heated, and bordered on the irrational. This is a pity, since neither position appears to have been well enough understood to merit such heated defence or competitive positioning. It is most likely that the record/message conflict will be resolved in the future, with the advent of standards which address both areas in an integrated way.

There are a number of dimensions in which a discussion on messages versus records can take place, as follows:

- Historical context
- The information lifecycle
- Domain scope
- Purpose of messages
- Level of abstraction

These are discussed below.

Historical Context

Let us consider briefly why standardised messages of the kind developed by HL7, EDIFACT etc exist.

Standardised receiver interface: receiver systems, typically health record systems, clinical applications etc, need to be able to receive data (lab results, etc) from disparate source systems. For this to be technically and economically feasible, the structure, content and semantics of received messages needs to be standardised, allowing the receiver system to treat all messages in a uniform way.

Standardised source system message model: many source systems are based on relatively simple, in-house information models, or information models which may be disciplined but concern only the information stored at the source (such as lab test data, which may be only a fraction of the total information of interest for a patient). A standardised message specification is one way of enforcing discipline on the information emitted by disparate source systems, without having to say anything about how they work internally. It also serves as a software specification for the maintainers of the source systems.

Source systems, and in most cases to date, receiver systems such as hospital health record systems, are not based on standardised information models, and consequently, message specifications are a natural way to deal with the requirement of communicating health information, without having to say anything further about how any participating system functions internally. This has been the historical position of organisations like HL7, and is entirely reasonable, since it is based on reality, and has proven itself in actual use.

In contrast, EHR proponents have historically concentrated on both the semantics of *in situ* information, i.e. information that is to be stored about patients inside a system and the semantics of *information sharing*.

In some cases they have wondered why message standardisation was necessary at all, since generic means (CORBA, COM, sockets, etc) existed to transport information to and from

EHR systems. While this is technically true, it ignores the fact that message standards were not primarily conceived for EHR systems based on standardised information models, from which well-formed extracts could be automatically obtained; they were conceived in an environment in which very little could be assumed about any of the participating systems. The developers of message standards quite rightly did not make any assumptions about participating systems, and concentrated on explicit messaging standards.

Today, with the advent of technologies such as standardised distribution frameworks (middleware); e.g. RPC, CORBA, DCOM, and more recently .NET. and SOAP, and also XML (a platform-independent information representation format) there is no longer a need to specify message content at the physical level. Instead, messages can be specified as logical compositions of objects defined by an information model.

Consequently, the meaning of the word “message” has effectively changed to mean “logical object composition”.

In HL7, this distinction was not visible in versions prior to version 3, but now clearly exists. Version 3 messages are now defined independently of their final implementation medium, which might be XML, CORBA or something else.

The Information Life-cycle

A justifiable observation which has been made by many people is that it is not really logical to define a standard that describes only messages, since for any reasonable information system, a specification of communicated information strongly implies a model of information inside the system. In other words, messages are just derivative artefacts emitted by systems, and we should just describe “abstract system architecture” standards.

While reasonable, we must be careful not to miss part of the picture. To describe everything properly, we need extra semantics for how communicated information is built and sent. The following simple taxonomy of information lifecycle phases is thus helpful in knowing how to define standards.

Content: semantics of information residing in a system. Typically expressed as an information-oriented class model.

Packaging: semantics of packaging information into extracts or messages, usually describing rules to do with information integrity (e.g. minimum granularity) and security. Typically expressed as extra classes and constraints.

Distribution: semantics of interactions between systems. Typically expressed as a service-oriented class model.

Clearly, EHR standards should consist of at least these three broad parts. Message standards to date have dealt with all three at once, with no clear separation of semantics. However, modern model-based message standards such as HL7v3 are beginning to separate out some of the semantics, potentially making it easier for software engineers to build more modular software.

In any case, the real difference between EHR and messaging standards is not the “systems/message” dichotomy, but *differences in scope and abstraction*, as described below.

What should also be made clear, accordingly, is that while models such as the HL7 RIM might be used to build software, it is not likely to be the same software as an EHR system.

Domain Scope

One of the more obvious differences between messages and health records is that messages sent between systems may relate to numerous things not of interest in the health record, for example:

- Billing
- Patient administration
- Operational (i.e. low-level) information about orders, materials, or demographic entities.

In general, the scope of the health record is smaller than that of messages as exemplified by the HL7 standards, as illustrated in FIGURE 4. However, there are other concepts in the health record not found in messages, such as “headings”, certain items of context identity of authorising health care practitioner, version control and so on.




FIGURE 4 Domain Scope of EHRs and Messages

An example of a information outside the EHR scope (“C” in FIGURE 4) is a billing message from a provider to a payor organisation (usually an insurer or the government). Information in the common scope (“B” on the diagram) might be an EHR entry for a prescription, and messages for the resulting order(s); however, while the scope is the same, the level of abstraction is not. EHR-only information (“A” on the diagram) might be higher-level information relating to care plans and guidelines.

Purpose of Messages

Another aspect of scope is to do with the type of communications between peer EHR systems and between non-EHR systems. Peer EHR systems communicate logically using something called an “EHR extract” (CEN, GEHR, etc), which is not the same as the kind of messages used when one of the systems is a non-EHR system. An EHR extract is a selection of the total EHR, of the same form and level of abstraction as the original EHR, and is intended for integration into the same patient’s EHR at a destination system, and eventual use by a human carer.

The contents of an EHR extract are likely to be significantly larger than a message, and consist of numerous items of information (e.g. the entire chain of entries for a problem); the model of access control will be different as well.

EHR extracts are more likely to be sent between providers, whereas many messages are likely to be between systems inside the same provider organisation.

Level of Clinical Abstraction

Health records differ from messages in another crucial aspect: they are *accumulators of clinical-level health information*. Messages deal with clinical and other information at an *operational* level; models describing messages do not describe a time-based accumulator of information, but rather discrete packets of information as sent at distinct points in time between cooperating systems.

As an example, consider the notion of “order for treatment”, e.g. a prescription, order for physiotherapy and so on.

The clinician's view of an order is really that of a stateful process, where the treatment goes through various states, such as: "proposed", "ordered", "executing", "completed" (with exception states as needed); new information relating to each state is successively incorporated into the record. The level of abstraction of the information is at the clinical level, that is to say, only to a level of detail required for the patient care process..

By contrast, in the messaging world, it is argued in [23.] (Unified Service Action Model section) that the various phases of an order should not be treated as states in the HL7 reference model, but as distinct message types. When one considers the dynamics and purpose of messages, this argument is entirely justifiable, since there will be separate messages for the phases “proposed” versus “ordered” versus “executing” of a treatment, and most likely they will not even be between the same systems. The state machine paradigm is therefore inappropriate.

In general, we can say that:

- Messages are discrete packets of information designed to transmit a quantum of information between two systems at a point in time, rather than descriptions of processes or things evolving in time.
- The type of information in messages is operational, and is primarily for consumption by computer systems, whereas the information in the health record view of an order is primarily destined for human carers, and omits most of the operational detail, which is of no interest to clinicians.
- Messages will usually be numerous but each contain only small, specific pieces of information.

Naturally, it is of crucial importance to align the models defining messages and those describing EHRs, so that common concepts are understood in a compatible way. However, it also implies that the approach for cooperation between such standards is one of “harmonisation of concepts”, rather than “convergence of models”.

