

A Distributed Electronic Patient Record System for the Mauritian Healthcare Service

Oveeyen Moonian

Abstract

Developed and developing countries have recognized the importance of ICT in healthcare services and most developed countries are investing important amounts in nation-wide healthcare information systems. At the heart of a nation-wide healthcare information system is the national Electronic Patient Record System.

As per the white paper published by the Ministry of Health in December 2002, Mauritius wishes to have a modern healthcare service of international standards. For such aspirations, Mauritius needs to give due attention to the development of a nation-wide EPR system. This paper briefly describes the Mauritius public healthcare service, motivates the need for a nation-wide EPR system and presents, from an Information Technology perspective, a suitable approach for partitioning and replicating the data and proposes an appropriate architectural model. It also proposes a suitable security scheme for such a system.

Keywords

Electronic Patient Record, Healthcare Service, Architectural Model, Replication, Security.

1. Introduction

The healthcare service is a highly information intensive area, however [Rindfleisch97] considered it as one of the few areas of society where computer access to information has had only limited success. The paper states, however, that the move to widely-accepted Electronic Patient Records (EPR) is accelerating and is inevitable because of many pressures. In an article published in the US News and World Report in August 2005, [Geron 05], Christopher J. Geron refers to the benefits (especially for emergency room (ER) treatments) provided by an online registry containing the healthcare records of 1.5 million patients in Indianapolis. Such a system, he says, provides an ER doctor essential information about his patient's recent hospitalizations and ER visits, medications, allergies, recent lab results, and previous diagnoses and may stop doctors from prescribing medicine that can harm or kill patients in certain situations. The article then emphasizes on the innumerable benefits of a nation-wide EPR system for the US.

by the U.S. Institute of Medicine that nearly 98,000 people die each year as a result of medical errors such as incorrect dosage due to poor legibility in manual records or delays in consolidating needed information to discern the proper intervention. Stephan Kudyba et al [Kudyba05], illustrate how using critical data with predictive modelling, to identify high-risk candidates, can enhance healthcare operational efficiency by initiating strategic management tactics to mitigate the potentially significant costs of fully developed illnesses.

A nationwide EPR can additionally facilitate statistical analysis and help to identify trends in diseases and, furthermore, provide valuable data beyond the boundaries of the healthcare services. More specifically they can prove extremely useful for the police department and insurance companies. Besides, electronic data have vast potential to improve the quality of healthcare audit and research. Thus many developed countries are pursuing efforts to develop EPR systems. In England, implementing EPR systems is one of the main aims of the 10-year National Programme for Information Technology which was launched in 2002 [HC07]. An article in the Psychiatric News, published in December 2006 [PN06], states that 98 percent of Dutch physicians use electronic healthcare records. In a technical report written for the UK National Audit Office in May 2005 [Anderson05], Ross Anderson overviews the

A nation-wide EPR has benefits far beyond emergency room treatment. Such systems generally provide one consistent set of data for each patient at whichever hospital the data is being used or updated. In an article published in 2008 [Abraham08], C. Abraham and his co-authors highlight the importance of having ubiquitous point-of-care access to complete current patient information. The article quotes information reported

Department of Computer
Science & Engineering
Faculty of Engineering
University of Mauritius

ovn@uom.ac.mu

Referencing this article

Moonian, O. (2009). A Distributed Electronic Patient Record System for the Mauritian Healthcare Service [Electronic Version]. *Journal of Health Informatics in Developing Countries*, 3(2), 37-43 from <http://www.jhdc.org/index.php/jhdc/issue/view/8>

healthcare IT infrastructure of eight developed countries apart from the UK, namely France, Germany, Italy, Spain, Sweden, the Netherlands, Poland and the USA. It is to be noted from the report that apart from Poland and Spain, all the other countries have a per capita health expenditure exceeding € 2,000, with that of the USA exceeding € 4,000. The report also states that for most of these countries an EPR is either one of the major projects or one of the priorities in the healthcare sector.

