Communication papers of the Federated Conference on Computer Science and Information Systems, pp. 199–205
DOI: 10.15439/2017F439
ISSN 2300-5963 ACSIS, Vol. 13

A model for context-sensitive usage control in healthcare information systems

Rodolfo Barriviera
Federal Institute of Paraná State,
Londrina PR – Brazil
Email: rodolfo.barriviera@ifpr.edu.br

Carlos Maziero
Federal University of Paraná State,
Curitiba PR – Brazil
Email: maziero@inf.ufpr.br

***Abstract*—The secure management of access to patient information in medical and hospital environments is a relevant and widely debated problem. Although the electronic patient record is envisaged in the current legislation, it is still a promise to the Brazilian health reality. The aim of this work is to define a continuous access control model for patient's medical information, which considers contextual information and that can be implemented in large scale environments, such as computational clouds. A bibliographic study was conducted related to the computational area of access control, electronic and paper medical records, and on the current health legislation in Brazil. Subsequently, a qualitative research was conducted in the BHUs of the city of Londrina - Paraná (Brazil). The results presented a great lack of computerization, the use of paper charts, and the need for a computer model of access control that considers the aspects of a real medical-hospital environment. From the results obtained in this research, it was possible to elaborate a computational model of access control to the patient's electronic medical record for health information systems using the UCONabc access control model and the XACML architecture. A prototype of the proposed model was implemented. The experiments conducted and their results demonstrated the proposal's feasibility.**

I. Introduction

COMPUTATIONAL advancement allows known technologies to be enhanced and is used in different ways in computing environments, especially in cloud computing. However, cloud computing brings new concerns, and new paradigms of data management, as well as worries about users that require improvement related to the data access.

Since the 1990s, health information systems have been developed to provide medical records in distributed environments and cloud computing. Nevertheless, due to concerns to the fact that only authorized people should have access to the system, certain security requirements were required, such as authenticity, non-repudiation, integrity, confidentiality, availability, and especially access control.

Research into health information systems in the cloud has been intensified in the pursuit of technologies that can consider not only the attributes of those involved in the information access but also the characteristics of the environmental context of a specific system.

According to Abowd [1], context is any information that can be used to characterize an entity. The entity can be the representation of a person, place, object or physical state that can be considered relevant in the interaction between the user and the system.

In general, health studies are based on the information contained in the current legislation and do not consider real health environments as Basic Healthcare Units - BHUs and Hospitals. Thus, there is a gap between proposals for health information systems and the current reality.

In this work, a qualitative research approach was used to investigate the real health environment and, under the circumstances, to serve as a basis for the development of a conceptual model that meets the demands of the server's access to the patient's medical records on the health environment. However, since the objective of this article is not to show a qualitative approach, these results can be appreciated in another opportunity.

The objective of this work is to propose a computational model of continuous access control for health information systems, which considers contextual information and can be introduced in large scale environments.

This article is structured in 5 sections. Section 2 presents the concepts of health information systems; Section 3 describes the access control model proposed in this work, who is based on the results obtained through bibliographical and qualitative research; Section 4 encompass the evaluation and results of the access control model proposed in this paper; And, finally, Section 5 concludes this work.

II. Hospital Information Systems

According to [2], the advancement of technology and its increasingly presence in the patient care process has increased the demand for collaborative sharing of patients' clinical data among health professionals. In 1997, the Institute of Medicine defined the Electronic Medical Record (EMR) as

> *An electronic record that resides in a system specifically designed to support users through availability of complete and accurate data, practitioner reminders and alerts, clinical decision support systems, links to bodies of medical knowledge, and other aids.*

The world-renowned architecture of a healthcare and hospital environments includes the concepts of Personal Health Record - PHR, Electronic Medical Record - EMR, and Electronic Health Record - EHR and Picture Archiving and Communication Systems - PACS. The most widespread architectures worldwide are user centric. Thus, all access must have

authorization from the owner user. In the Healthcare Cloud [3] software the user has autonomy to control the accesses to its electronic medical record through the mechanisms of control of access.

In Brazil, there is no system pattern, only a guideline of the Conselho Federal de Medicina (CFM, Portuguese for "Federal Council of Medicine"). The CFM Resolution No. 1,331/89, Ordinances No. 1,638/2002 and No. 1,639 / 2002, informs how the temporality, custody, and handling of the medical records have to be processed, as well as how the computerized systems have to be used. These Brazilian details allow the exploration and research in the health area aiming the computer vision of the access control area.

