

An Environment for the Creation of an Integrated Electronic Health Record in HYGEIAnet, the Regional Health Telematics Network of Crete*

Dimitrios G. Katehakis¹, Panayiotis Lelis¹, Eleni Karabela¹, Manolis Tsiknakis¹,
Stelios C. Orphanoudakis^{1,2}

¹ Center of Medical Informatics and Health Telematics Applications (CMI-HTA), Institute of Computer Science (ICS), Foundation for Research and Technology – Hellas (FORTH), P.O. Box 1385, GR 711 10, Heraklion, Crete, Greece, Tel: +30 (81) 391600, fax: +30 (81) 391601, E-mail: {katehaki, lelis, karabela, tsiknaki, orphanou}@ics.forth.gr, URL: <http://www.ics.forth.gr>.

² Department of Computer Science, University of Crete, P.O. Box 2208, GR 714 09, Heraklion, Crete, Greece.

Abstract

National and international healthcare networks are increasingly used to facilitate the sharing of health-related information and meet the healthcare needs of an increasingly mobile population. In this context, the vision of the Integrated Electronic Health Record (I-EHR) enables autonomous healthcare facilities to operate in a cooperative working environment and the delivery of seamless care. The Patient Clinical Data Directory (PCDD), as a central element of the overall architecture is a middleware component that provides clinical information on the distributed EHR segments maintained by autonomous information systems in a regional health-telematics network such as HYGEIAnet, the regional health-telematics network of the island of Crete. The PCDD maintains a distributed registry of feeder systems, patient key demographics, and references to clinical objects (medical encounters) inside patient EHR segments. Feeder systems, supporting diverse access methods ranging from human-mediated to CORBA object references and HTTP URLs, are the clinical information systems that qualify as PCDD information sources. The PCDD middleware enables an I-EHR service that provides encounter-centered views of patient clinical data.

The PCDD has been designed in the frame of the HYGEIAnet Reference Architecture (HRA), which provides a generic and evolving frame of reference for the reuse of services, components, and interfaces. This infrastructure can be utilized for the extraction of primary information and appropriately defined epidemiological indices. The technological approach followed in the currently installed pilot implementation includes CORBA interfaces (for data acquisition, patient identification, and semantic mapping), X.500/ LDAP (for security services, naming services, user profiles, patient clinical information, and healthcare resources), and data mining techniques (for the extraction of healthcare indicators). Dedicated gateways (e.g. SQL/ ODBC-LDAP) have been implemented for scheduled directory updates, and XML is used for maintaining collected clinical information in a consistent way.

Keywords: Healthcare Information Infrastructure, Component-Based Architecture, Adaptive User Interface, Federation of Clinical Information Systems, Population Screening.

* Proceedings of the 16th Annual Towards an Electronic Patient Record Conference and Exhibition (TEPR 2000), Your Connection to Electronic Healthcare, San Francisco, California, USA, May 9-11, 2000, Vol. 1, pp. 89-98. (On-line version: <http://www.ics.forth.gr/~katehaki/publications/tepr2000.pdf>)

1. Introduction

When a patient moves inside a healthcare system, several segments of the personal health care record are left behind in electronic or other form. All these segments together form what is called the patient's Healthcare Record (HCR). Although each healthcare facility is autonomous and devoted to the delivery of a particular set of services, continuity of care requires that different healthcare facilities, offering complementary services or different levels of expertise, exchange relevant patient data, and operate in a cooperative working environment. In this environment, diverse user groups require secure, customizable access and sharing of information residing at geographically distributed autonomous information systems. This sharing of information resources is generally accepted as the key to substantial improvements in productivity and better quality of service. Therefore, the HCR is the cornerstone for the provision of the continuity of care, and is the point of reference for any exchange of information related to the patient under consideration. However, when someone tries to implement it, a number of practical problems arise: (a) clinical data are composite information that comes from various sources and exists in many different forms. For example clinical data produced in the intensive care, are completely different from those needed in a pediatrics surgery clinic or the ones produced by a PACS, (b) the exchange of clinical data between heterogeneous information systems is not feasible and in the majority of the cases, this is also true for the exchange of data among homogeneous information systems too, (c) precise medical information for a patient is extremely difficult to be obtained, and this is because even when recorded, clinical information is sometimes inadequate or even inaccurate or anonymous, (d) in addition, the perfect world where all information systems are continuously interconnected and accessible over fast networks is far from being reality, even for the most developed countries of the world. And in most cases this is not due to lack of infrastructure but rather due to certain restriction policies related to accessing personalized information. Nevertheless, under certain circumstances even the knowledge of the existence of non-directly-accessible HCR segments can be of vital importance.