The importance of EPR or Electronic Medical Records (EMR) systems for developing countries is also being recognized by researchers and health authorities. Thus H. Fraser and his colleagues [Fraser05] discuss the potential benefits of EMR for developing countries and describe a number of pilot projects on EMR carried out in different developing countries, namely Kenya, Peru, Haiti, Uganda, Malawi and Brazil. The World Health Organization has published an electronic health record manual for developing countries, in 2006 [WHO06].

Unfortunately, so far, the Mauritius public healthcare service does not possess an integrated EPR system. Although some healthcare institutions are computerized, the data is generally only used internally and their electronic accessibility does not cross the boundaries of the healthcare institution where they are generated. The private clinics also have computerised information systems for internal uses. This paper presents, from a technical perspective of Information Technology, a model for a national patient records system for the Mauritian Healthcare service. After presenting an overview of the Mauritian public healthcare service in section 2, it discusses data partitioning and presents a suitable replication model in section 3 leading to the required architectural model in section 4. Section 5 then looks at security considerations, and section 6 concludes the discussions.

2. The Mauritius Healthcare Service

Mauritius has a free public healthcare service accessible to the entire population. A number of private clinics also provide (paid) healthcare services. The public healthcare service consists of a number of regional hospitals and district hospitals that provide curative healthcare services, and a number of specialized hospitals for services such as Eye treatment, ENT (Ear, Nose and Throat) treatment, Psychiatric treatment [MoH02]. The regional and district hospitals provide services which include, accident and emergency services, general medicine, general and specialised surgery, gynaecology and obstetrics, chest medicine, orthopaedics, traumatology, paediatrics and intensive care services. One of them also provides radiotherapy services. As per statistics available on the website of the Ministry of Health, at end of 2007, there were 5 regional hospitals, 3 district hospitals (essentially similar to the regional hospitals in type of services), 4 specialized hospitals and one high-tech specialized cardiac centre.

To reduce the load on the regional and district hospitals, more than hundred smaller healthcare units have also been set up to provide primary care at the local level. These consist of community health centres, area healthcare centres and mediclinics. Patients usually start their treatments at the local health centres or regional hospitals. Patients who cannot be treated at the local health centres are referred to regional hospitals. For specialized treatments the patients are referred to the specialized hospitals, by the regional hospitals. The specialized hospitals are national level ones. Patients having been treated at the specialized hospitals can have follow-up of their treatments at regional or district hospitals.

As can be seen from the above there is need for a lot of movement of patient data from one healthcare institution to another. In addition to the benefits presented in section 1, a nation-wide Electronic Patient Record system will improve the day-to-day service provided by the public healthcare system. It will enable Mauritius to have a modern healthcare service that can allow easy access to patient data when a patient has been transferred from one healthcare institution to another. Such access will allow doctors to have personalized treatment of patients and easy follow-up on post treatments. It can also allow patients to schedule their own appointments and can allow private clinics to have access to the required information when patients get recourse to their services. Such a system will, furthermore, lead itself to easy statistical analysis and can help the Mauritian healthcare service identify trends of diseases.

3. Data Distribution

The design of a nation-wide EPR should ensure easy accessibility and sharing of data, reliability of the data, distributed management and security of the data. The architectural model to be proposed will depend on the required placement of data, which is determined by the locations of creation and access of the data. This section presents how the data should be partitioned and replicated. These will lead to the required architectural model which will be presented in section 4.

3.1 Data Partitioning

Data for electronic patient records will be essentially of two types: the personal data of patients, which are generally static, and treatment data, which grows with each visit of a patient to a healthcare institution. Since the system will involve a large number of patient records of both types there is need to partition the data. Also, as the data will originate from each of the healthcare institutions, there is need for distributed management of the data. The geographical location of the different healthcare institutions and the nature of services provided by each lead to a natural partitioning of the data. There is first a need of partitioning the static personal data to avoid a bottleneck on the access of the personal information. A horizontal partitioning of data can be performed based on locality of patients' addresses. The proposal is for the personal data at national level to be partitioned

among a number of regional servers handled by the ministry of health but placed at different locations of the country based on addresses of patients. The record for each person can be created at a registration time. This can be either at birth or at a later date, such as a first visit to a hospital, which typically occurs within the first month following the birth for babies born after the system is in place. For the existing population a registration can be performed at a first visit to a hospital or the data can be obtained from existing piecemeal systems.

