Interactive analysis of geographically distributed population imaging data collections over light-path data networks

Baldur van Lew[*a], Charl P. Botha[bc], Julien R. Milles[ad], Henri A. Vrooman[e],
Martijn van de Giessen[ad], Boudewijn P. F. Lelieveldt[ab]
[a] Division of Image Processing, Leiden University Medical Center, Leiden, The Netherlands; [b] Delft Bioinformatics Lab, Delft University of Technology, The Netherlands; [c] vxlabs (Pty) Ltd, 20 Gardner Williams Avenue, Somerset West 7130, South Africa; [d] Philips Healthcare Benelux, Boschdijk 525, P.O. Box 90050 5600 PB Eindhoven, The Netherlands; [e] Biomedical Imaging Group Rotterdam, Depts. of Medical Informatics & Radiology, Erasmus MC – University Medical Center, Rotterdam, The Netherlands

ABSTRACT

The cohort size required in epidemiological imaging genetics studies often mandates the pooling of data from multiple hospitals. Patient data, however, is subject to strict privacy protection regimes, and physical data storage may be legally restricted to a hospital network. To enable biomarker discovery, fast data access and interactive data exploration must be combined with high-performance computing resources, while respecting privacy regulations. We present a system using fast and inherently secure light-paths to access distributed data, thereby obviating the need for a central data repository. A secure private cloud computing framework facilitates interactive, computationally intensive exploration of this geographically distributed, privacy sensitive data. As a proof of concept, MRI brain imaging data hosted at two remote sites were processed in response to a user command at a third site. The system was able to automatically start virtual machines, run a selected processing pipeline and write results to a user accessible database, while keeping data locally stored in the hospitals. Individual tasks took approximately 50% longer compared to a locally hosted blade server but the cloud infrastructure reduced the total elapsed time by a factor of 40 using 70 virtual machines in the cloud. We demonstrated that the combination light-path and private cloud is a viable means of building an analysis infrastructure for secure data analysis. The system requires further work in the areas of error handling, load balancing and secure support of multiple users.

Keywords: Epidemiological studies, light-path data sharing, light-path, interactive analysis, cloud computing, data protection, image processing, omics

1. INTRODUCTION

The ability to recognize the first signs of disease early has enormous socio-economic benefits by supporting a shift from treatment to prevention. Population imaging-genetics studies have shown great promise for identifying such biomarkers from combined imaging and genetic data. They provide unique databases that can be exploited to investigate patho-physiological substrates of disease in a pre-symptomatic phase, with the aim of prediction and early diagnosis of disease.

The cohort size required in epidemiological imaging genetics studies often necessitates pooling data from multiple hospitals to increase the statistical power of the findings. In both image analysis, for example voxel-based morphometry (VBM)[1], and in -omics data analysis, for example genome-wide association studies (GWAS)[2], a sufficiently large sample size is crucial to detect relevant disease biomarkers. In addition, biomarker discovery from imaging and genetic data often requires compute intensive processing, therefore fast and scalable compute resources are required to accelerate the processing of large data sets. However, patient data is subject to strict privacy protection regimes and physical storage may be legally restricted to the hospital. Within the European Economic Area (EEA) the originating institution must retain control over the data as specified by directive 95/46/EC[3]. A patient or study participant can potentially object to

---

[*] b.van_lew@lumc.nl; phone +31 715265049

Medical Imaging 2015: PACS and Imaging Informatics: Next Generation and Innovations,
edited by Tessa S. Cook, Jianguo Zhang, Proc. of SPIE Vol. 9418, 941806

· doi: 10.1117/12.2081985

processing and the responsible institution must be able to revoke access when required (see directive 95/46/EC[3] section 25). As such, simply transferring the data to centrally accessible network storage or sharing via a physical medium is not feasible.

The objective of this work is to explore a novel approach to secure pooling of data from multiple hospitals governed by the data protection directive 95/46/EC[3], while at the same time providing fast and interactive access to compute intensive data analysis resources for biomarker discovery.

1.1 Related work: secure processing of medical imaging data in grid and cloud

A number of related efforts have been described in recent literature. Systems combining data sharing between medical institutions and high performance computing (HPC) have been created in a number of projects based on the European DataGrid[4]. For example in the context of the German e-Science D-Grid there is MediGRID[4] and its successors: Services@MediGRID[5], PneumoGRID and MedInfoGrid[6]. MediGRID features a PACS interconnect component and the facility to transfer image data to grid nodes (using the gridDICOM protocol). The connected PACS machines are located in a “demilitarized zone” (DMZ), an area of the hospital network with more relaxed firewall rules. These contain anonymized images, not the actual hospital PACS machine.