The goal of this work is to demonstrate the technological approach that will enable knowledge discovery in a regional health-telematics network such as HYGEIANet. This can be achieved by defining all the middleware enabling services that are required for the provision of basic services in a consistent, reusable and expandable way. This can be done by developing the necessary interfaces for accessing clinical information, and by implementing a number of interfaces adapted to the needs of a number of healthcare-related domains. A promising approach to this integration problem is to gain control of the organization's information resources at a meta-data level, while allowing autonomy of individual systems at the data instance level. However, achieving integration at the semantic level is a challenging problem mainly because the logic, knowledge, and data structures used in various systems are complex and often incompatible. Thus, a realistic solution should hide heterogeneity at the top level, while making the individual sources of information appear to end users as a large collection of objects that behave uniformly.

The rest of the paper is structured as follows: Section 2 (Background Work) presents an overview of already existing technologies and work that is considered reference for the healthcare domain. HYGEIANet, the regional health telematics network of Crete is introduced in Section 3 (HYGEIANet) and its Healthcare Information Infrastructure (HII) software components are described in Section 4 (Components of the Architecture). Section 5 (I-EHR Services) describes how the overall infrastructure of HYGEIANet is utilized in order to provide end-user services. The currently available implementation status is presented in Section 6 (Status of the Work) together with conclusions and considerations regarding future work.

2. Background Work

The concept of the Patient Meta-Record (PMR), as presented in [Tsiknakis96] could turn a large, heterogeneous and distributed set of Computerized Patient Record (CPR) components into a seemingly integrated and homogeneous healthcare record and permitted the integration of EHR segments in order to assist personalized utilization of the EHR by the respective functional units or healthcare providers. The PMR managed, at a single logical point, references to all of the physical information related to a patient, therefore, among it was an indexing system, providing access to all stored data for a particular patient, utilizing naming services, directory ser-

vices, and the Health Reference Data Model (HRDM) to deliver a semantically consistent view of the geographically distributed healthcare record segments [Leisch97].

Apart from these initial efforts to approach the issue of highly distributed and multi-format segments of personal information, several other important attempts have been made worldwide both as part of serious standardization effort or by means of EU financed projects. For example, in the Synapses Project [Hurlen98] the objective was to facilitate the sharing of health records among information system consistently, simply, and securely. The approach followed was that of middleware servers that enabled healthcare professionals to access information from feeder systems. This approach bared some similarity with the PMR approach in that in its deployment it involved the definition of basic query types and the clinical constructs they may be involved. However, the deployment cost of such a solution was really high. GEHR [GEHR95] provided specifications for an exchange format, and supported the idea of systems exchanging healthcare data in a standard exchange format. Lately, the advance of XML has swept away the GEHR approach. According to the XML approach, different systems exchange data using documented Document Type Definitions (DTDs).

In the TeleMed project [Kilman97], the virtual patient record forms the basis for a CORBA-based collaboration environment in which multiple physicians, and ultimately the patient, can engage in interactive electronic discussions. Multiple physicians at remote locations can simultaneously view, edit, and annotate multimedia patient data. Each physician can see the data another physician has entered as well as monitor some of the other physician's interactions with various user interface windows. In this way, TeleMed may be used to support referral and teleconsultation sessions. TeleMed lately employs CORBA technology by adopting fully all available OMG specifications [OMG].

