Technology Development for Building Distributed, Scalable, and Reliable Healthcare Information Store

Technical Annexure

R&D Partners

Centre for Development of Advanced Computing (C-DAC), Pune
&
Swedish Institute of Computer Science (SICS), Sweden

Contents

1 Introduction

With the extensive use of computerized medical equipment in healthcare domain, health data of individuals is often available in electronic form. Several information systems developed to deal with healthcare parameters also generate and handle individuals' healthcare information in electronic form. A few examples of such systems are Hospital Information Systems (HIS), Telemedicine Systems, and Tele-Follow-Up Systems. Several of these systems have been operational for the last few years and have generated healthcare data in electronic form for the treated patients.

The Department of Information Technology (DIT), Ministry of Communication and Information Technology (MCIT) have funded several projects during the past decade for development of IT-based Healthcare solutions. During the course of these projects, the need for a standard EHR for the nation has been strongly felt for interoperable health-care solutions. During the meetings of the National Knowledge Commission for creating the national health informatics vision for India, the need for standard EHR with secure storage and access of EHRs in a storage system spanning national geo-political hierarchy has been discussed. Hence the current proposal for developing technologies for a distributed health care information store is in-line with the vision of DIT and National Knowledge Commission for better health-care services for citizens of India.

1.1 Current Scenario and Associated Problems

The current way of developing and using healthcare information store systems has led to a chaotic state of affairs due to following reasons:

- They have been developed independently and do not easily interoperate with each other.
- They follow their own convention of creating, maintaining, and storing Electronic Health Records (EHRs) of patients.
- If a patient is treated at different hospitals at different instances, different EHRs are generated and stored for the same patient by the two different information systems in use at the two hospitals. A single EHR for an individual is desirable irrespective of his/her time and place of treatment.
- Each system has its own way of creating and managing its storage of EHRs. Obviously, such an information store is based on both relational database technology, due to its ubiquity and maturity in managing large volume of information, and media storage software, for example X-ray picture archiving. Different systems use different relational databases and different media storage software making data transfer across systems impossible/inconvenient.

- Most of the existing systems use centralized storage, leading to limited scalability and poor reliability (single point of failure).

As a result, it has become difficult to exchange EHRs across different systems and to have a unified information system to deal with one EHR per individual, irrespective of the time and place of treatment of an individual.

1.2 Requirements of a Future EHR System

To overcome the problems mentioned above, there is an urgent need to create a framework for building a distributed, scalable, and reliable healthcare information store system that can have a single EHR for every individual of a nation. Building such a system is a challenging task involving several research and development problems. Some of the key issues that will need to be addressed are:

- Evolving a standard for common EHR
- Evolving policies and mechanisms for a single EHR for each individual
- Evolving policies and mechanisms for secure storage and access of EHR
- Building a national distributed, scalable, and reliable healthcare information store for storage and access of all EHRs of a nation

1.3 Current Initiatives at C-DAC and SICS

C-DAC has already initiated efforts towards addressing some of the issues mentioned above. In particular, it is currently working towards:

- Evolving a standard for common EHR
- Evolving policies and mechanisms for a single EHR for each individual
- Evolving policies and mechanisms for secure storage and access of EHRs in a storage system spanning national geo-political hierarchy

C-DAC has extensive experience in developing distributed computing systems and Grid based networks, emulation of Grid, Parallel Processing and Programming architectures. C-DAC is an active contributor to the telemedicine and health informatics domain in India. The organization has developed multiple products servicing the health sector in various nations.

SICS has actively been doing research in the context of distributed, reliable, and scalable information systems. Within the area of distributed computing, the group's focus has been on overlay networks, content-distribution networks, self-management, gossip-based algorithms, middleware, emulation/simulation of large scale distributed systems, and analytical methods for analyzing the performance of these systems. The group has, over the years, built and refined an event-based middleware for building large-scale fault-tolerant distributed systems. This middleware has been the basis of

the group's overlay networks, and will be shared with all consortium partners for the development of EHR products.

