Security of Data of EHR in a grid computing Environment

Kulvinder Singh Mann
Associate Professor,
Department of IT, GNDEC, Ludhiana

Ankit Bansal
Department of CSE, GIET, Khanna

ABSTRACT

Data protection in EHR application needs to be enforced more strictly than in other application areas [1, 2]. Since the 1960s, data protection of personal health information was and is still of high concern. The details of where information flows, who has access to the data and for what purpose are of major importance. Considering this and bearing in mind that in grid computing it is not only simple data sharing, but rather it is the sharing of distributed resources like algorithms, storage, computing power, etc., it is necessary to study closely the aspects of grid security and to find suitable solutions to enforce it. Phenotypic data in a patient electronic medical record can lead to identify for example whether the subject has a particular infection or not. Analyses methods for such identification can be simple statistical procedures, like in [4], or machine learning systems – artificial neural network – like in [5]. Anonymizing procedures are not enough for protecting the data without loosing the scientific value of this data. For instance, sharing high-resolution imaging datasets online may be risky; a full reconstruction of the face using computerized 3D techniques is indeed possible [6-8]. Sharing of medical 3D imaging datasets has already been reported in some pre grid environment applications [9, 10], anyhow, the risk of this sharing is not well studied so far. The major problem is that we still cannot enforce a dependable security policy in grids; i.e. we cannot assure that administrators, developers, or other staffs do not have an access to the medical data.

Keywords - EHR, Grid Computing, SOA, distributed system

1. Introduction

Grid computing is becoming the solution for researchers looking for vast storage and computing capacity, for sharing programs and algorithms. “A computational grid is a set of computing elements and data storage elements, heterogeneous in hardware and software at geographically distant sites which are connected by a network and use mechanisms to share resources as computing power, storage capacity, data” [30] and algorithms. Comparing it with the internet, grid computing goes one step farther in sharing also computing power, storage, applications and algorithms beside sharing information. Similar to semantic web, there is the semantic grid. The semantic web is an extension of the current web in which information is given a well-defined meaning, better enabling computers and people to work in cooperation, i.e. embedding knowledge alongside information. The semantic grid is an extension of the current grid in which information and services are given a well-defined meaning, better enabling computers and people to work in cooperation [12, 31]. Grids are virtual pools of resources rather than computational nodes. Although current systems focus on computational resources (CPU cycles and memory) [32], grids operate on a wider range of resources like storage, network, data, software but also on graphical and audio input/output devices, sensors and so on [33, 34]. All these resources typically exist within nodes that are geographically distributed in multiple administrative domains. Precisely spoken, “the grid is a virtual hypothetical concurrent machine, which is constituted of a set of resources taken from the resources pool” [35].

2. Comparison of grid computing with other distributed systems

In traditional resources/services access control, the user has a direct connection to the resources, which will authorize the user or block her. What distinguishes grids is that, unlike conventional distributed systems (like cluster computing), users and resources appear differently at the virtual and at physical levels. This requires an appropriate mapping to be established between them [35]. Therefore, grid technology uses agents for user and resources mapping. The resource mapping agents (brokering systems) are responsible for mapping a user’s request to a suitable resource, which is available at the request time point. The user mapping agents will in turn map the requests arrived from the brokers to the local nodes as jobs from local users. Semantically spoken, “the inevitable functionalities that must be present in a Grid system are resource and user abstraction” [35].

Nemeth and Sunderam characterized grid computing and presented formal definition expressed in Abstract State Machine (ASM) to describe grid computing and to distinguish such a computing environment [35]. The existence of resource mapping agents and user mapping agents is the main characteristic, which distinguish grid systems from other systems.
To capture the notion of these two kinds of agents, Nemeth and Sunderam defined two processes: *CanUse* and *CanLogin*. If *CanLogin: USER* × *NODE* → *{true,false}* evaluates to true, it means that user has a credential that is accepted by the security mechanism of the node. It is assumed that initiating a process at a given node is possible if the user can log in

to the node. *CanUse: USER* × *RESOURCE* → *{true,false}* is a similar logic function; if it is true, the user is authentic and authorized by an authorizing mechanism to use the abstract given resource defined in the request. While *CanLogin* directly corresponds to the login procedure of an operating system, *CanUse* is a new concept of grid systems and corresponds to an authorizing process to determine what the user is allowed to do and redirect her request to the suitable nodes and resources [35]. In other words: *CanUse* is a brokering function and *CanLogin* is a local account access control function.

