Application Engineering in Health Care

John Barkley

May 9, 1995

1 Introduction

CSL has undertaken a number of projects which seek to further develop Information Technology (IT) that supports the goal of reducing the cost of health care services. The goal of the CSL Health Care projects is to advance the applicability of IT in the health care environment. The more effective IT is for health care applications, the lower the cost of health care services and the greater the advancement of IT itself.

More specifically, the CSL Health Care projects have the following goals:

- Identify the IT requirements for health care applications.
- Identify IT areas which can be most effectively used in health care applications.
- Promote development of those areas and participate in their development.
- Advance IT in those areas so that IT can better provide solutions to health care problems.
- Improve the synergy between the health care and IT industries as well as between the technologies.

More information about the CSL Health Care projects is available on the World Wide Web (WWW). The Universal Resource Locator (URL) is: *http://hissa.ncsl.nist.gov/rbac/prog/*

Among the basic technologies required to support health care applications are:

- Application Engineering - making the development of reliable applications more efficient.
- Multimedia Information Representation, Storage, and Retrieval - reducing the time to access information and the amount of storage needed to store information.
- Interoperability - enabling the rapid reliable exchange of information electronically.
- Security and Integrity - minimizing the possibility of unauthorized access and modification/destruction of information.
- Human Machine Interface - making the use of IT more accessible.
- Intelligent Agent - automating knowledge discovery and business processes.

This paper[1] describes two of the ongoing projects in Application Engineering. These are "Software Engineering Environments for Distributed Applications in Health Care" and "The Use of Role Based Access Control in Health Care Information Security."

2 Distributed Applications

This project assesses the capabilities of several technologies designed for developing distributed applications with respect to generally accepted requirements for health care application development environments. Included in the technologies to be studied are the Object Management Group's (OMG) Common Object Request Broker Architecture (CORBA), Microsoft's Object Linking and Embedding (OLE), Remote Procedure Call (RPC), and the Protocol Independent Interfaces (PII, the IEEE POSIX.1g draft standard). The output of this project includes a report and a demo illustrating the capabilities of each technology studied.

RPC and PII are the traditional tools used to develop distributed applications. These tools are particularly useful for applications requiring high performance. CORBA is becoming the technological descendent of RPC and provides RPC's basic functionality. Consequently, the demo for this project will not use RPC as a development tool although NFS, the remote file system access method, which was implemented using RPC, may be used. For more information about RPCs, see: *http://hissa.ncsl.nist.gov/nistir/5277/*

CORBA and OLE are examples of more integrated environments. Both CORBA and OLE are now mature and both are continually being enhanced by industry consortia. NIST participates in the OMG meetings.

These technologies for developing distributed applications are compared with regard to health care application requirements. In particular, the following capabilities are compared:

- Ease of use by the developer - How effectively does the technology contribute to the development process?
- Class of applications for which the technology is particularly effective in developing - Is the technology more effective in developing any specific classes of applications?
- Security capabilities - Does the technology provide the capabilities needed to ensure the confidentiality, integrity, and availability of the application's processes and data?
- Protocols utilized - What communications support does the technology provide and how interoperable is that support?
- Performance - What are the inherent performance limitations of the technology and is the technology suitable for real-time applications?

[1] Because of the nature of this report, it is necessary to mention vendors and commercial products. The presence or absence of a particular trade name product does not imply criticism or endorsement by the National Institute of Standards and Technology, nor does it imply that the products identified are necessarily the best available.

