Towards a Deconstructed PACS-as-a-Service System

Roberto Álvarez[1,2(✉)], Jon Haitz Legarreta[1,2], Luis Kabongo[1,2], Gorka Epelde[1,2], and Iván Macía[1,2]

[1] eHealth and Biomedical Applications, Vicomtech-IK4, Donostia-San Sebastián, Spain
{ralvarez, jhlegarreta, lkabongo, gepelde, imacia}@vicomtech.org
[2] Biodonostia, Donostia-San Sebastián, Spain

Abstract. Traditional Picture Archiving and Communication Systems (PACS) were designed for vendor-specific environments, dedicated radiology workstations and scanner consoles. These kinds of systems are becoming obsolete due to two main reasons. Firstly, they don't satisfy the long-standing need in healthcare to put all the resources related to the patient into a single solution rather than a multitude of partial solutions. And secondly, communication, storage and security technologies have demonstrated that they are mature enough to support this demand in other fields. "Vendor Neutral Archives" are becoming the new trend in medical imaging storage and "deconstructed PACS" goes one step beyond proposing a totally decoupled implementation. Our work combines this implementation with the scalability and ubiquitous availability of cloud solutions and internet technologies to provide an architecture of a PACS-as-a-service system that handles a simple enterprise workflow orchestration of tele-radiology. **AQ1**

Keywords: Cloud-based technologies · Deconstructed PACS · DICOM · RSNA · Tele-radiology · VNA

1 Introduction

Traditional medical image archiving systems have been closed, vendor-dependent solutions, requiring non-negligible initial investment and additional maintenance costs. Such setups are commonly designated by the general term Picture Archiving and Communication Systems (PACS) [1], they typically comprise an infrastructure of data acquisition, storage and visualization systems integrated into a digital network, as well as the necessary software to provide such services.

According to [2], one of the biggest problems with traditional PACS arises when the hospital needs to migrate from one PACS vendor to another, typically motivated by the exponential growth and increased complexity of information. Since each solution has a different proprietary implementation of archiving strategies, migrations entail that the complete earlier data needs to be adapted to the format of the newly procured PACS, which is resource consuming. Furthermore, despite that these solutions are

Y.-W. Chen et al. (eds.), *Innovation in Medicine and Healthcare 2017*, Smart Innovation, Systems and Technologies 71, DOI 10.1007/978-3-319-59397-5_25

designed to handle large amounts of data, storage demand is growing enormously, which implies a continuous and increasing number of upgrades in storage capabilities and configurations.

Statistics report that the average number of radiology scans is over 1 per year per inhabitant (different years reported), and trends show increasing figures [3]. Also, it was estimated that 1 billion diagnostic imaging procedures were performed in the United States in 2014, adding up 100 petabytes of volume data [4].

Medical image requirements are heterogeneous across medical specialties and clinical centers. In this context, structured and consistent digital data availability to physicians potentially avoids data loss, scan repetitions or delayed diagnosis. The use of modern web-based data storage, communication and image visualization technology, along with the integration of the Electronic Health Record (EHR) as key element, has the potential to transform and improve new solutions beyond PACS capabilities. Therefore, a pressing need for increasingly flexible and efficient solutions is arising.

This paper proposes an implementation that aims to avoid problems derived from the departmental focus of traditional approaches, which tends to create silos of information. The remainder of the document is structured as follows: Sect. 2 presents a summary of related research and development works, as well as existing commercial products, and other relevant aspects; Sect. 3 presents the specific context the work is aimed at; Sect. 4 provides an overview of the technical approach adopted in order to provide with an evolved solution in such a scenario. Finally, Sect. 5 summarizes the most notable contributions of this work, as well as future research avenues.

2 Related Work

Several approaches have been proposed to overcome the limitations of PACS systems. On a theoretical approach, Pohjonen et al. [5] revisited the grid computing and streaming concepts applied to a PACS system. Vossberg et al. [6] adapted the DICOM protocol to use a grid infrastructure, distributed across different institutions. Yang et al. [7] also used an analogous approach. Although they reportedly improved performance over other web PACS approaches, the grid infrastructure required a complex setup in both cases, and none of these works developed a DICOM viewer. Furthermore, data grids are usually employed when resources are required for solving large-scale, data-intensive scientific applications, which may be beyond medical imaging storage and access requirements in many cases.

