Medical Informatics Europe 2002
G Surján et al. (eds)
IOS Press 2002

An Architecture for a Virtual Electronic Health Record

H. van der LINDEN, J. TALMON, H. TANGE, G. BOERS, A. HASMAN

Dept. Medical Informatics, Maastricht University, PO Box 616, 6200 MD Maastricht

Abstract. The Healthcare Domain Taskforce of the Object Management Group has specified standards for secure access and retrieval of demographic and medical data. This paper discusses the strengths and weaknesses of an electronic healthcare record that implements these specifications.

1. Introduction

We define a virtual Electronic Health Record (vEHR) to be *"an EHR that contains all, primarily medical, information on a patient, stored in a variety of systems in a variety of locations over a long period of time, secured against illegal access, provided with an audit trail, and presented to the reader as one dossier"*. We will use the term vEHR to focus on the fact that there will probably not be one location, at any period of time, where the full record of one patient is stored.

We present an architecture of a vEHR that can serve as a base architecture for future generation EHRs of various kinds and sizes. We will discuss the strengths and weaknesses of the proposed architecture, taking results from various related projects into account.

2. Background

The PropeR project studies the use of decision support software integrated with an EHR and its effect on the quality of care. The architecture was developed for one of the subprojects (Proper Transmuraal), which will support a team of care providers involved in the rehabilitation of CVA patients in a home care environment.

In home care only conventional communication methods (telephone, fax and paper) are currently used. So the first objective in the project was to implement an EHR.

Literature survey showed there is no EHR that would meet our needs. This made us decide to build our own system, but with the additional requirement that it should conform to internation standards.

3. System Requirements

- It must be capable of processing data from different sources and different disciplines. The team members have different specialties and therefore require different sets of data.

- It should be accessible from different locations, i.e. users should have access at their office and at the patient's home. The system should also be usable from behind a firewall.
- It should be based on middleware and components technology to reach the flexibility and scalability necessary for building the complex systems needed in health care. [1-4]
- It should meet current security standards and privacy regulations.
- It should be open, i.e. modifiable, to be able to integrate decision support software.
- It should be reusable for other medical domains and settings.
- It should be flexible enough to be updated without rewriting major parts.
- It should meet (inter-)national standards and adhere to state-of-the-art techniques.

4. Architecture

One of the standards that meet our requirements are the HDTF specifications of the OMG [5] and the open source implementation OpenEmed by the Los Alamos National Laboratories [6]. These specifications provide a solid base but are not sufficient to build a complete EHR. Most importantly they lack a description of the actual data stored and/or retrieved.

In the following we present an overview of the HDTF specifications and an EHR architecture based on them.

5. HDTF specifications

The OMG formed a Healthcare Domain Taskforce (HDTF) in 1996. One of their goals is *"To improve the quality of care and reduce costs through the use of interoperability technologies throughout the global healthcare community"* [7]. They try to reach this goal through the production and maintenance of "the Healthcare Component-based Reference Architecture". *"The purpose of the Component-based Reference Architecture is to delineate and describe the interfaces and interactions between the various logical components in health care systems. The interactions and interfaces between the components will then serve as a reference against which the issuance of future health care related RFIs and RFPs can be considered."* [7]

The former name of the specifications, CORBAmed, was abandoned to remove the suggestion of a dependence on CORBA. These specifications could also be implemented using other middleware e.g. SOAP. Currently the following specifications and prototypical implementations exist:

- PIDS - Person Identification Service [8]. A specification that addresses the problem of identifying a person labeled with different IDs in various medical systems.
- RAD - Resource Access Definition [9]. A specification of an interface to a server that handles the authorization profiles and provides security on a detailed level.
- COAS - Clinical Observation Access Service [10]. This specification provides a generic interface to medical data stored in a variety of systems.
- TQS - Terminology Query Service [11]. This specification provides an interface to terminology servers and mapping information to map one terminology to another.

6. Strengths

- HDTF specifications build upon the advantages of CORBA. It specifies services for health related problems with the same properties as the CORBA specification: scalability, platform independency, language independency, stability and component-based. Using the security services of CORBA, makes security transparent.
- One of the goals of the PIDS specification is to remove patient identifying data from the medical data and therefore enhancing security, privacy and anonymous research.
- The COAS specification allows for data access with different granularity. A resident on the ICU ward can be interested in all data concerning a patient gathered in the ICU, while the patient's GP is only interested in the outcome of the treatment. The query functionality of COAS is flexible enough to handle these different requests.
- Another advantage of COAS is the feature to store and retrieve context-rich information. Not only are values supplied with units, valid ranges and other information that allows correct interpretation of the value, but every observation can be supplied with qualifiers that state relations to other values (e.g. 'pertains to problem X', 'result of test Y').