A vertical partitioning of the data is also required between personal data and treatment information. Treatment data for a patient will originate at each hospital where the patient is undergoing treatment. The treatment data for a patient is expected to keep growing over his lifetime and may originate from different healthcare institutions (primary care centres, regional hospitals, specialized hospitals), at different times. The vertical partitioning can be based on location of treatment. Thus the proposal is that the data for treatment at each regional, district and specialised hospital should be stored with that hospital. The system should adopt a client-server architecture, with each hospital at regional or higher level hosting one or more servers for data originating from that hospital. Primary healthcare institutions do not need to store data. Instead, each regional or district hospital will be responsible for the management and storage of the data of a number of primary healthcare centres, geographically located close to it.

3.2 Replication Scheme

Due to the importance of the data and importance of the sector there is need for continuous availability of the data. If one server is temporarily inaccessible (eg. suffered a crash), its data should continue to be available. To ensure continuous availability, it is important to replicate both the personal and treatment data. For the data to survive a server crash, it is sufficient to have an additional server at each point of storage, mirroring the data. However, such replication cannot guarantee continuous availability in case of a network failure, causing a network partitioning. Besides, replication can be exploited for improved performances, specially, for data that are only read and not modified. Mirroring data as above does not help much in improving performance, since a set of data is only available from one location and the two copies of the same data have to compete for the same network bandwidth. Thus although each regional server will be the primary storage of data for a region this data should be replicated at other locations. For good performance, a copy of the data should be available as closely to the point of use as possible. Considering these two features, the proposal is for the data from each server to be replicated at two other servers, geographically distant from the first one. Thus data for a server in the south may be replicated in the north and west, whereas data from a server located in the north may be replicated in the east and centre. The country will be divided into a number of regions, with each region containing a regional or district hospital. For the purpose of data storage servers at specialised hospitals will be treated in a similar way to the servers at regional and district hospitals. The personal data servers will also use a similar partitioning and replication

scheme for the personal data.

One important issue with replication is ensuring data consistency. Each time a new patient registers with the system, new data will be created at a server. Each time a registered patient goes for treatment, again new data will be generated causing an update of the database at the required server. Since there are multiple copies of each data set, the system should ensure that all copies of a data set are updated in a consistent way, whenever an update occurs. Additionally, since one of the main aims of replication is for fault-tolerance, such as tolerating server crashes, the system should ensure consistency when bringing data at servers up-to-date after a crash. A number of update schemes are available for replicated systems and described in [Coulouris05]. The two main categories are active and passive replication. In active replication, whenever an update occurs at one server, the update is immediately communicated to all other servers. In passive replication, updates occur at one server and are then slowly propagated to other servers. Among the active replication schemes are Quorum Consensus Method, the Available Copies Algorithm and the Virtual Partition Algorithm. Among the passive replication schemes are the Gossip architecture, the Bayou approach and the Coda File System approach.

An important advantage of active replication is that all updates occur at all servers in real time, but the associated cost is that for any update or any data access a number of servers have to be contacted to avoid possible conflicting access/updates. In the Quorum Consensus Method, for example, each update requires the agreement of a majority of servers.

In a distributed healthcare system, it is not crucial that an update be immediately seen at all servers, but it is important that updates be not lost and that any client sees updates consistently. For example, if a client has performed an update of data and that update has been performed by one server, next time the same client obtains data, it expects to at least see the updates it had performed even if the data comes from another server. However, the client need not see immediately updates performed by other clients. Thus a passive replication scheme is sufficient provided that it has the property of ensuring a consistent data view to each client.

