Volume : 4 | Issue : 6 | Special Issue June 2015 • ISSN No 2277 - 8179

Scalable and Secure Sharing of Personal Health Records in Cloud Computing using Attribute-Based Encryption

Engineering

KEYWORDS: Cloud computing, electronic health record, personal health record, personalized medicine, radiology.

D.Vikkiramapandian	Assistant Professor , Department of Information Technology, K.L.N. College of Information Technology, Sivagangai Dt, Tamilnadu, India
A.Midhunkumar	Assistant Professor , Department of Information Technology, K.L.N. College of Information Technology, Sivagangai Dt, Tamilnadu, India
R.T.Subhalakshmi	Assistant Professor , Department of Information Technology, K.L.N. College of Information Technology, Sivagangai Dt, Tamilnadu, India

ABSTRACT

Personal Health Records (PHRs) should remain the lifelong property of patients, who should be able to show them conveniently and securely to selected caregivers and institutions. In this paper, we present My PHR Machines, a cloud-based PHR system taking a radically new architectural solution to health record portability. In My PHR Machines, health-related data and the application software to view and/or analyze it are separately deployed in the PHR system. After uploading their medical data to My PHR Machines, patients can access them again from remote virtual machines that contain the right software to visualize and analyze them without any need for conversion. Patients can share their remote virtual machine session with selected caregivers, who will need only a Web browser to access the pre-loaded fragments of their lifelong PHR. We discuss a prototype of My PHR Machines applied to two use cases, i.e., radiology image sharing and personalized medicine.

I. INTRODUCTION

Health record (PHR) as "a set of computer-based tools that allow people to access and coordinate their lifelong health information and make appropriate parts of it available to those who need it" [1]. PHRs should be portable, i.e., remain with the patient, contain lifelong information, and should not be re-stricted by file formats or other local issues [2]. In other words, they are electronic health records (EHRs) that are owned by pa-tients. These are usually opposed to hospitals' electronic medi-cal records (EMRs), which only contain medical data generated within one specific care institution.

Sustainability in this context refers to the financial and political aspects of the health care and software industries. Point

(1) focuses on raw PHR data since care institutions may not be able or willing to provide their EHR data in "one" standardized PHR format. Such standards are useful and slowly emerging, but we argue that regardless of such evolution, patients should already be em-powered with the ability to manage their own (potentially raw) data. With point

(2) we aim at the so-called functional interoperability (i.e., "the ability of two or more systems to exchange information so that it is human readable by the receiver" [4]). Concretely, we aim at providing patients (and their trusted care-givers) remote desktop or tablet computer access to all their PHR data, and support this access by the software that matches the data format. Since we do not tackle semantic data integration in this paper, one can more specifically label this as health record mobility and portability.

Cloud computing offers unique opportunities for support-ing long-term record preservation. we present My PHR Machines, a cloud-based PHR system that answers our research question.

One of the agreed key requirements for share-ability of the EHR is to break the nexus between the EHR and the EHR system [4]. The My PHR Machines architecture clearly separates PHR data from the software to work with these data. This paper demonstrates how this creates novel opportunities for the market of PHR software services without compromising patient privacy.

Commercial PHR systems positioning themselves within the cloud computing paradigm are emerging. For example, See-MyRadiology [6] enables patients to upload their medical im-ages and then selectively share these with caregivers. Unfor-tunately, such so-called software-as-a-service (SaaS) systems are typically (1) specialized for one medical function and (2) specifically programmed for web browsers. The SeeMyRadiol-ogy example indeed consists of a DICOM viewer that has been programmed in HTML 5 and related technologies. MyPHRMa-chines is an academic prototype that is more generally applica-ble since it exposes to its users the so-called infrastructure-as-a-service (IaaS) tier of cloud architectures [7]. In a nutshell, the system provides infrastructure to (1) store and share (subsets of) patient data and (2) deploy and use specialized software in remote virtual machines (Vms).

My PHR Machines allows patients to build PHRs which are robust across the space and time dimensions.

Space: Patients relocating or simply traveling across different countries during their lifetime will always be able to reproduce their original health records and the software required to analyze/visualize those. This is often currently not possible because of the high functional and architectural heterogeneity of health care information systems across different countries/states [8].

Time: As technology evolves, application software typically becomes obsolete.