III. A USAGE CONTROL MODEL FOR HOSPITAL INFORMATION

The objective of this proposal is to elaborate a model that allows the continuous access control to medical information and considers external information (such as data of the entities context and the identities of the subjects and objects involved) besides considering the results of the qualitative research.

Studies starting from the reality of the subject through a previous qualitative research, involving the concepts of continuous control of access, using the contextual information of the environment, and applying this information to the processing of rules policies were not found in the area of access control and cloud computing. This aspect shows the relevance of this research not only in the conceptual aspect but also in the real applicability of the model.

In the next sections, the necessary requirements for the development of the proposed model will be discussed as well as the related works, the technologies involved, and the architecture of the model and its interaction form.

A. Requirements analysis

The objective of the qualitative research was to understand how the Basic Healthcare Units - BHUs control the access to patients' medical records, to verify the occurrence of the use of the medical record in paper and in digital format and to compare the reality practiced with the current legislation. Thus, through its results, serve as a basis for the development of a computer model of access control related to the real aspects of the health environments.

According to the results of the qualitative research, context-sensitive health information, such as patient address, blood pressure, and location of the care service, is essential for the proposed model. Interview reports, such as "Medical records are taken from the storage room. Particularly in medical visits", emphasize the importance of conducting qualitative research and the application of contextual information in the development of the model.

The contextual information identified by the qualitative research is processed along with other information through the rules policies in the proposed model. For example, if a patient does not reside in the BHU coverage area he will not have home care, the patient will have to request the service from the BHU of its area.

The health environments present several characteristics identified through the qualitative research that can be translated to the proposed model, for example: Subject and its attributes - the one that requests access to the available resource in the system; Object and its attributes - the resource available in the system in this model represented as the patient's medical record; Rights - the actions that the subject can execute on the object; Obligations - characterizes what the subject must do; Conditions - defines the characteristics to be respected related to the system environment; And authorizations - the policies that represent the rules that will be processed to determine whether or not the subject can access the object.

Because of the results of the qualitative research, it was verified that during the process of the patient care and access to its medical records, there was a continuous behavior: The patient's medical record was constantly used by health servers and, in some cases, shared with others workers to clarify cases.

From this result, the concept of continuity and delegation was applied in the proposed model. Thus, the proposed model was able to consider the continuous accesses to patients' medical records and also to provide the possibility of having several health professionals acting in the same clinical case through the delegation of access.

Continuous accesses of patients' medical records were associated with a constant change in patient characteristics (defined as contextual information) and should also be considered in the proposed model. The contextual information covers, for example, blood pressure, service location and health status.

To encompass the behavior identified by the results of the qualitative research, the concept of mutability was applied. Hence, the changes occurred during the patient care can be considered in the proposed model.

As an example, one of the situations considered in the model tests was the authorization request to access the patient medical records within a BHU. This request is referred to the moment the patient and the nurse are in the anamnesis room. The nursing auxiliary needs the access to the patient's medical record for a period of time to take the pertinent notes, so the auxiliary requests authorization.

The model considers for this authorization request: the subject and its attributes, the object and its attributes, rights, obligations, conditions, contextual information, and authorization request. Combining all this information, the model does the processing and issues the access authorization response.

B. Technologies involved

After a bibliographic study, the technologies considered essential for the success of the proposal were selected for the development of the suggested model. The technologies chosen are described in the following sections.

1) Usage Control - UCONabc: In the proposed model the concepts of continuous access control to medical information are essential because it is necessary to consider the dynamicity

of real health environments. The UCONabc technology application has in its specification this behavior and has proved to be efficient.

The UCONabc is a Usage Control Model proposed in [4] that has as objective the definition of rigorous theoretical concepts for Usage Control in which it considers the dynamicity of the actions occurred during the use of the system resources. This model contains the concepts relevant to current computational applications and introduces in its architecture the concepts of Subject (S), Subject Attributes (ATT(S)), Objects (O) and Object Attributes (ATT(O)), Authorizations (A), Rights (R), Obligations (B), Conditions (C), as well as the concepts of mutability and continuity.

The application of UCONabc allows the proposed model to use information related to the subject and object in a wider and more dynamic way. In addition, it allows the model to apply obligations and conditions to the subjects in the environments and control them from their rights and authorizations. Health environments are active, shared and in constant change, therefore it is important to consider the changes that are occurring in this environment as well as ensuring its continuity and delegation.