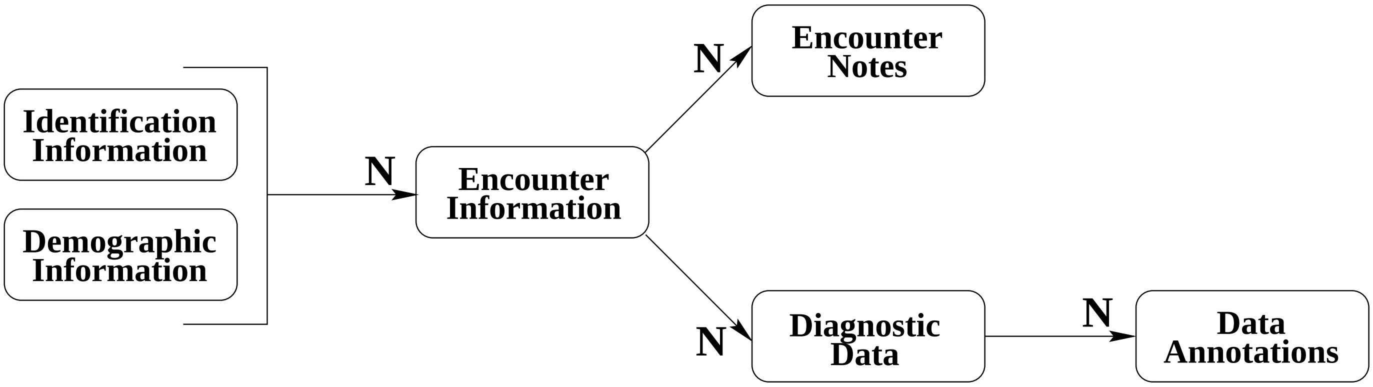

Figure 1: Patient Record Data Base Object entity relationships

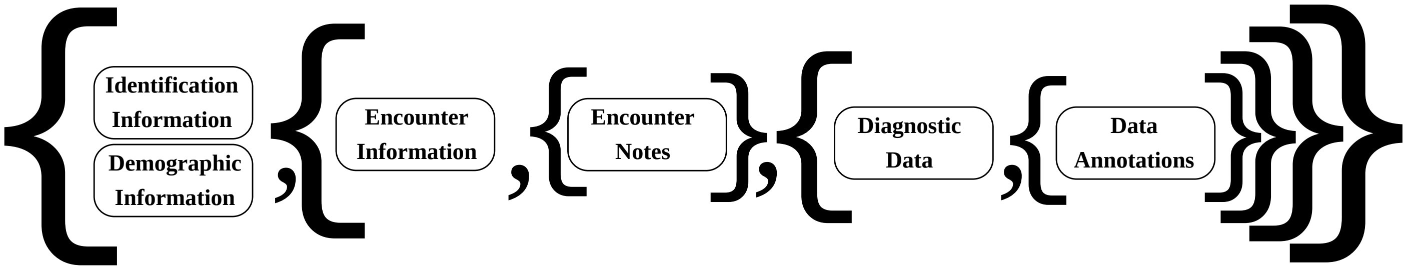

Figure 2: Patient Record Data Base Object set notation

The demo illustrating the capabilities of each technology studied in the project consists of a distributed application for clinical and administrative patient data access. For this demo, a patient record data base object (PRDBO) is defined. This object provides a consistent view of the patient information. The concept is to access patient data through this object whose methods provide a consistent specification for accessing the data. How the data is actually stored is independent of how the object client accesses the data. The methods in the object implementation access the data however and wherever the data is actually stored. CORBA is being used as a means of implementing the PRDBO.

The PRDBO organizes patient information into groups. Figure 1 shows the information groups of the PRDBO and how they relate to each other. The Identification Information Group contains information like name, address, and patient ID. The Demographic Information Group contains information like birth date and sex. The Encounter Information Group contains information like encounter date, physician seen, symptoms, and diagnosis. The Encounter Notes Group contains physician notes on the encounter. The Diagnostic Data Group contains the results of diagnostic procedures (e.g., x-rays) associated with the encounter. The Data Annotations Group contains annotations to the diagnostic data, such as, notations on an x-ray highlighting abnormalities. The Diagnostic Data and Data Annotations Groups usually contain multimedia information such as images and sound.