Purpose of Reference Models

Both EHR and message standards have reference models - object-oriented semantic models describing entities which are either stored in records, or transmitted in messages. However, in the EHR approach, the reference model defines semantics of the health record only, leaving other semantics - such as for demographics and billing - to related models and specifications. In other words, EHR standards fit more or less into the “landscape” illustrated at the beginning of this paper.

In contrast, the reference model used in the HL7v3 standard describes semantics for everything which is relevant to any message in the scope of the standard. For example, it includes a flexible model of demographic entities such as persons, organisations, places, and the relationships which can occur between them. The reason for this is that since the message is the

only place in which information can be stored in the messaging paradigm, it must be capable of recording the full richness of all information.

In the EHR world, sophisticated demographic models do indeed exist (the Corbamed PIDS specification is probably the best-known example), but they are designed for a demographic service, not an EHR service *per se*. The reference model in EHR standards includes only a modicum of demographic semantics required for minimum identification of demographic entities *in case the demographic service is not available*.

Neither of these approaches is necessarily wrong, but as with other aspects of EHR and message standards, they serve different purposes. The HL7 reference model for example would not be a good basis for an EHR design, since the semantics for demographics and financial entities are too sophisticated for the kind of information stored in the health record, as opposed to stored in other repositories in an EHR environment; in distributed environments, the whole point of separation of concerns is that different kinds of data can be stored, processed and accessed according to their uses and users. This is difficult to achieve if all kinds of information are stored in one place.

On the other hand, EHR reference models may not provide all the semantics needed for messaging, depending on the level of detail and scope required in messages.

That said, there appears to be no reason why EHR and messaging reference models such as found in ENV 13606 and HL7v3 could not one day be converged to a “super” model, probably in the form of a core model with extensions. Doing so would rely on both standards using the same methodology for two-level modelling, since this has significant impact on how reference models are designed.

Use of EHRs and Messages

A final point of comparison between EHRs and messages is the observation that EHRs are created by and for human users (and additionally other computerised systems such as decision support), while messages are targeted at systems.

Clinical authors of EHR content accordingly add entries to the record in logical navigational structures, appropriate to the task at hand. For example, information may be classified under the problem/SOAP headings commonly used in general practice.

Additionally, a requirement of the EHR is that the contributions made by clinicians (or other systems) need to include auditing information describing the details of who, where and when information was added. In contrast, this kind of context information is not generally required for messages, since they are intended simply to carry content.

In practice, it is not obvious that there will be such a clear distinction between the kind of navigational and contextual information that will be used in messages and that used in EHRs.

However, in principle we can expect that EHR standard models will contain a certain number of formal concepts relating to organisation and context not found in message standards models.

Summary

The remarks made above are summarised in Table 1.

	Message	EHR_EXTRACT
Purpose	transmit quantum of information	transmit selection of clinical information accumulated over time
Abstraction level	operational	clinical
Time-space characterisation	small, numerous	large, probably few
Domain Scope	clinical, demographic, financial, pathology, ICU, ...	clinical
Representation of process	Message per transition	State machine, with version per transition
Reference Model scope	Semantics of all possible message content, including clinical, demographic, billing, orders, etc, etc	Semantics of EHR.

Table 1 Messages versus EHRs

From this section, it should be clear that:

- Message and EHR standards serve different purposes.
- Both kinds of standards express content, packaging and distribution semantics, but for different domain scopes and at different levels of abstraction.
- Message and EHR standards should probably be harmonised, not converged.

Design Methodology

Overview

When considering how to construct or review good health information standards, insufficient attention is typically paid to the consequences for software construction and runtime systems. Many of the major problems of the past for information-intensive systems, including most EHR and related systems, have to do with the inability to deal with change.

This section reviews design methodologies used to build systems based on a “reference model” and a second level of domain models.

The Archetype Methodology

A recent methodology which has been proven in independent work by GEHR Australia and University College London, to avoid the problems of system maintainability, and to address a number of others is described in [7.]. Standards which use this methodology can be said to subscribe to the following basic principles:

- Separation of concepts into a small technical model, usually known as a “reference model” (or “reference information model”, as in HL7v3), and another layer of formal domain concept models (each describing a distinct domain concept, such as “blood pressure” or “microbiology results”), which may be quite numerous.
- Software is formally based on the reference model and EHR data are technical instances of it.
- Systems have the property of being able to be deployed before domain concept models are created; or in other words, domain concept model creation is not bound to the software development of systems.

Software Meta-architecture

The archetype software meta-architecture is shown in FIGURE 5. This figure illustrates a meta-architecture based on two models.

Reference model (RM): the model from which software can be built, and of which all data are instances.

Archetype model (AM): a related model whose instances are domain concepts, or *archetypes*, which are directly processable by information systems.

The important addition to systems engineering offered by this meta-architecture is that of archetypes, which are formal domain concept models. Each archetype is an instance of the archetype model, and is essentially a structural constraint model for valid constructions of reference model instances. Archetypes are used:

- To enable domain experts to construct formal models of their own working concepts, at a domain user level.
- To enable systems to formally validate data during input or batch processing.
- To enable systems to be interoperable at the knowledge level, not just the data level, but exchanging data generated by archetypes.
- To enable intelligent querying and decision support, based on archetype-aided automatic processing

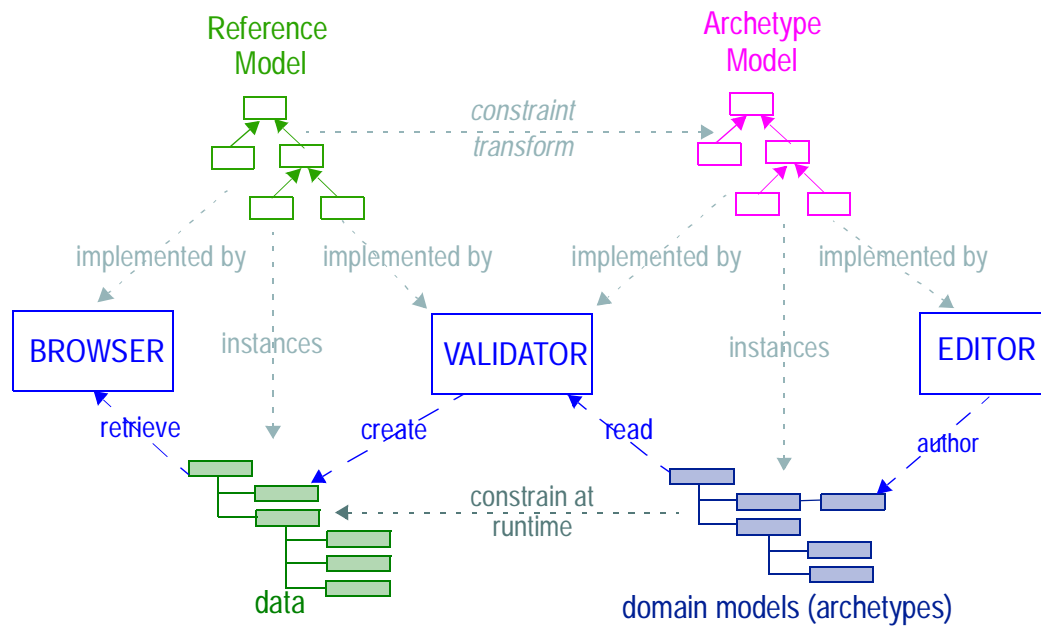


FIGURE 5 Archetype Meta-architecture

In FIGURE 5, several processors (implemented as software) are shown, being based on the reference and/or archetype models. These are:

Domain Model Editor: a GUI application for creating new domain concept definitions, based on the constraint model.