2) Context Aware Computing: During qualitative research, health environments presented quality contextual information, such as current location and blood pressure. Due to the relevance of this information they were applied in the model. Using the Context-Aware Computing technology it is possible to apply this information to the model.

The development of context-aware information systems brings new challenges. According to [5], the context architecture is the basis for the abstraction of a context application. In information system that uses context, it is necessary to obtain relevant environment information that has effects on the interaction of the user with the system. Contextual information is constantly obtained from the environment while using the system.

Context-Aware Computing is an essential factor in the proposed model. There are significant contextual characteristics in the UBS environment that were identified in the qualitative research, such as physical status of the patient and physician's current location. This type of information is very relevant and can be used to control access.

3) Extensible Access Control Markup Language - XACML: XACML[6] is a markup language, an XML extension (eXtensible Markup Language) that allows the modeling, storage, and distribution of descriptive Usage Control policies across entire systems, specific resources, proprietary or public environments. In this language, the policies are defined to provide forms of requests and responses to Usage Control and authorization. Because of this standard, it is possible to control access and authorize the use of certain system contents.

According to [6], the XACML architecture is composed of: Access Requester, Policy Enforcement Point - PEP, Obligations Service, Context Handler, Policy Decision Point - PDP, Policy Administration Point - PAP, Policy Information Point - PIP e Environment. XACML is an XML-based request-response language that, after processing requests and rules, evaluates it to "permit", "deny", "indeterminate" or "not-applicable".

C. Related work

The proposed model presented in this paper was developed based on theoretical studies and from the results of a qualitative research in order to follow the reality of the healthcare environment. The related works described in this section are linked to this paper because they partially apply the technologies used here.

Anastasi in his work [7] presents a federated cloud access control framework based on the UCONabc model and on the XACML language. The framework extends the XACML technology and also proposes a way to manage the authorization session thus creating a control of continuity or revocation of access to the federated cloud environment.

In his study, Almutairi[8] applies the concepts of the UCONabc model and the contextual information to provide continuity of access to changes in the environment. The suggested model called CA-UCON uses contextual information in conjunction with UCONabc concepts to identify changes that occur during access.

The related works presented the development of models from the UCONabc model, contextual information, and processing with the XACML language. However, in the references cited and in the literature, it was not identified works that have proposed continuous access control models for health information systems that applied the concepts of the UCONabc model, obtained contextual information of the environment in a constant way and made the access session management. Also, no work found applied to real health environments on a large scale, considered the results of the qualitative research, or that was applied through of XACML.

D. Architecture

The proposed model applies the concepts of the UCONabc Usage Control Model mentioned in section III-B1 as well as the concepts of continuity or revocation of access mentioned in the work of Anastasi [7]. In addition, it applies the concept of Context-Aware Computing discussed in section III-B2 and the application of the results of the qualitative research. The model uses the XACML markup language, cited in III-B3, as a policy processing language. Thus, the attributes, the characteristics of the environment, the context information of the health environment, the attributes mutability, the continuity and the delegation of access are considered.

The model considers that the subjects requesting access to the object are already authenticated in the system. The architecture is concerned about the authorization process.

The model architecture, as can be seen in figure 1, is structured through several components with specific actions. There are nine components: *(i)* **Subjects and Attributes of the subjects**: The subjects (user and/or system) are the entities that want to have access to the objects. The attributes represent the characteristics of the subjects. Examples of

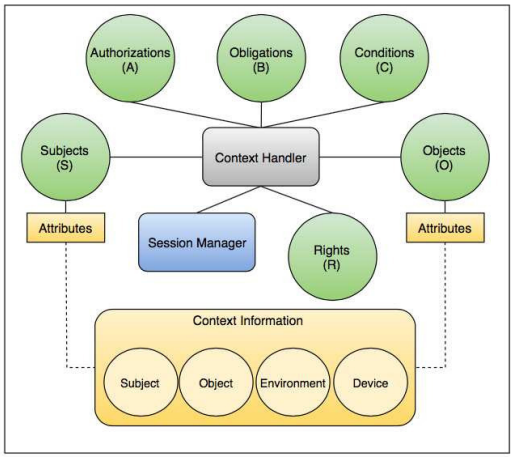

Fig. 1. Architecture of the model.

