Issues in the Security Architecture of the Computerized Patient Record Enterprise*

Konstantin Beznosov†

April 13, 1998

Abstract

We discuss issues in CPR enterprise security architecture. The main goal is to provide a security environment where a user will be viewed the same across all enterprise systems, and access control decisions will be consistent across all components of the CPR enterprise.

1 Introduction

The Computerized Patient Record (CPR) enterprise is and will be a heterogenous environment for a long time if not forever. Stovepipe systems are going to co-exist with new component-based systems as well as with CORBA services, facilities and vertical domain services. The enterprise will always have to accommodate emerging technologies of processing and delivering information to physicians and the staff with old disappearing technologies. The main goal for the CPR security architecture is to provide a security environment where the view of an enterprise user will be consistent across all its components, and access control decisions will be made according to one enterprise-specific model. This paper discusses various issues that make the goal difficult to achieve and maintain. We present our vision on how a CPR enterprise architecture can be designed so that the described problems can be addressed in the realm of existing constraints. The problems discussed in this paper are based on experiences from the ongoing project of designing the CPR security architecture[1] at Baptist Health Systems of South Florida[2](BHS).

In order to facilitate understanding of issues in the CPR security architecture, we will provide background information on CPR enterprise and its specifics next.

2 CPR Enterprise

The Computerized Patient Record is a long-term initiative at BHS. Wreder et al [1] describe its ultimate goal as "to provide the mechanism to capture, manage and present information required throughout the continuum of care in a manner that optimizes the business process" by taking advantage of distributed object computing technologies. BHS's CPR can be viewed as a set of object services and clients distributed across a healthcare enterprise. Since all clinical and some business services are eventually expected to be integrated into the CPR infrastructure, the CPR is considered as an enterprise itself. The CPR architecture is being constructed utilizing the Object Management Architecture described in [2]. CORBA-compliant ORBs constitute the backbone for the CPR components.

All deployed application systems are selected according to the criteria of the best fit for a particular business process they serve and according to the mandatory requirement to comply with the CPR architecture. Particularly, application systems and services are required to provid CORBA-compliant interfaces to their main functionality and to use services available within the CPR enterprise to avoid redundancy. For example, any application system and service, which has a notion of patient, is required to utilize a CORBA-compliant Patient Identification Service (PIDS)[3] and expose any data related to clinical observations via interfaces compliant with a future Clinical Observation Access Service (COAS)[4] standard from the OMG. The very first CORBA-based CPR service was deployed at BHS in February 1998. The service provides access to clinical transcription records. BHS is in the process of deploying a Master Patient Index service that will provide PIDS among other services. An anatomic pathology system that will be using PIDS and will also provide access to its data via COAS-compliant interfaces is expected to be deployed within the next 12 months.

Even though all new components deployed in the

*This document is available in electronic form at http://www.bhssf.org/IT/Projects/cpr/security/architecture-issues/

†Information Technology Department, Baptist Health Systems of South Florida, 6855 Red Road, Coral Gables, FL 33143 <beznosov@baptisthealth.net>

[1]The project web site is at http://www.bhssf.org/IT/Projects/cpr/security.

[2]More information about BHS can be found at http://www.baptisthealth.net

CPR enterprise are based on CORBA technology, there are legacy systems that have to be integrated in the CPR architecture at some point. Also, some new non-CORBA-compliant services will be deployed within the CPR enterprise. Such systems and services have to be integrated in the CPR enterprise including its security infrastructure. In the next sections, we will discuss the issues of designing the CPR security architecture.

2.1 Characteristic features of the CPR enterprise

- Many different application systems (Recent inventory for Y2K showed we have about 200)
- Some products come from narrow niches with few vendors
- Heterogenous operating system environments
- Vendors are oriented towards numerous more conservative customers
- Outside visitors have the potential for physical access to desktops and network infrastructure
- Different departments have different levels of urgency and different requirements for confidentiality and service availability
- No in-house development

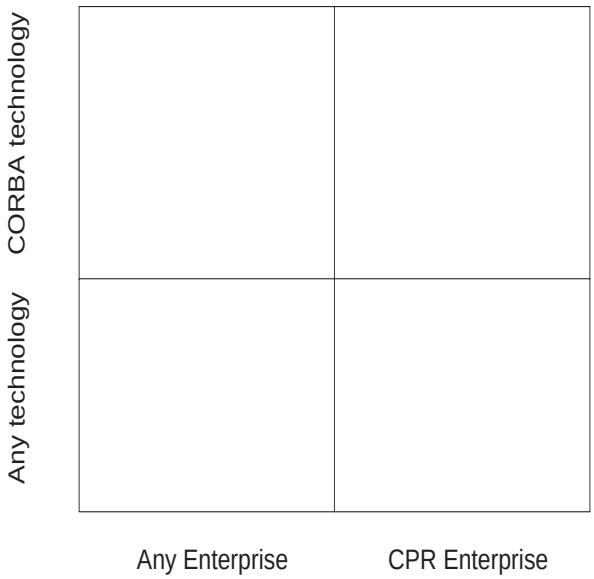

Figure 1: CPR security issues space

3 Security Architecture Issues

We will present four groups of issues related to the security architecture of the CPR enterprise. To ease the understanding of how the described groups relate to each other, we place them on a discrete 2-dimensional space depicted in Figure 1on page 2. The horizontal dimension identifies if the issue can be found generally in any information enterprise or only in a CPR enterprise. The vertical dimension identifies if the issue is related to any technology or it is specific to CORBA-based enterprises.

