National Committee on Vital and Health Statistics

Subcommittee on Standards and Security

July 13-14, 2000

Hubert H. Humphrey Building
Room 705A
200 Independence Avenue, SW
Washington, D.C. 20201

Statement by:

Solutions for Distributed Business^{SM}

Thomas C. Culpepper
Senior Architect
2AB, Inc.
www.2ab.com
Technical Co-Chair of CORBAmedÔ
tculpepper@2ab.com
(801) 619-4161
1700 Highway 31
Calera, AL 35040

11 July, 2000

Table of Contents

1 Biography

Thomas C. Culpepper received the B.S. degree in Computer Science with a minor in Mathematics from Weber State University in 1986. He joined the School of Epidemiology at the University of Utah and concluded work on the Down Winders Project. Moving to the University Hospital Information Systems at the University of Utah he spent several years in a technical management and liaison role for the integrated Pharmacy Systems. At that time he was involved with many projects devoted to parsing, encrypting and extracting clinical information from the Patient Medical Record System.

In late 1989 he took a Senior Programmer role at Intermountain Health Care (IHC) who was at the time in joint development with 3M Health Information Systems (HIS) in moving the HELP Systems into a commercial venture. While at IHC he designed and development a template language used for creating decision support logic for the Order Communications subsystem as well as a code-generator used in conjunction with the template language to automatically generate decision support logic which allowed clinicians to order and modify laboratory, pharmacy, respiratory care, radiology and nursing orders through the Order Communications subsystem. He also sat on several technical committees:

- The Quality Circle/Development Process Committee which developed a set of procedures and methods that ensured the effective performance involved in the definition, design, implementation and testing of the software development life cycle which were used throughout the development environment.
- The Ad-hoc Committee that was dedicated to providing a mechanism whereby the user could query stored medical information.

Accepting a position as Senior Software Engineer with 3M Health Information Systems (HIS) he pursued and engaged in the object-oriented paradigm of computer science and assisted in moving the proprietary legacy system into a n-tired open system architecture. Many of the projects he worked on include:

- RW-Report Writer
- CW-Clinician WorkStation
- LDR-Lifetime Data Repository – Electronic Patient Record
- DAS-Data Access for Clinical Observations
- EMH-Error Message Handler
- HDD-Health Data Dictionary

He also held a position on the 3M HIS Technical Operations Committee and the 3M HIS Standards Committee where he was instrumentally involved with the Object Management Group® (OMG™) Medical Domain Task Force (DTF) called CORBAmed™. He was a contributor on the Person Identification Service (PIDS) and the Terminology Query Service (TQS) as well as the primary lead on

the Content Access Service (COAS) formally Clinical Information Access Service (COAS). At this time is was also a US Delegate to the ISO/JTC1 SC32 WG2 and author of the Metadata Query Service (MQS) New Work Item (NWI).

January 2000 he accepted a position with 2AB, Inc. as a Senior Architect and currently resides as the Technical Co-Chair of CORBAmed™. He is a Senior Architect and object-oriented expert for the Government Computer-based Patient Record (GCPR) Framework Project (FP) for the Department of Defense Health Affairs. He is also the current Chair of the Reference Architecture and Common Services Working Group (RACS WG) for GCPR FP.

He has been published in:
- Object World
- Health Care Informatics
- AMIA
- Advanced Technology Program for the National Institute of Standards Technology
- Washington Post Interview

He is a guest lectures for the School of Medical Informatics at the University of Utah on the:
- OMG™
- Component Object Request Broker Architecture (CORBA®)
- CORBAmed™

He has presented at the:
- Object-Oriented Programming Systems, Languages and Applications (OOPSLA) Conference
- OMG™ Technical Meetings
- Health Level 7 (HL7) Vocabulary Special Interest Group
- HL7 Executive Board
- Objects 6000 - Brazil

2 Abstract

The healthcare domain's business has evolved and is becoming increasingly more difficult in providing the right information to the users based on who needs it, what they need and when they need it, notwithstanding, that this information needs to be secured as it makes it way through the healthcare environment.

3 A Standards-Based Approach

The Object Management Group (OMG™) was formed in 1989 and has grown into a consortium of over 850 companies. OMG™'s mission is to promote the theory and practice of object technology for the development of distributed computing

systems. The goal is to provide a common architectural framework for object-oriented applications based on widely available interface specifications.

The model of the OMG™ infrastructure is similar to the model of the technology that they are working on – a distributed "object" technology. The infrastructure developed by the OMG™ is referred to as the Object Management Architecture (OMA™). It has grown into a powerful toolkit of tools and services, including a set of lower layer interoperability tools known as the Common Object Request Broker Architecture (CORBA®). It also includes a rich set of services and facilities, which serve as the infrastructure and glue to make distributed computing possible.