The software to create the idealized environments on con- temporary hard- and software is maintained by big vendors, regardless of the My PHR Machines-specific extensions. On the client-side, My PHR Machines does rely on contemporary web technologies, but only to realize a generic remote desktop client. Hence, also client software maintenance is decoupled from the number and complexity of PHR software services. PHR systems typically offer functionality to share, visualize, and analyze PHR data [10]. My PHR Machines also enables its users to share software to work with the health-related data, keeping data and software clearly separated in the system architecture. Having separate data and functionality also allows a finer grained delegation of access to different stakeholders. Specifically, My PHR Machines allows patients to selectively reveal health information to other stakeholders and it guarantees (spine curvature of less than 20) and discopathy (intervertebral disk fracture) due to physical traumas. The diagnosis and treatment of such conditions is not an easy task and physicians often tend to waive intensive and expensive treatment referring the patient to physiotherapy or even commercial fitness clubs for palliative therapy. The condi-

tion, however, may remain latent for years and reappear in the long run. The decision to start a professional, long-term revalidation program may be postponed too long especially when caregivers lack access to prior scans and analyses. This use case concerns the medical history of a real patient of the Belgian health care system affected by the previously mentioned condition. For reasons of privacy, the case has been made anonymous. The medical history of the patient can be synthesized as follows.

Before discussing the implementation and application of My PHR Machines, we introduce two use cases exemplifying the potential for innovation in health care brought about by our prototype. The first use case concerns radiology image sharing, showing how My PHR Machines can be used to build and main-tain efficiently a lifelong PHR of radiology images. The second use case concerns genomic data analysis in the context of per-sonalized medicine, showing how the separation between PHR data and PHR functionality allows finer grained privacy-related control over PHR data access and utilization.

III. DESIGN AND IMPLEMENTATION OF MY PHR MACHINES

The main idea behind My PHR Machines is to leverage the cloud for allowing patients building their own personal health data repository and share these data with different care institu-tions. In the current implementation, patients have to manually upload the data they obtained from care institutions, e.g., in a DICOM CD, in the repository. In a near future, we envision that care institutions could directly push patient data to the repos-itory. Once stored in My PHR Machines, patients can flexibly share these data with any other care institution or interested stakeholder. Access to My PHR Machines, in fact, requires only a Java-enabled browser, and access to a selected part of the repository can be easily granted by patients to any care institu-tion, e.g., a GP, a hospital, or an insurance company, health-related data. In this way, caregivers need not be able to run specialist software, since they can get access to this software directly from the cloud.

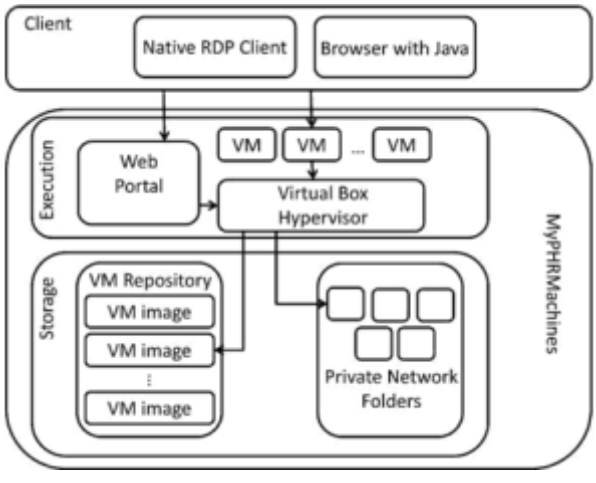

Fig. 1. Technical architecture of My PHR Machines.

A. Technical Architecture

Fig. 1 shows the technical architecture of My PHR Machines, identifying, besides the components constituting MyPHRMa-chines, also the components of the front-end Client.

The prototype reuses parts of SHARE [15], a mature system for making computational research results more accessible and reproducible. The key technological components have, there-fore, undergone various development cycles, which adds to the robustness of My PHR Machines technical architecture. On the one hand, My PHR Machines excludes functionality developed for the SHARE-specific use cases (e.g., generating BibTeX code for conveniently citing a VM image from a research paper). On the other hand, My PHR Machines requires the development of new functionality specific to the PHR context (e.g., access del-egation to a VM session). We also redesigned the user interface of the Web portal to become simpler and coherent to facilitate access by non expert users.

