Low-cost nationwide medical imaging storage

Ariel Sabiguero Yawelak
Dirección de Informática - ASSE —— Instituto de Computación - FING - UdelaR
ariel.sabiguero@asse.com.uy | asabigue@{ieee.org|fing.edu.uy}
Montevideo, Uruguay

***Abstract*—ASSE is the largest Uruguayan health services provider. Being State owned, it has national coverage, serving 38% of our population. This work describes ASSE's initiative to provide a nationwide PACS. This paper represent ongoing work, with a clear roadmap and interesting intermediate results. Levels of fault tolerance, availability and reliability of the solution are outstanding while assembling entry level hardware components, open/free software and existing communications infrastructure. Different architectural options are discussed and key practical lessons are presented. Short and mid-term challenges for successful organization imaging success are presented too.**

Keywords - DICOM, PACS, GNU/Linux, DCM4CHEE, COTS, FOSS

I. INTRODUCTION

ASSE[1] (Administración de los Servicios de Salud de Estado - National's Health Services Administration) is the largest national effector of health services. ASSE serves 1.264.643 Uruguayans, which represent 38% of our national population. ASSE administers about 800 different facilities, ranging from hundreds of polyclinics to 61 hospitals. About five million medical consultations are performed yearly.

Until 2009 almost all X-Ray machines were analogical, thus, storage, retrieval and distribution of produced images was based on the possession of the films, but since then, digital-hybrid machines started being deployed. The technological upgrade was inevitable as former technology was being deprecated mostly for international ecological reasons. Straightforward usage of digital equipment through acquisition and printing costs roughly 50% more than previous technology, making the distribution of digital images a must.

The rest of the document is organized as follows: Section II present key technological building blocks used in the assembly of the solution, together with Uruguayan and ASSE details. In Section III ASSE imaging requirements, together with some working hypothesis are presented. With the described problem in mind, different possible architectures are analyzed and discussed in Section IV. Section V describes with further details the selected architecture and provides figures regarding project evolution, current and expected costs. The work concludes on Section VI.

II. STATE OF THE ART

Despite the fact that digital imaging was new to ASSE, it is not new as medical technologies. Academia, industry and standard organizations have been working for almost 20 years in the definition of medical grade storage of digital images and standard formats for image interchange. Two core technological elements considered in current project evolution are the Picture Archiving and Communication Systems, PACS [2] and Digital Imaging and Communications in Medicine, DICOM [3], respectively.

A. DICOM

DICOM is a versatile, adaptable standard for the transmission of digital images addressing in a vendor-independent approach. It has been adopted as preferred image format by radiology, cardiology, endoscopy, MRI, among others. DICOM standard is layered, aligned with the seven layer model of ISO/OSI. DICOM defines an upper layer protocol on top of TCP/IP, being independent of physical, lower layer, networks. Adhesion to DICOM was straightforward. Further analysis of DICOM standard is beyond the scope of this paper.

B. PACS

The PACS concept comprises modalities -different image sources-, visualizers, communication networks and digital archives for information storage and retrieval [4]. A PACS archive requires enough storage to cope with archiving needs, an information system to allow gathering of statistics, analysis, search capabilities and retrieval of specific images and an implementation of DICOM that allows it to interact with other imaging devices. Early works in PACS technology and its definitions started during the '80.

C. Digital Imaging in Uruguay

Previous digital imaging experiences did exist in Uruguay. Most notably, Hospital Maciel[1] and Hospital Dr. Manuel Quintela[2] worked on PACS systems for some areas inside them. Both are teaching hospitals, have deep relations with UdelaR, our national University, and invest on research and innovation. These two examples provides us with local experiences that addresses intra-hospital acquisition, distribution and archiving of images.

Recently our government started making initial announcements of a national health initiative named SaludUY [5]. One of the long term goals is to make the government a keeper of patients medical records and images. Announcements are made that they would build a large PACS system that would interconnect all health effectors and store all images centrally. Their schedule is not firm yet, but only prototype systems would exist until 2016. Despite similarities and differences in our approaches, time frame is not adequate for ASSE. We did need a solution and a road-map sooner than that.

[1] http://www.hmaciel.gub.uy/
[2] http://www.hc.edu.uy/

We introduced key components and national situation that determines the broad scenario where the project take place. The next Section provides a better insight on ASSE's imaging scenario.

III. Imaging and it's requirements in ASSE

It is important to note that IT staff approach to imaging became as a response to an organizational urgent request to decrease costs, more than a consequence of it's evolution to new technologies, opportunities and process re-engineering. By the time IT divisions started interacting with other groups, several computed radiography (CR) devices were already deployed and being used for acquisition and digital printing. Key concern to our work were CR modalities, as they have been recently renewed. The scope of the initiative covers 31 hospitals and their imaging divisions. Approximately half a million CR studies are performed every year, yielding almost 8TB of images.