The companies involved in the OMG™ work together to produce an infrastructure and set of standardized interfaces that make distributed, object-oriented computing possible on a global scale. Each company participates, contributes and profits within their area of expertise.

3.1 What is the OMG™'s role in the Medical Domain?

The Medical Domain Task Force (DTF), known as CORBAmed™ was chartered in 1996 by the Domain Task Force Committee (DTC) of the OMG™ to produce a standard set of object interface specifications within the medical domain that can assist in bringing secure information to the right person, at the right time. The following CORBAmed™ specifications have been adopted by the OMG™:

- CORBAmed™ PIDS (Person Identification Service)
- CORBAmed™ TQS (Terminology Query Service)
- CORBAmed™ COAS (Clinical Observation Access Service)
- CORBAmed™ RAD (Resource Access Decision)
- CORBAmed™ CIAS (Clinical Image Access Service)

Several others are currently being worked on:

- CORBAmed™ SLIMS (Summary List Management Service)
- CORBAmed™ MTS (Medical Transcription Service)
- CORBAmed™ HILS (Health Information Locator Service)
- CORBAmed™ OETS (Order/Entry Tracking Service)

Of course, there are also many CORBA® standards that can be utilized within the healthcare domain; they will not be addressed here, but are referenced here:

- CORBA® Security
- CORBA® Naming
- CORBA® Trader
- CORBA® Transaction

With these technologies you can integrate products to begin taking advantage of the strengths each provide in order to bring a complete solution to the medical information requirements of today and tomorrow.

The following briefing describes the adopted CORBAmed™ Standards.

3.1.1 PIDS - Person Identification Service

Throughout an individual's lifetime they may have episodes of care provided by hundreds of healthcare providing organizations (e.g. hospitals, medical centers, Dr. offices, etc.). These organizations maintain medical records for the patients they have cared for. When a patient comes into a healthcare organization for care there is a need to find the records for any previous care that patient had with the institution. Each healthcare provider may have used a different scheme (e.g. numbering system) to identify the patient. The system used for identifying a patient is called a Master Patient Index (MPI).

In addition it is desirable to combine the medical records from multiple institutions in order to show a complete picture of a person's health record. This need to combine records from different organizations has increased dramatically in the last few years due to consolidations and collaborations between providers.

Because of the rapid change in the healthcare environment within the last few years the systems and standards needed to satisfy this need to share patient records that do not yet exist. One of the major impediments to this sharing of patient records between organizations is a lack in the ability to identify a patient in a consistent manner. Due to this inability there is no standard way today to combine a patient's records from multiple institutions.

This PIDS specification defines common features of a patient identification system that allows multiple of these patient identification systems to interoperate.

The complete Patient Identification Services (PIDS) RFP can be found on the OMG™ web server as document:
  http://cgi.omg.org/cgi-bin/doc?corbamed/98-02-29

3.1.2 TQS - Terminology Query Service

There is discussion occurring in the Lexicon Query Service (LQS) Revision Task Force (RTF) about changing the name from Lexicon Query Service to Terminology Query Service. Here TQS is referenced.

The TQS specification defines interfaces for the common features of a set of lexicon query services. TQS also describes the requirements for services to support lexicons (controlled terminology resources) in a distributed object system.

Despite many efforts over the years, the ability to consistently and precisely represent information, such as observational and historical data in healthcare, has eluded the industry. This ability to represent a concept in an unambiguous machine-readable format is key to the better management of clinical processes within a healthcare organization, and between a healthcare organization and its various trading partners. The ability to support a discrete coded lexicon is of critical importance within the healthcare business segment. It is the first step towards being able to:

1. Better manage the communication of information between disparate organizations
2. Support the collection and analysis of clinical processes and outcomes as a result of consistent and clinically specific encoding
3. Enable the use of sophisticated rule-based `decision support' tools, which require consistent data representation to be effective. For example, the rule:
    a. If the order is for any drug in the category antibiotics and there is a history of allergy to any antibiotic, send an alert regarding possible cross-allergic reactions. This requires the ability to classify all antibiotics under a single `parent' in a specified hierarchy to assure that no matter what drug is ordered, if it is in the category antibiotics, this rule is triggered.
4. Assist in the reporting of information to various interested parties in a consistent manner

It is important to make the distinction between the lexicon content (i.e., the "vocabularies" themselves), and the methods to support lexicon queries and functions. In fact, we should not assume that the lexicon query services defined through this effort are necessarily limited to support of a health lexicon/domain of content. It may be the case that these services are a requirement across other domains/task forces within OMG™. It is anticipated that responses could be received from vendors who provide similar services outside of the healthcare arena. However, since the primary interest and critical, near term need resides within the healthcare domain, CORBAmed™ has taken the lead the effort to define these services.