General issues are propagated into more specialized areas. For example, those problems that exist in any information enterprise are propagated also into a CPR enterprise. To illustrate it, we represent the same issue space in the propagation pyramid shown on Figure 2 on page 2. More general problems at the foundation of the pyramid, if not addressed, would propagate upward.

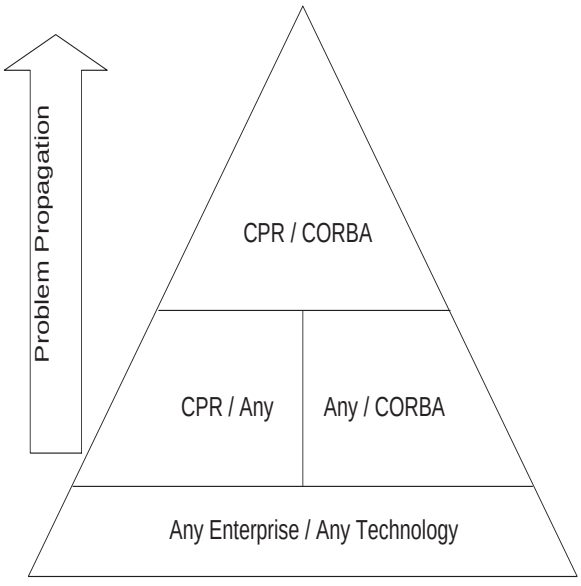

Figure 2: Propogation of problems from more general domains into specific ones

3.1 Any enterprise based on any distributed computing technology

Increasing complexity and size – Due to the increasing rate that an information enterprise grows

with and due to replacement of conventional monolithic solutions by component-based ones, maintenance and administration complexities are rapidly increasing. Increasing size and complexity exacerbates all other factors.

Business gets faster – Business workflows change much faster (18 months) than they used to (5 years) some 10-15 years ago [5]. This means that the information enterprise configuration has to be adjusted at the same rate. For a security architecture, this means decentralized administration and extensive delegation of administration privileges, as well as more frequent changes of access control decision logic driven by business workflows.

Multiple user repositories – Conventional application services have their own user data repositories, which are used to authenticate users and find out user credentials (such as *userid*, and groups). Having multiple user repositories brings inconsistency into the user image across the enterprise. The same user has different *userids*, passwords and group memberships from application to application within the same enterprise. Multiple user security data repositories also create a heavy administrational burden. An administrator has to track changes about the same user across multiple systems. Not only does much more work have to be done to perform changes per user or per logical change, but also it leads to a much higher degree of human error, and it annoys users by making them remember multiple IDs and passwords.

Coupled access logic – Conventional application services have their own access control decision logic, which is coupled tightly with an application itself. The enterprise security administrators end up having to configure such access logic on an application-by-application basis, which brings tremendous administration overhead and highly increases chances of human error as well.

Decisions about which users can have what access to what assets of the information enterprise should ideally depend only on the following factors:

- User security credentials
- Enterprise security policies
- Business workflow constraints

All listed items are properties of a particular enterprise and not of a particular application. Also, access control models must have a common denominator to map enterprise security policies and business workflow constraints uniformly into particular access control rules. Therefore, all access decisions should be foreign to an application service and native to the enterprise security infrastructure as well as the enterprise business workflow.

No standard administration interface – Each application system has its own proprietary interface to administrate access control logic if it has such an interface at all. This makes it impossible to administrate access control and other security policies for multiple applications using a single administration environment.

Inconsistent security models – Due to multiple representations of the same user and access decision logic being tightly coupled with an application system itself, multiple inconsistent security models co-exist in the same information enterprise. In this case, it is highly difficult to insure consistency of access control rules across the enterprise. Most of the time, security administrators end up having no guarantee, whatsoever, that access rules and, especially, changes to them are consistent across all application systems as well as with required company policies.

3.2 CPR enterprise based on any distributed computing technology

YES/NO access control – It is hard to draw exact borders between what a healthcare provider, as an enterprise user, is supposed to have access to and what he/she is not. Some scenarios are clear (e.g., a registration clerk trying to change lab test results of a patient) and some are not (e.g., emergency room physician browsing encounter history of John Smith). There is a need for so called "soft" access control when a principal is granted access; however, audit and (maybe even) non-repudiation "alarms" go off for later investigation. Meanwhile, the user is warned that they are accessing information they are not supposed to. Such a "soft" access control notion is missing from most access control models including CORBASEC. Additional abstraction is needed in security administration solutions to accommodate "soft" access control.