Costa et al. [8] developed a PACS solution based on peer-to-peer (P2P) communication models, and utilized document-based indexing techniques within a DICOM network. The system is marketed by BMD Software, Lda. (Aveiro, Portugal). Although P2P may be a solution to increase the dynamics of DICOM nodes, or to reduce the latency, it may not be the most appropriate solution from a security point of view.

Works bridging mobile and PACS technologies are polarized into two distinct approaches with advantages and drawbacks to each. Early attempts to bridge portable devices and PACS technologies required implementing PACS functionality on the devices [9]. This approach was aimed at providing wireless PACS access within a given institution. This type of solution benefits from a seamless integration with the AQ2

PACS. However, it falls short when requiring external access. In such cases, web protocols, which are platform independent and can be leveraged to provide a common interface for multiple platforms, need to be used. Valente et al. [9] proposed a RESTful architecture, employing the well-known open source dcm4che library for DICOM communications. Although a medical imaging viewer was also proposed, their approach required some set-up, even if minimal, on the client side.

Finally, the advent of mobile devices provides flexibility of work, anytime, any-where. Although mobile environments have inherent constraints (such limited display size compared to traditional radiology viewers or limited computing capabilities), the FDA cleared in 2011 the first diagnostic radiology app [10] and the number of cleared apps has grown ever since.

The so-called zero-footprint, web-based DICOM viewers consist of applications that do not require end-users to install any software. Since only a browser is required and run on any platform, the solution is very cost-effective. Different commercial products [11, 12] offer such DICOM viewers. However, although some PACS solutions incorporating web-viewers claim to be zero-footprint solutions, yet many require plugins to run, only run on specific browsers, or do not support tablet or mobile device capabilities without further specialization. Hence, the true advantage of a zero-footprint, such as decreasing cost of ownership and deploying data with minimal or virtually no support, is diminished.

3 PACS Evolution

Transitioning to a new PACS vendor or platform has traditionally involved complex and expensive software and hardware migration efforts, caused by the proprietary mechanisms that have been at the core of such systems. Within recent years, alternative technologies have experienced an unprecedented development that has made the emergence of new actors in the medical imaging industry possible. At the same time, the complexity of the infrastructure required to develop and maintain medical image management systems has been reduced with respect to the amount of stored data. This has allowed novel services to be proposed to healthcare centers that would not have been possible otherwise.

To address these issues, the concept of Vendor Neutral Archives (VNA) [2] has emerged in the last few years. VNA solutions provide hospitals with a decoupled solution for image archiving, so traditional PACS have started to be shifted to these enterprise imaging solutions. With the introduction of VNA and cloud-based digital archives, data from all departments can be put into one pot and accessed using a universal method. This enables the data to be managed by the healthcare system's information technology (IT) staff, rather than individual departments or radiology units.

More recently, in late 2014, the notion of "deconstructed PACS" architecture was introduced as an evolution of VNA systems during the Radiological Society of North America (RSNA) meeting. The underlying idea behind a deconstructed PACS (also referred to as PACS 3.0), is not only decoupling the archive system, but also the diagnostic and clinical viewing of images, the enterprise workflow orchestration and the imaging analytics systems. Deconstructed PACS enable institutions to gain control of their images and optimize workflow, but perhaps most importantly, they can choose the best viewer available for interpretation and also the best analytical processes [13, 14].

4 Proposed Architecture

We have designed and developed an implementation for a deconstructed PACS relying on its hardware-agnosticism and modularity principles. Besides these features, other key aspects for its delivery "as-a-service" have been taken into consideration, such as the adoption of cloud-based solutions, data protection or cost-related issues. The sum of all these elements led our development into a new concept of "Deconstructed PACS-as-a-service" solution. The proposed architecture is depicted in Fig. 1.