The complete Terminology Query Services RFP can be found on the OMG™ web server as document:

http://cgi.omg.org/cgi-bin/doc?corbamed/98-03-22

3.1.3 COAS - Clinical Observation Access Service

There is discussion occurring in the COAS RTF about changing the name from Clinical Observation Access Service to COntent Access Service.

The COAS specification provides the necessary interface specification necessary for accessing information within the medical domain. Although the COAS specification is called "Clinical Observation Access Service" it has the ability to retrieve all types of information within the medical domain not just clinical observations. This is important because within the medical domain there are a number of stakeholders i.e., administrators, clinicians, educators, patients and researchers all of which need varying types of information to project futures, provide care, instruct others, make informed decisions and create new solutions to problems.

Clinical observations do however; constitute a significant proportion of the information recorded about any patient. Examples of clinical observations include the following: laboratory results; vital signs; subjective and objective observations and assessments; observations and measurements provided by a specialist such as radiologist or pathologist who interprets images and other multi-media data. Information that may also be retrieved but not considered clinical include such things as patient demographics, billing information, medical supplies, dietary information, administrative information such as census, etc.

The complete Clinical Observations Access Service can be found on the OMG™ web server as document:

http://cgi.omg.org/cgi-bin/doc?dtc/00-01-01

3.1.4 RAD - Resource Access Decision

The problem addressed by the Resource Access Decision (RAD) facility is not a new one. The catalyst for this facility was a realization that access control was becoming increasingly unmanageable in enterprise application integration environments. Vendors are spending an increasing percentage of their development time building access control into their applications. This is accomplished in a variety of ways with the obvious problem that each time an enterprise purchases a software component, they are also purchasing an access control mechanism that must be deployed and administered as part of their security infrastructure. This fact has made it impossible for enterprises to design and implement consistent application resource access control policy. The requirements of many government regulations, such as the U.S. Healthcare Information Protection & Privacy Act (HIPPA), make it mandatory that this problem be solved.

The lack of a framework to support fine-grain access controls required by "application-level" security is a well-known problem in distributed computing. The problem is not specific to healthcare but the complexity of the problem in the

healthcare environment is escalated by the need to ensure privacy and confidentiality of clinical information.

Today's commercial authorization products need to address the sophisticated access control policies required by the healthcare industry. For example, a security policy may need to be based on transient relationships such as "attending physician" or individual elements of patient records such as "HIV test results."

This has forced healthcare software vendors to develop proprietary access control mechanisms, known as security policy engines, as part of their healthcare products. This has several implications:

1. It puts the responsibility for defining the capabilities of security policy for an enterprise with the healthcare vendor and not the customer. This requires all customers of a vendor product to use a common security policy model to control secured resources.
2. The proliferation of healthcare organizations, products, policy engines, and security policy implementations makes it difficult to administer and maintain an enterprise-wide security policy.
3. It forces healthcare software providers to develop security architectures that may not be their core competency. This, in turn, detracts from their primary business mission.

Security is a complex problem. The commonality of business domain tasks and security requirements across healthcare computing environments promotes and requires exercising fine-grained access control policies in a uniform and standard way. Access control is only one aspect of the security domain and to fully address the requirements of healthcare industry solutions that integrate auditing, non-repudiation, and notification of security breaches.

Healthcare vendors are increasingly asked to be security vendors, driving up the cost of solutions. The healthcare industry must integrate existing security architectures, technologies and products and not continue to develop proprietary security solutions.

RAD addresses these problems, providing a uniform way for application systems to enforce resource-oriented access control policies. RAD was designed by security specialists to address the requirements of the healthcare industry. By standardizing this service, we enable the healthcare organization to define and administer an enterprise security policy consistently across systems.

The RAD service provides:

1. The ability to leverage existing security standards, such as Kerberos, SSL, PKI.
2. Secure interoperability across distributed technology environments (OSF/DCE, CORBA, LDAP)
3. For the identification of caregivers' privileges.
4. Commercial vendor support by companies specializing in security solutions.
5. To maintain security policy without the need to modify healthcare applications. This allows the security administrators to dynamically maintain and enforce security policies as they evolve.
6. Leverage the power of "componentization": the separation of security policy from the healthcare applications.
7. A simple, standardized interface that product vendors use to request access control decisions.

The RAD specification standardizes the interface but leaves implementation details to the vendors. RAD supports the use of a wide variety of access control policies and provides a framework to plug-in diverse policy enforcement engines.

The complete Resource Access Decision (RAD) can be found on the OMG™ web server as document:

> http://cgi.omg.org/cgi-bin/doc?corbamed/99-05-04

3.1.5 CIAS – Clinical Image Access Service

CIAS is a set of interfaces and data structures with which a server can provide clinical images. It also provides a means to transforming images from non-Internet formats such as DICOM to standard Internet formats such as GIF and JPEG.