Validator: any component which creates or manipulates valid data using archetypes. This is based on the reference and archetype model classes.

Browser: a generic browser can be built, based solely on the RM, although a smarter browser can be built using the archetype model as well.

Almost all applications in a real system are instances of the “validator” component; that is, they have the property of being able to manipulate data in the presence of the constraints expressed by domain models.

The most important property of systems based on this scheme is that instance data (shown at the bottom left) are not only technically conformant to the RM (as per the usual object-oriented class/instance relation), but are also conformant to one or more archetype instance (bottom right). That is, they are both *valid RM instances*, and *logical instances of domain models*. Further, the variability expressed in archetype constraints enables more than one data instance to be identified as instances of the same domain model.

The constraint transform relationship between the reference and the archetype models is a new kind of formal relationship between models, and is not typically treated in the object-oriented literature. However, it is not technically difficult to devise such a relationship, and it has been implemented in the GEHR (Australia) and SynEx (UCL) projects.

This meta-architecture can be used to create a family of reference models, which collectively describe the semantics of the EHR, terminology services, demographics, and other areas shown in FIGURE 1. For each reference model, instances of the archetype model are domain level specifications for concepts in that area. For example, an instance of the archetype model

for an EHR reference model might be an archetype for the concept “blood pressure” or “biochemistry results”; similarly, an instance of the archetype model for a demographics reference model might be an archetype for the concept “dutch person name” or “basic UK NHS patient details”.

Messages and Legacy Systems

An important side-effect of the two-model meta-architecture is that they can be used to develop other specifications. In particular, message specifications can be automatically derived from archetypes, by generating the allowed permutations of each archetype.

Archetypes can also be used to model the data available in legacy systems, enabling it to be extracted in a manner whereby it can be processed generically. One way to deal with legacy system data is to use both a custom archetype modelling the legacy data, and a standard archetype as a target. Legacy data can then be massaged from one form to the other, based on formal conversion techniques and XML/XSLT technology.

This separation of concerns is the key to having enough flexibility to deal with diverse requirements, systems, users, and changing technologies.

See Appendix A for more details on the archetype methodology.

The HL7v3 Methodology

Interestingly, the imminent HL7 V3 specification has recently evolved to a two-level model as well, and from an entirely different starting point. HL7v3 contains what are effectively formal domain models, in the form of its R-MIMs (refined message information models) and their derivative HMDs (hierarchical message descriptions) and CMETs (common message element types). Although originally conceived for messages, the HL7 effort has discovered that the only feasible way to express message structures as would be used in real systems is not to rely solely on a reference model (known as a “reference information model”, or RIM in HL7 parlance), but to use discrete models based on the reference model.

However, the methodology differs in a number of important aspects from the archetype methodology described above. See Appendix A for details.

Comparison

Table 2 summarises the differences between the archetype and HL7v3 domain modelling approaches.

	Archetype	HL7v3
Purpose of reference model	Normative type definition for all EHR data.	Informative definition of the language of R-MIMs
Formalisation of clinical concept	Archetype	HMD [and CMETs]
Clinical concept authored by...	Domain experts	Technically proficient users

Table 2 Archetype versus HL7v3 R-MIM methodology

	Archetype	HL7v3
Clinical concept authoring method	Tool-based	Manual, using diagramming tool (future: automated tool?)
Relationship of data and reference model	EHR data are instances of the RM classes.	Message data are instances of HMD and CMET classes.
Software based on ...	RM only	R-MIMs / HMDs / CMETs

Table 2 Archetype versus HL7v3 R-MIM methodology

A Taxonomy of Standards

A good design paradigm provides a qualitative basis for constructing good models for standards; the archetype method described above thus shows how reference models can be limited in size, and successfully used as the basis for high-quality software.

However, there is more to the story: we still need to know how to scope models, both thematically (what concepts do they describe?) and computationally (how to relate attributes, functional interfaces and constraints).

Separation of Concerns

FIGURE 1 shows a “health information landscape” in which the various information functions of an environment are separated out into different services. This kind of separation is crucial for making reference models, and consequently, software, manageable and maintainable.

Once a thematic separation is described, the methodology for developing standards should be decided.

For many areas including the EHR, demographics, terminology, and guidelines, the archetype methodology is suitable (although not the only approach possible), and would thus result in a reference model, archetype model, and archetypes.

How the domain is divided up in the first place is probably as much an art as science, but basic criteria for distinguishing areas are:

- Who are the creators and maintainers of the information? (compare the authors of terminologies with the authors of prescribing reference data.)
- Who are the users of the information/service? (compare the users of decision support systems and administration systems)
- Usage patterns of the service (compare the fast, constant interrogation of a terminology service to the relatively infrequent but potentially complex interactions of EHR users and EHR systems).
- What technologies are used for the service (e.g. imaging requires specialist computing systems).

The OMG HDTF separation of concerns in the domain is probably still one of the best general statements on the matter.

ISO RM/ODP

The ISO Reference Model for Open Distributed Processing (RM/ODP; see [2.]) offers a way of teasing out the requirements as well as the different layers of technical abstraction (information versus service versus implementation), in order to see how to structure reference and knowledge models with respect to each thematic area of the domain.

The ODP model consists of five viewpoints, of which three are tabulated in Table 3 against a number of levels of enterprise requirements relating to health records.

Cells in the table contain concepts found in some of the major existing standards and work efforts.

The classification used in Table 3 is inspired by work done at the Veteran's Health Administration (see [15.]), and at the University of Magdeburg (see [10.]).

	Enterprise (purpose, scope, policies)	Information (semantics & processing)	Computational (service model)
Data Management	Clinical EHR requirements e.g. GEHR requirements	Reference Model: Content part (1) e.g. GEHR Aust. RM; CEN pt 1; HL7 RIM core 6 classes	Service Model: fine-grained API e.g. HDTF COAS
	peer - peer EHR communication	Reference Model: Packaging part (2) e.g. GEHR Aust. EHR_EXTRACT; CEN pt 4	Service Model: Distribution part (3) e.g. CEN pt 3 HDTF HILS
	point - point communication (non-EHR)	Reference Model: Messages HL7 RIM message classes	Service Model: Messaging service interface
Knowledge Management	Requirement to formally model domain concepts	Archetype Model e.g. GEHR Aust. AM, SynEx AM	
	Domain concept model library		Archetype Repository Model service model of archetype libraries
Knowledge	Clinical / Admin processes	Clinical / Admin workflow archetypes , e.g. workflow archetypes, HL7 RIM Referral class, HL7 R-MIMs	Service Model: Workflow e.g. HDTF Orders
	Clinical / Admin knowledge	Clinical /Admin content archetypes e.g. clinical content, demographic archetypes, CEN demographic classes, HL7 RIM Patient_encounter, Diet classes, Entity, Role, Financial_Act, some Act subtypes HL7 R-MIMs	Service Model: Knowledge e.g. HDTF PIDS

Table 3 RM/ODP classification of EHR and Messaging Standards

What this table indicates is that well-formed models neatly fall into one or a small number of boxes in the same viewpoint. For example, well-formed reference information models should fall into the information viewpoint. There should also be only one way to express a concept in

a given viewpoint. For example, in the GEHR system, the only way to express demographic entities is using archetypes (see last row, informatin viewpoint), whereas both CEN and HL7 include demographic classes in their reference models, and HL7, numerous other domain concepts. This might be reasonable in a single-level modelling situation, but where archetypes or other two-level modelling mechanism is being used, e.g. HL7 RMIMs, confusion is likely.