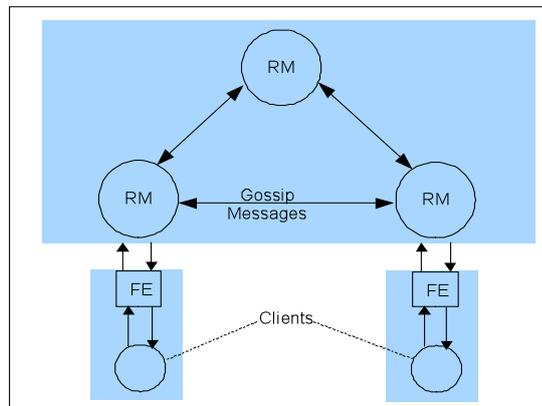
A suitable replication scheme having the above features is the *gossip architecture* [Ladin92]. Figure 1, below, briefly illustrates the type of interactions involved in the gossip architecture. A set of servers holding copies of the same data is viewed as a set of *replica managers* providing the service for these data. All servers holding a copy of a given data have the same importance. A key feature of the gossip architecture is that each client communicates with the set of servers through a service known as a *front end*. Clients are not really concerned about which replica they use as long as the front ends see a consistent view.

The name 'gossip architecture' stems from the fact that updates are communicated among the replica managers through periodic exchange of 'gossip' messages. Whenever independent updates

take place at more than one sever, the gossip architecture ensures that all updates are finally propagated to all servers in a consistent way and that any client always sees at least the last update that it had performed.

and the gossip architecture will ensure that the update is propagated to all servers concerned.

Figure 1
The Gossip Architecture



In the distributed healthcare system, the proposal is for each hospital to have its front end, and thus each hospital be viewed as one client. Although the front end for a hospital will generally communicate with the server at that hospital, in case of server crash it should be able to communicate with other servers. However, to improve reading performance, a cache can be used at each hospital so that if the local server fails, the cache can play the role of the server for data currently being used. Such caches will typically require much less powerful computers than the main data servers but they can communicate with the servers at any hospital. Front-ends also should be replicated to avoid the system of a hospital depending on one component. Since front ends generally deal with a low amount of data, all front ends at a hospital can have concerted actions. Front-ends will generally see two sets of servers, one for treatment data and one for personal data. For any update of data, a front-end will contact the closest available server for the data,

4. Architectural Model

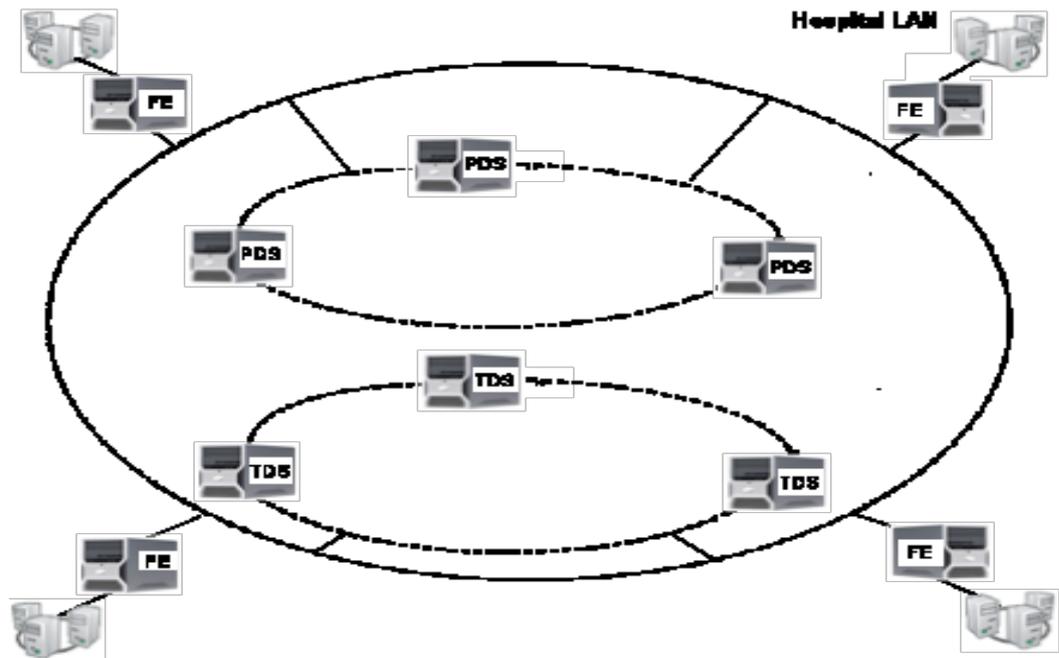
The data partitioning discussed in section 3 leads to two sets of servers which, among them, hold all national patients data. These servers will, for the purpose of the current discussion be referred to as the Personal Data Servers (PDS) and Treatment Data Servers (TDS). The PDS will form one set of servers partitioning and replicating (static) patients' personal data. The TDS partition and replicate treatment data. Figure 2, below, proposes a suitable EPR architectural model for the Mauritian Public Healthcare Service.