Clinical images are sets of two-dimensional arrays of picture elements (pixels) that may be related by anatomy or time series, or both. From the image format point of view, it can be divided into two categories: DICOM images and non-DICOM images.

CIAS is only for clinical image access. Real-time or broadcast streaming video information is specifically excluded from this specification.

The scope of CIAS is to provide a clear, unambiguous mechanism to access medical image information and related metadata in a manner that is useful to general practitioners and clinical specialists where diagnostic image quality is not a consideration.

The complete Clinical Image Access Service (CIAS) can be found on the OMG™ web server as document:

4 Catalyst Project Using CORBAmedÔ Standards

The following is excerpts from a paper I helped author at the onset of the Government Computer-based Patient Record (G-CPR) Framework Project and is presented here to further articulate the need of interoperability and to identify a project that is addressing many of the concerns of the NCVHS Committees.

The Government Computer-based Patient Record (G-CPR) Framework is an effort by a consortium of federal agencies; the Department of Defense, Veterans Affairs, and the Indian Health Services to create an environment where comprehensive clinical patient information can be shared across points of (clinical) service and as though the patient's lifelong medical record was a single view of information.

The overall problem addressed is the lack of sufficient information transfer and research access in disparate data systems. These problems exist in a number of industries, but the healthcare industry is an ideal environment in which to pursue an initial solution due to the critical nature of the data involved and the high costs associated with the industry.

Healthcare enterprises are currently capable of storing comprehensive data about each patient, including medications, clinical procedures, and lab tests. Likewise, medical researchers are continually adding to a knowledge base that would be extremely useful to clinicians—if they were able to master and remember it all. However, this huge amount of patient data and medical research doesn't help the clinicians treat patients, or help enterprises cut costs, unless the right information gets to the right person in a form that is useful and understandable at the right time.

This difficulty is compounded by several factors. As enterprises merge, they find that the databases storing their vital information have incompatible data structures. Therefore, it is difficult and expensive to share data among their systems. In addition, the rapid technology growth in many industries means a rapid growth of new terminology's and concepts. As these vocabularies grow, there is often the duplication of variant terms with the same meaning or single terms that evolve to have variant meanings. These and a number of other linguistic factors create a "Tower of Babel" for systems that attempt to share information. Data sharing or research over distributed systems becomes

extremely challenging. An excellent example is the difficulty posed by enterprise consolidation and the growth of computerization in the healthcare industry.

The information revolution has had a profound impact on medicine. In 1995 alone, the healthcare industry spent over $3 billion on systems to store healthcare data.

Medical researchers are continually adding to a huge knowledge base about illnesses and the effectiveness of various treatments. However, this escalating ability to collect information has not led to the dramatic improvements in healthcare quality and efficiency that one might assume. It's clear that mountains of patient data and medical research don't improve treatment or reduce costs, unless the right information gets to the right person in a form that is meaningful and useful.

To enable this vital access to information, diverse computer systems must be able to share patient data with no loss of meaning or usefulness, and be able to cooperate in the joint execution of tasks—that is, systems must achieve interoperability. In addition, to find information focused on a specific patient problem and the best methods to treat it, systems must be equipped with tools to sift through records collected on an enterprise's entire patient population, as well as access the knowledge provided by research-focused clinical trials with patient populations.

The GCPR Framework is utilizing CORBAmed™ Standards to address these issues and those concerning HIPPA.

5 Conclusion

The OMG™ is continuing to advance and develop the tools and technology necessary to make distributed object computing a reality. The OMG™ has also become a central forum for discussion between and across various groups within the domains.

CORBAmed™ also serves as a central forum for discussion between those in the dictation, transcription, imaging, data repository, decision support, coding and terminology, instrumentation and many other arenas. These discussions are leading to the discovery of common patterns and interest and will continue to serve to reduce the cost while increasing the quality of computerized medicine.

As Technical Co-Chair of CORBAmed™ I urge you as the statutory public advisory body to the Secretary of Health and Human Services in the area of health data and statistics:

1. To strongly promote medical domain interoperability based on open interfaces in an effort to provide the right information, to the right people at the right time in a secure fashion. Which, in turn, will allow the consumer to buy solutions that are based on service, reliability, and best of breed rather than on an interface that binds them to proprietary solutions.

2. To promote a government-sponsored center for the management and dissemination of structural medical information standards (Medical Domain Templates – MDT)) so that unambiguous medical information can be exchanged. For example, if I wanted to exchange a patients laboratory information there would exist a standard structure that provided the necessary structural, attribute and relationship semantics to convey information from one source to another without intermediate intervention. These can be done in a way to allow extensibility while at the same time provide the basic elements of information that do not need to be considered intellectual property or be used in a competitive manner.