Vanilla security administration – A low-level generic security administration model, where access control (and other) rules are expressed in terms of subjects' security attributes and (groups/domains of) objects/interfaces, is not much useful. A domain-specific environment that will abstract the access model to the level of business workflow is needed.

Non-configurable authentication – Most of the application systems and services that come with a CPR enterprise need authentication mechanisms to be replaceable. Depending on business workflow (whether it is an emergency room or a registration desk) and company security policies (that depend directly on increasing legal and liability requirements), either stronger, or based on different principles (what you are – biometric properties – instead of what you know – passwords), or yet more convenient (smart cards with X.509 public-key certificates instead of passwords) authentication mechanisms will be required.

3.3 Any enterprise based on CORBA technology

Heavy-weight desktop – Today implementations of the CORBA Security service require preinstalled heavy-duty SESAME or Kerberos environments on each user's desktop. If the business process requires functionality that any network computer with a web browser and JVM downloaded to it can provide, then installing and maintaining a full-blown desktop with at least Windows NT on it for each of BHS 1,700 users is financially unjustifiable. The goal is to have an underlying security technology environment to be downloaded with the client application itself in the similar way to how a user can download a Java applet, including ORB implementation or be pulled during the desktop boot phase, as it happens during the boot process of JavaOS or any other storage-less network entety.

3.4 CPR enterprise based on CORBA technology

Interoperability of security services – No two CORBA Security services are known as of April 1998 to be interoperable. This is becoming the main obstacle of deploying a CORBA Security environment in the CPR enterprise.

"Heavy" security domains – Ideally, we want to use the notion of security policy domains actively to leverage the CORBA Security service access control model. All information about a particular patient can be represented as a collection of objects that belongs to the same access control policy domain. So, when a new patient walks to a registration desk and that patient record is created, all data about the patient is accumulated into objects belonging to the patient's domain and the access control (as well as other) policies are instantiated appropriately. Take into account that a healthcare enterprise serves thousands of patients. We do not have empirical knowledge, but it seems that the current underlying security technologies, like SESAME and Kerberos, would not scale to scenarios with thousands of security policies domains.

Coarse-grain access control – Preliminary modeling of a CPR access control model [6], [7] shows that the basic CORBA Security service access control model does not take into account such important for a healthcare enterprise factors of authorization decisions as the content of requests and replies, and the context of client/server interactions. Hopefully, Healthcare Resource Access Decision Facility requested in [8] will resolve this issue.

4 Prioritization of the Issues

Not all problems are as urgent in the shirt term period or as important in the long term period as others. Some of them are highly critical for the CPR enterprise success. Below, we state the goals that we believe will impact significantly the way the CPR enterprise security architecture will evolve.

4.1 Long Term Most Important Goals

1. Central user security attributes repository that will allow a single view of a user no matter what underlying security technology is used
2. Fine grain uniform access decision model across all application services
3. Ability to "plug" various authentication mechanisms
4. Domain-specific security administration abstraction

4.2 Short Term Critical Goals

1. Interoperability of CORBA Security service implementations
2. Light-weight downloadable CORBA security services along with underlying technologies

5 Conclusions

In this paper, we outlined the main issues in constructing a security architecture for the CPR enterprise at BHS. We grouped them into four categories according to the type of information enterprise (general or healthcare) they can appear in, and the type of distributed computing technology they characterize (any

or CORBA-specific). We hope the paper discussion will help other security architects of information enterprises to clarify outstanding issues they face. We also believe the paper will help application vendors to prioritize functional and non-functional properties of their systems designs.

References

[1] Kent Wreder, Konstantin Beznosov, Alan Bramblett, Eric Butler, Alicia D'Empaire, Eddie Hernandez, Eric Navarro, Alex Romano, Mimi Tortolini-Taylor, Enrique Urzais, and Rita Ventura. Architecting a computerized patient record with distributed objects. In *Proceedings of Health Information Systems Society Conference*, pages 149–158, February 1998.

[2] Richard Mark Soley and Christopher M. Stone. *Object Management Architecture Guide*. John Wiley & Sons, 3 edition, June 1995.

[3] Object Management Group. Patient Identification Service RFP. OMG document number: corbamed/96-11-02 http://www.omg.org/library/schedule/Patient_Identification_Service_RFP.htm, November 1996.

[4] Object Managment Group. Clinical Observations Access Service RFP. OMG document number: corbamed/97-12-28 http://www.omg.org/library/schedule/Clinical_Observations_RFP.htm, December 1997.

[5] Larry R. DeBoever. Concept of "highly adaptive" enterprise architecture. Enterprise Architecture Conference keynote address, December 1997.

[6] Konstantin Beznosov. Applicability of corba security to the healthcare problem domain. OMG document number: corbamed/97-09-11, September 1997.

[7] Wayne Wilson and Konstantin Beznosov. CORBAmed security white paper. OMG document number: corbamed/97-11-03, November 1997.

[8] Object Management Group. Healthcare Resource Access Control RFP. OMG document number: corbamed/98-02-23, http://www.omg.org/library/schedule/Healthcare_Resource_AC_RFP.htm, February 1998.