Cloud computing for medical imaging is now beginning to attract the attention of the medical community[7]. A practical example is illustrated by the CloudDRN project[8]. Here the focus is on using commercially available systems, CloudDRN uses Microsoft Azure. Individual data aggregation queries can be authorized using certificates based on a public key infrastructure (PKI). Moore et al.[9] examine the cloud as a possibility to outsource hospital IT resources as well as for clinical research.

Both MediGRID and CloudDRN are based on existing open or semi-open grid and cloud systems. The systems have taken a largely regulatory approach to security instead of applying physical security. However, there are three open issues in these current approaches, which we will discuss consecutively.

First, *secure storage and network transfer of patient data* is required to enable pooling data across institutions. MediGRID based in Germany, must deal with the problem that current grid middleware does not address the data security requirements imposed by the EU data security directive (95/46/EC)[3] and related German regulations. The CloudDRN project notes that compliance certification is allowing enterprises to store more data in the cloud. In our case, however, this is complicated by the different legislation between the U.S. and Europe. In particular the functioning of data Safe Harbor agreements[10], self-certification of cloud providers and governmental access to data are all items of inter-jurisdictional conflict.[11,12,13].

Second, *providing and revoking subject specific data access* must be addressed: Within the European Economic Area (EEA) the originating institution must retain control over the data as specified by directive 95/46/EC. A patient can potentially object to processing and the responsible institution must be able to revoke access when required (directive 95/46/EC section 25). Data revocation is further complicated by the existence of data copies. MediGRID implements a complex data auditing mechanism to address this. However, protection of data copied to the grid remains problematic as Krefting et al. describe in their conclusion[4]. CloudDRN offers the choice of placing data in Cloud local storage or accessing in the originating institution via a networked server. The first option requires that an explicit copy be made, the second does not, but it offers no guarantee that the accessor will not retain a copy of the data. While CloudDRN does retain a data access audit trail there is no mention of a data removal mechanism.

Finally, *scalable computing*, should be provided to apply compute intensive processing pipelines on pooled data within the aforementioned constraints. A typical grid infrastructure such as d-Grid or any sub-grid of the European Grid Infrastructure is a federation of smaller clusters. These may be based on different hardware and OS versions are also not guaranteed to be the same[14]. It is known that certain software is sensitive to the exact versions of hardware and software, for example Freesurfer[15]. This means that to get consistent results in a grid the user must either restrict worker nodes on which jobs can run, or alternatively test the software on all the combinations they intend to use to allow the greatest possible exploitation of grid resources. While the effort will repay itself for commonly used software for infrequently used tools and pipelines it is problematic.

While it is possible for clouds to be implemented on top of heterogeneous grid systems[14], cloud systems in general are homogeneous in their hardware and the complete software stack, the virtual machine itself, is chosen by the end user. This then greatly simplifies the process of adding new tools and pipelines. However testing that they work as expected in the distributed system without wasting resources remains a challenge.

1.2 Our contributions

The objective of this work was to develop a novel approach for pooling data from secure sources and processing this in a secure and scalable computing system. Our area of interest is data from hospitals governed by the data protection directive 95/46/EC. The solution presented here uses a closed intrinsically secure, nationally based, and compliant cloud service and delivers a way to exploit cloud services in a manner which is transparently compliant with the data regulations in the EEA.

Specifically, the contributions of this work are:

- We demonstrate that the combination of a private, compliant cloud with a physically secure light-path meets the requirements for connection to a hospital network, and provides a natural solution to data revocation, with minimal overhead.
- We show that directly mounting disks over the 1Gbit light-path gives acceptable performance for the order of 100 simultaneous jobs and data visualization without making explicit data copies. We demonstrate this in a pilot experiment involving the registration of brain images.
- We show how open source software allows the construction of simple but functional cloud processing interface, the EYR-Garlic system, and that this system is readily expandable with new image processing pipelines

2. METHODOLOGY

2.1 System overview

The core of the proposed system is the availability of private optical fiber connections between a secure and scalable compute resource and multiple hospitals acting as data providers. Access to a secure private cloud infrastructure through physically secure light-paths allowed us to assemble a closed system that met the design goals of functionality and security. This system can be most easily described as a "secure, scalable, distributed PC" (see Figure 1) in which the disks, processing and visualization are connected by secure light-paths. In this section we examine the architectures and implementation we used to achieve the desired functionality.