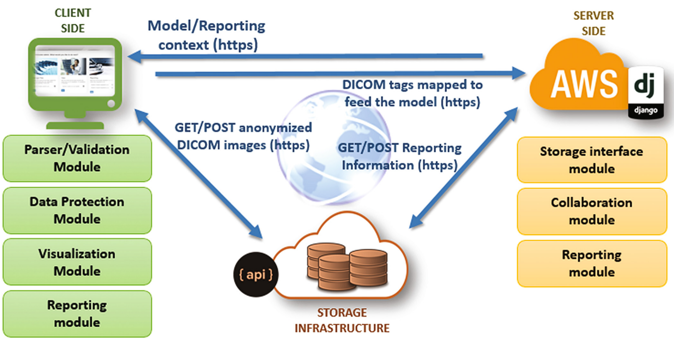

Fig. 1. Proposed architecture.

4.1 Hardware-Related Strategy

The proposed deconstructed PACS-as-a-service relies on a three-tier architecture: a client-side for end-user data access; a storage infrastructure as data repository and a server side to place the backend actions.

4.1.1 Client side

The developed interfaces have followed a responsive design, so that the application can be accessed from both PC or mobile devices. Although according to the IT Reference Guide for the Practicing Radiologist [15], higher resolution displays do not necessarily translate into better diagnostic quality, a set of minimal requirements related to aspects such as calibration or luminance are required to maximize diagnostic accuracy. AQ3

4.1.2 Storage infrastructure

All files managed by the application are stored into a cloud-based storage provider that can be selected by the user, provided that an API to communicate with third parties exists. With this strategy, all storage specific tasks are delegated to the cloud storage service, thus ensuring the right scalability of the system.

4.1.3 Server side

The system can be deployed into any "Platform-as-a-Service" provider. These kinds of services provide a platform that allows customers to develop, run and manage applications without the complexity of building and maintaining the infrastructure. Some of the common tasks they provide are capacity provisioning, load balancing or auto-scaling. We have successfully tested the deployment of our solution into an AWS Elastic Beanstalk environment.

4.2 Software-Related Strategy

The logic of the system is divided into the following modules: an acquisition module in charge of the image gathering process; a visualization module for image inspection; a searching module for locating tasks and finally a collaborative module where we have implemented several simple enterprise workflows.

4.2.1 Acquisition module

Parser/validation module

Digital Imaging and Communications in Medicine (DICOM) [16] is the standard for the communication and management of medical imaging information and related data. One of the goals of this standard is the definition of the Data Dictionary that shall be supplied alongside the images, to achieve a seamless data interchange between digital imaging computer systems.

We have integrated into our solution an adapted version of the dicomParser library [17] which is a lightweight tool for parsing DICOM byte streams in modern HTML5 based web browsers. The system discards non-DICOM-compliant files, and also verifies all mandatory DICOM fields.

Data protection module

The fast growth of digital media solutions has provided great advances in healthcare, but also has brought increased risk related with data protection. According to the European Parliament and Council 2016/679 Regulation of 27 April 2016 [18], data protection issues must be taken into consideration from the design phase of any digital solution. Pursuant to this legal resolution, we have implemented a data protection module in charge of anonymizing all sensitive DICOM tags. This module is based on a JavaScript cross-compilation of DCMTK [19], so it is executed on the client-side, minimizing the risks associated to personal data transfer to third-party cloud storage.

In addition to this data protection module, other security-related strategies have been also developed, such as the adoption of the HTTPS protocol, a 3-step registration design or a password expiration policy.

Storage interface module

We have developed a common interface capable of working with any public API exposed by cloud-based storage providers. The set of files sent to these storage providers comprise the DICOM images, the textual reporting information and the audios created for reporting (explained in more detail within the *Reporting module.* section).

This interface has been successfully tested against the Box Inc. [20] storage provider using the Box Python SDK [21].

4.2.2 Visualization module

We have implemented both 2D and 3D visualization (Fig. 2) modalities. Images are sent directly from the storage provider to the client side, without consuming resources from the backend server which helps to reduce costs and waiting times. Relevant DICOM metadata are presented alongside the image stacks. Furthermore, the system allows users to enrich this metadata with descriptors and comments.