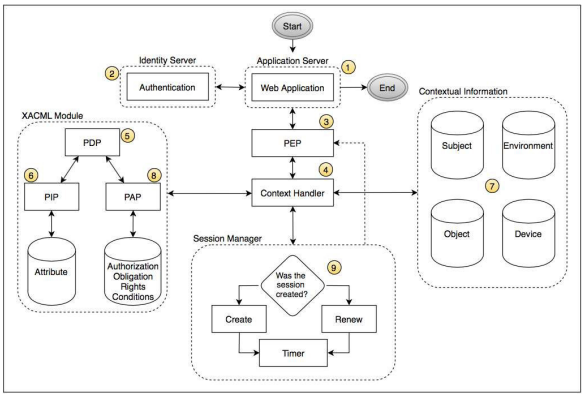

Fig. 2. Interaction of the modules.

this component are name, address, and location; *(ii)* **Objects and Attributes of objects**: The objects (patient's electronic medical record) are the entities that the subjects want to access for a certain period of time. Attributes represent the characteristics of objects. Some examples are: ID, date, and location; *(iii)* **Rights**: Rights are the entities that represent the privileges that a certain subject can obtain during access to the object in a certain environment. Some examples are: the patient can be vaccinated in the health environment; *(iv)* **Authorizations**: Authorizations are the entities that represent the location of all the policies that will be processed are stored. The policy processing (containing the rules) results in the access authorization; *(v)* **Obligations**: Obligations are the entities that represent the obligations and/or requirements that must be fulfilled before or during the use of the authorization; *(vi)* **Conditions**: The conditions are the entities that represent the factors related to the environment and that are relevant to the authorization process; *(vii)* **Contextual information**: Contextual information is the entities that represent the attainment and updating of the relevant contextual information used in the authorization process. Any contextual information applied to the model is obtained and updated by that entity. The contextual information used is obtained and updated by this entity in parallel to the policy processing; *(viii)* **Context Handler**: Module responsible for integrating the entire model. It interacts with each architecture module in order to manage the operation of requests and responses; *(ix)* **Session manager**: The session manager is the entity that represents the control of the access period of the subject to the object. With this entity, it is possible to define the subject access time and the renewal (continuity) of its access or revocation. The proposed model reflects the application of the defined concepts and also the real environment of the UBSs.

E. Interaction Model

Based on the UCONabc model, a context-sensitive access control software was developed in health information systems for the computational environment.

As described in the figure 2, the prototype was implemented with the following modules: Application Server (1), Identity Server (2), PEP (3), Context Handler (4), PDP (5), PIP (6), Context Information (7), PAP (8) and Session Manager (9).

The interaction between the modules follows the order: **Step 1**. The web application receives the authentication request of the subject (containing login and password) through the web application and forwards the obtained data to the authentication module; **Step 2**. The authentication module checks the credentials of the subject and sends the authentication process response to the web application. If there is any problem with the authentication process, the user is refused in the web application and there will be no authorization request for access to the patient's electronic medical record; **Step 3**. After successful authentication, the Policy Enforcement Point - PEP module intercepts the subject's authorization request in its native language and converts it to XACML language. After the conversion, the PEP sends a request to the Context Handler; **Step 4**. The Context Handler receives a request in XACML language sent by the PEP. It is required to provide a response to the PEP. The response can be: "permit", "deny", "not applicable" or "indeterminate". To obtain a response, the Context Handler needs to request the processing of the policies that are associated with the request received. Thus, the Context Handler sends to the Policy Decision Point - PDP the XACML language request to be processed; **Step 5**. The Policy Decision Point - PDP receives the request processing of the Context Handler request. Then, the PDP requests the Policy Information Point - PIP the attributes related to the request received. The attributes refer to the subject, object, environment, and device; **Step 6**. PIP searches the stored attributes on its basis. The attributes not found are requested for the Context Information module, through the Context Handler. After receiving the requested attributes, PIP sends them to the PDP; **Step 7**. Contextual Information represents a module that acts in parallel with the model. This module is constantly obtaining the relevant contextual information through the environment sensors. A recording is performed in the corresponding database from the contextual information