Message / EHR

A general scheme of systems and communications is illustrated in FIGURE 6, showing the role of EHR extracts and messages with respect to EHR and non-EHR systems.

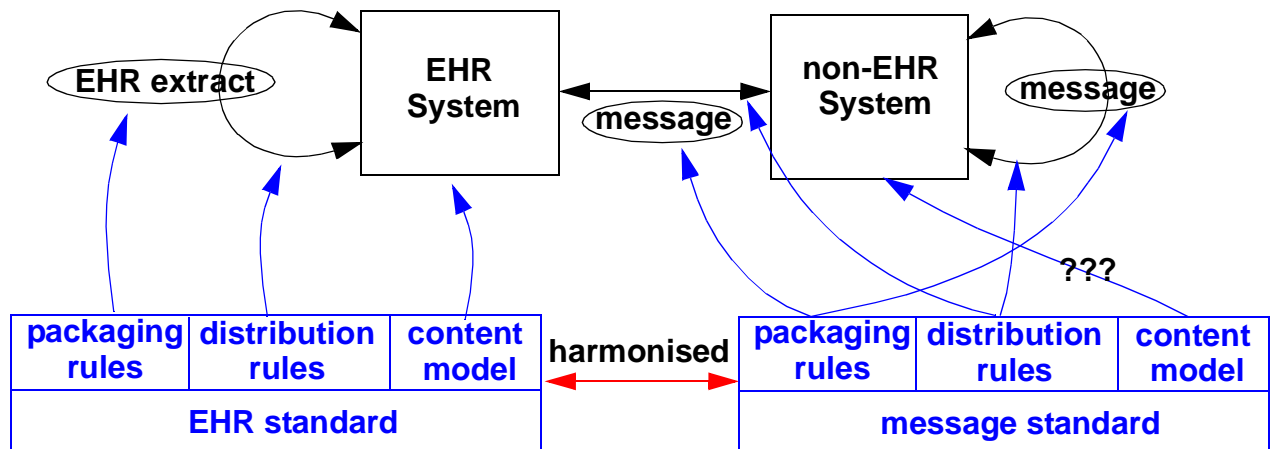


FIGURE 6 EHR and Message Standards Relations

Both FIGURE 6 and Table 3 lead to a general formulation of the structure of standards required in the future, which essentially follows the form of Table 3.

In concrete terms, it is therefore highly desirable to make tomorrow’s version of specifications such as HL7v3, CEN ENV 13606, GEHR and OMG HDTF, all based on the same underlying set of models. Doing so would require a widening of perspective by each party to accommodate the historical interest areas of the others.

Criteria for Evaluating Standards

This section describes a set of criteria by which standards efforts and other work efforts can be judged. It is based around the critical attributes of real software systems, since the only point of HIS standards in the end is to make it possible to build good quality, interoperable health information systems. These are broadly classified under two headings: quality of service (the user perspective) and economic value (the payor perspective). We can identify the stakeholders in more detail as follows:

- Healthcare Consumers:
 - Patients
 - Well healthcare consumers (preventative strategies, including lifestyle etc)
- Clinicians:
 - General practitioners
 - Specialists
 - Allied health professionals
 - Alternative & complementary medicine practitioners
- Health provider organisations
- Public health organisations
- Health researchers
- Health educators
- Administrators
 - Managers
 - Health policy makers
 - Regulators
- Software Developers:
 - Software systems builders
 - Systems integrators
 - Product/Component developers
- Payors:
 - The government, via public health funds and publicly funded providers
 - Private health funds
 - Privately funded providers

For all of these Stakeholders, there are a number of materially important perspectives, which are described below.

Critical Attributes

Separation of Concerns

Standards should consist of a family of models which address each of the areas of the landscape (as understood by the standards - not necessarily the same as the one presented here). Separation can easily be achieved using e.g. UML packages.

Why: The separation of concerns is helpful in understanding and maintaining the model, and essential for modular software engineering.

Test:

To Be Continued:

Explicit Domain Concept Models

Systems should understand domain concepts explicitly.

Why: all systems process information representing concepts in the domain. If such concepts are not described by formal models which are standardised (within their scope of use), their semantics are ambiguous. Distinct implementations will create incompatible implicit models of these concepts, and there will be no hope for reliable interoperability or standardised automatic processing.

Test: clinical concepts such as “blood pressure” or “biochemistry results” can be found explicitly modelled somewhere in systems.

Test: decision support systems are able to inspect information at the granularity of domain concepts, and do not have to resort to brute-force methods to find data items of interest.

Test: queries can be constructed in terms of domain models rather than low-level technical information concepts.

Future-proof Software

Deployed software systems and databases need to be immune to domain-level changes in requirements, such as the need to adjust or add new domain concepts.

Why: the primary cost factor in software systems is in maintenance and enhancement, typically accounting for 70% of the cost of the system over its lifetime. Having to constantly change software and databases, test them and redeploy them constitutes a significant part of health provider budgets. Secondly, constantly modifying software systems increases the risk of errors, data corruption, and reduces the quality of service.

Test: The software itself cannot include concrete models of domain concepts, since domain concepts are the source of change. Software and databases which concretely model these concepts will constantly have to be modified and redeployed. Domain concept models must exist outside of the software, in a form consumable by systems.

Domain Empowerment

Domain stakeholders should have the technical possibility to author models of domain concepts themselves, and introduce them directly into their systems, without having to involve software developers, or requiring software or database modifications to be carried out.

Why: the health domain is changing all the time, and new and changed information concepts are constantly being defined. If user bodies can define these concepts in a form directly usable by systems, the systems will learn in time, and not become obsolete. Secondly, healthcare professionals are best placed to define domain concepts and should be able to do so without recourse to IT specialists. Thirdly, users need to be able to introduce new concepts into systems post-deployment.

Test: domain-oriented tools (i.e. ones which do not expect knowledge of the underlying technical information models or semantics) for authoring domain concept models exist, which create formal definitions directly processable by health information systems.

Knowledge-level Interoperability

Systems need to be able to communicate at the level of domain level concepts rather than just generic informational concepts (e.g. “Quantity”, “Coded_text”).

Why: if common models of domain concepts can be assumed by sender and receiver systems, they can make powerful assumptions regarding the structure and semantics of the data, enabling sophisticated display, data entry and batch processing.

Test: systems are able to request information by identifying it by its domain model identifier (e.g. “prescription”, “family history” etc)

Test: systems can process received information assuming particular domain models.

Appropriate Use of Terminologies

Standards should support systems which make the best use possible of externally defined terminologies (also known as “controlled vocabularies”), specifically avoiding hard-coded terms or enumerated types, and using domain models which also avoid inflexible uses of terminology as well.

Why: terminologies represent a set of facts and relationships in the domain, which changes from time to time. .

Test: systems automatically take account of terminological changes. For example, when the reclassification of hepatitis which has occurred in the last decade was introduced into terminologies, systems should automatically be able to permit the new classifications to be used without changes to software or domain models.

Technical Interoperability & Future-proof Information

Health information must remain available regardless of changes of computing platform, software vendors and so on.