2 Further Efforts Needed for a Complete Solution

Although C-DAC's current efforts are targeted towards conceptualizing, evolving, designing, and developing the core framework for a prototype demonstration of a national healthcare information system, several additional issues need to be addressed to transform it into a complete operational solution. That is, to build a nation-wide distributed, scalable, and reliable healthcare information store, further research needs to be carried out towards the following:

2.1 Building a Common Standard EHR

The current way of developing and using healthcare information store systems has led to a chaotic state of affairs due to following reasons:

- Independently developed solutions do not easily interoperate with each other.
- Follow own convention of creating, maintaining, and storing Electronic Health Records (EHRs) of patients.
- Treatment at different hospitals at different instances results in generation of different EHRs for the same patient.
- Own way of creating and managing its storage of EHRs. Different systems use different relational databases and different media storage software makes data transfer across systems impossible/inconvenient.
- Existing systems use centralized storage, leading to limited scalability and poor reliability (single point of failure).

As a result, it has become difficult to exchange EHRs across different systems and to have a unified information system to deal with one EHR per individual, irrespective of the time and place of treatment of an individual.

There is an urgent need to create a framework for building a distributed, scalable, and reliable healthcare information store system that can have a single EHR for every individual of a nation. Building such a system is a challenging task involving several research and development problems.

2.2 Scalable Distributed Storage

Current relational database technology has matured to the level where systems can store terabytes of data in a database cluster using Storage Area Networks. However, centralized information storage architectures provide impediments to scalability and high availability, including a central point of failure in the system, and the need to use expensive high availability hardware.

Recently, an increasing number of companies operating in extreme computing environments (where the rate of updates and access are very high), such as Telecom

and Internet companies, have developed geographically distributed information store architectures to provide both high scalability and high availability. Examples of such systems include telecom products such as Home Location Registers and Internet services such as Google Mail. Typically, these systems partition data geographically, generally using a user-specific identifier, so that user-specific data is located on the same partition, preventing data queries that span multiple, geographically distributed partitions. The partitioning of user-data enables relational database and image stores to handle vastly increased data sizes, and to massively increase read and write throughput by removing the centralized bottleneck in the system.

Obviously, the National EHRs repository will be so large that it will have to be modeled on scalable distributed storage mechanism mentioned above. Suitable database partitioning algorithms will have to be devised for geographical distribution of EHRs, while ensuring that data access is efficient and within acceptable time limits.

2.3 Reliable Distributed Storage

The EHR database also needs to have high availability feature to ensure that the data is always available. Suitable data replication mechanisms need to be researched for making this possible. Since the EHR database will be huge in size, implicit data replication mechanisms are desirable in which the entire replication process is automatically controlled by the system without users' knowledge. That is, the number of replicas of an EHR and their locations should be transparent to the users.

2.4 Need-based Consistent Distributed Storage

Since the EHR database will use data replication technologies for high availability feature, it is important to also research suitable consistency algorithms to guarantee consistency of replicated data.

Existing database systems support the geographic replication of data to ensure that data is always available. However, existing geographic replication support for MySQL and Oracle is based on asynchronous replication: they provide no data consistency guarantees for replicated data. So, in the event of the crash of a partition, there is potential for data loss. This is not acceptable in the health care domain.

Although strict consistency of multiple copies of replicated EHRs is intuitively the most desirable consistency model, this may unnecessarily degrade performance. Moreover, since an EHR is a huge set of information, not all information will require strict consistency of data for the information system to function correctly. Hence, there is a good scope of research to identify and implement weaker consistency models to enable the system to function both efficiently and correctly. The idea is to define a set of consistency models with varying degrees of consistencies and to associate them with different sets of information in an EHR depending on the consistency needs of the system.

2.5 Efficient Distributed Access of EHRs

Client software that accesses data in the distributed repository of EHRs will need to provide routing and lookup functionality to enable data updates and queries to be sent to the partition where the EHR of current interest is stored.