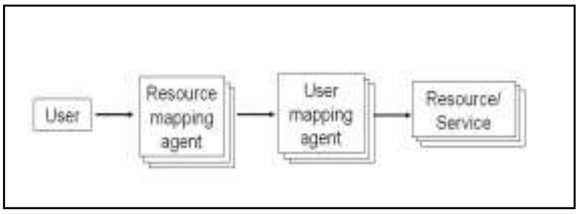

Figure 1: A simplification of the resources access process in grids. The arrow means "communicates".

This results in particular difficulties for authorization, namely:

- The service does not know the user; it only knows the local account to which the user is being mapped to. This leads for example to an auditing problem since the log files and activity protocols refer to the local accounts. It also results in fine granular access rights problems because the service grants access according to local accounts, not the real users.

- The grid user cannot choose explicitly the provider of a service; typically, the user's request includes only an abstract description of the needed service, not the address (location, administrative domain) of the service. monitoring of her work through the service provider. Similarly, the user cannot decide what information about herself will be communicated to the service provider. This may result in a problem about maintaining the "Informational self- determination" right.

- The user, the service and the agents may have conflicts of interest; all entities in the grid have their own policies and interests. The grid community is talking about the problem of (un)suitable mapping: "The Grid system should follow each domain's security policies and also may have to identify users' security policies." [40]. The new "buzz" in the grid community about developing the grid Service Level Agreements [41] is an attempt to solve this problem.

- The usage of a resource/service is temporary and not static; i.e. the users' rights to use the resource/service are in turn temporary. This is also applicable for services from a specific provider in the grid. In this regard, there is a practical problem in isolating usages of different users regarding time (for example to avoid a memory espionage [42-44]). Problems appear also regarding saving profiles for later usages. In other words: time is not considered in the authorization decision.

3. Access control models

Authentication-based access control

One effort to solve the authorization problem in the new environment was to provide authentication in distributed systems. Once a remote user is securely authenticated, ACLs on the server-side can be used to provide authorization. Needham and Schroeder protocol and Kerberos are examples for these efforts [93, 94].

Capability-based access control

Anyhow, the early work on an access control model for distributed computing systems goes back to the 1970s as Wulf et al. published the design of the HYDRA operating system [95]. Beside Wulf's work, later works on distributed operating systems like the Cambridge Distributed Computing System [96] and the Tanenbaum's work on Amoeba [97] employed the idea of capability based access control for distributed systems. For the protection of the capabilities in the distributed environment suitable cryptographic mechanisms were used (hashing the capabilities [46][4]). Capability-based models did not provide a real solution. The origin of weakness was that "the right to exercise access carries with it the right to grant access" [98]. Critiques include that such systems cannot detect stolen capabilities, nor prevent duplication of capabilities, and that revocation invalidates all capabilities for the target object [46].

Credential-basedaccess control

The introduction of credential-based systems solved the shortcomings of capability- based approaches. With the notion of credential-based authorization, trust management was born and authorizations for strangers became possible [99-101]. Credential-based systems utilize user's capabilities for authorization in the form of digital credentials or certificates (signed for example by a trusted issuer). Early approaches go back to the late 1980s when Gong modified the semantics of the traditional capabilities to incorporate the user's identity [102][5]. Different approaches enrich the idea of using credential-based access control. Most notably is the Public Key Infrastructure.

4. Access Control Models in Distributed Systems

Access control models have focused on protecting digital resources on the server- side and do not deal with client-side controls for locally stored digital information. In order to control the usage of already disseminated digital objects, the Digital Right Management (DRM) gave the access control problem a new perspective [103, 104]. DRM technologies attempt to control the use of disseminated digital media by preventing unauthorized access (copy or conversion to other formats) by end users. They have emerged in mid-1990s and gained attention because of their applications in the commercial sector. Current DRM solutions focus on commercial use cases, mainly on intellectual property rights (IPR) protection [25, 105, 106]. Lately, DRM is being addressed for medical applications [107, 108] and general DRM use scenarios include different medical use cases [109].

To sum up, the models for access control in distributed systems address two issues: server-side and client-side control enforcement. They concentrate on protecting digital objects within an environment that consists of a user that interact with a service provider to access the digital object. Third parties (like Certificate Authorities in Public Key Infrastructure) play a passive role in delivering any needed information in a trusted way, i.e. assuring the identity of the user to the service. They do have certificating policy (i.e. an authentication policy), but they do not have an access control policy regarding the users nor can they enforce such a policy. Next subsection (2.2.3) discusses these aspects in a grid computing environment and shows how third parties (authorization authorities) have an active role in defining authentication as well as authorization policies.