Why: changes of software vendor and technology in systems are inevitable, as are heterogeneous computing environments. But health information must remain usable for far longer periods than the typical technology cycle of 2 - 5 years; therefore it must be immune to such differences.

Test: information can be moved from one vendor product to another, with no loss of meaning or other errors

Test: information can be shared between systems built from different technologies.

Test: users do not suffer from “vendor lock-in”.

Test: information remains available in the long term.

Security and Privacy

Systems need to implement a flexible model of security which can be modified as needed in specific contexts, including taking account of varying legislation in different countries. The guiding principle should be patient-based consent for the use of information, balanced by an acceptance by patients of the responsibility for the consequences of particular decisions made with respect to privacy. Systems also need to respect privacy and security even when information is transferred outside of its original environment, to places where security and privacy definitions may be different.

Why: Electronic health systems will not be accepted if they violate basic rights and preferences of citizens as patients.

Test: Systems correctly deny access to users whose identity or time / space context does not match the preferences consented to by the patient.

Test: Systems are able to share information across different security environments, while respecting the patient's privacy specifications.

To Be Continued:

Possibilities for Standards Convergence

Global Standards?

We may now (finally) be on the verge of viable standards-based EHR systems, based on the following major work around the world:

- The European GEHR project (AIM project 2014; 1992 - 1995) produced what remains the most comprehensive requirements statement about health records.
- The European standards agency CEN has produced the third iteration of its EHCRA standard, ENV 13606; while it still contains some implementability problems, it serves as a formal model embodying many of the GEHR, CEN and other EHR requirements identified to date.
- Various projects at University College London (1995 -), including EHCR Support Action, Synapses, and SynEx have resulted in a prototype implementation of a 2-level model-based architecture for extracting data from non-EHR source systems into a standardised health record.
- UK efforts such as Prodigy and GP2GP communication.
- Corbamed-based systems in the US and Brazil
- The Australian GEHR project (1997-) has a prototype implementation of a 2-level model based EHR architecture. The second level contains formal clinical models which have a formal relationship to the reference information model, enabling significant improvements in the way health systems are built and health information is modelled.
- Recent (Nov 2001) agreements between *openEHR*, CEN TC/251, PROREC, and the European Records Institute to collaborate on convergence of models and methodologies.

Given the directions taken in recent work efforts, including GEHR Australia, UCL (SynEx) and numerous open source projects, it is suggested that global standards for EHRs and messaging will be possible under the following circumstances:

- A harmonised understanding of the RM/ODP framework is adopted
- A harmonised design paradigm is adopted, ensuring compatible descriptions of the archetype concept
- Reference models are harmonised

Here, “harmonised” means that entities describing the same concept (possibly at different levels of detail, or in different views) should be modified so that the common subset of semantics is the same.

An EHR standard should thus consist of:

- A reference model in three parts: content, packaging rules, distribution rules
- A corresponding archetype model
- Technical expressions of the reference model and archetypes suitable for use in:
 - EHR extracts. XML schema might be used for example.
 - System development. Expressions in this case might be in (for example) UML, IDL, .net form.

- Example archetypes, i.e. formal models of domain concepts related to the area described by the reference and archetype models.

A messaging standard should consist of similar deliverables, harmonised with their counterparts in the EHR standard.

Players

Standards organisations: ISO/TC 215, CEN TC/251, HL7, ASTM, OMG (HDTF), DICOM.

Non-profit organisations: *openEHR* Foundation, Open Source Health Care Alliance.

Work efforts including:

- G-CPR project (US) and its successor, the Federal Health Information Exchange (FHIE).
- Open source efforts, e.g. FreeMed, openEmed, Gnumed, Vista

Technical Methodology

Methodologies featuring two levels of models are starting to appear in HL7 and CEN. Other aspects of the methodology need to be discussed.

Methodologies need to be heavily reviewed for suitability for software implementations, not just their prescriptions for information structures or services.

EHR Reference Models

Overview

Sources of semantics for a converged EHR reference model include:

- CEN ENV 13606 and its various implementations in Europe
- GEHR Australia GOM ; UCL SynOM (soon to be merged as the openEHR reference object model)
- DICOM models
- HL7v3 Reference Information Model

Each of these tends to have a number of identifiable aspects, namely, data types, terminology, interface to guidelines, protocols and reference data, and content and structure, and messages or extracts.

Data types

Sources: CEN (theoretical, various implementations), GEHR (implementation-tested), HL7v3.

Expressions: UML, IDL (derived), MSIL (.net; derived), XML-schema (derived)

A converged set of data types would probably include the semantics of most of the HL7v3 data types and CEN DATA_VALUE subtypes, contained in a single model.

Terminology

There are two categories of terminology: reference terminologies, which relate to domain concepts (such as found in SNOMED, ICPC, etc), and controlled vocabularies which define allowable values for attributes in constructed models such as CEN and HL7.

It is desirable to standardise such “internal” term sets, which requires that the meanings correspond to an attribute with the same meaning in domain models used in each type of system. The HL7 internal terminologies and CEN categorial terms appear to be primary candidates in this area, and should be aligned for harmonised attributes.

Representation of coded terms needs to be compatible in all reference object models.

Guidelines, Protocols, and Reference Data

A standard approach to interfacing to reference data including terminology, protocols, clinical guidelines, prescribing data needs to be defined.

Structured Content

Requires semantically equivalent information building blocks, else receiving systems cannot understand messages. In other words, reference models (HL7 RIM, GEHR OM, CEN model, SynOM) need to be at least semantically equivalent.

Reference models should only include:

- Internationally acceptable concepts
- Future-proof concepts
- Domain-generic information concepts

Consequently, HL7 should remove financial types and some other non future-proof concepts from the RIM, and express these as RMIMs/HMDs.

Messaging / Extracts

Agreement on packaging required between standards organisations, i.e. on:

- Meaning of EHR extract (CEN, GEHR, SynEx) versus message (HL7), document (HL7 CDA)
- Alignment of Transaction (GEHR), EHR contribution (CEN, SynEx), special Act Context (HL7)
- Sufficient contextual information to allow information to be added to EHR; i.e. to satisfy EHR requirements
- Agreed rules about integrity and granularity of messages. Requires a semantic specification of concepts such as “transaction” and “message”.

The Domain Model Level

For knowledge-level interoperability (needed for any automatic processing), convergence is needed at the domain level - i.e. between HL7’s RMIMs / HMDs / CMETs and GEHR’s and SynEx’s archetypes. Other formal domain models such as Odyssee (France) and other open source systems (e.g. OIO and FreeMed) should also be considered.

Archetype Models

To have interoperability of domain concepts, the formal models describing the semantics of the domain level need to be aligned. In GEHR and SynEx, this is an archetype model for EHR semantics, i.e. a formal class model expressing constraint semantics on all aspects of the reference model.

In HL7v3, it is not yet clear what the formal model of RMIMs, HMDs etc is, even though these artifacts clearly exist; however, if the formal model can be described, a technical basis for converging domain models should be available.

This alignment is crucial for the sharing of data between systems built on different standards, and also for the sharing of the domain concepts created by different groups.

Archetypes

Archetypes or other constraint models from different standards groups need to be aligned in such a way that wasteful replication is avoided. This means all archetype authors having access to online libraries of archetypes (and their equivalents), no matter who creates them. It also means that there should be ways of creating specialised versions of these models, to allow for localisation and specialisation.