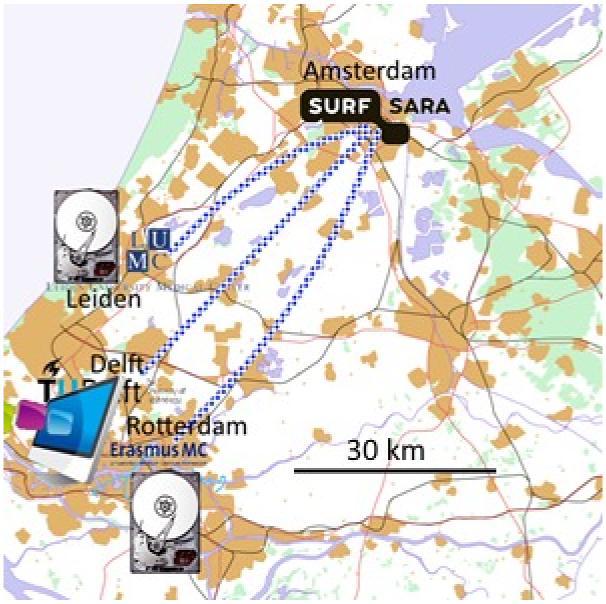

Figure 1. The secure, scalable, distributed PC superimposed on the map of South Holland in the Netherlands showing location of the data and the visualization unit. The processing unit is located at the center of the star network in Amsterdam

2.2 Network architecture: security and data sharing

The architecture (see Figure 2) involves coupling two hospital networks (ErasmusMC and LUMC), a research institution (TUDelft) and cloud compute and network resource provider (SurfSARA) for the Dutch academic institutions via physically secure light-paths. Data from the hospitals is NFS mounted to a private network range in the virtual private cloud (VPC). The VPC is implemented using VLAN[16] technology. The VPC ensures that the virtual machines in this system are isolated from others using the SARA HPC-Cloud infrastructure.

Virtual machines in the SurfSARA cloud receive addresses in the range 10.0.X.10.to 10.0.X.240. Addresses above 240 can be fixed to connected machines, in practice these are assigned to the external endpoints in the hospital or research institute. This address range was directly compatible with the ErasmusMC network and in the LUMC and TUDelft endpoints NAT address mapping was used to resolve IP conflicts with the institution internal addresses. The coupling between these end nodes and the cloud takes place over secure light-path connections.

From both the LUMC and ErasmusMC access-points the data directories were NFS shared using the /etc/exports table to the VPC VM address range. This allowed data from both ErasmusMC and LUMC to be directly mounted on all cloud VMs. Virtual machines are configured to mount the NFS directories from the end nodes to a /datadir/<DataResourceID> directory where <DataResourceID> is a unique identifier for the NFS shared directory. Data is never copied in this system though it may be temporarily cached on the VMs due to NFS caching.