Within My PHR Machines, we distinguish between the Execution and Storage layers. Each VM in the execution layer represents the virtualization of specific application software (or a software bundle) serving the purpose of either viewing or analyzing patients' health data. Patients can log into

My PHR Machines and decide which VM to load in a given session using a standard Web portal. The Hypervisor is a generic piece of software to start, stop, clone VMs, and control their Internet access. For our prototype, we decided to use VirtualBox, an off-the-shelf hypervisor. Being heavily used in several industries, Virtual Box benefits from periodic functionality updates and security reviews. Note that, as discussed more in depth later, the VMs for specialist software are stateless and deprived of Internet access.

The storage layer includes the repository of VM images, i.e., virtual disks containing a bootable operating system and addi-tional applications. Patient-specific VMs are simple instances of these VM images. In order to publish new VM images, software vendors go through the following procedure: first, they clone an existing VM containing the right operating system and perhaps some additional libraries of interest through the MyPHRMa-chines Web portal.

Finally, the vendor "publishes" the VM image for other users of My PHR Machines. Users can-not change the published VM image since any personal instance of a VM image is stateless. By keeping VM instances stateless, one can deploy updates at the VM image level, which is much more scalable and secure than trying to do this at the level of patient-specific Vms.

The labor cost of requesting a VM clone via the MyPHRMa-chines portal is negligible. Other labor costs relate to (1) up-loading application executables to My PHR Machines and (2) configuring them in an instance of the new VM image. The first cost is unavoidable since by definition it becomes relevant any time a software vendor wants to deploy software to a cloud-based system. About the second cost, any IaaS-based approach would provide the same level of flexibility as My PHR Machines.

However, in more general IaaS platforms (e.g., Amazon EC [2]), VM images would have to be cloned explicitly for each end-user. Also, end-users would be able to change the VM images, intro-ducing huge maintenance costs. Instead, the My PHR Machines approach of using stateless VM sessions (i.e., sessions that do not affect the VM image involved) avoids that cost problem by design.

The PHR data are stored into network folders, which remain private folders within the My PHR Machines domain. Put differ-ently, the VM-based architecture ensures that all patient data can remain on the server-side, on a trusted infrastructure. The latter feature, combined with stateless VMs deprived of Internet ac-cess, guarantees the privacy of the patient's health-related data. In particular, even if software in one of the VMs is programmed with some sort of malware, this will not be able to push PHR data outside the network domain of My PHR Machines.

Clients can view remote VM sessions using the remote desk-top protocol (RDP). Therefore, VM sessions can be viewed in any Java-enabled Web browser without installing any addi-tional software, by using a simple applet-based RDP viewer. For operating systems not supporting Java, e.g., iOS, a na-tive RDP client is required. All communications between the Web portal and the hypervisor are delivered via SSH, a se-cure and stable communication protocol that, among others, provides measures to prevent man-in-the-middle (MITM) attacks. MITM attacks can also be prevented for the HTTPS traf-fic between a client browser and the My PHR Machines web server Since My PHR Machines is currently deployed only as a prototype, we have

not yet invested in the required certifi-cates issued by a major certification authority. MITM attacks can also be prevented for the RDP traffic: the My PHR Machines hypervisor supports RDP over TLS [18] so both RDP client and server identity can be protected using certificates. Again, for the prototype deployment, certificates are not yet used. Cer-tificate configuration needs to be handled once at the level of the web server and once at the level of the hypervisor. VM-level certificates are not needed, so the scalability of the architecture is safeguarded.

One important downside of sending around URLs that pro-vide direct VM access is that, without additional security mea-sures, the access delegation messages could be intercepted by malicious Internet users. Fortunately, care institutions are likely to have secure messaging tools in place and, therefore, the access delegation message can be sent securely from the MyPHRMa-chines web server to the inbox of the caregiver. Therefore, we do not consider this as a major threat.

These images can be offered to other software vendors, who want to specialize in offering end-user oriented VM images. Steps 16 to 18 involve publishing a VM image to a library. The library concept is not only important both to separate the developer-oriented images from the end-user oriented ones, but also to organize end-user images in various more fine-grained categories (e.g., per medical condition or per insurance plan). After having presented the technical architecture of our pro-totype, we can now go back to the requirements listed in Sec-tion II and discuss how My PHR Machines addresses explicitly all of them.