CORBAmed as the Healthcare Domain Task Force (DTF) of the Object Management Group (OMG) defines standardized object-oriented interfaces that serve to promote interoperability between a variety of platforms, operating systems, languages and applications, in the healthcare domain. Well-documented interfaces expressed in the Interface Definition Language (IDL) associated with the integration framework of CORBA, provide basic support for interoperability among computer systems. This is particularly important in large hospitals, where many kinds of different computers have been installed and cannot be changed.

In Europe, CEN/ TC 251 [CEN] develops standards that enable compatibility and interoperability between independent systems in healthcare, and has gained strong support by the European Commission, healthcare authorities, suppliers of Information and Communications Technology (ICT) solutions, and users as well. Such standards are needed for the efficient use of information and communications technology allowing new ways of interaction between different actors promoting cost-effective procedures giving better quality of care for the patients. In the United States, ISO/ TC215 also works on standardization in the field of information for health, and health ICT to achieve compatibility and interoperability between independent systems.

It is evident that a number of projects around the world address the problem of locating, indexing, and accessing the distributed segments of a patient's EHR. The level of integration supported by each of these efforts varies, and so is the complexity of the effort. In essence, there is a trade-off between the diversity of clinical objects stored in centralized repositories and the generality and expressiveness of models supporting it. The more information placed inside centralized repositories, the richer query model is supported. This, however, limits the range of clinical information systems that can be federated in a non-trivial way.

3. HYGEIAnet

HYGEIAnet (Figure 1) is currently the point of reference and also a long-term objective of the Center for Medical Informatics and Health Telematics Applications (CMI-HTA) of the Institute of Computer Science (ICS), Foundation for Research and Technology - Hellas (FORTH). It is the integrated health telematics network that is under development in Crete, which is expected to serve as a model for national and transnational healthcare networks. The principles followed during HYGEIAnet's design phase, came from currently existing status in the regional healthcare system that is applied for Greece and aims towards servicing the local population, no matter whether they are patients, healthcare professionals, researchers, or managers.

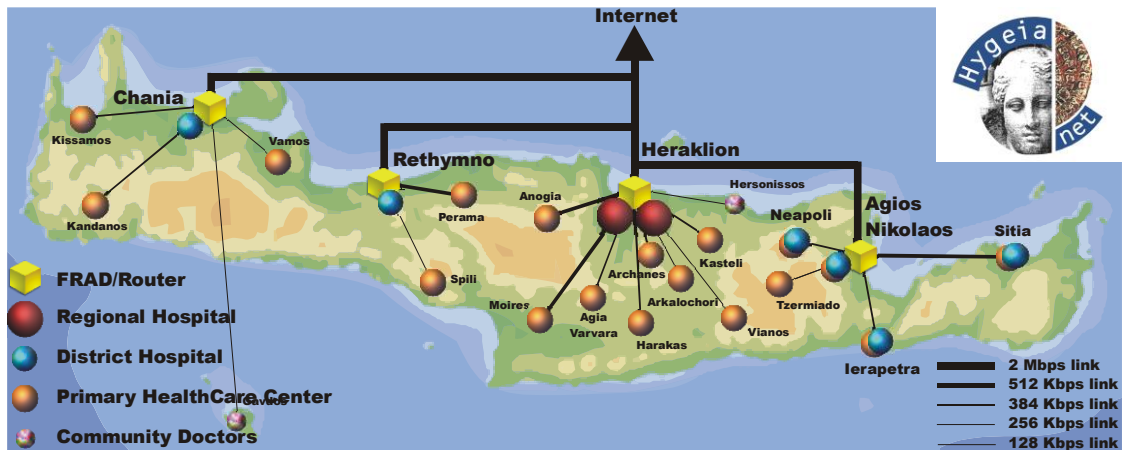


Figure 1: HYGEIAnet, the Regional Health Telematics Network of Crete.