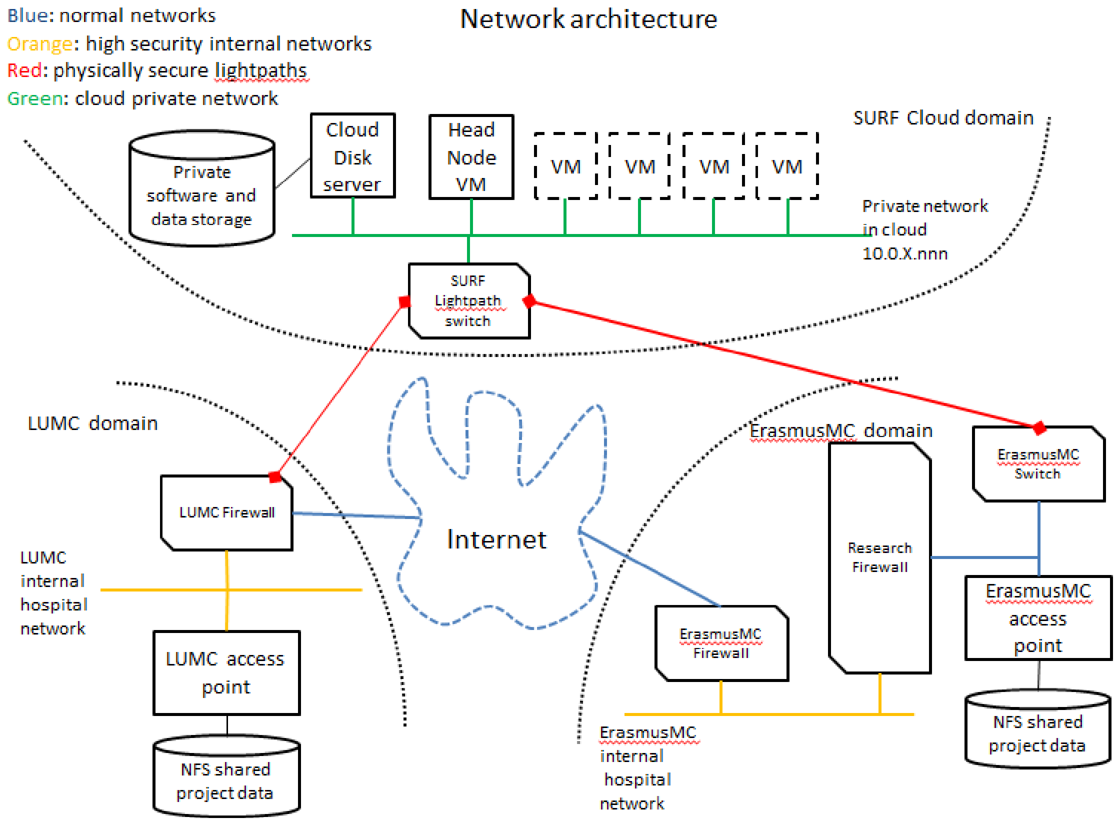

Figure 2. EYR-Garlic network architecture showing security domains, networks, machines, and data shares (for brevity TU Delft has been omitted). The light-paths bypass the regular internet and connect the VPC and hospital networks in a secure domain.

2.3 Software architecture

The architecture of the EYR-Garlic software addresses a number of practical issues in development and functionality. Algorithms were developed in an open source architecture with new components developed in Python to enable rapid prototype development. Secondly performance should be sufficient to execute real research pipelines consisting of thousands of individual jobs running on tens to one hundred virtual machines. Thirdly the system should shield the user from cloud resource management. Virtual machines should be started and stopped on demand to ensure that pipelines begin executing in a timely manner and that core-hour resources are not wasted. Fourthly the system should be capable of supporting new pipelines which should be testable, with minimal consumption of core hours. Finally the system should allow the development of result visualization prototypes.

Where possible we used proven open source software in preference to developing our own solutions. The base operating system for the VMs is a Linux Ubuntu 12.04 LTS 64-bit server. Standard Python library components were used to build web-server, and logging. The Nipype[17] library, also Python based, provides a ready-made means of sub-dividing processing pipelines into individual steps and, by means of plugins, these pipeline steps are executed in a distributed computing environment. Figure 3 shows the modular approach adopted in the development of the central system software. A main EYR-Garlic server running on a single head node virtual machine provides the web interface. In addition it manages jobs, and starts and stops worker VMs in response to job requirements. The same head node also runs the HTCondor processes needed to distribute jobs and the mongodb that acts as both repository for job status, and as result cache (together with the common cloud storage).