obtained; **Step 8**. After all necessary attributes were received from the PIP, the PDP requests to the Policy Administration Point - PAP the policies regarding the software environment. Then, the PAP searches the pertinent policies (authorizations, obligations, rights, and conditions) to the current requisition and sends them to the PDP to process. The PDP receives the policies of the PAP and does the processing of the policies. The result is sent to the Context Handler in XACML format; **Step 9**. The Context Handler receives the response from the PDP and checks it. If the response is "permit", it prompts to the Session Manager to create or renew the session for the last request. The Session Manager receives the request from the Context Handler to create or renew the session of the subject. Initially, it is checked if the session requested by the Context Handler has already been created (in the cases of continuity of access). If there is no session regarding the subject of the active request, then a session is created and its timer activated. If the session is active, it is refreshed and the timer restarts. When the subject's use time expires, the Session Manager sends a renewal request to the PEP (continuity process). From this request, all procedures for authorization requests are executed again; **Final step**. After the steps described above, the Context Handler receives the message from the Session Manager that the section was created. Next, the Context Handler sends to the PEP the response of the authorization request in XACML language. PEP converts the XACML language authorization result received from the Context Handler into native language and sends the result to the web application. Consequently to "permit" response, the web application authorizes the subject to access the electronic medical record of the patient.

The delegating access process is done by a user who is using the system. In other words, its access request has already been authorized. This user sends an access request to the system and as a result, a certain employee can access the medical record that the first user is accessing at that moment. This request follows the same steps described above.

IV. Evaluation of the proposed model

The objective of this section is to describe the implemented prototype that represents the model proposed in this work as well as its characteristics and results obtained.

A. Prototype implemented

For the evaluation of the model, a MacOS Sierra 10.12.3 (3.4 GHz Intel Core i5 processor, memory module holding 16GB RAM and 395 Gb Hard Drive capacity) was set as the server machine. The programming language used was Java SE 1.8.0 version and the free software framework WSO2 Application Server version 5.3.0 (WSO2 AS 5.3.0) as a web application server. The free software framework WSo2 Identity Server 5.3.0 was utilized as the authentication server. And, finally, MySQL 5.6.2 version 14.14 was used as a database and the Apache Axis2 module was chosen as SOAP message interceptor.

The proposed model was developed through Java SE language oriented to object with the Model View Controller - MVC structure integrated with Java Server Pages - JSP, a total of 1,960 lines of code. The free software framework WSO2 Application Server is used to integrate web technologies such as Apache TomCat, Web Services (Apache Axis2), RESTful, cluster and log extensions. The framework makes it possible to host and manage web services and the integration with WSO2 frameworks. It was responsible for hosting the developed prototype as a web service.

The WSO2 Identity Server software framework is an identity manager. The LDAP is used as a native module to authenticate users. It also has the XACML engine for versions 1.0, 2.0 and 3.0 in order to process the rules. This framework was responsible for authenticating users and processing the rules defined in the proposed model. The Apache Axis2 module was used as the request interceptor. This module acts as the PEP in the XACML structure.

The MySQL database was used to store all the contextual information obtained from the environment. It has been integrated into the WSO2 Identity Server as an external data source. Therefore, the developed prototype was hosted on the WSO2 Application Server which was integrated with the WSO2 Identity Server to authenticate users and process the rules.

B. Conducted Experiments

The conducted experiments using the developed prototype were a load test. This test aims to overload the web service to observe its behavior from a certain volume of data. The volume measured in the prototype is from the access requisitions. For each test, all the services involved (the web server and the XACML authentication and processing server) were restarted. The restart of services avoids server caching and performance problems that occurred after long periods.

For the simultaneous execution of access requests, the free software Apache JMeter version 3.1 was used. The software was configured for each test with the variation of users between 8, 16, 32, 64, 128 and 256, and variation of connections per second between 8, 16, 26, 41, 55 and 64. The rules configuration was defined on the WSO2 Identity Server with XACML rule processing. The combinatorial algorithm used to execute the rules was deny-unless-permit. This rule forces the processing of all active rules in the system to check the permit or deny condition. The rules were varied between 8, 16 and 32 users. The attributes were configured in the external database MySQL. The attributes varied between 8, 16 and 32 users as well. Initially, for configuration and adjustment of the tool, 50,000 tests were generated to verify the system stability.