2.6 Technique, Specification, and Framework for Conversion Tools

Several conversion tools also need to be developed for migration of existing healthcare data from existing data stores to the proposed national data store. The project will evaluate, consolidate, refine, and evolve common techniques, standard approaches, and framework for development of conversion tools.

3 Project Proposal

The Joint Project Proposal for Sweden-India Cooperation focuses on addressing the issues mentioned above under the section on "Further Efforts Needed for a Complete Solution". The focus of the proposed project is to:

- Incorporate distributed information store to make the solution scalable.
- Devise policies and mechanisms for creation, unique identification, and proper management of a single EHR for every individual in distributed environment.
- Design and build a generic silo framework for partitioning EHRs. In particular, we propose the use of Distributed Hash Table (DHT) technology to model a nation-wide scalable distributed storage for efficient storage and retrieval of EHRs. The DHT technology has proven successful in building partitioned distributed storage systems, such as OceanStore. Hence, its application to build a partitioned distributed storage system for national EHRs with suitable customization would be researched. Routing and lookup functionality will also be developed to enable data updates and queries to be sent to the partition where EHR of current interest is stored.
- To ensure high availability of EHRs, it is proposed to develop scalable geographic replication algorithms with implicit data replication mechanisms. Need-based consistency algorithms will also be developed to ensure that multiple replicas of an EHR are kept consistent to the extent required. This will enable the system to function both efficiently and correctly.
- Incorporate suitable security policies and mechanisms to prevent corruption or theft of EHR data in distributed environment.
- Design and develop several conversion tools to enable smooth migration of existing healthcare data from existing data stores to the proposed nation-wide distributed data store.

3.1 Roles of SICS (Swedish Institute of Computer Science)

- SICS and C-DAC will jointly work towards evolving an overall framework design for a Distributed EHR store to make the solution scalable.
- SICS will carry out research towards the following:
  - Designing and building a generic silo framework for partitioning EHRs based on Distributed Hash Table (DHT) technology
  - Developing replication algorithms for high availability of EHR data
  - Developing routing and lookup functionality to enable data updates and queries

3.2 Roles of C-DAC (Centre for Development of Advanced Computing)

- SICS and C-DAC will jointly work towards evolving an overall framework design for a Distributed EHR store to make the solution scalable.
- C-DAC will carry out research towards the following:
  - Devising policies and mechanisms for a single, unique EHR for every individual in distributed environment
  - Developing need-based consistency algorithms for replicated EHRs
  - Incorporating security policies and mechanisms to prevent corruption or theft of EHR data in distributed environment
  - Developing technique, specification, and framework for conversion tools

3.3 Estimated Resource Need

To perform the research described above and develop a functional prototype of the distributed information store, following requirements are estimated:

- Project duration of three years
- Approx. 3 people per year (starting with 2 and ramping up to 4) at SICS
- Approx. 10 people per year (starting with 4 and ramping up to 14) at C-DAC

3.4 Other Applications of the Results

The proposed project has a clear focus on developing enabling technologies for a future EHR system. However, the developed technologies would also be applicable to many other distributed applications. For example, large scale distributed databases are relevant for applications in telecommunications (e.g. HLRs) and Internet services (e.g. web mail), while reliable distributed databases are relevant for applications in e.g. disaster relief and robust web services in general.

4 Work Program

Note: M0 signifies Month-0, M1 signifies Month-01 and so on, from start of the project.

4.1 WP 1: State-of-the-art Review

Task Description: The review will cover an analysis of existing replication, consensus algorithms, as well as standard, policies, and mechanisms for a common single EHR, and secure storage and access of EHR.

Time frame: Intensively M0 to M3, thereafter ongoing coverage of related research which results in document updates as needed.

Contribution: SICS 2 person months (PMs), C-DAC 40 person months (PMs).

Dependencies: None.