As soon as our first prototype was deployed, other digital imaging devices, as CT, were connected to our PACS archive solution, forcing us to re-design our solution to cope with possible existing legacy digital imaging devices. The following is a list of requirements, constraints and facts that have to be considered for the solution:

- The PACS system needs to address the distribution of images inside the hospital. In order to minimize printing of images, it is required that visualization workstations are available everywhere. Medical grade workstations were acquired and distributed nationwide together with CR devices, intended for radiography reporting by imaging technicians. Apart from these workstations, visualization grade workstations have to be deployed everywhere where access to images is a must.
- The requirements for the solution are 24x7, as emergency needs are. Availability and robustness must be considered in the design of the solution.
- Serve as the basis for nationwide imaging reporting services that are being defined currently.
- Images should outlive the particular study and be integrated to patient's medical history. This requirement extends life of the digital image in our systems, possibly, for all patient's life.
- In some rural areas of our country our population is served by traveling physicians. They send patients to hospitals for different studies and meet patients again later in time. For those medical consults, images, together with their report, have to be delivered over the 3G network used by MD.
- Long-term openness of the PACS. We know that any solution we deploy today will be replaced sooner than later. For that reason, we must take every possible consideration in order to ease the transition to unknown future systems.
- For a few years now, ASSE evolved to a GNU/Linux organization, thus, visualizers must be Linux compatible.

Other non functional requirements, like transmission and processing times, are present too, but, as it is a new area they are not strongly imposed in this stage. Internally we require the PACS only a few seconds for LAN image distribution times.

IV. Analyzed architectures

How we deploy archive elements has a deep impact in PACS solution as a whole. Different storage options were analyzed in order to provide required archive space [6], [7], [4], [8]. Deep analysis of every option is beyond the scope of this work. Analyzed architectures were: centralized, distributed, hybrid and peer oriented. We will describe the hybrid one, the one being implemented.

A. *Hybrid distributed/centralized architecture*

A purely distributed architecture could be easily improved by adding copies of the studies in a different facility. Combining the centralized and the distributed architecture we can benefit from both strengths. Distribution of mini-PACS solves distribution of images inside hospitals. Adding a forward rule to each mini-PACS that forces them to replicate each study to a central PACS maintains existing workflows and adds an extra layer of redundancy to all studies. There would be a copy inside the hospital where the study is performed and a copy in a central location. The central copy will be done in background, in a few minutes after the image is stored in the mini-PACS. The communication time for this copy do not affect the work in the hospital and with a minimal delay, it is replicated geographically. In this scenario, at least two systems holds every study, thus, when one system is being serviced, the other is still accessible, minimizing even more global PACS downtime.

The availability of a central archive of all images simplifies a nationwide reporting service, as they have access to all required material. After reports are made, they should be delivered to both mini-PACS and central-PACS in order to maintain coherency.

The existence of central copies also eases 3G distribution of images. The central-PACS can be used as a single point of distribution of lossy-compressed images over the 3G network, easing integration with the provider network and simplifying the enforcement of firewall rules.

There is a drawback on a single central-PACS that is the capacity of the system. As systems grow beyond certain point, scaling them up makes us move on super-linear cost curves. This fact has to be considered closely in order not to get into extremely expensive solutions. Properties like availability, survivability and reliability of the PACS are a consequence of its distributed and redundant architecture. We should also consider scale in order to maintain adequate cost parameters.

V. State of the project

The architecture selected is the hybrid one. In the following paragraphs we provide a brief description of key elements.

A. mini-PACS

There was a recent acquisition of general purpose servers, a hundred of them, equipped with 6 core processors, 8GB of RAM and two 500GB hard drives, GNU/Linux compatible hardware. These servers, complemented with enough hard drives, became the mini-PACS hardware.

The software selected for the mini-PACS had to run on top of GNU/Linux, be as free as possible, and count with a minimum local experience. Both local experiences (Hospital Maciel and Hospital Dr. Manuel Quintela) shared a similar platform and run DCM4CHEE [9]. Other health services providers, privately owned, use DCM4CHEE too, like Hospital Británico, Médica Uruguaya, IMPASA and Asociación Española. That was our straightforward option, combined with the knowledge that some local companies provide support and develop extensions for DCM4CHEE.

Each of the mini-PACS is equipped with three 2TB, 3TB or 4TB HDD, RAID-5, depending on the storage needs of the hospital. Each server has dual gigabit connectivity, allowing a physical separation of imaging network from regular hospital network. Gigabit switches are deployed among imaging equipment and a separate numbering space is defined for all imaging networks. Table I resumes relevant figures of the three scales of mini-PACS being deployed.

	small	**medium**	**large**
capacity (GB)	3.725 GB	5.589 GB	7.452 GB
capacity (img)	248.000	372.000	496.000
mini-PACS cost	U$S 960	U$S 1.110	U$S 1.440
cost per GB	U$S 0,26	U$S 0,20	U$S 0,19
cost per image	U$S 0,0039	U$S 0,0030	U$S 0,0029

TABLE I. COMPARISON OF DIFFERENT MINI-PACS.