The Role of Standards Organisations

An important question is what changes are needed in the role of standards organisations in order to accommodate the challenges of RM/ODP viewpoints, archetypes and component-based software. The following sections contain a few thoughts on this subject which may well be controversial!

Technical Models and Methodology

Clearly standards bodies should continue their current role in the formulation of technical models and specifications.

An area of improvement for standards processes such as those found in CEN and HL7 would undoubtedly be an improved capacity for using implementation experience to inform future revisions of standards. The general situation to date has been one in which standards are published without the benefit of open, well-studied implementations, with the consequence that they contain even quite simple problems for implementation. This in no way reflects on the high quality of work that has gone into the specifications, but rather a basic human short-coming of being unable to see far enough down the myriad paths of complex development processes, in order to know the consequences of choices which appear logical in the discussions which take place during standards development.

A recommendation therefore would be for standards groups to find ways to identify implementation groups which are actively testing a current version of the standard, and to find ways of accepting feedback from them. This needs to be done in such a way as not to compromise the proper process of standards development, while taking into account learning experiences which might make crucial differences to standards.

It should be noted that untested assumptions early on in standards development can be extremely problematic, since if they turn out to be wrong, the quality of all subsequent work is seriously compromised. Conversely, early engagement with external implementation efforts can ensure that all major assumptions are sound.

The eventual outcome of such a collaboration is standards which have effectively been “road-tested” on the day of release, and which can confidently be picked up by industry and the public sector.

Domain Concept Models and Archetype Management

The situation for the domain concept models is less clear. While it has been suggested that the formal domain models should be developed by domain stakeholders, there is clearly a potential role for the standards bodies in:

- Helping to logistically organise archetype development
- Developing or reviewing methodologies for developing archetypes (e.g. a workshop process which enables domain people to crystallise formal expressions of their everyday information concepts).
- Developing a quality assurance methodology for archetypes, and performing actual QA on archetypes.
- Certifying archetypes.

It is suggested that it should be possible for domain stakeholders to perform a large proportion of archetype development without having to directly take part in standards meetings etc; in this way work can proceed without the bottleneck of over-formal process being applied right from the start.

However, at some late stage of development of each archetype, it seems reasonable that the development should move into a forum managed by a standards organisation, followed by a relatively quick quality assurance, testing and certification process.

The key is to find a balance whereby the weight of standards process does not stifle the development of archetypes during the “creative” phase, but can be used to ensure quality in the final stages.

Archetypes intended for experimental or purely local use probably do not need formal certification.

As long as archetypes are marked in a standard way as to what their status is, disciplined archetype management should be possible, even if in some parts of the health domain, their creation is quite anarchic, while in others it may be heavily controlled.

Registration Authorities

To Be Continued:

Terminology

Archetypes

Guidelines

Reference data (prescribing, demographic)

Appendix A

The GEHR Archetype Methodology

GEHR Australia has developed and refined the archetype methodology, based on the general principles described earlier.

A microbiology results archetype example is illustrated in FIGURE 7, showing that archetypes are quite comprehensible for domain users.

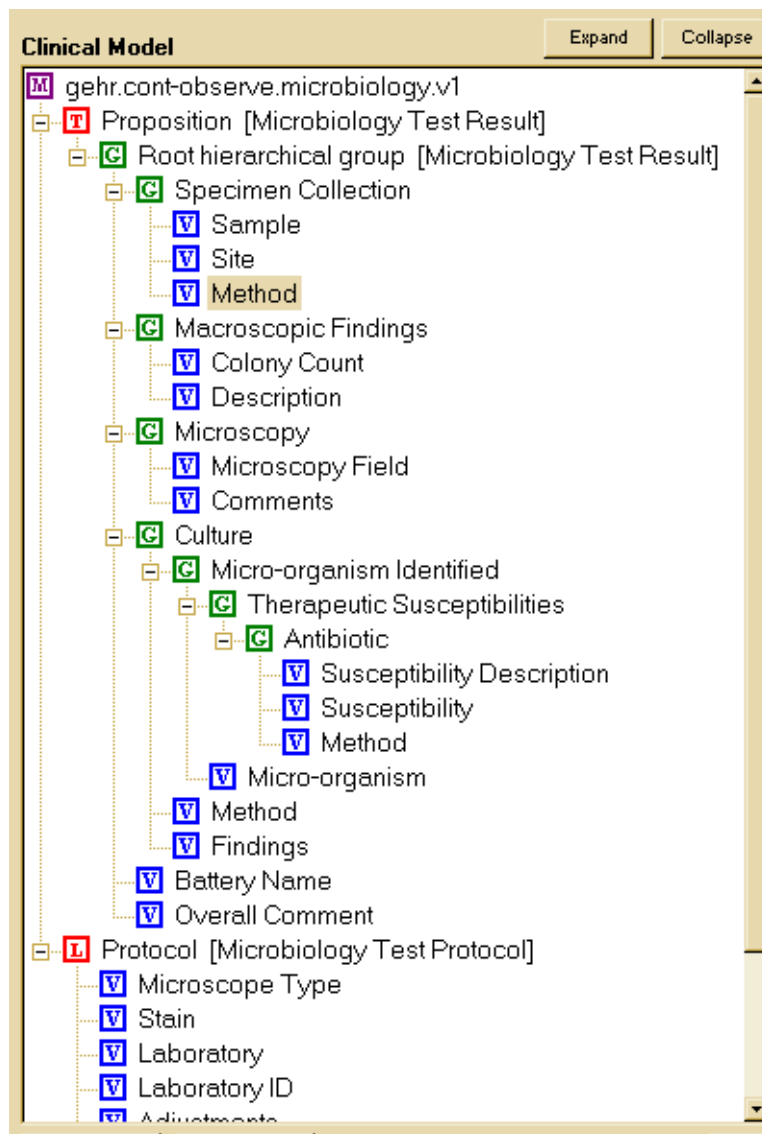


FIGURE 7 Archetype structure for Microbiology results

The essence of a methodology for developing not only systems, but standards, is to recognise three threads of work:

- The technical work of developing reference models and archetype models. This is carried out by a mixture of domain experts and IT specialists.
- The development of domain level models, i.e. archetypes.

- The development of domain level structured vocabularies.

FIGURE 8 illustrates these, and shows how standardised models, vocabularies and archetypes relate to operational systems.