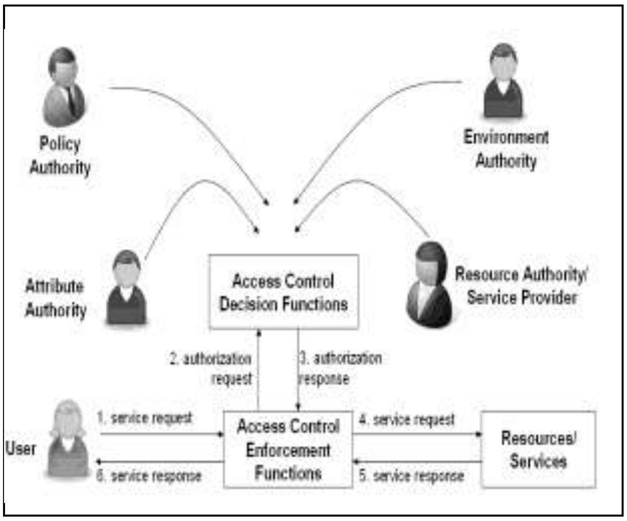

Fig 2 Represents an overview of a grid authorization system based on the pull scenario. It also shows various authorities that may be involved in issuing and determining the authorization parameters, attributes and policies. Similar diagrams could be drawn for the push scenario [69].

Comparison between the different grid computing authorization systems

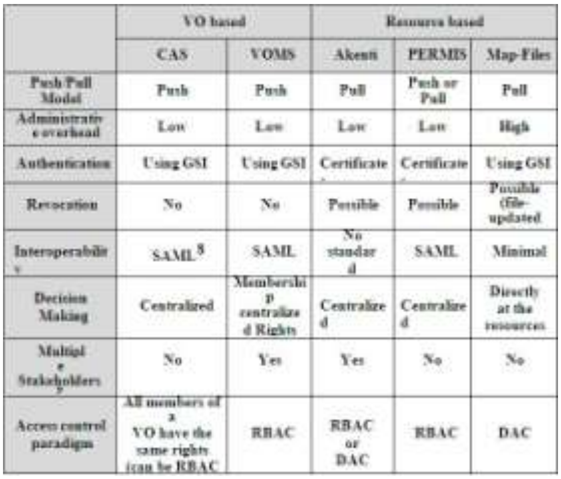

	VO based		Resource based		
	CAS	VOMS	Akenti	PERMIS	Map-Files
Push/Pull Model	Push	Push	Pull	Push or Pull	Pull
Administrative overhead	Low	Low	Low	Low	High
Authentication	Using GSI	Using GSI	Certificate	Certificate	Using GSI
Revocation	No	No	Possible	Possible	Possible (file-updated
Interoperability	SAML[8]	SAML	No standard	SAML	Minimal
Decision Making	Centralized	Membership centralized Rights	Centralized	Centralized	Directly at the resources
Multiple Stakeholders	No	Yes	Yes	No	No
Access control paradigm	All members of a VO have the same rights (can be RBAC	RBAC	RBAC or DAC	RBAC	DAC

Role based access control as the HL7 access control standard

HL7 adapted RBAC as an approach for access control for medical documents and application in 2003 [130]. In May 2007, HL7 balloted the Role-Based Access Control (RBAC) Healthcare Permission Catalogue as a standard and presented a normative language to the permission vocabulary in constructing permissions {operation, object} pairs. In this context, permission is an approval to perform an operation on one or more RBAC protected objects. An operation is an executable image of a program, which upon invocation executes some function for the user. Within a file system, operations might include read, write, and execute. Within a database management system, operations might include insert, delete, append, and update. An object is an entity that contains or receives information. The objects can represent information containers, e.g. files or directories in an operating system, and/or columns, rows, tables, and views within a database management system [131].

Virtualization using gridcomputing

Grid computing opens new opportunities to create a virtual record; i.e. to connect data from different locations and resources. Through the distributed storage and processing of data, grids offer also the possibilities of handling large datasets and performing complex analysis and data mining on the various records as well as real-time data updates. In this context, a true problem in the use of personal data (in medical applications) is the possibility of unauthorized re-identification of individuals, known as the disclosure risk.