Thus, for each test, we defined an attribute value, rules value and the variation of users with the variation of connections per second. For example, the first test was performed with 8 attributes, 8 rules, with the variation of 8, 16, 32, 64, 128 and 256 users combined with the variation of connections per second of 8, 16, 26, 41, 55 and 64. Each user value corresponded to a value of connections per second, such as: 8 users with 8 requests per second, 16 users with 16 requests per second, 32 users with 26 requests per second, 64 users with 41

requests per second , 128 users with 55 requests per second, and 256 users with 64 requests per second. The next test was done with 8 attributes, 16 rules, and the same variation of users and requisitions per second. The tests were then successively conducted up to 32 attributes and 32 rules. In total, 4,536 executions were conducted.

C. Results and discussion

The results obtained with the configuration of the environment and experiments performed show the viability of the proposed model. These results had the objective of evaluating the prototype performance.

The figure 3 (a) shows the response time as a function of the requisition rate per second. It is verified that the lowest response value for the requests occurs when there are 8 attributes and 8 rules, whereas, when there are 8 attributes and 32 rules the response time is higher. This result shows that the variation of rules directly impacts the results. The same behavior is indicated in tests with 32 attributes and 8 rules compared to 32 attributes and 32 rules.

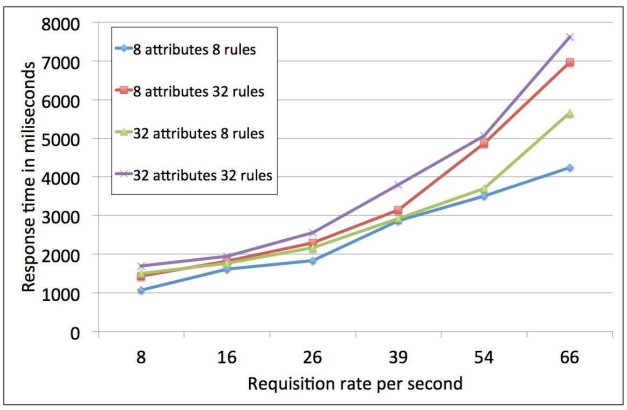

(a) Response time according to the requisition rate per second.

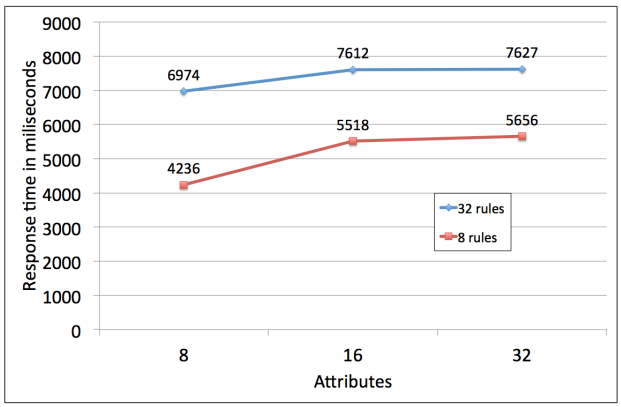

(b) Response time according to the number of attributes.

Fig. 3. Response time.

When a better case, the smallest number of attributes (8) and the lower number of rules (32), is observed and compared with the worst case, the highest number of attributes (32) and the biggest number of rules (32), there is also an increase in response time for requests.

This behavior was expected as a prototype performance because the prototype needs to process a larger number of rules and more requisitions per second, leading to an increase in the response time for each request.

The figure 3 (b) shows the performance of the response time in function of the variable attributes and rules. As the attributes increase, the response time also gradually increases. However, after 16 attributes, it turns out that the response time increases more slowly, representing a better performance in the search time in the database, since the attributes need to be collected in the external MySQL database.

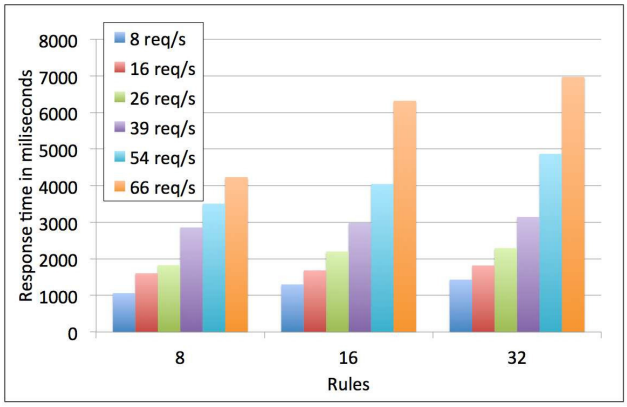

(a) Relation between rules and requisitions per second - 8 attributes.