In the framework of an environment that provides advanced protection of personal life, the infrastructure built has as an ultimate target the provision of quality services to be offered to patients, as well as the development of integrated healthcare related services. It is expected to provide the set of information and services that is considered to be necessary, for the continuous development of feasible services and the better understanding of the status of the population health.

Users are primarily interested in information processing applications, which they may own or gain access to as end-users via communications networks. These services are ‘enabled’ by other underlying, transparent services provided by information and network service providers. Applications and enabling services, in turn, employ certain information processing services and systems for data transport, which can be geographically distributed throughout the region. Thus, the reference architecture model refers to three basic components: applications, enabling (or middleware) services, and physical infrastructure.

The components of the information infrastructure in the health-telematics network of Crete are part of an execution architecture that fits within the reference architecture model presented in [Tsiknakis97] and [Orphanoudakis98]. Throughout the region, clinical information systems operate as autonomous networked information sources. Appropriate extensions or modifications to these information systems may be required so that they exploit middleware components and become part of a cooperating information environment that enables information integration. Healthcare-specific middleware services support applications with services related to activities of the healthcare domain. These middleware services are associated with information and procedures that are considered to be of paramount importance to the correct functioning of the healthcare organization as a whole. These services include resource, authorization, naming, messaging, terminology, semantic mapping, and other meta-data services, as well as services for the management of medical acts, patient identification, and clinical data location services.

4. Components of the Architecture

Based on the work of on-going CORBAMED activities as well as the existing real-world situation in Greece, CMI-HTA has managed to identify the most important middleware that are currently required in order to support I-EHR. Primarily it is essential for a regional network to know its own resources. Therefore a registry where all the available regional healthcare resources can be found must exist. Regional resource services provide information on the availability of physical resources such as hospital departments, diagnostic modalities, mobile emergency units and their characteristics. Based on that, a registry for people having parts of their personal clinical information needs to be created, and on top of that all the relevant clinical information links need to be mapped on a clinical encounter registry. The Patient Clinical Data Directory (PCDD) does exactly this work. It is a middleware enabling service that indices the EHR segments of patients health records and facili-

tates global access to their clinical observations in feeder systems, which are heterogeneous, autonomous, and geographically distributed clinical information systems. The PCDD maintains a distributed registry of feeder systems, patient key demographics, their EHR segments, clinical meta-data, and references to clinical objects. The data model of the PCDD is based on the Subjective Objective Assessment Plan (SOAP) model [Bainbridge96] that originated from the primary healthcare domain. Further access to detailed information on particular healthcare encounters is delivered on a case-by-case basis.

For the PCDD to be able to offer secure services it needs to be able to authenticate each client's principal identity, role and sensitivity, and to transmit information confidentially and with integrity. It is evident that the existence of a trust infrastructure needs to evolve in parallel as part of the HYGEIANet. This way security and confidentiality services can be based on a regional certification authority, which will provide digital certificates to healthcare facilities and human resources. The combination of digital signatures for authentication, public key cryptography for recipient authentication, and Secure Socket Layer (SSL) protocols for secure data-transfer, provide the necessary technological framework for secure communication of healthcare related information across the Internet. This way support for open standards can be achieved and provide end users with confidence that they will be able to communicate within heterogeneous computing environments.

At the moment, access control is achieved by means of user profile information. In this context any patient is able to have complete access to personal information. A physician has access to all information that has been provided by him, as well as to his/ her referral data. In addition patients are able to grant and restrict access to their personal information. Of course there exists the clear danger that very tight security constraints are capable of making ineffective the use any I-EHR environment. Hence more emphasis is paid on auditing "who accesses what type of information at what time," instead of trying to enforce very tight security constraints. This is one of the reasons for the very tight coupling between user profiles and auditing services as a more generic set of services that enable user information collection, capable of modifying/ customizing user preferences on the way of both restricting and presenting information.