5. CONCLUSIONS

Gartner Inc., the known information technology research and advisory company, predicted in 2003 that “by 2008, SOA will be a prevailing software-engineering practice, ending the 40-year domination of monolithic software architecture” [283]. The grid computing activities were reformed in 2003 by introducing the third version of Globus Toolkit in order to adapt the Web Services (WS) and Service Oriented Architecture (SOA) technologies [32]. This adaptation was a strategic step for the developing of the grid technology to be an infrastructure for science [284]. In April 2006, the first grid web services standards – the Web Services Resource framework (WSRF) – were adopted and are still evolving [169, 222]. In the same year Gartner predicted also that “the lack of working governance mechanisms in midsize-to-large (greater than 50 services) post-pilot SOA projects will be the most common reason for project failure" [285]. In 2006 also, Manes, the vice president and research director in Burton group pointed out that “many organizations don’t start to think about governance until things are completely out of control” [286]. In this context, SOA governance is the process of defining and enforcing organizational policies and standards. Hence, it was clear that governance would be a problem in the different approaches to a SOA environment including grid computing. Nevertheless, limited efforts were invested to be ready to face this problem.

Authorization is comprised in IT governance. Weill and Ross define governance shortly as "specifying the decision rights and accountability framework to encourage desirable behavior in the use of IT" [287]. IT governance implies control and measurement in the computing environment. It also provides the framework, mechanisms and methodology for involving all interacting people, from those being supported to those who provide support [288]. Governance in grid computing lacks not only a well defined framework that facilitates the enforcement of needed policies; moreover it lacks a method for defining needed policies (like defining policies using computer logic in order to check their correctness). Simply put, in regard to security, grid computing lacks correct functional authorization mechanisms, but more important, it lacks a functional authorization model.

6. References

[1] Anderson, R., *Security engineering: a guide to building dependable distributed systems*. 2001, New York: Wiley. xxviii, 612 p.

[2] Anderson, R., Isaac Newton Institute for Mathematical Sciences, and British Medical Association, *Personal medical information: security, engineering, and ethics, Proceedings of the Personal Information Workshop, Cambridge, U.K., June 21-22, 1996*. 1997, Berlin; New York: Springer. x, 250 p.

[3] Martin-Sanchez, F., V. Maojo, and G. Lopez-Campos, *Integrating genomics into health information systems*. Methods of Information in Medicine, 2002. **41**(1): p. 25-30.

[4] Diero, L., et al., *Can data from an electronic medical record identify which patients with pneumonia have Pneumocystis carinii Infection?* International Journal of Medical Informatics, 2004. **73**(11-12): p. 743-750.

[5] Haraldsson, H., L. Edenbrandt, and M. Ohlsson, *Detecting acute myocardial infarction in the 12-lead ECG using Hermite expansions and neural networks.* Artificial Intelligence in Medicine, 2004. **32**(2): p. 127-36.

[6] Nakajima, S., et al., *3D MRI reconstruction for surgical planning and guidance*, in *Advanced Navigation in Neurosurgery*, E. Alexander and R.J. Maciunas, Editors. 1999, New York: Thieme. p. 137-145.

[7] Jones, M.W., *Facial reconstruction using volumetric data*, in *Proceedings of the Vision Modelling and Visualization Conference 2001.* 2001, Berlin: AKA GmbH. p. 135-150.

[8] Evison, M.P., *Computerised 3D facial reconstruction.* assemblage - the Sheffield graduate journal of archaeology, 1996. **1**(1).

[9] *Medical data storage and processing on the GRID.* 2002 [Accessed 2008 March 13]; Available from: http://creatis-www.insa-lyon.fr/MEDIGRID.

[10] Benkner, S., et al., *GEMSS: Grid-infrastructure for Medical Service Provision.* Methods of Information in Medicine, 2005. **44**(2): p. 177-181.

[11] Purdam, K., et al., *Grid computing and disclosure control.* New Review of Information Networking, 2004. **10**: p. 161-176.

[12] Banatre, J.-P., et al., *Future for European Grids: grids and service oriented knowledge utilities. Vision and research directions 2010 and beyond.*, D.D. Roure, Editor. 2006, Next Generation GRIDs Expert Group, Information Society Technologies (IST), European Commission, Office for Official Publications of the European Communities, Luxembourg. p. 60.

[13] Tomita, Y., et al., *Artificial neural network approach for selection of susceptible single nucleotide polymorphisms and construction of prediction model on childhood allergic asthma.* BMC Bioinformatics, 2004. **5**(1): p. 120.

[14] Bao, L. and Y. Cui, *Prediction of the phenotypic effects of nonsynonymous single nucleotide polymorphisms using structural and evolutionary information.* Bioinformatics, 2005. **10**(21): p. 2185-2190.