The proposal is for each hospital to contain a LAN with computers (generally PCs) in the different wards and other sections of the hospital, such as casualty, X-ray department, blood laboratory. These computers will be clients to the PDS and TDS. All communications between the client machines/processes and the servers will occur through the front-end service which will be provided by a set of machines, working in close coordination so that the failure of any one of them does not affect the data access capability of the hospital LAN.

It is proposed that the ministry of health host a number of PDS distributed in the different regions of the country. To distribute the task of management of treatment data, each regional, district and specialized hospital should be given the task of managing data of its patients' treatment data and thus will have its own TDS.

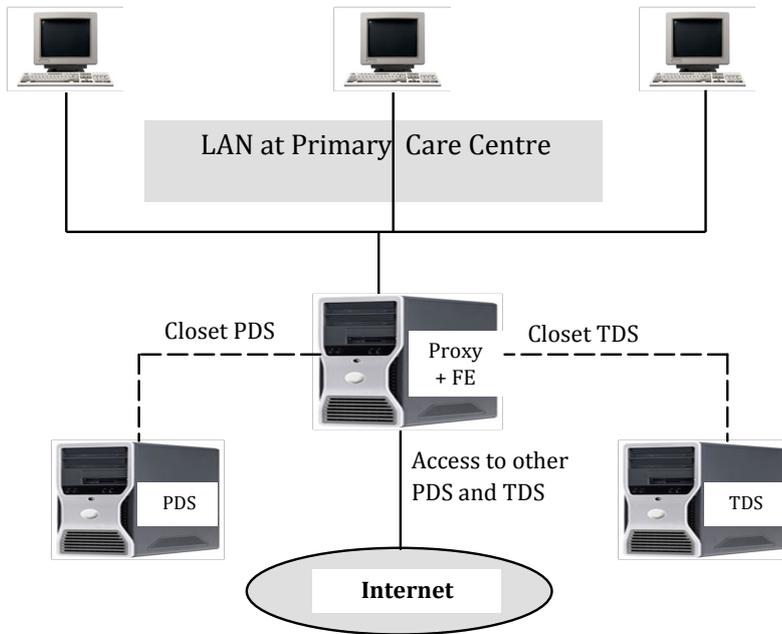
Primary care centres will not host expensive servers for permanent storage of data. Instead, as mentioned in section 3, each regional/district hospital will be made responsible for the data of a number of primary care centres. Each such centre will have relatively inexpensive computers that can

Figure 2
The Proposed Overall IT Architecture for the Public Mauritian Healthcare Service



act as proxy servers. These proxy servers can serve as cache for data of patients receiving treatments in these centres. Whenever a centre requires patient data, the data will be pulled from the main storage at the main server at the hospital responsible for that centre. If that hospital is not accessible, the replication scheme (the gossip architecture) will take the required steps to obtain the data from other relevant servers. The proxy server will also contain the front-end software component and will be the main connection point between the internal LAN at the primary healthcare centre and the rest of the system as shown in Figure 3, below.

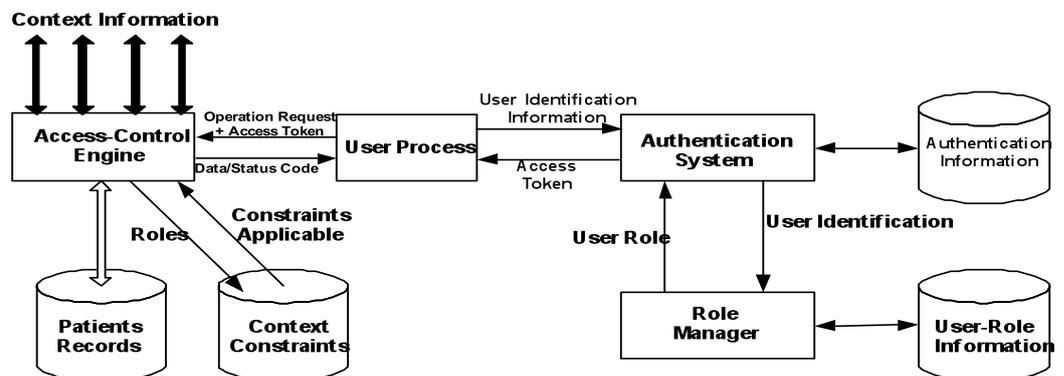
Figure 3
Connection of Primary Care Centres to the Overall IT Architecture