5. I-EHR Services

Challenges in the area of human-computer interaction have always focused on the design of User Interfaces (UIs) accessible by different user groups and the propagation of existing guidelines into user-adaptable UI implementations [Stephanidis97]. This has contributed to an ever-increasing number of computer users, characterized by diverse abilities, requirements, and preferences. In addition, context has become more complex and knowledge demanding. The fact that, in the majority of cases, domain experts use computers in a completely different manner than simple users makes necessary the existence of a number of adaptive user profiles to individual user needs requirements. In addition, due to the fact that users do not share the same level of experience, interaction techniques should be suitable for all the range of users including both computer-literate and computer-illiterate.

At present, the objective of the I-EHR environment is to deliver an encounter-centered view of the patient's EHR. It utilizes the existing CORBA interfaces to provide a consistent way to locate, access and transmit secure information about a patient's EHR segments. The I-EHR is a front-end to an EHR indexing service, managed by PCDD which indices both structured and unstructured information that is provided by co-operating information systems, without imposing any constraint on their internal operation or their interface, beyond the medical encounter level. The Usability Context Analysis (UCA) methodology that has been followed [Thomas96] helped collecting specifications on the number and special characteristics end users have, as well as the tasks they need to carry out. Awareness of contextual factors is important for the development of appropriate and usable human-computer interfaces. Figure 2 provides a view of how a user perceives and can adapt the presentation offered to him by the currently existing user interface, which has been implemented entirely by means of the Java programming language, and utilizes the existing middleware infrastructure, described in the previous section. Apart from the lifeline view of all the available encounters of the patient, two more views are currently supported: a per-clinical system view of the encounters' history, as well as the traditional tabular view of old

generation GUIs. Certain filters are supported for the better isolation and location of information based on time and examination type criteria.

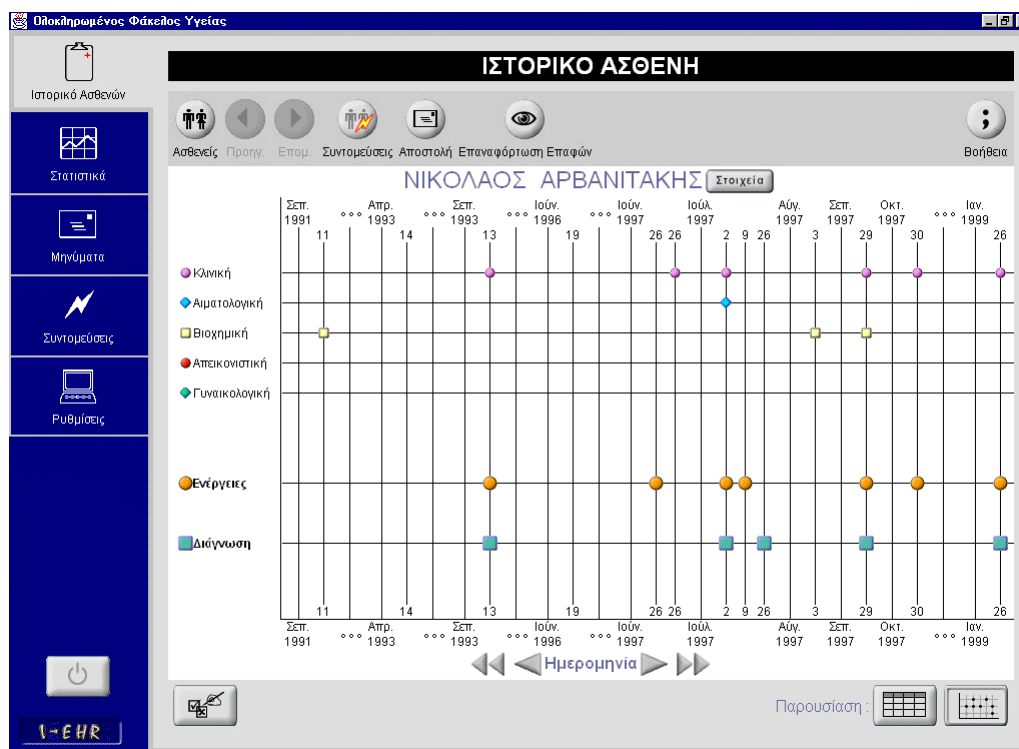


Figure 2: Human Computer Interaction with the I-EHR environment.