Deliverable: Report on State-of-the-art on replication, consensus, common standard/policies/mechanisms for single EHR, and secure storage.

4.2 WP 2: Routing to Cluster

Task Description: This task involves the routing of information store read and write operations to the appropriate cluster. We will investigate both DHT and consistent hashing approaches.

Time frame: M0 to M6.

Contribution: SICS 6 person months (PMs) , C-DAC 2 person months (PMs).

Dependencies: None.

Deliverable: Routing Algorithms.

Milestone: Review of Work to Date. First look at the Middleware.

4.3 WP 3: Defining Common EHR Standard

Task Description: This task involves evolving a standard for common EHR for each individual.

Time frame: M3 to M9.

Contribution: SICS 2 person months (PMs), C-DAC 65 person months (PMs).

Dependencies: None.

Deliverable: An EHR standard document.

4.4 WP 4: Cluster Failover Algorithms

Task Description: This task involves developing a suite of algorithms for managing the failover of clusters. It involves developing failure detectors for WANs, data recovery from failed clusters, replication event conflict detection, and replication event conflict resolution.

Time frame: M0 to M18.

Contribution: SICS 11 person months (PMs), C-DAC 2 person months (PMs).
Dependencies: None.

Deliverable: Implementations of the algorithms for Cluster Failover.

4.5 WP 5 Client-Side Disconnected Operation

Task Description: This task involves developing caching algorithms for client-side storage of EHRs during network disconnection, the synchronization of client-side caches with the Information Store.

Time frame: M0 to M24.

Contribution: SICS 11 person months (PMs), C-DAC 2 person months (PMs).

Dependencies: None.

Deliverable: Disconnected operation algorithms and need-based consistency replication.

4.6 WP 6: Distributed Consensus and Replication Algorithms for WANs

Task Description: This task involves the development of distributed consensus algorithms to support synchronous replication of data to multiple clusters. This task

involves designing consensus and replication algorithms to handle delay-tolerant networks that may be common in parts of India.

Time frame: M0 to M30.

Contribution: SICS 29 person months (PMs) , C-DAC 3 person months (PMs).

Dependencies: None.

Deliverable: Design and Evaluation of distributed consensus algorithms, Design and simulation of consensus and replication algorithms for delay-tolerant networks.

Milestone: Review of Work to Date. Set of test algorithms simulated for the system.

4.7 WP 7: Devising Policies and Mechanism for a Common Single EHR and Implementation of Distributed EHR

Task Description: This task involves devising policies and mechanism for a single common EHR, and implementing the common distributed EHR.

Time frame: M9 to M19.

Contribution: SICS 5 person months (PMs), C-DAC 75 person months (PMs).

Dependencies: WP3.

Deliverable: Report on single EHR policies and mechanisms and evaluation of Common Distributed EHR.

4.8 WP 8: Secure Storage and Access of EHR

Task Description: This task involves evolving policies and mechanisms for secure storage and access of EHR.

Time frame: M11 to M24.

Contribution: SICS 6 person months (PMs), C-DAC 60 person months (PMs).

Dependencies: WP3.

Deliverable: Design and Evaluation of secure storage and access of EHR.

Milestone: Implementation of Secure Storage and access mechanisms.

4.9 WP 9: Researching and Implementing Replication and Need-based Consistency Algorithms

Task Description: This task involves researching various replication and need-based consistency algorithms and implementing most efficient algorithms.

Time Frame: M11 to M24.

Contribution: SICS 12 person months (PMs), C-DAC 90 person months (PMs).

Dependencies: None.

Deliverable: Efficient algorithms for replication and need-based consistency.

4.10 WP 10: Prototype Development for Integration of EHR and Algorithms and Validation

Task Description: This task involves developing a prototype for integrating replication and consistency algorithms and Electronic Health Record.

Time Frame: M25 to M36.

Contribution: SICS 3 person months (PMs), C-DAC 110 person months.