The proposed caching scheme for the proxy servers is the NFS client caching scheme [Coulouris05]. In this scheme whenever an update needs to be performed at a proxy server, the updated information will be immediately be communicated to the main server. Any client (proxy server) holding cached data will regularly check with the main server to ensure that the cached data is up-to-date. System administrators will be responsible for setting the update frequency.

The model proposed in this work is for the public healthcare service. An important and interesting issue will be how private clinics contribute data to the system and use available data. Architecturally, they can have their own LANs with their own

Figure 4
The Essential Components of the Security Model



servers. These servers can act as clients to the main servers, contributing data to the main servers and also pulling information from these servers. However, the legal issues of ownership of data will have to be sorted out among the concerned authorities.

5. The Security Model

Security consideration is an important issue in the system for three main reasons. First, medical records are sensitive information and thus the data access should be closely controlled. Second, the vast amount of information and large number of users involved require an access control system with easy management. Third, since the data is distributed and required to provide remote access it becomes vulnerable to attack over the network. Role-Based Access Control (RBAC) [Sandhu96] has proven to be an efficient way of assigning access rights to users, while ensuring security of the data. To cater for the sensitive nature of the data and the various attacks that can occur, researchers have tried to combine RBAC with various kinds of context constraints [Bertino01, Covington02, Zhang03, Ardagna06, Damiani08].

[Moonian08] presented a security model for the Mauritian healthcare service. The model is based on RBAC, supported by context-awareness. The type of access allowed being linked to roles, it allows users to access data according to the responsibilities associated with their occupations. Thus, users registered as doctors will have different data access from those registered as nurses. To ensure that data access by various kinds of employees are restricted to their duty locations and schedules, the model also include context constraints based on location and time of access.

The essential components of the model are as shown in Figure 4, below. The system contains a Role Manager that handles the assignment and verification of roles. Users access the system through some identification information, such as an RFID tags or login names and passwords. The identification information is passed to the Authentication System which validates the user and contacts the Role Manager to obtain the role of the user. The Authentication system then builds an access token consisting of the user id and the role id and returns this token to the user process. This token will be used for all data accesses.

All data accesses are performed through the Access-Control Engine (ACE). For each data access operation, the user process passes to the ACE the requested operation together with the access token. The ACE goes through a number of steps verifying the kind of context constraints applicable to the requested operation and the required data object(s). If the requested operation is allowable for the given role and the required context constraints are satisfied, the operation is performed and the result is returned to the user application. Otherwise the operation is denied and the user process is informed accordingly.

To provide access control in the national EPR system, the model is to be applied at each point of access of the data. Thus each of the PDS and TDS should implement the access control scheme. Each hospital LAN needs to have a Role Manager to perform Role to User assignment and an ACE to perform access control verification for each requested operation. Although, the primary care centres also need to assign roles, perform authentication and access control, each centre can use these components from the hospital to which it is attached, and thus need not have its own Role Manager and ACE.

6. Conclusion

This paper has motivated the need for a nation-wide Electronic Patient Record system for Mauritius and has provided the technical models for such a system from a perspective of Information Technology. It has described the benefits of a nation-wide EPR in terms saving lives and costs through precision in treatment and avoidance of medical errors by having all information on a patient available at the finger tips.

The paper has presented suitable horizontal and vertical partitioning of patients data and proposes the gossip architecture as the replication model to be applied. The paper proposes an architectural model that can suitably accommodate the proposed data partitioning and replication and discusses a data access control model that can be applied to such a system.