A further extension of the architecture allows for the execution of more dynamic queries by providing as parameters the geographic region, date range, and target citizen group (age, gender, etc) along with the type of information to collect. For the information collection engine to proceed with the data collection process, it needs to cooperate with all the available and existing services through stable interfaces. For example, in order to locate all the relevant information sources, it needs to communicate with the healthcare resource registry, and as a result to get all installed clinical information systems that reside in the geographical area under consideration. Then, the output of the previous step is filtered according to the type of information required, by means of a Terminology Query Service (TQS) implementation that can support this kind of functionality. Then, the information collection engine tries to trace all clinical encounters, which matches the given parameters (date range, citizen characteristics and system installations), by using PCDD. References to recorded data are obtained, and are used to retrieve actual data through Clinical Observation Access Service (COAS) implementations. Finally, the engine communicates with TQS to transform data to the desired format, and records output to the Data Repository as an XML document instance (by means of a predefined DTD). During the process, the engine uses Security Services to grant access to resources, User Profile Services to record the most important queries provided by the individual user, and Auditing Services to record all interaction with the services and components of the HII.

Beyond the process of accessing patient clinical information, dispersed throughout the healthcare facilities of the island of Crete, users should be gradually allowed to pose global queries into the available information space. This is extremely important since it will also allow healthcare professionals to access epidemiological information that is currently hidden within the individual clinical information systems. Work under progress is currently focused on the development of the necessary elements of technology that will co-operate towards the envisioned process of knowledge discovery (Figure 3).

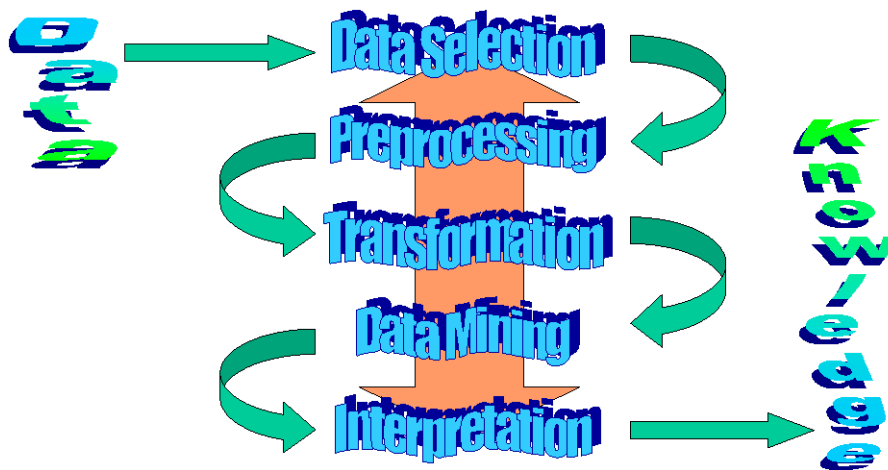


Figure 3: Knowledge Discovery Process.

The deployment of the HII environment in Crete facilitates for the design, development, and evaluation of an architecture that can exploit common components of the regional infrastructure to retrieve health information, homogenize, and make them available for knowledge discovery. This infrastructure can be utilized for the extraction of primary information and appropriately defined epidemiological indices. The technological approach followed in the currently installed pilot implementation includes dedicated gateways (e.g. SQL/ ODBC-LDAP), CORBA interfaces (for data acquisition, patient identification, and semantic mapping), and XML for maintaining collected clinical information in a structured way.