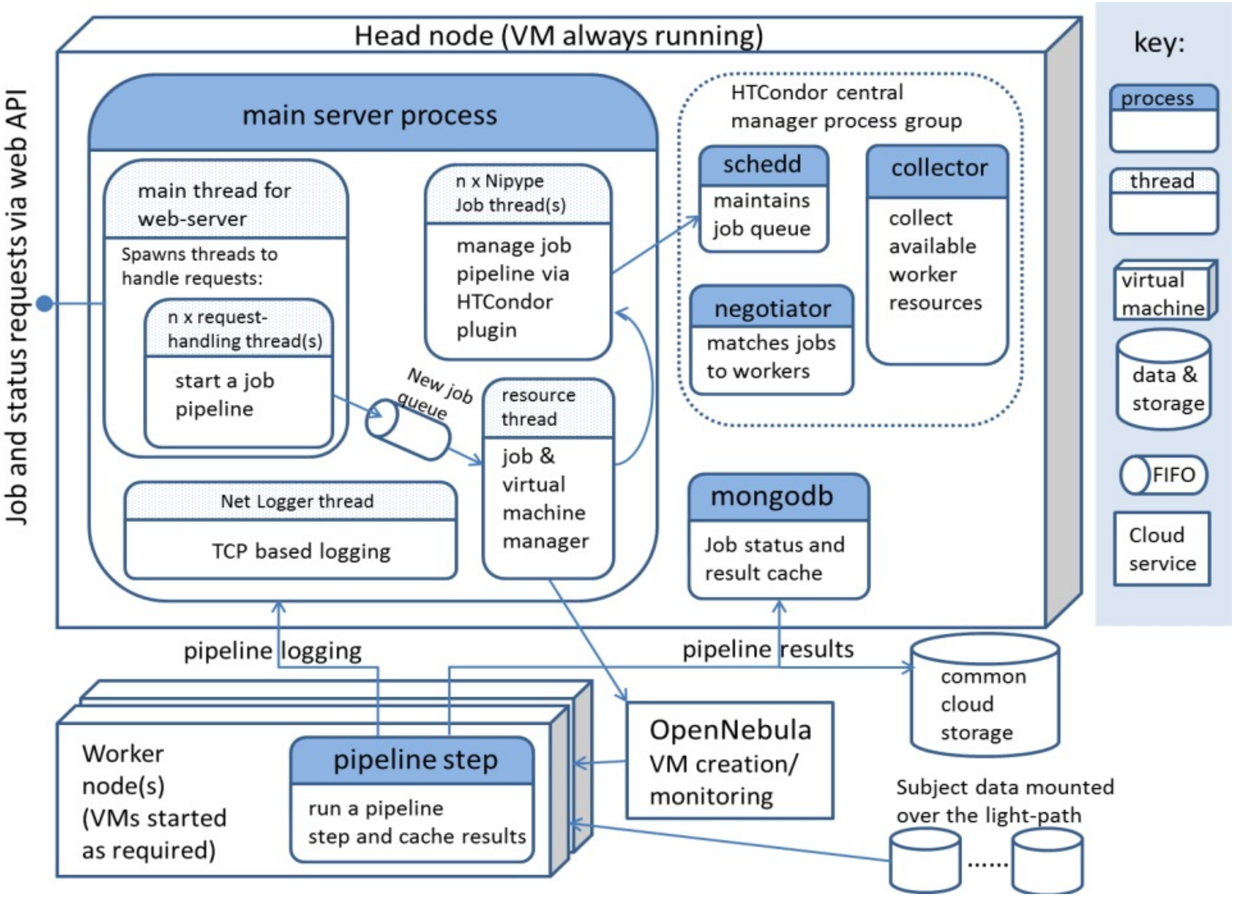

Figure 3. Overview of the EYR-Garlic system. The head node architecture is highly modular. The main server process includes several sub threads and the head node includes other components, HTCondor and mongodb, which have their own process trees. All pipeline processing is performed on the worker nodes and these are created on demand

System performance was improved using multi-threading/processing and result caching. The architecture of the head node involves several cooperating processes and can benefit from being run on a multi-core VM if extra performance is needed.

Users may repeat pipeline runs using overlapping input data or parameter sweeps. This can result in the same pipeline step being executed with the identical data and input settings. These overlaps can be exploited to speed up execution if individual pipeline steps cache their results based on input data and settings. Before a pipeline step executes the cache is checked for a pre-existing result, or if none was found then after the step completes the step results are added to the cache. The cache itself is a combination of mongodb documents and file storage.

Job and virtual machine management rely on a number of common interface functions implemented in each pipeline.. The pipeline interface includes a function that allows the Job and VM management thread to determine the maximum number of parallel steps that can be run. This number and the actual number of free cores in the cloud system are used to allocate a number of VMs to the job, and VMs are started as needed. When a pipeline completes these VMs are shutdown in a controlled fashion. VM control takes place via the OpenNebula xml-rpc interface[18].

Testability of the system was designed such as to minimize consumption of cloud core-hours. This was achieved by designing the architecture to be configurable to run outside the cloud/light-path environment. This configuration is done automatically by central Python config file that checks the host name and sets a number of predefined paths and variables in the main EYR-Garlic process. The system can therefore be executed on our development system, a local 16-core blade server, as well as in the cloud environment.