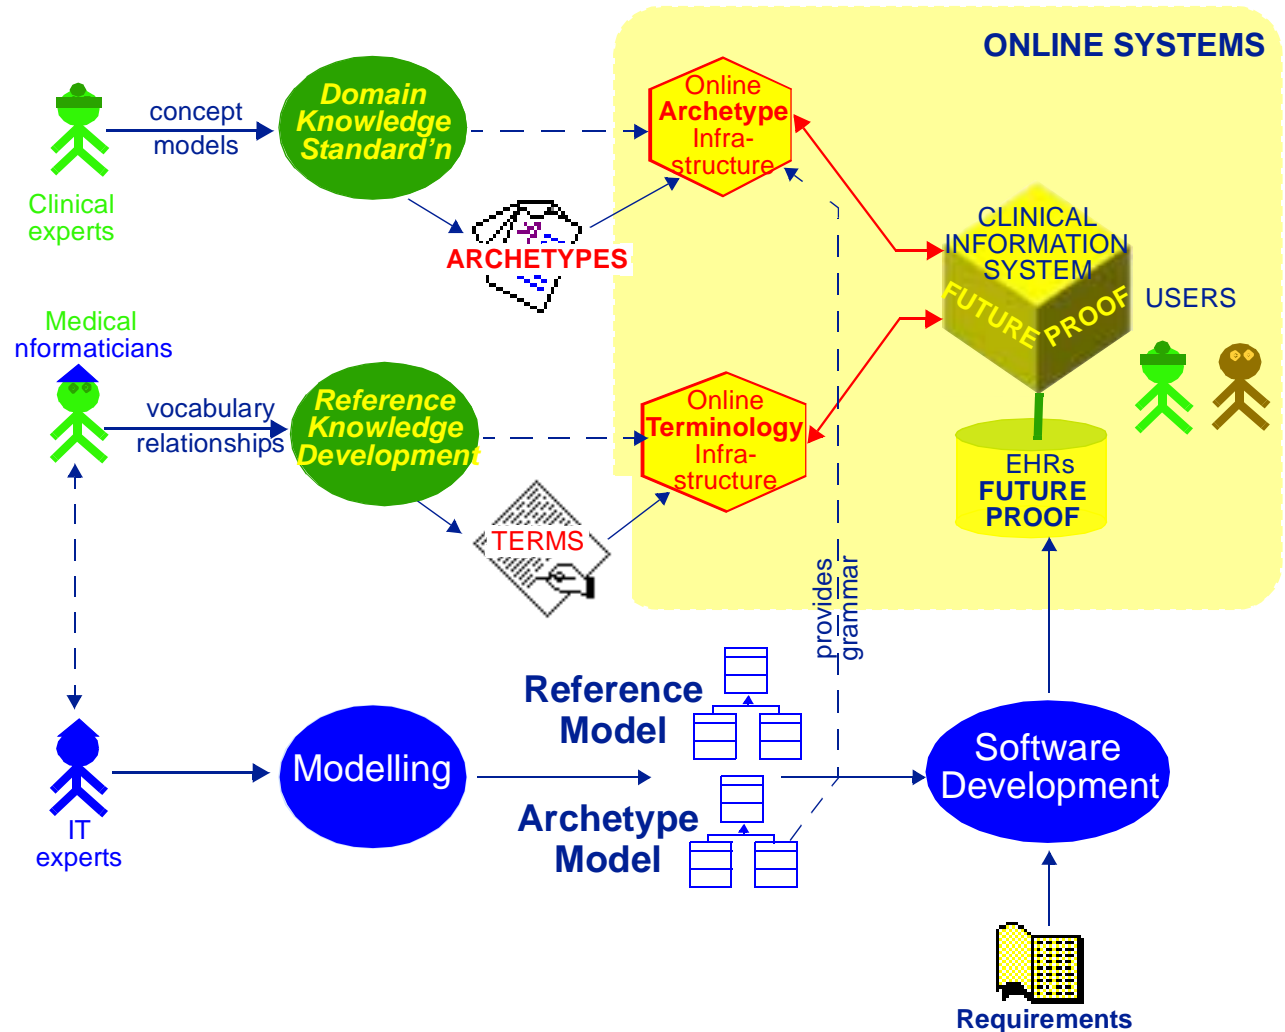


FIGURE 8 Processes in an Archetype-based Methodology

This separation of concerns is the key to having enough flexibility to deal with diverse requirements, systems, users, and changing technologies.

The HL7v3 Methodology

Interestingly, the imminent HL7 V3 specification has recently evolved to a two-level model as well, and from an entirely different starting point. HL7v3 contains what are effectively formal domain models, in the form of its R-MIMs (refined message information models) and their derivative HMDs (hierarchical message descriptions) and CMETs (common message element types). Although originally conceived for messages, the HL7 effort has discovered that the only feasible way to express message structures as would be used in real systems is not to rely solely on a reference model (known as a “reference information model”, or RIM in HL7 parlance), but to use discrete models based on the reference model.

However, the methodology differs in a number of important aspects from the archetype methodology described above. Instead of archetypes, the HL7 approach uses R-MIMs, or

“restricted message information models”, based on the reference model. Each R-MIM is in fact a specialised schema, or model, in its own right, derived from the RIM by:

- Removal of classes and attributes irrelevant to the particular message being defined
- Cloning of classes to provide compositional replication
- Addition of some constraints

The result is that each R-MIM is like a small RIM in its own right, with its own class names and namespace. R-MIMs are expressed using a diagrammatic language, as shown in the example in FIGURE 9.

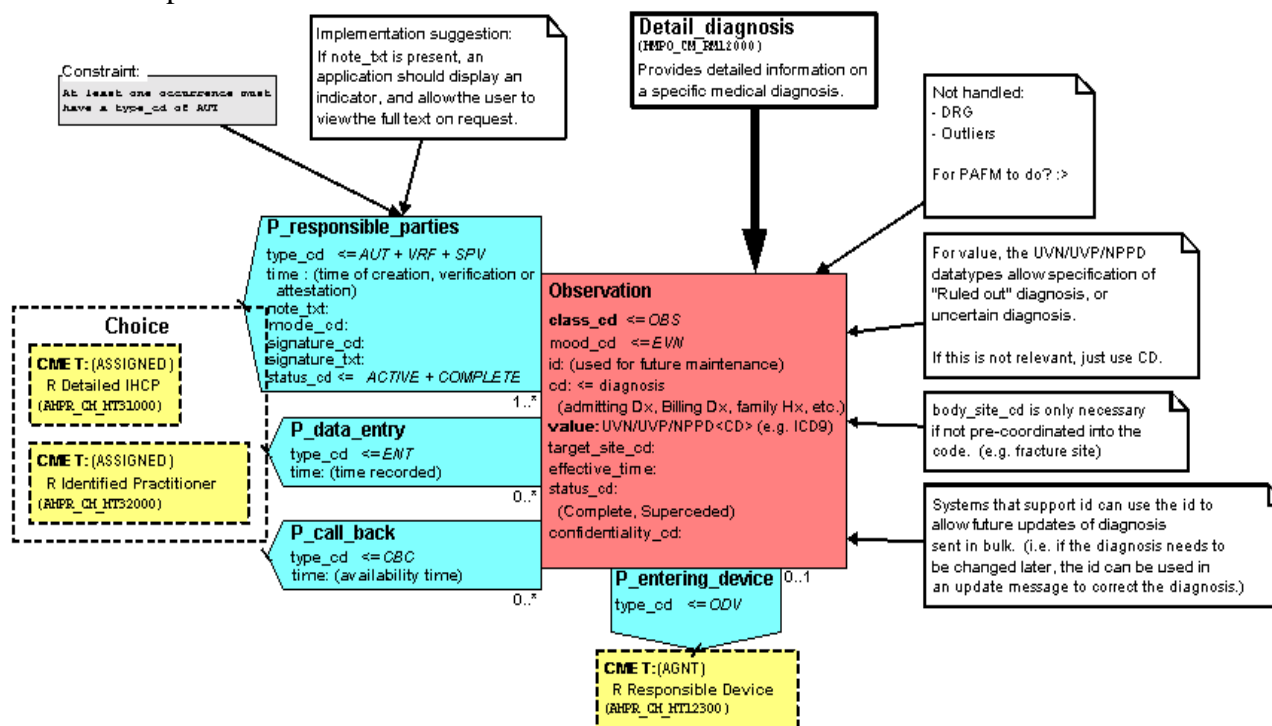


FIGURE 9 R-MIM: detailed information on a specific medical diagnosis

R-MIMs are further refined to produce HMDs and CMETs; the former is effectively a template for messages, while the latter is a template for re-usable subparts of messages which are likely to occur in more than one message.