6. Status of the Work - Conclusions

The public healthcare model for the region of Crete consists of one University Hospital, two prefecture hospitals, five general hospitals, fifteen healthcare centers, and one hundred and seventeen community doctors, without counting a number of Health Emergency Units, as well as the private sector. Figure 4 shows the currently available physical infrastructure that is required for providing this initial I-EHR service. A directory server supporting a number of X.500 services, web, and CORBA servers, the individual clinical information systems, as well as the auditing server, which is at the moment accessed through Java RMI constitute the installed infrastructure. Currently this infrastructure provides with access to almost the entirely Primary Health Care Center (PHCC) information space, to a number of selected clinical information systems (for pathology, cardiology, pediatrics and lab exams), as well as to IASO, the information system of the Health Emergency Coordination Center (HECC) in Heraklion [Leisch99].

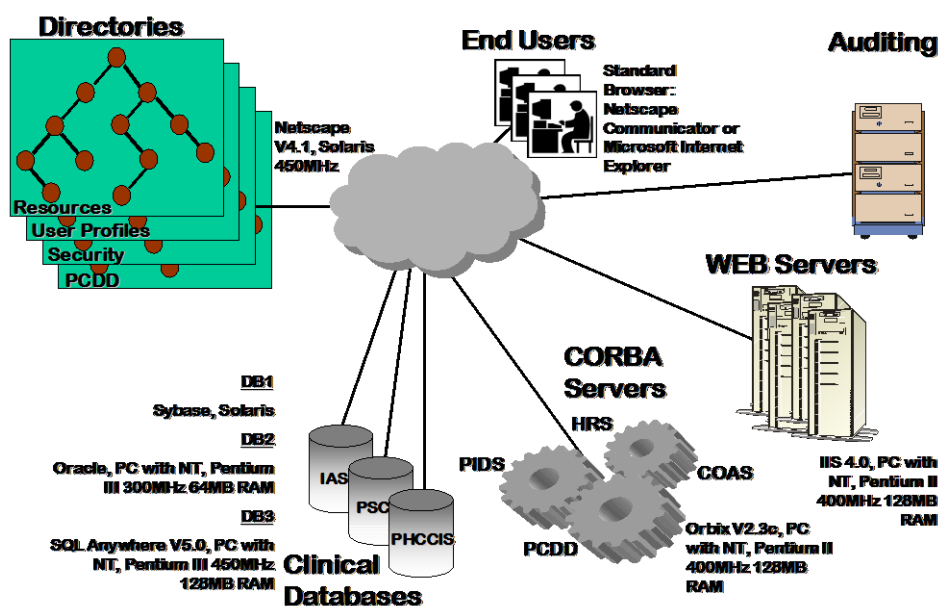


Figure 4: Schematic Description of the available infrastructure for providing mid-ware-enabling services on the island of Crete.

The I-EHR, as used in the current context is “virtual” in the sense that it provides a uniform view of data (meta-data) possibly configured to work differently at different locations. Due to the fact that users seek selective information, following specific paths depending on their personal preferences, it is expected that the I-EHR concept will eventually lead to a uniform applications and services environment. Since electronic records can provide much easier navigational facilities, navigational issues will become even more important in the future, mainly because of the end-user requirements to have similar interfaces in terms of look and feel. In addition any such interface should be able to work both on secure and insecure environments. This is the main reason why the existence of cached information is not desirable in any sort of UI implementation.

From all the above it is quite evident that although clinical systems may be designed with particular health professionals in mind, any EHR architecture should meet the needs of all health professionals. Access and use of the primary medical data must be maintained by the individual healthcare information systems in different authorization domains, which is subject to strict confidentiality policies. The enforcement of these policies necessitates the synergy of healthcare-specific services, such as authorization services, with generic services such as security and certification services. Flexible domain-specific authorization services based on generic security and certification services, conformant to the evolving legislation regarding telemedicine, when enforced, they will promote the acceptance of health-telemedicine services.