An *xml-rpc web interface* supports a number of methods to query information about the supported pipelines. Running the pipeline returns a unique identifier for the job that can then be used to monitor the job status and finally to download the job results. The results include the paths to the original data in the VPC. This allows the user to perform the “drill down” to the original data as shown in Figure 4

3. RESULTS AND DISCUSSION

The following *pilot experiment* was carried out with the system described. All data processed was stored locally in the hospital sites. From the publicly available OASIS dataset[19] 66 brain volumes were chosen, and the set was divided into two equal parts. Half of the data was locally stored at LUMC within the hospital network, and the other half at Erasmus MC. The registration pipeline of Klein et al.[20] was run on the Cloud-based VM, but with data mounted via the local ErasmusMC and the LUMC access points. In this case we performed 2,145 jobs (dissimilarity calculations for 66 volumes) on 70 VMs in an elapsed time of approximately 10 hours or roughly 20 minutes per calculation. The time per job was approximately 50% longer when compared with running on a local system, however cloud-based processing the entire data set was an order of magnitude faster.

Job execution was performed by accessing the web-interface from an IPython notebook. Multidimensional scaling was applied from inside the notebook in order to visualize results. Two different hospital infrastructures were successfully coupled to the system, one using an access point in a special research ‘DMZ’ (ErasmucMC), and one located directly on the hospital network (LUMC). The low latency of the light-path network (ping time < 1 ms) enabled instantaneous interactive inspection of analysis results, where image data was retrieved over the light- path on the fly from the source disks in the hospital. This provides users with the responsiveness of processing and viewing the data on a local machine:

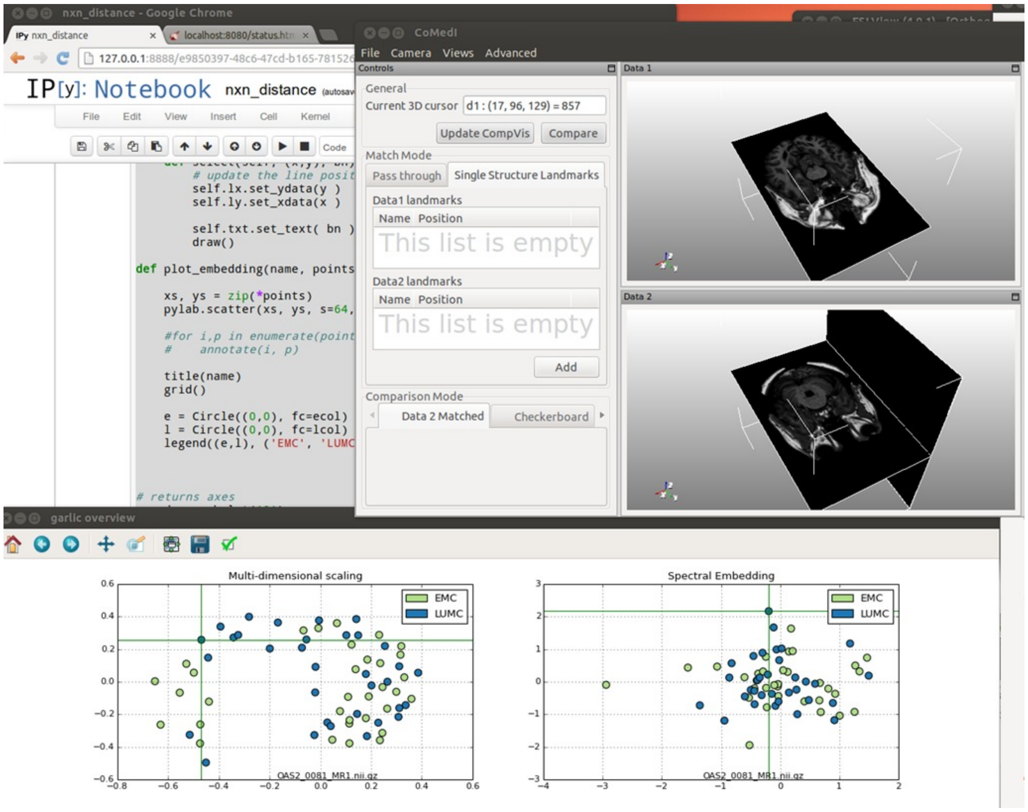

```
        # update the line posit
        self.lx.set_ydata(y )
        self.ly.set_xdata(x )

        self.txt.set_text( bn )
        draw()

def plot_embedding(name, points

    xs, ys = zip(*points)
    pylab.scatter(xs, ys, s=64,

    #for i,p in enumerate(point
    #    annotate(i, p)

    title(name)
    grid()

    e = Circle((0,0), fc=ecol)
    l = Circle((0,0), fc=lcol)
    legend((e,l), ('EMC', 'LUMC


# returns axes
```

Current 3D cursor d1 : (17, 96, 129) = 857

Figure 4. A screen shot of the IPython notebook driving the pilot experiment. It displays two low dimensional embeddings of the experimental result. Selecting a point in either graph displays the original data.