Dependencies: Completion of research over replication and consistency algorithms (WP5).

Deliverable: Prototype for Integration process of EHR and algorithms.

4.11 WP 11: Techniques, Specifications, and Framework for Conversion Tools

Task Description: Several conversion tools are to be developed for migration of existing healthcare data from existing data stores to the proposed national data store.

Time frame: M25 to M31.

Contribution: SICS 5 person months (PMs), C-DAC 24 person months (PMs).

Dependencies: None.

Deliverable: Conversion Tools Framework with documentation.

4.12 WP 12: Dissemination of Results and Exploitation of IP

Task Description: The software output of the project will be made available under an open-source license. Commercialization and exploitation opportunities for project Intellectual Property will be realized through regular meetings with the Business Development partners in Sweden and India.

Contribution: SICS 3 person months (PMs), C-DAC 3 person months (PMs).

Dependencies: Other WPs

Time frame: M0 to M36. Ongoing for the duration of the Project.

Deliverable: Publications will be produced throughout the duration of the project.

5 Exploitation and Utilization of Project Results

SICS and C-DAC will exploit the results of their research by publishing their results in high quality journals and conferences. Before publication, the partners will discuss whether the core contribution of the publication should be patented or not. Thus, results which give technological competitive advantage will be patented before public publication. These results should open up new research directions and opportunities to pursue further research funding in the fields of peer-to-peer networks and decentralized optimization of distributed systems.

C-DAC is in the process of setting up a corporate entity for commercialization of its research output. The technologies developed as part of this project will be commercialized through this corporate entity. C-DAC will also make efforts to identify other parties in India who are interested in commercializing the research outputs of this project. Similarly, SICS will make efforts to identify suitable parties in Sweden who are interested in commercializing the research outputs. Together both C-DAC and SICS will also make efforts to identify parties for commercialization of the research outputs internationally.

The common EHR standard which will be developed can be put forward for consideration as an International Standard for defining health records. In particular, the nascent area of decentralized optimization offers the opportunity for SICS to provide global leadership in a growing research area.

6 Risk Analysis

The main risks SICS may face that could affect the success of the project are summarized below in Table 1:

Risk Factor	Probability of Occurrence	Consequences
The design and implementation of client-side routing algorithms and disconnected operation may not be successful.	Low	SICS has extensive experience in building robust P2P and routing systems; instability is unlikely but would affect system performance.
Key personnel will leave the project.	Low	There is a large enough team to handle loss of a person.
The design and implementation of cluster failover algorithms may not be successful.	Low	SICS has extensive experience with cluster-based systems, but failure

		would increase user involvement in data recovery after cluster failure.
The system may be self-managing in normal operation, but has not been designed to handle unexpected situations.	Low	Extensive testing should ensure that this does not happen.
Distributed Consensus Algorithms not as efficient as Asynchronous Replication for Cluster Replication.	Medium	Revert to asynchronous replication algorithms for cluster replication.

Table 1: Summary of Risk Factors for the EHR project – SICS

The main risks C-DAC may face that could affect the success of the project are summarized below in Table 12:

Risk Factor	Probability of Occurrence	Consequences
The definition of national level common EHR and access mechanisms may not be successful	Low	Ample know-how and experience is present in the organization for development of this part of the project as the organization has built a Telemedicine Solution, healthcare solutions, Garuda grid, etc.
Need-based consistency algorithms for replication may not be efficient	Low	There is ample know how as C-DAC has been involved in developing distributed applications over Grid Computing networks
Breach of security in Storage and Access of EHR may occur.	Low	This can be mitigated by using Strong encryption algorithms and devising secure access mechanism. Security can never be guaranteed it can only be strengthened.
Key personnel may leave the project.	Low	There is a large enough team to handle loss of a person. Transfer of Technology among persons will be done on a continuous basis to

	maintain competency in the team

Table 2: Summary of Risk Factors for the EHR project – C-DAC

NOTES