HYGEIAnet builds on a regional healthcare information infrastructure to improve the quality and accessibility of health care and to enable the delivery of integrated health care services. It provides the information and services that are the foundation for accountability, continuous improvement to health care and better understanding of the determinants of the health of the population. The open issues that remain to be resolved are many, but experience has shown that any carefully planned and designed component-based architecture can support advanced health telemedicine services of a wide range, as long as open interfaces are available, and currently both CORBAMED and the recently introduced Microsoft Health Users Group have gained a quite wide range of acceptance. High priority issues for CMI-HTA in the near future are: common terminology so that information can be represented in a semantically unified way, the establishment of a more flexible security framework that will allow the co-existence of various user profiles having different authorization levels, and knowledge extraction over the population of the island of Crete, in the form of healthcare indicators.

7. References

- [Bainbridge96] M. Bainbridge, P. Salmon, A. Rappaport, G. Hayes, J. Williams, S. Teasdale, "**The Problem Oriented Medical Record - just a little more structure to help the world go round?**". Proceedings of the 1996 Annual Conference of PHCSG, Downing College, Cambridge, UK, September 1996 (<http://www.ncl.ac.uk/~nphcare/PHCSG/conference/camb96/mikey.htm>).
- [CEN] European Committee for Standardization - Technical Committee for Health Informatics (<http://www.centc251.org/>)
- [GEHR95] Good European Health Record Project (EC AIM GEHR A2014), "**GEHR Architecture**", deliverables 19, 20, 24, 30.6.95.
- [Hurlen98] P. Hurlen, K. Skifjeld and E. P. Anderson, "**The basic principles of the Synapses FHCR server**", International Journal of Medical Informatics, vol. 1998(52), 123-132, 1998.
- [Kilman97] D. G. Kilman, D. W. Forslund, "**An International Collaboratory Based on Virtual Patient Records**", Communications of the ACM, Vol. 40, No. 8, pp. 111-117, August 1997 (<http://www.acl.lanl.gov/TeleMed/Papers/p110-kilman.pdf>).
- [Leisch97] E. Leisch, S. Sartzetakis, M. Tsiknakis, and S.C. Orphanoudakis, "**A Framework for the Integration of Distributed Autonomous Healthcare Information Systems**", Medical Informatics, Special Issue, vol. 22(4), 325-335, 1997 (http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/papers/1997/mi97int/mi97int.html).
- [Leisch99] E. Leisch and S.C. Orphanoudakis, "**The HECTOR Pilot System of Crete as Part of a Regional Health Telematics Network**", 1st International Conference on Health Emergency Telematics, Seville, Spain, March 15 - 17, 1999 (http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/papers/1999/hector_pilot_system/hector_pilot_system.html).
- [OMG] The Object Management Group, (<http://www.omg.org/>).
- [Orphanoudakis98] S.C. Orphanoudakis, "**Integrated Telemedicine Networks and Added-Value Services**", Proc. VIII Mediterranean Conference on Medical and Biological Engineering and Computing (MEDICON'98), Lemesos, Cyprus, June 14-17, 1998.
- [Stephanidis97] C. Stephanidis, D. Akoumianakis, "**Preference-Based Human Factors Knowledge Repository for Designing User Interfaces**", International Journal of Human-Computer Interaction, Vol. 9, No. 3, pp. 283-318, 1998.
- [Thomas96] C. Thomas, N. Bevan, "**Usability Context Analysis: A Practical Guide Version 4.02**", Crown Copyright 1996.
- [Tsiknakis96] M. Tsiknakis, D.G. Katehakis, and S.C. Orphanoudakis, "**Intelligent Image Management in a Distributed PACS and Telemedicine Environment**", IEEE Communications Magazine, vol. 34(7), pp. 36-45, 1996 (http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/papers/1996/ieec96/ieec96.html).
- [Tsiknakis97] M. Tsiknakis, C.E. Chronaki, S. Kapidakis, C. Nikolaou, and S.C. Orphanoudakis, "**An Integrated Architecture for the Provision of Health Telematic Services based on Digital Library Technologies**", International Journal on Digital Libraries, Special Issue on "Digital Libraries in Medicine", vol. 1(3), 257-277, 1997 (http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/papers/1997/dglib97/dglib97.html).