FIGURE 10 illustrates the refinement methodology.

The differences between this approach and the archetype approach include the following:

- The reference model (RIM) is only an informative definition of the language of R-MIMs, rather than being a model of information instances. To find out what model data are instances of requires the particular R-MIM / HMD / CMETs to be available.
- Message data items are not instances of RIM classes, but of the classes found in HMDs and CMETs.

Each new R-MIM (and its HMDs and CMETs) causes new class definitions to be created. If message-processing software is based on classes corresponding to the R-MIM classes, the addition of new R-MIMs will continually require update of software. Since each type in an R-MIM is a subtype of a RIM type (+/- attribute modifications), the links from the R-MIM classes to the generating RIM classes are also needed, if any general validation capability is required.

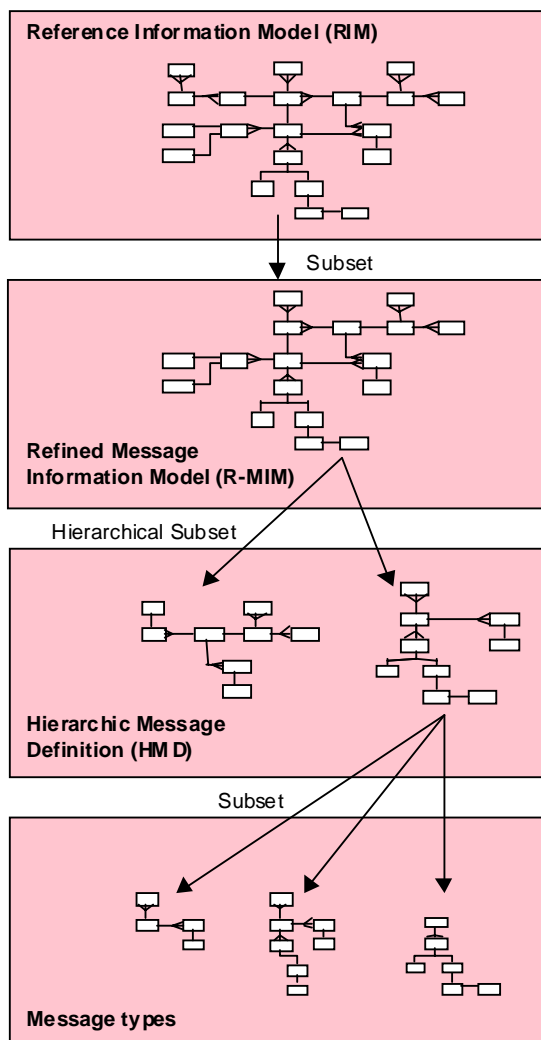


FIGURE 10 HL7v3 Message Development Methodology

- R-MIMs are not intended for domain experts to understand (according to some HL7 practitioners), rather they are created and maintained by technical personnel. (The current diagramming format is certainly challenging to understand for most people, whether from the domain or a technical background. However, improvements in the formalisation of R-MIMs can be imagined such that tools could be built which display R-MIMs in a much more domain-friendly way).

References

General

1. Feyerabend, Paul. *Against method*. (3rd Edition). Verso, 1993.
A classic in the philosophy of science. Discusses the way in which new hypotheses are often compared to the extant theory of something, which is unreasonable, since there is nothing to say that the extant theory is the better one. All theories should be judged against evidence, experiments, and requirements (being the evidence of user needs).

Information Technology

2. ISO/IEC: Information Technology. *Open Distributed Processing, Reference Model: Part 2: Foundations*.
3. Raymond, K. RM/ODP XXXXXXXX full ref required
4. Sowa, John F. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Brooks/Cole, California, 2000.

Health Informatics

5. Ameline, Philippe and Brûlet, Jean François. *La Ligne de vie: Continuité et coordination des soins*. (In French). 2001 (to be published).
6. Beale, Thomas, Goodchild Andrew. *openEHR Modelling Principles*. See <http://www.deep-thought.com.au/health/XXXX>.
7. Beale, Thomas. *Archetypes*. 2001 (available at <http://www.deepthought.com.au/it/archetypes.html>).
The technical basis of two-level modelling.
8. Beale, Thomas and Heard, Sam. *The GEHR Object Model Architecture*. 1999, the GEHR project (available at http://www.gehr.org/technical/model_architecture/gehr_architecture.html).
9. Beale, Thomas. *The GEHR Kernel Architecture*. 1999, the GEHR project (available at http://www.gehr.org/technical/kernel_architecture/kernel_architecture.html).
10. Blobel, Bernd. *Application of the component paradigm for analysis and design of advanced health system architectures*. International Journal of Medical Informatics 60 (2000) pp 281–301.
11. National Health Records Task Force, Australia. *The Health Information Network for Australia*. July 2000. <http://health.gov.au/healthonline/healthconnect>.
12. Kalra, Dipak et al. *Software Components developed by UCL for the SynEx Project & London Demonstrator Site*. UCL, 2001. See <http://www.chime.ucl.ac.uk/HealthI/SynEx/>.
This EC-funded project looks for solutions for federating legacy systems, and uses the CEN ENV 13606 model as a Reference Object Model, and a “dictionary” of configuration concepts very similar in purpose and form to the GEHR archetypes.
13. Rector, A. L. *Clinical Terminology: Why Is It So Hard?* Yearbook of Medical Informatics 2001.

14. Rossi-Mori, A. and Consorti, F. *Assembling clinical information*. Note sent to HL7 discussion list, 2001-04-19.
15. Rubin, Ken. *Using a Taxonomy Matrix as a Communications Instrument*. Version 1.01 April 14, 2000

Resources

16. Arden Syntax. See <http://www.cpmc.columbia.edu/arden/>
17. ASBRU/ The Asgaard Project. See <http://smi-web.stanford.edu/projects/asgaard/>
18. CEN TC 251 EHCRA standards. See <http://www.centc251.org/>.
19. OMG Health Domain Taskforce (was Corbamed). See <http://www.omg.org/>
20. GEHR (Good Electronic Health Record). See <http://www.gehr.org>.
21. GEHR (Good European Health Record). See <http://www.chime.ucl.ac.uk/HealthI/GEHR/Deliverables.htm>.
22. GLIF - Guideline Interchange Format - <http://www.glif.org/>.
23. HL7 version 3. See <http://www.hl7.org>
24. ICD (International Classification of Diseases). See <http://www.who.int/whosis/icd10/>.
25. Odyssee Project. See <http://www.nautilus-info.com/odyssee.htm>.
This project from France has developed a system of "fils guides" (guide paths) and constraint definitions which is similar to the archetype concept, and is effectively used for the same purpose in the Odyssee product.
26. ProForma language for decision support. See <http://www.acl.icnet.uk/lab/proforma.html>.
27. SNOMED (Systematized Nomenclature for Medicine). See <http://www.snomed.org/>.

Acknowledgements

Review comments have been provided by:

- Denis J. Protti, Professor, School of Health Information Science, University of Victoria, British Columbia, Canada.
- Jane Curry, Sierra Systems, HL7, Canada
- Dr. Marion Lyver, Smart Systems for Health, Ontario, Canada
- Ken Rubin, EDS.