About requirement RA1, PHR data can be stored by patients using the cloud storage provided by My PHR Machines. When required, PHR data can be easily exposed to care institutions, e.g., a physician, as long as an Internet connection and a Java-enabled browser are available. In particular, My PHR Machines enables the virtualization of any type of operating systems and application software. These will remain available to patients and care institutions even when they are no longer in use or accepted in practice.

The requirement RB1 is implemented as a feature of the Web portal. When launching a VM, in fact, patients can select only part of their PHR data currently available within MyPHRMa-chines to be shared with a given care institution. Finally, the requirement RB2 is forced by design because, as we discussed before, VMs do not have Internet access and, therefore, the PHR data used by them cannot be pushed outside the domain of My PHR Machines to pursue improper use. Having VMs without Internet connection may represent a limitation of our prototype. My PHR Machines ensures that even when end-users or applications tamper with the firewall settings of a VM, no harm can be done.

A second radiology related VM contains a specialized DICOM viewer that can also visualize DICOM data in case a viewer has not been embedded in the hospital-provided DICOM CD (or DVD). This is the case for example for CDs containing DICOM data of cone beam computed tomography scans.

The VM for the personalized medicine use case combines the DNA data of an anonymous patient available on the Inter-net with the open source Promethease software as application software, in this case to analyze the PHR data. Fig. 10 shows an example report generated by Promethease within MyPHRMa-chines. Note that the information generated by VMs cannot by default be pushed out of the My PHR Machines network do-main by a possibly malicious implementation of the applica-tion software because Internet access is by default disabled for VMs. For such cases, MyPHRMa-chines VMs can be given Internet access through a virtual net-work proxy. My PHR Machines platform administrators can de-fine fine-grained policies, e.g., to give the trusted Promethease VM access to the Internet address of the genomic expert rule repository. An alternative solution would be for the application software vendor to routinely update its provided VM image with the latest expert rules for genome interpretation.

VI. CONCLUSION

Leveraging virtualization techniques, My PHR Machines allows patients to build lifelong PHRs. The records can be shared by the patient with any stakeholder interested in those. My PHR Machines allows also the controlled sharing of application software that is required to view and/or analyze health records. Moreover, as technology evolves, patients will always be able to use original software to view and analyze data, even when that software becomes obsolete and possibly no longer supported by the stakeholder that produced the data.

Besides a clinical experimentation, to fairly assess patients' propensity in using such an innovative PHR system as My PHR Machines, we are currently working on extending our prototype in several ways. One of the major extensions regards creating an open App market for application software, through which medical software providers could compete to provide the best suited functionality required by patients. We are currently studying the issue of how various security techniques can be employed to protect data in My PHR Machines at various levels, such as encryption techniques at the level of VM instance logs, private key transfers between RDP clients and remote VMs, and encryption at the level of mounted network folders. Further-more, we are surveying practitioners to understand more broadly and deeply the specific use cases for which My PHR Machines forms a unique enabler. Finally, we will deploy data translation services to My PHR Machines. Such services will enable a smooth transition from the already provided functional inter-operability to the deeper system interoperability. The private network folders will be used as the blackboard for exchanging data between different VMs.

REFERENCE

[1] D. C. Kaelber, A. K. Jha, D. Johnston, B. Middleton, and D. W. Bates, "Viewpoint paper: A research agenda for personal health records (PHRs)," J. Amer. Med. Inform. Assoc., vol. 15, no. 6, pp. 729–736, 2008. | [2] AHIMA e-HIM Personal Health Record Work Group, "Defining the personal health record," J. AHIMA, vol. 76, no. 6, pp. 24–25, Jun. 2005 | [3] . P. C. Tang, J. S. Ash, D. W. Bates, J. M. Overhage, and D. Z. Sands, "White paper: Personal health records: Definitions, benefits, and strategies for overcoming barriers to adoption," J. Amer. Med. Inform. Assoc., vol. 13, no. 2, pp. 121–126, 2006. | [4] Health Informatics—Electronic Health Record—Definition, Scope and Context, International Standards Organization, ISO/TR 20514:2005, Jan. 2005.