7. Weaknesses

- There might be a performance penalty involved that could make this setup unsuitable for time critical environments like ICU wards and Emergency rooms. This is especially important when used with continuous monitoring devices. Other projects [3] have demonstrated that an agent-based solution to this problem in an ICU ward, using CORBA as the middleware, is effective. Technological developments have taken away much of this disadvantage with the release of real-time CORBA [12]. From there it is only a small step to integrate continuous data monitoring into an HDTF based EHR.
- COAS does not specify the actual data structure, i.e. there is no predefined COAS structure for a lab test or a blood pressure. This allows for flexibility (every kind of structure can be composed), but also allows variations in the construction of similar information. This problem can be solved by the use of an information structure definition.

8. Information structure definition

An information structure definition (ISD) is a structured machine-readable definition of context-rich medical data. This ISD defines all available medical data and their structure and constraints (i.e. a blood pressure is represented by two values, one is smaller than the other) as well as the context of the medical data (i.e. blood pressure is found under "vital signs").

This model has a loose coupling with the other components of the EHR. It is therefore possible to build ISDs complying with various standards and approaches such as CEN/TC251 13606 [13], HL7 RIM [14] and GEHR archetypes [15]. Substituting one ISD with another would make the entire system compliant with one standard or the other.

The ISD is connected to the system through a TQS interface, which provides a uniform query interface to it. TQS can handle various versions of the ISD. Content management has to be performed with an ISD repository editor.

We want to build this ISD based on a full record structure (e.g. Synapses SynOD/ SynOM [16, 17] or CEN/TC 251 13606) complemented with GEHR Archetypes. This should provide a generic meta data model to describe the structure of medical data.

9. The full picture

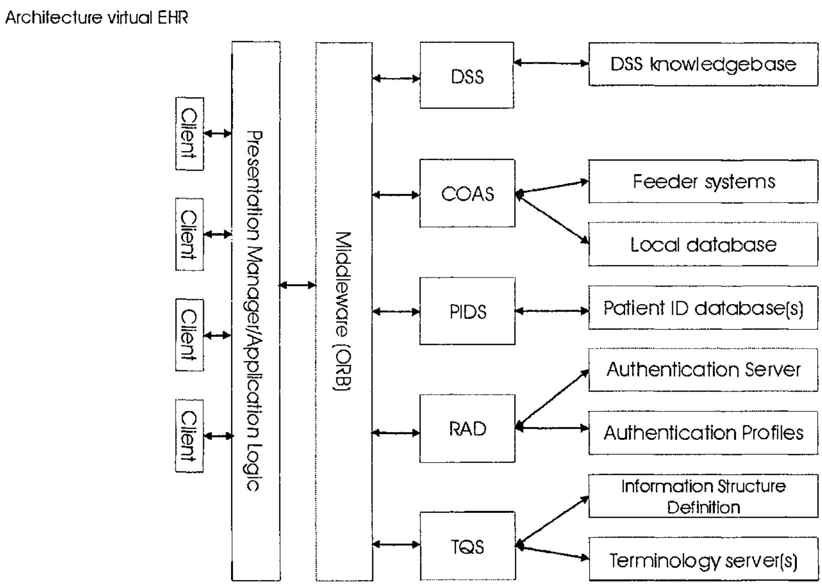

Fig. 1: Architecture of vEHR, explanation in par. 4.5

Figure 1 gives an overview of how the various components of the vEHR operate together. Looking at the individual components we see the following picture (from right to left):

- The rightmost column shows the various "databases" used. Database is used here as a generic term to describe a collection of type information that can be accessed and edited (where relevant and allowed). E.g. a feeder system, a system external to this EHR, is considered to be a database, while it might actually be a complete HIS.
- The column left of the database column contains various services that provide standardized access to their own database. These services communicate to each other and to the presentation layer through an ORB.
  - DSS – the decision support module
  - COAS is the interface to systems containing relevant medical data.
  - A PIDS server handles all requests for a patient ID from it's own database to provide enhanced security and privacy.
  - A highly secured authentication server interfaced with RAD implements the necessary authorization profiles. [2]
  - TQS provides interfaces:
    - To the various terminology servers and the information to map one terminology to another;

  - To the Information Structure Definition. An ISD specifies various views on the medical data as well as the structure of the data items.
- As middleware we use a CORBA ORB, although it is not required to be CORBA.
- A presentation layer handles the requests on behalf of the clients and returns the results in a format appropriate for the client.
- Clients are web or application based, whatever is called for.