Pieces of information within a group have a one-to-one relationship to each other. For exam-

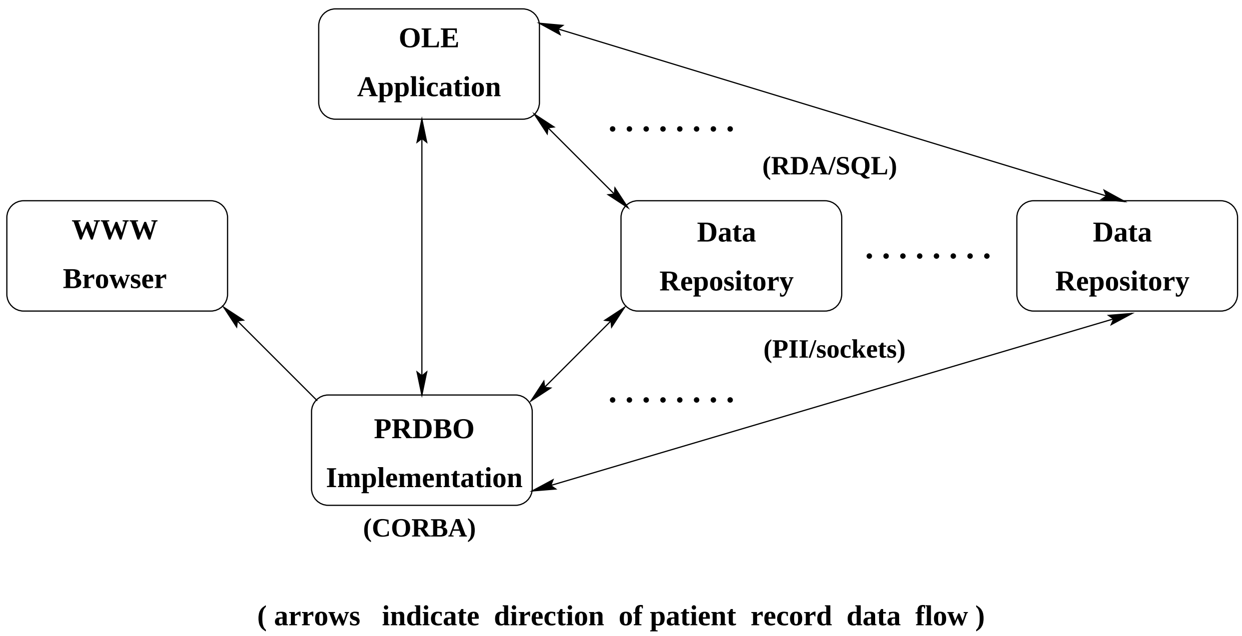

Figure 3: Distributed application architecture

ple, within the Demographic Information Group, each patient has only one birth date and is of only one sex.

The information groups can relate to each other in either a one-to-one relationship or a one-to-N relationship. For example, the Identification Information Group and the Demographic Information Group have a one-to-one relationship. Each patient has only one group of identification and demographic information. However, for each patient, there may be several visits to a physician. Consequently, there may be several Encounter Information Groups associated with each patient.

The information groups may be thought of as elements of sets. Figure 2 illustrates the PRDBO when viewed from this perspective. The PRDBO is a set whose elements are information groups or sets of information groups for each patient. Each element of the PRDBO set has as elements: the Identification Information Group, the Demographic Information Group, and a set of information about each encounter.

Figure 3 shows the architecture of the distributed application. Two clients of the PRDBO are being developed. One illustrates access from within the organization that created the information. This client is being developed using Object Linking and Embedding (OLE) on the PC. The other illustrates access from outside of the organization that created the data. This client is being developed for use with World Wide Web browsers. The arrows indicate the direction of patient information flow. The OLE application is capable of both reading and writing information to the data repositories within the organization which created the information. WWW browsers are capable only of reading information and they provide access to information created within an organization to those external to that organization.