The *network security* by design was sufficient to allow the connection of the EYR-Garlic system to a live hospital network. However the system as currently implemented only supports a single security domain in which users can access all data. In the future we would like to extend the system to support multiple projects where the data is not shared, a situation known as multi-tenancy.

Performance of the pilot experiment fell within the desired range, the calculation in question could be run overnight compared to taking up to a week on one of our in-house systems. However, it is clear that with larger numbers of jobs and increasing web requests bottlenecks may appear in the system. The head node architecture is designed to be scalable over multiple cores and machines, however remaining areas of concern are the common cloud storage used for results and the NFS mounted data disks across the light-path. One possible means of addressing both of these is to use the local storage of intermediate VMs as caches. For the shared data this will not violate our data local policy since the caching VMs can be run on non-persistent images which are guaranteed to clear on VM shutdown.

Virtual machine management was complicated because of the lack of an offline facility. In practice this was also the least reliable part of the system. Although VMs were started as required, occasionally there was a failure to stop all the VMs at pipeline end. This issue can be addressed by improving the off-line testing by running an OpenNebula Sandbox[21] as a stub for the VM management. An additional weak point in VM management is that there is no mechanism to track varying VM demand in serial and parallel sections of a pipeline. This was not an issue for the pipeline we tested but we anticipate that it should be addressed in the future.

Apart from the pilot experiment reported here, we have successfully integrated other tools in to the system and created several different pipelines one of which was used in the analysis of Alzheimer's disease[22]. By design, the testing of pipelines could be done on our blade server and they worked as expected in the cloud environment

4. CONCLUSIONS

Our work shows that cloud systems coupled with light paths offer a secure and flexible resource for central processing of distributed data. The system presented is conformant to a regulatory framework within a single EEA country the Netherlands. The existence of cross-border fiber links, to Germany and Belgium, indicate the potential for further expansion of the system, however before this is a reality the issue of differing interpretations of the EEA data protection directive[3] will have to be addressed[23].

We have not yet tackled the issues that arise in multi-tenancy, that is how to run several separate projects simultaneously in this system and the problems of controlled access to different shares. In addition further experimentation is required to discover where the performance issues lie. At the moment we do not know if the central storage or the light path form bottlenecks at high loads or if we need to look more closely at load balancing. Nonetheless we believe that this secure, scalable, distributed PC model demonstrates a simple but effective approach to increasing the power of clinical studies on a national basis.

ACKNOWLEDGEMENTS

This research has received partial funding from the Dutch Technology Foundation STW, as part of the STW project 12720 ("VANPIRE") under the IMAGENE Perspective program, and from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project). The authors gratefully acknowledge the SURFSARA, and the staff of the ICT directorates at LUMC, TU Delft, and ErasmusMC.

REFERENCES

[1] Whitwell, J. L., "Voxel-Based Morphometry: An Automated Technique for Assessing Structural Changes in the Brain," J. Neurosci. **29**(31), 9661–9664 (2009).

[2] Xie, J., Cai, T. T.., Li, H., "Sample size and power analysis for sparse signal recovery in genome-wide association studies," Biometrika **98**(2), 273–290 (2011).

[3] "Data Protection: DIRECTIVE 95/46/EC OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL.", 1995, <http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:31995L0046&from=en>.

[4] Krefting, D., Bart, J., Beronov, K., Dzhimova, O., Falkner, J., Hartung, M., Hoheisel, A., Knoch, T. A., Lingner, T., et al., "MediGRID: Towards a user friendly secured grid infrastructure," Future Gener. Comput. Syst. **25**(3), 326–336 (2009).

[5] Dickmann, F., Falkner, J., Gunia, W., Hampe, J., Hausmann, M., Herrmann, A., Kepper, N., Knoch, T. A., Lauterbach, S., et al., "Solutions for biomedical grid computing—Case studies from the D-Grid project Services@MediGRID," J. Comput. Sci. **3**(5), 280–297 (2012).

[6] Dickmann, F., Kaspar, M., Löhnhardt, B., Kepper, N., Viezens, F., Hertel, F., Lesnussa, M., Mohammed, Y., Thiel, A., et al., "Visualization in Health Grid Environments: A Novel Service and Business Approach," [Grid Economics and Business Models], J. Altmann, R. Buyya, and O. F. Rana, Eds., Springer Berlin Heidelberg, 150–159 (2009).