Although it might not be obvious from the picture this EHR not only integrates the information of various systems for viewing, but also supports a standardized way of storing new information in databases for which the user is authorized to update the data.

10. Current state and future plans

At this moment we have a rudimentary client that can connect to PIDS and COAS servers we have up and running. Our next goal is to add the support for the ISD which can then be used to implement all necessary medical data.

11. Summary

New generation EHRs should be based on an architecture flexible enough to last several generations and allowing for improvements due to experience gained, but without starting from scratch.

We presented an architecture for a vEHR that wil allow incremental development. Due to the component-based approach temporary shortcuts can be used for functionalities that are either not fully understood or are less relevant for the time being. These shortcuts can later be replaced by better solutions for the issues at hand.

References

All references to URLs are last accessed 24 April 2002.

[1] B. Blobel, Application of the component paradigm for analysis and design of advanced health system architectures, Int J Med Inf 60 (2000) 281-301.

[2] B. Blobel, F. Roger-France, A Systematic approach for Secure Health Information Systems, Int J Med Inf 62 (2001) 51-78.

[3] G.C. D.G. Katehakis, M. Tsiknakis, S. C. Orphanoudakis,, A Distributed, Agent-Based Architecture for the Acquisition, Management, Archiving and Display of Real-Time Monitoring Data in the Intensive Care Unit, FORTH-ICS/TR-261, Heraklion October 1999.

[4] D.D.W. Forslund, Building a Virtual Medical Record Capability with Low-cost, Open-source, Open-architecture Components, www.pharmabriefing.com/businessbriefing/pdf/GHealthcare2001/Book%20Section/Forslund.pdf

[5] Healthcare DTF, http://www.omg.org/healthcare

[6] OpenEmed, http://www.openemed.org

[7] W. Ruh, T.C. Culpepper, Healthcare DTF Overview/Tutorial, ftp://ftp.omg.org/pub/docs/health/01-04-04.pdf

[8] Person Identification Service (PIDS) Specification, Version 1.1, http://cgi.omg.org/docs/formal/01-04-04.pdf

[9] Resource Access Decision Facility Specification, Version 1.0, http://cgi.omg.org/docs/formal/01-04-01.pdf

[10] Clinical Observations Access Service Specification, Version 1.0, http://cgi.omg.org/docs/formal/01-04-06.pdf

[11] Lexicon Query Service Specification, Version 1.0, http://cgi.omg.org/docs/formal/00-06-31.pdf

[12] Common Object Request Broker Architecture (CORBA), Real-Time CORBA, v2.6, http://cgi.omg.org/docs/formal/01-12-28.pdf

[13] prENV 13606: Health informatics - Electronic healthcare record communication - Part 1 - 4, http://www.centc251.org/TCMeet/doclist/TCdoc99/N99-040.pdf, N99-041.pdf, N99-042.pdf and N99-043.pdf
[14] HL7 Reference Information Model, Version 1.12, http://www.hl7.org/Library/data-model/RIM/C30112/rim0112l.zip
[15] T. Beale, Archetypes Constraint-based Domain Models for Future-proof Information Systems, Revision 2.2.1, http://www.deepthought.com.au/it/archetypes/archetypes.pdf
[16] J. Grimson, W. Grimson, D. Berry, G. Stephens, E. Felton, D. Kalra, P. Toussaint, O.W. Weier, A CORBA-Based Integration of Distributed Electronic Healthcare Records Using the Synapses Approach, IEEE Transactions on Information Technology in Biomedicine 2 (1998) 124-138.
[17] W. Grimson, D. Berry, J. Grimson, G. Stephens, E. Felton, P. Given, R. O'Moore, Federated healthcare record server - the Synapses paradigm, Int J Med Inf 52 (1998) 3-27.