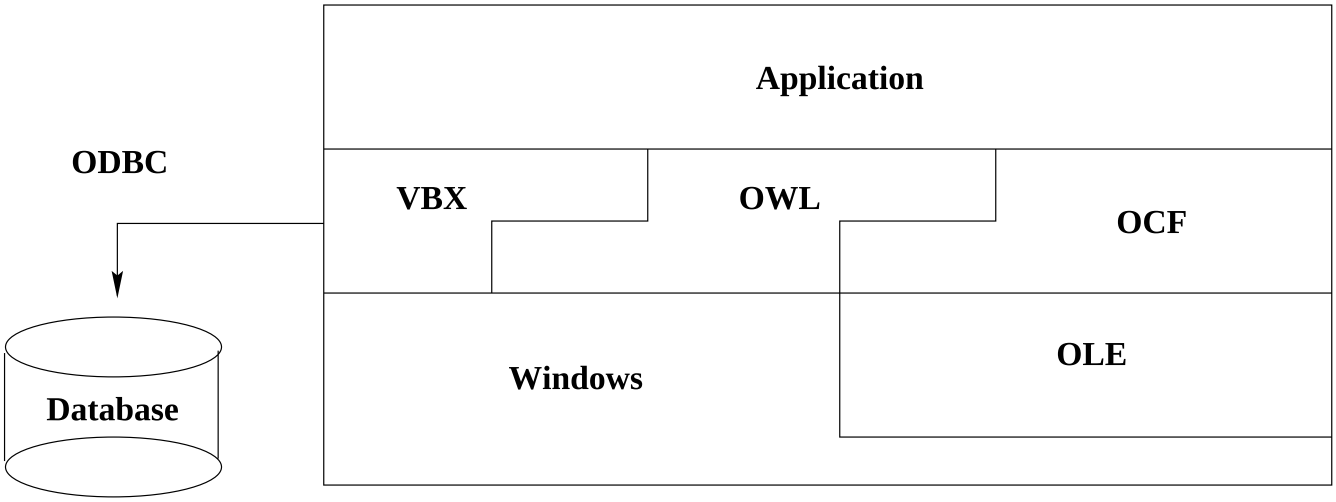

Figure 4: OLE application architecture

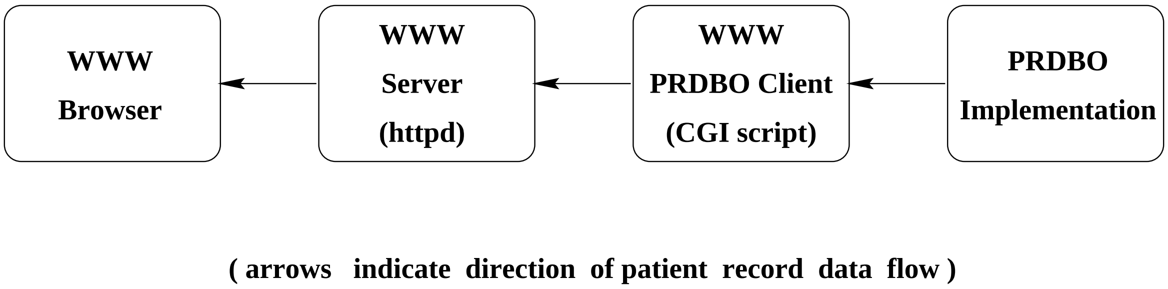

Figure 5: Access to patient information from the World Wide Web

The data repositories contain data suitable for traditional relational databases and multimedia data. The PRDBO Implementation and the OLE client are capable of performing SQL queries on relational databases. Where these databases are remote, the Remote Database Access (RDA) protocol with the SQL specialization (RDA/SQL) is used. For more information see: *http://hissa.ncsl.nist.gov/ ~len/rda_info.html*

Multimedia data is transmitted by means of the sockets interface of the Protocol Independent Interfaces IEEE Standard (PII/sockets). The PII/sockets interface is derived from the Berkeley sockets interface and is useful for transmitting large amounts of data.