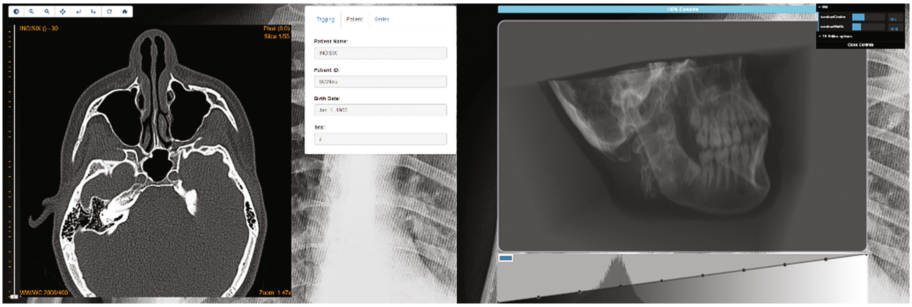

Fig. 2. Partial captures of the 2D (left) and 3D (right) visualizations of a dataset.

2D Visualization

We have integrated the Cornerstone JavaScript library [22] in order to provide 2D visualization capabilities in browsers that support the HTML5 canvas element. The Cornerstone WADO Image Loader [23] is used as the engine for retrieving DICOM images over HTTP. In addition, we have included alongside the display a set of common medical imaging interaction tools [24], such us window-level, pan, zoom or rotate. To allow a smoother navigation, datasets are delivered and rendered accordingly in cached blocks of 13 slices (this number was empirically set). Table 1 depicts caching and rendering times of one of these blocks with different resolutions.

Table 1. 2D visualization caching and rendering time measurements.

Resolution (size)	128 × 128 × 13 (0.5 MB)	256 × 256 × 13 (1.63 MB)	512 × 512 × 13 (6.53 MB)	1024 × 1024 × 13 (26 MB)
Cache/Render time	3510 ms	3755 ms	4660 ms	6919 ms

3D Visualization

Some applications rely on the ability of the viewer to offer 3D visualization for a better appreciation of patient anatomy and location of possible pathologies. For that purpose, we have included a web implementation of Direct Volume Rendering (DVR) [25, 26] using the X3DOM Library that offers implementation for most common representation use cases of volumetric medical imaging data [27]. Table 2 depicts mean preparation times for datasets with different resolutions before interactive rendering.

Table 2. 3D visualization rendering time measurements.

Resolution (size)	128 × 128 × 30 (1 MB)	256 × 256 × 13 (3.77 MB)	512 × 512 × 30 (15 MB)	1024 × 1024 × 30 (60 MB)
Render time	6301 ms	6681 ms	7740 ms	18202 ms

4.2.3 Search module

We have developed a retrieval tool for medical images capable of performing queries within a pre-defined set of fields including the name of the patient, the description of the study, the description of the series, the descriptors assigned to the images or the annotations made by the reporter.

4.2.4 Collaboration module

This module oversees ruling the three collaboration actions that users can perform with their images: (1) **Share**: users can share images with read-only permissions (2) **Request for reporting**: users can request a report for their cases. This action will emit a request for reporting to the recipients whom will be able to accept or reject the request (3) **Transfer ownership**: users can transfer the ownership of their series.

Reporting module. We have developed a specific module to manage the reporting requests. The main tasks performed by this module are:

- Manage the workflow of the request, providing state-aware information to both involved users (requestor and recipient).
- Provide with a set of 2D image interaction tools specific for reporting needs, such as Region of Interest (ROI) delineation, measurements (distances, angles), probe, annotation and screenshot.
- Provide users with three different, non-excluding ways of reporting cases:
  - RSNA templates: this module is able to communicate with the Radreport API [28] in order to assign to the requested case a "structured" report template which the recipient is expected to fill in. According to the RSNA Reporting Initiative, these templates: (1) Improve communication between radiologists and referring providers. (2) Enable radiology practices to meet accreditation criteria and (3) Help radiology practices earn pay-for-performance incentives.
  - Free text: recipients can report requested cases using free text. We have integrated the Web Speech API [29] to facilitate this task. This allows users to dictate the reports using a microphone; the Web Speech API automatically performs the speech-to-text process.
  - Audio files: cases can be reported using audio files. Once the speech has been recorded, the audio clips are available to be listened to within the application at a later time, using an embedded player.
- Export the reported information to a PDF file.