[7] Kagadis, G. C., Kloukinas, C., Moore, K., Philbin, J., Papadimitroulas, P., Alexakos, C., Nagy, P. G., Visvikis, D.., Hendee, W. R., "Cloud computing in medical imaging," Med. Phys. **40**(7), 070901 (2013).

[8] Humphrey, M., Steele, J., Kim, I. K., Kahn, M. G., Bondy, J.., Ames, M., "CloudDRN: A Lightweight, End-to-End System for Sharing Distributed Research Data in the Cloud," Proc. 2013 IEEE 9th Int. Conf. E-Sci., 254–261, IEEE Computer Society, Washington, DC, USA (2013).

[9] Moore, K. L., Kagadis, G. C., McNutt, T. R., Moiseenko, V.., Mutic, S., "Vision 20/20: Automation and advanced computing in clinical radiation oncology," Med. Phys. **41**(1), 010901 (2014).

[10] "U.S.-EU Safe Harbor Overview.", <http://www.export.gov/safeharbor/eu/eg_main_018476.asp>.

[11] WP29., "Letter from the Article 29 Working Party to Mr. Martin Schulz, President of the European Parliament,on BCR for processors" (2014), http://ec.europa.eu/justice/data-protection/article-29/documentation/other-document/files/2014/20140612_wp29_bcr-p_general_ep_president.pdf.

[12] "Ireland voices 'serious concern' over U.S. order on Microsoft emails.", Reuters, DUBLIN, Sept 4 (2014) http://www.reuters.com/article/2014/09/04/usa-microsoft-emails-ireland-idUSL5N0R529J20140904.

[13] "US cloud providers' Safe Harbor guarantees do provide adequate data protection to EU businesses, US Government says.", 22 April 2013, <http://www.out-law.com/articles/2013/april/us-cloud-providers-safe-harbor-guarantees-do-provide-adequate-data-protection-to-eu-businesses-us-government-says/> (11 April 2014).

[14] Foster, I., Zhao, Y., Raicu, I.., Lu, S., "Cloud computing and grid computing 360-degree compared," Grid Comput. Environ. Workshop 2008 GCE08, 1–10, Ieee (2008).

[15] Gronenschild, E. H. B. M., Habets, P., Jacobs, H. I. L., Mengelers, R., Rozendaal, N., van Os, J.., Marcelis, M., "The Effects of FreeSurfer Version, Workstation Type, and Macintosh Operating System Version on Anatomical Volume and Cortical Thickness Measurements," PLoS ONE **7**(6), e38234 (2012).

[16] "802.1Q VLAN — OpenNebula 4.10 documentation.", <http://docs.opennebula.org/4.10/administration/networking/hm-vlan.html> (7 November 2014).

[17] Gorgolewski, K., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L.., Ghosh, S. S., "Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python," Frontiers in Neuroinformatics **5**(13), 1-15 (2011).

[18] "XML-RPC API — OpenNebula 4.10 documentation.", <http://docs.opennebula.org/4.10/integration/system_interfaces/api.html> (13 November 2014).

[19] "OASIS: Home Page.", <http://www.oasis-brains.org/> (4 August 2014).

[20] Klein, S., Loog, M., van der Lijn, F., den Heijer, T., Hammers, A., de Bruijne, M., van der Lugt, A., Duin, R. P. W., Breteler, M. M. B., et al., "Early diagnosis of dementia based on intersubject whole-brain dissimilarities," 2010 IEEE Int. Symp. Biomed. Imaging Nano Macro, 249–252 (2010).

[21] "OpenNebula | VirtualBox SandBox.", <http://opennebula.org/tryout/sandboxvirtualbox/> (12 November 2014).

[22] Doan, N. T., van Lew, B. A., Lelieveldt, B., van Buchem, M. A., Reiber, J. H. C.., Milles, J., "Deformation texture-based features for classification in Alzheimer's disease," Proc SPIE 8669 **8669**, 866927–866927 – 7 (2013).

[23] Hon, W. K.., Millard, C., "Data Export in Cloud Computing – How Can Personal Data Be Transferred Outside the Eea? The Cloud of Unknowing, Part 4," SSRN Scholarly Paper ID 2034286, Social Science Research Network, Rochester, NY (2012).