Figure 4 shows the architecture of the OLE application. The application is developed using Borland C++ and uses three libraries. The Object Windows Library (OWL) and Object Components Framework (OCF) libraries provide a high level interface to Microsoft Windows and OLE. The Visual Basic Control (VBX) library provides direct access to databases by means of the Object Data Base Connection (ODBC) interface specification. The OLE application also can be a client of the PRDBO to access information from outside the organization and, optionally, from within the organization. As a practical consideration, applications which access information created by an organization from within that organization may forego using the PRDBO.

Figure 5 shows in more detail how patient information is accessed over the World Wide Web (WWW). The arrows indicate the direction of patient information flow. A WWW browser, such as, *mosaic*, connects to the Web Server *httpd*. The Web Server initiates the WWW PRDBO Client as a Common Gateway Interface (CGI) script. The PRDBO Client makes requests to the PRDBO Implementation which accesses the Data Repositories.

3 Role Based Access Control

Within the health care industry, there are continuing problems associated with how to ensure the security and integrity of health care information, in particular, patient information. These problems will only become larger in the future with the increasing automation and integration of health care information.

There are two basic types of access control mechanisms used to protect information from unauthorized access: discretionary access controls (DAC) and mandatory access controls (MAC). Because DAC places the decision of who can access information at the *discretion* of the creator of the information, DAC is not applicable to the majority of health care information. Because MAC requires all those who create, access, and maintain information to follow rules set by administrators, MAC is the kind of access control mechanism required of health care information.

The most commonly used MAC is the multi-level security mechanism used by the Department of Defense (DOD). This is the mechanism which associates information with such labels as *TOP SECRET*, *SECRET*, and *CONFIDENTIAL*. It has become apparent that this type of MAC is not sufficiently flexible for industry use. This type of MAC is also not adequate for the needs of health care.

Role Based Access Control (RBAC) is a MAC which has been developed at NIST to meet the needs of industry. Rather than labeling information, it associates *roles* with each individual who might have a need to access information. Each role defines a specific set of operations that the individual acting in that role may perform. The operations may be broad or very specific, e.g., when a diagnosis is entered into a patient record, the symptoms leading to that diagnosis must also be entered. Once an individual has been properly identified and that identification authenticated, the individual chooses a role that has been assigned and accesses information according to the operations assigned to the role. For more information on RBAC, see:
http://hissa.ncsl.nist.gov/rbac/

This project determines the applicability of RBAC to health care information. While it is generally accepted that RBAC is more suited to health care than others, the question remains as to whether RBAC meets all of the requirements for the security of health care information. Moreover, there are several variations on the RBAC model and there is the question of which variations are most suitable for health care information.

In order to illustrate the usefulness of RBAC to health care, this project also produces a demonstration of the use of RBAC with patient records. The demonstration suggests different roles that are appropriate with patient records and defines sample operations associated with those roles.

Role ID	Access
Patient	All information for the patient
Doctor	All information
Voluntary Caring Agency	name, address, clinical data
Researcher	age, sex, clinical data
Epidemiologist	age, sex, clinical data
Environmental Health Officer	name, ID, address
Organizational Staff	name and ID

Table 1: Simplified Version of the UK RBAC Policy

A sample RBAC policy related to clinical and administrative patient data has been identified. This draft specification, entitled "A Strategy for Security of the Clinical Record and its Transfer" by Dr. Antony Griew of the Institute for Health Informatics in the UK, represents some degree of consensus on a policy for patient information access. The UK policy is RBAC with the addition of the capability of labeling information that is only available to the patient and the doctor. It specifies roles and the level of access permitted by each role.

The UK policy is being incorporated into the demo for the health care project "Software Engineering Environments for Distributed Applications in Health Care." See section 2 for a description. At this point in the development of the demo, the UK policy has been somewhat simplified by eliminating the labeling and limiting the number of roles. Table 1 lists the roles and the patient information that each role may access in this simplified form of the UK policy.