All these modules, as well as the wrapper that orchestrates them (the Python web framework 'Django') are based on open-source products, which facilitates the flexibility and maintainability of the solution.

5 Conclusions and Future Work

We have achieved to use existing web-based technologies like HTML5, JavaScript and WebGL to build a decoupled, scalable and secured tele-radiology solution based on extensively used cloud platforms like Amazon Web Services for web servers and Box for cloud storage. This work could define a new product concept called deconstructed PACS-as-a-Service to provide easy access of medical image content between institutions, their professionals and their patients.

According to the annual title "Imaging & Oncology" published by The Society of Radiographers [30], the adoption of this kind of solution would have a wide range of benefits for the 3 actors involved in a radiology process.

Patients can have immediate access to expert opinion, they benefit from enhanced speed and turnaround time of experts, from immediate availability of appropriate emergency investigations and immediate availability of a second opinion. Second, hospital management can supplement the lack of local expertise, it can face shortfalls of radiology provisions, balance demand variations, provide continuous service and availability of expert opinion. And finally, radiologists can review images from anywhere, they can manage more effectively their workload and provide their services across the world, along with the benefits of underlying stable working patterns and improved work life balance.

We have covered three aspects that typically compose this kind of hardware-agnostic solutions: image management and archiving, simple enterprise workflows orchestration and clinical viewing of images for diagnostic. Our work is modular and proves the maturity of existing technologies to support tasks like tele-radiology and remote collaborative exploration of medical images. Other aspects of deconstructed PACS approaches, like imaging analytics (along with all the necessary middleware), are being considered for future work. First, we will carry out a thorough study of more complex enterprise workflows that lead these kinds of processes in order to implement mechanisms to facilitate their integration. Secondly, we will work on improvements interoperability-related to our system, focusing on integrating information coming from the EHR using the HL7 FHIR (Fast Healthcare Interoperability Resources) [31] standard. DICOM tags will mapped to a FHIR resource instead of using an ad-hoc model. Another aspect to be considered is the fast deployment of automatic or semi-automatic image analysis processes (e.g. deep learning algorithms for detection, annotation, segmentation or screening) making use of actual technical products currently mature, to ship isolated processes fast (e.g. Docker containers). Finally, we will advance towards a unification of reporting criteria, implementing a comparison feature based on semantic analytics able to detect disparities between two different reports belonging to the same case.

Acknowledgments. This research work has been achieved in collaboration with the companies Bilbomática S.A., Derten Sistemas S.A. and Centro Médico Udalaitz S.A. It is part of the research project named "RADMOVE" (file reference # IG-2014/01264, IG-2015/00247, ZL-2016/00048) under HAZITEK support programs for the experimental development and industrial research projects of the government of the Basque Country (Spain).