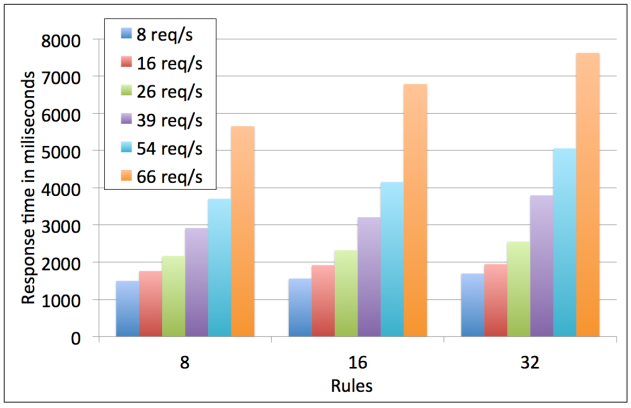

(b) Relation between rules and requisitions per second - 32 attributes.

Fig. 4. Rules in relation to 8 and 32 attributes.

Figures 4 (a) and (b) represent the response time with the relationship between the increase in the number of rules and the rates of requests per second. The chart confirms the expected prototype behavior. As more rule processing and requisitions per second are required, there is an increase in response time for each request.

V. Conclusion

The qualitative research allowed the proposal to be closer to the reality of the user. The UCONabc model granted through its concepts the development of a strict control of

continuous access to the health environments that have several characteristics. The use of contextual information was essential to obtain from the environment the relevant data considered in decision making. The results obtained demonstrated good performance when simulating situations close to reality with several requisitions per second. We have not found any work in the literature with a proposal equivalent to the one presented in this paper. Future research opportunities include studying the impact of the model on higher requests rates per second and evaluating the prototype with more complex rules in a more realistic scenario.

REFERENCES

[1] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, and P. Steggles, "Towards a better understanding of context and context-awareness," in *Proceedings of the 1st International Symposium on Handheld and Ubiquitous Computing*, ser. HUC '99. London, UK, UK: Springer-Verlag, 1999, pp. 304–307. [Online]. Available: http://dl.acm.org/citation.cfm?id=647985.743843

[2] J. a. Filho, M. Figueiredo, D. Santos, D. A. Pizzol, L. C. D. Medeiros, A. Fernanda, F. Bezerra, G. Henrique, and M. Bezerra, "Infraestrutura de segurança para comunicação , autenticação e autorização transparentes em hospitais federados (in Portuguese)," *Journal of Health Informatics*, vol. 3, no. 2, pp. 58–63, 2011.

[3] R. Zhang and L. Liu, "Security models and requirements for healthcare application clouds," *Proceedings - 2010 IEEE 3rd International Conference on Cloud Computing, CLOUD 2010*, pp. 268–275, 2010.

[4] J. Park and R. Sandhu, "The UCON ABC Usage Control Model," *ACM Trans. Inf. Syst. Secur.*, vol. 7, no. 1, pp. 128–174, Feb. 2004. [Online]. Available: http://doi.acm.org/10.1145/984334.984339

[5] A. Dey, G. Abowd, and D. Salber, "A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications," *Human-Computer Interaction*, vol. 16, no. 2, pp. 97–166, 2001.

[6] D. Ferraiolo, R. Chandramouli, R. Kuhn, and V. Hu, "Extensible access control markup language (xacml) and next generation access control (ngac)," in *Proceedings of the 2016 ACM International Workshop on Attribute Based Access Control*, ser. ABAC '16. New York, NY, USA: ACM, 2016, pp. 13–24. [Online]. Available: http://doi.acm.org/10.1145/2875491.2875496

[7] G. F. Anastasi, E. Carlini, M. Coppola, P. Dazzi, A. Lazouski, F. Martinelli, G. Mancini, and P. Mori, "Usage Control in Cloud Federations," *2014 IEEE International Conference on Cloud Engineering*, pp. 141–146, Mar. 2014. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6903467

[8] A. Almutairi and F. Siewe, "Ca-ucon: A context-aware usage control model," in *Proceedings of the 5th ACM International Workshop on Context-Awareness for Self-Managing Systems*, ser. CASEMANS '11. New York, NY, USA: ACM, 2011, pp. 38–43. [Online]. Available: http://doi.acm.org/10.1145/2036146.2036153