In all cases, fault tolerant per gigabyte effective cost is adequate and affordable for the requirements of the service. Mini-PACS are dimensioned in order to provide several years of service to each hospital. The distribution of mini-PACS would cost approximately U$S 35.500 for the whole organization.

B. central-PACS

A central PACS was assembled, in order to provide replication, thus, further availability to the solution, and, basis for 3G distribution and reporting services. The central server is equipped with twelve 4TB hard drives, RAID-6. The central-PACS offers 36.248 GB of storage, roughly, 2.417.000 studies, about four years of studies. The cost per gigabyte is of approximately U$S 0,29 and the average cost of a study is U$S 0,0044.

There is not known limitation on the ability to migrate the images from one PACS to another. By the time the archive capacity of the central-PACS is exhausted, so will be the warranty of the hard drives, which is of 5 years. It would be possible to build a bigger one, with future drives of higher capacity, or to add another central-PACS, direct all new studies to the new one and maintain current mini-PACS for historical studies. Everything suggests that it would be possible to build a new one that doubles current capacity and would be able to hold data for a similar period of time. With this approach we can provide a road-map that outlives the first technological migration of deployed system.

C. Deployment status

The project is scheduled to be fully deployed by the end of 2013. Currently only 11 mini-PACS were deployed and all of them are replicated to the central one, providing digital access to more than 30.000 online images. Each image is stored both in the local and central PACS, each of them, with reasonable levels of data redundancy. Local mini-PACS are tolerant to single drive failures and central-PACS is tolerant to double drive failure. Five different, separated drives have to fail simultaneously in order to produce a non-recoverable data loss. Assuming perfect independence of drive failure, the probability of data loss is of $4,6 \times 10^{-15}$ each year. Each replicated image has a theoretical cost of less than U$S 0,0083, using described infrastructure. The approximate budget for the whole PACS should not exceed U$S 50.000.

From the central PACS we are delivering with 3G lossy-compressed images and expect to start tests with physicians in few weeks. Authorization is addressed globally, based on LDAP infrastructure. DICOM authorization is being analyzed for workstations.

VI. CONCLUSIONS

During this project, we managed to provide a FOSS solution combining GNU/Linux and DCM4CHEE that meets current organizational imaging storage and distribution needs with a moderated cost. Provided solution exceeds the characteristics of a centrally, naïve, implementation based on a single central high-end PACS solution. Levels of replication, availability and fault tolerance are extremely high, while keeping bounded overall costs. The solution being deployed not only provides local access and distribution of images, but, 3G access, which was not a must during the beginning of the project.

REFERENCES

[1] (2013) ASSE - Administración de los Servicios de Salud del Estado. [Online]. Available: http://www.asse.com.uy

[2] R. L. Arenson, "Pacs: current status and cost-effectiveness," *European Radiology*, vol. 10, no. 3, pp. S354–S356, 2000. [Online]. Available: http://dx.doi.org/10.1007/PL00014092

[3] (2001) DICOM - Digital Imaging and Communications in Medicine, Strategic Document. [Online]. Available: http://medical.nema.org/dicom/geninfo/dicomstrategyv105/dicomstrategyv105.htm

[4] L. Beird, "The importance of a picture archiving and communications system (pacs) manager for large-scale pacs installations," *Journal of Digital Imaging*, vol. 12, no. 1, pp. 37–37, 1999. [Online]. Available: http://dx.doi.org/10.1007/BF03168750

[5] "Programa Salud.uy - Presidencia MEF MSP AGESIC," http://www.agesic.gub.uy/innovaportal/file/2423/1/presentacion_programa_salud_uy.pdf, Dec. 2012.

[6] J. Edwards, J.C., S. Wong, B. Glicksman, and F. Prior, "Pacs as part of the enterprise: impact of corba," in *Medical Technology Symposium, 1998. Proceedings. Pacific*, 1998, pp. 63–67.

[7] S. Peer, R. Vogl, R. Peer, and W. Jaschke, "Sophisticated hospital information system/radiology information system/picture archiving and communications system (pacs) integration in a large-scale traumatology pacs," *Journal of Digital Imaging*, vol. 12, no. 1, pp. 99–102, 1999. [Online]. Available: http://dx.doi.org/10.1007/BF03168769

[8] J. Carrino, P. Unkel, I. Miller, C. Bowser, M. Freckleton, and T. Johnson, "Large-scale pacs implementation," *Journal of Digital Imaging*, vol. 11, no. 1, pp. 3–7, 1998. [Online]. Available: http://dx.doi.org/10.1007/BF03168246

[9] "Open Source Clinical Image and Object Management," http://www.dcm4che.org/, 2013.