References

1. Duerinckx, A.J.: Introduction to two PACS '82 panel discussions edited by André J. Duerinckx, M.D., Ph.D.: 'Equipment Manufacturers' view on PACS' and 'The Medical Community's View on PACS'. J. Digit. Imaging **16**(1), 29–31 (2003)
2. Agarwal, T.K., Sanjeev, : Vendor neutral archive in PACS. Indian J. Radiol. Imaging **22**(4), 242–245 (2012)
3. European Commission: Radiation Protection, Medical Radiation Exposure of the European Population, Part 1/2. Frost & Sullivan, Luxembourg, 180, November 2012
4. Frost & Sullivan: Prepare for Disasters & Tackle Terabytes When Evaluating Medical Image Archiving. Frost & Sullivan (2008)
5. Pohjonen, H., Ross, P., Blickman, J.G., Kamman, R.: Pervasive access to images and data - the use of computing grids and mobile/wireless devices across healthcare enterprises. IEEE Trans. Inf. Technol. Biomed. **11**(1), 81–86 (2007)
6. Vossberg, M., Tolxdorff, T., Krefting, D.: DICOM image communication in globus-based medical grids. IEEE Trans. Inf. Technol. Biomed. **12**(2), 145–153 (2008)
7. Yang, C.-T., Chen, C.-H., Yang, M.-F.: Implementation of a medical image file accessing system in co-allocation data grids. Future Gener. Comput. Syst. **26**(8), 1127–1140 (2010)
8. Costa, C., Ferreira, C., Bastião, L., Ribeiro, L., Silva, A., Oliveira, J.L.: Dicoogle - an open source peer-to-peer PACS. J. Digit. Imaging **24**(5), 848–856 (2010)
9. Valente, F., Viana-Ferreira, C., Costa, C., Oliveira, J.L.: A RESTful image gateway for multiple medical image repositories. IEEE Trans. Inf. Technol. Biomed. **16**(3), 356–364 (2012)
10. Press Announcements - FDA clears first diagnostic radiology application for mobile devices. http://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm242295.htm. Accessed 06 Feb 2017
11. HTML5 Zero-footprint viewer for DICOM and PACS. LEAD Technologies, inc., 1927 South Tryon Street Suite 200 Charlotte, NC 28203 USA (2015)
12. Vue Motion: Carestream's Enterprise Access Viewer Connecting Technologies, Physicians, and Patients. Carestream Health, Inc. (2012)
13. Bassett, M.: PACS 3.0: the next iteration of radiology's reading platform. Radiol. Bus. J., 40–44, October/November 2016 AQ4
14. Winthrop Resources: Finding ROI in Deconstructed PACS (2015)
15. Krupinski, E.A., Flynn, M.J.: IT Reference Guide for the Practicing Radiologist (2013)
16. DICOM Homepage. http://dicom.nema.org/. Accessed 30 Jan 2017
17. Hafey, C.: dicomParser - JavaScript parser for DICOM Part 10 data. GitHub. https://github.com/chafey/dicomParser. Accessed 30 Jan 2017
18. The European Parliament and the Council of the European Union: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016
19. dicom.offis.de - DICOM Software made by OFFIS - DCMTK - DICOM Toolkit. http://dicom.offis.de/dcmtk.php.en. Accessed 30 Jan 2017
20. Box Inc.: Box Content API. https://docs.box.com/docs. Accessed 30 Jan 2017
21. box/box-python-sdk. GitHub. https://github.com/box/box-python-sdk. Accessed 13 Mar 2017
22. chafey/cornerstone. GitHub. https://github.com/chafey/cornerstone. Accessed 30 Jan 2017
23. chafey/cornerstoneWADOImageLoader. GitHub. https://github.com/chafey/cornerstoneWADOImageLoader. Accessed 30 Jan 2017
24. chafey/cornerstoneTools. GitHub. https://github.com/chafey/cornerstoneTools. Accessed 30 Jan 2017

25. Arbelaiz, A., Moreno, A., Kabongo, L., García-Alonso, A.: Volume visualization tools for medical applications in ubiquitous platforms. In: eHealth 360°, pp. 443–450. Springer, Cham (2017)
26. Congote, J., Segura, A., Kabongo, L., Moreno, A., Posada, J., Ruiz, O.: Interactive visualization of volumetric data with WebGL in real-time. In: Proceedings of the 16th International Conference on 3D Web Technology, New York, NY, USA, pp. 137–146 (2011)
27. Arbelaiz, A., Moreno, A., Kabongo, L., García-Alonso, A.: X3DOM volume rendering component for web content developers. Multimed. Tools Appl., 1–30, July 2016
28. Developers | RadReport.org. http://radreport.org/dev/. Accessed 31 Jan 2017
29. Web Speech API Specification. https://dvcs.w3.org/hg/speech-api/raw-file/tip/speechapi.html. Accessed 31 Jan 2017
30. Dubbins, P.: Teleradiology – how much flatter is the world now? Imaging Oncol. **2009**, 56–59 (2009)
31. Index - FHIR v1.0.2. https://www.hl7.org/fhir/index.html. Accessed 31 Jan 2017