For such a model to be implementable, there is need for a detailed analysis to develop specifications of server configurations, bandwidth, equipment required and the costs involved. Following that, a detailed design must be developed. The implementation will require the setting up of LANs at each hospital, with support for internet connections. However such a system will also involve a number of legal issues that have to be investigated. Future works will also need to consider the use of wireless networks with user mobility.

References

- [Abraham08] Abraham C., Watson R.T., Boudreau M., "Ubiquitous Access: On the Front Lines of Patient Care and Safety", Communications of the ACM, June 2008, Vol. 51, No. 6, pp 95-99.
- [Anderson05] Anderson R., "Healthcare IT in Europe and North America – Foundation for Information Policy Research", Technical Report of the UK National Audit Office, published on 5 May 2005.
- [Ardagna06] Ardagna C.A., Cremonini M., Damiani E., De Capitani di Vimercati S., Samarati P., 2006, "Supporting Location-Based Conditions in Access Control Policies", Proceedings of the ACM Symposium on Information, computer and communications security.
- [Bertino01] Bertino E., Bonatti P.A., Ferrari E., "TRBAC: A Temporal Role-Based Access Control Model", ACM Transactions on Information and System Security, Volume 4, No. 3, August 2001, pp. 191-223.
- [Covington02] Covington M.J., Fogla P., Zhiyuan Z., Ahamad M., "A Context-Aware Security Architecture for Emerging Applications", Proceedings of IEEE 18th Annual Computer Security Applications Conference, 2002, pp 249-258.
- [Coulouris05] Coulouris G., Dolimore J., Kindberg T., "Distributed Systems- Concepts and Design"- 4th Edition, Addison Wesley, 2005.
- [Damiani08] Damiani M.L., Silvestri C., "Towards Movement-Aware Access Control", Position Paper, Proceedings of the SIGSPATIAL ACM GIS 2008 International Workshop on Security and Privacy in GIS and LBS ,Irvine, California, USA, November 04, 2008, pp 39-45.
- [Fraser05] Fraser S.F.H., Biondich P., Moodley D., Choi S., Mamlin B., Szolovits P., "Implementing Electronic Medical Records Systems in Developing Countries", Informatics in Primary Care 2005, Volume 13, pp 83-95, PHCSG, British Computer Society.
- [Geron 05] Geron C.J., Wired hospitals: safety net. Creating a national patient record system is a daunting but worthy goal, US News World Report, Aug 1, 2005, 139(4).
- [HC 07] House of Commons Health Committee, "The Electronic Patient Record", Sixth Report of Session 2006–07, Volume I, Published on 13 September 2007 by authority of the House of Commons, London.
- [Kudyba05] Kudyba S., Hamar G.B., Gandy W.M., "Enhancing Efficiency in the Health Care Industry", Communications of the ACM, December 2005, vol. 48, No. 12, pp 107-110.
- [Ladin92] Ladin R., Liskov B., Shrira L., Ghemawat S., "Providing Availability using Lazy Replication", ACM Transactions on Computer Systems, Vol. 10, No. 4, November 1992, pp 360-391.
- [MoH02] White Paper on "Health Sector Development & Reform", Ministry of Health and Quality of Life, Republic of Mauritius, December 2002.
- [Moonian 08] Moonian O., Khedo K.K., Cheerkoot-Jalim S., Doornun R., Nagowah S.D., Cadarsaib Z. "HCRBAC – An Access Control System for Collaborative Context-Aware HealthCare Services in Mauritius" - Journal of Health Informatics in Developing Countries, Vol. 2 No. 2, 2008.
- [PN06] Psychiatric News, December 1, 2006, Vol. 41, Number 23, published by the American Psychiatric Association.

[Rindfleisch 97]

Rindfleisch T.C., "Privacy, Information Technology and Healthcare", Communications of the ACM, August 1997, Vol. 40 Number 8, pp 93- 100.

[WHO06]

"Electronic Health Records: Manual for Developing Countries", ISBN 9290612177, World Health Organization, 2006.

[Zhang03]

Zhang G. and Parashar.M., Dynamic context-aware access control for grid applications. In IEEE Computer Society Press, editor, *4th International Workshop on Grid Computing (Grid 2003)*, pages 101 – 108, Phoenix, AZ, USA, November 2003.