Cloud Federation for Sharing Scientific Data

Spiros Koulouzis[1], Reginald Cushing[1], Dmitry Vasunin[1], Adam Belloum[1], Marian Bubak[1,2]

[1]Informatics Institute, University of Amsterdam, The Netherlands
[2]AGH Krakow, Department of Computer Science, Krakow, Poland

8th IEEE international conference on eScience 2012

System and Network Engineering

Introduction

The **objective** of this work is to **address** the **issues** facing the VPH community regarding **large datasets**. **VPH** aims at **understanding** the **physiological processes** in the human body across multiple length and time scales.

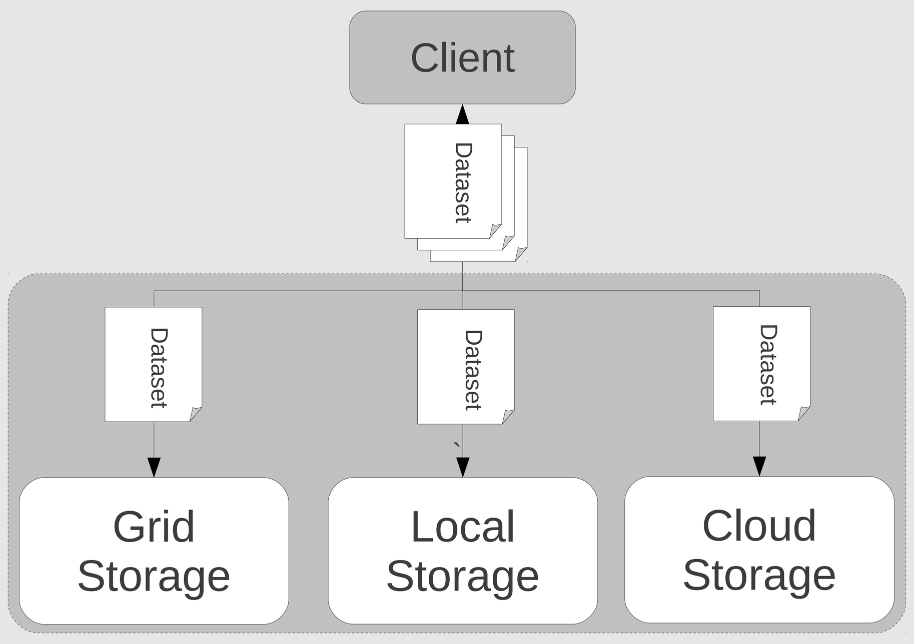

Fig. 1. Aggregating a pool of data resources in a client-centric manner in order to effectively use available storage.

Basic requirements :

- **datasets** are **not** located in a **single storage infrastructure**
- **data owners** prefer to **use their own storage** infrastructure.
- **privacy** on some datasets is essential

Cloud federation: the **detachment** of individual **services** in order to provide a more efficient and **flexible overall system.**

Storage federation: aggregation a pool of resources in a **client-centric** manner to reduce the cost of using cloud storage and to **avoid vendor lock-in.**

Existing storage resources outside the cloud must also be utilized while maintaining:

- **concurrency control,**
- **network transparency**
- **standard protocols**

Workflows are valuable for the transformation of data to knowledge. **VPH workflows** consists of three tasks: 1)obtain clinical data, 2)analyze them with models, 3)produce the clinical output.

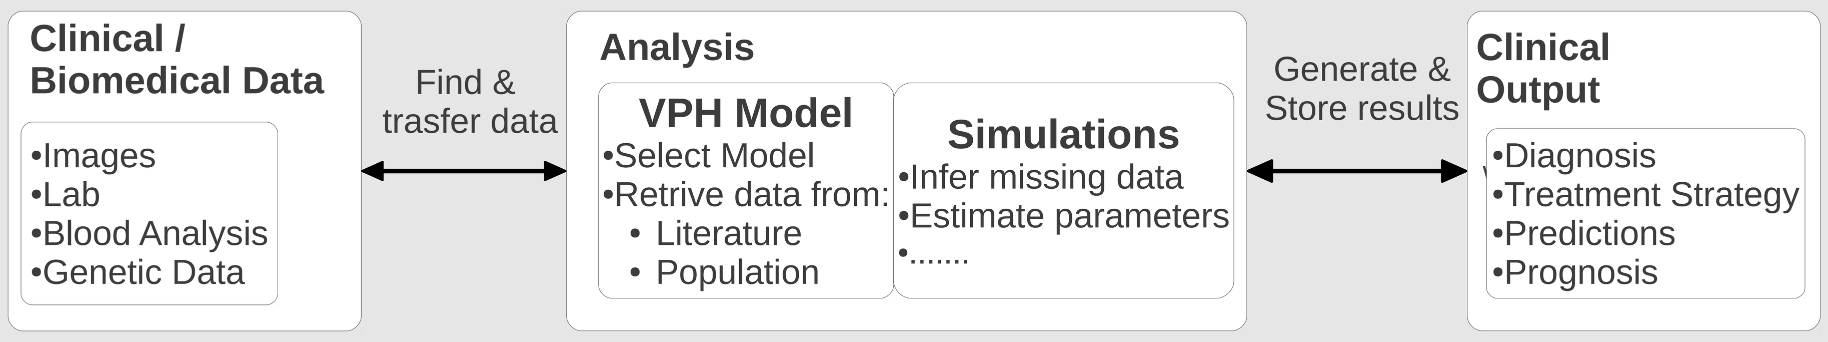

Fig. 2. The components of a generic VPH workflow. The first step always requires access to some dataset that is further analyzed

The **objective** of our investigation **is a large scale collaborative storage environment** able to **federate multiple storage resources** and present them as a **unified storage space**.

Design Considerations & Architecture

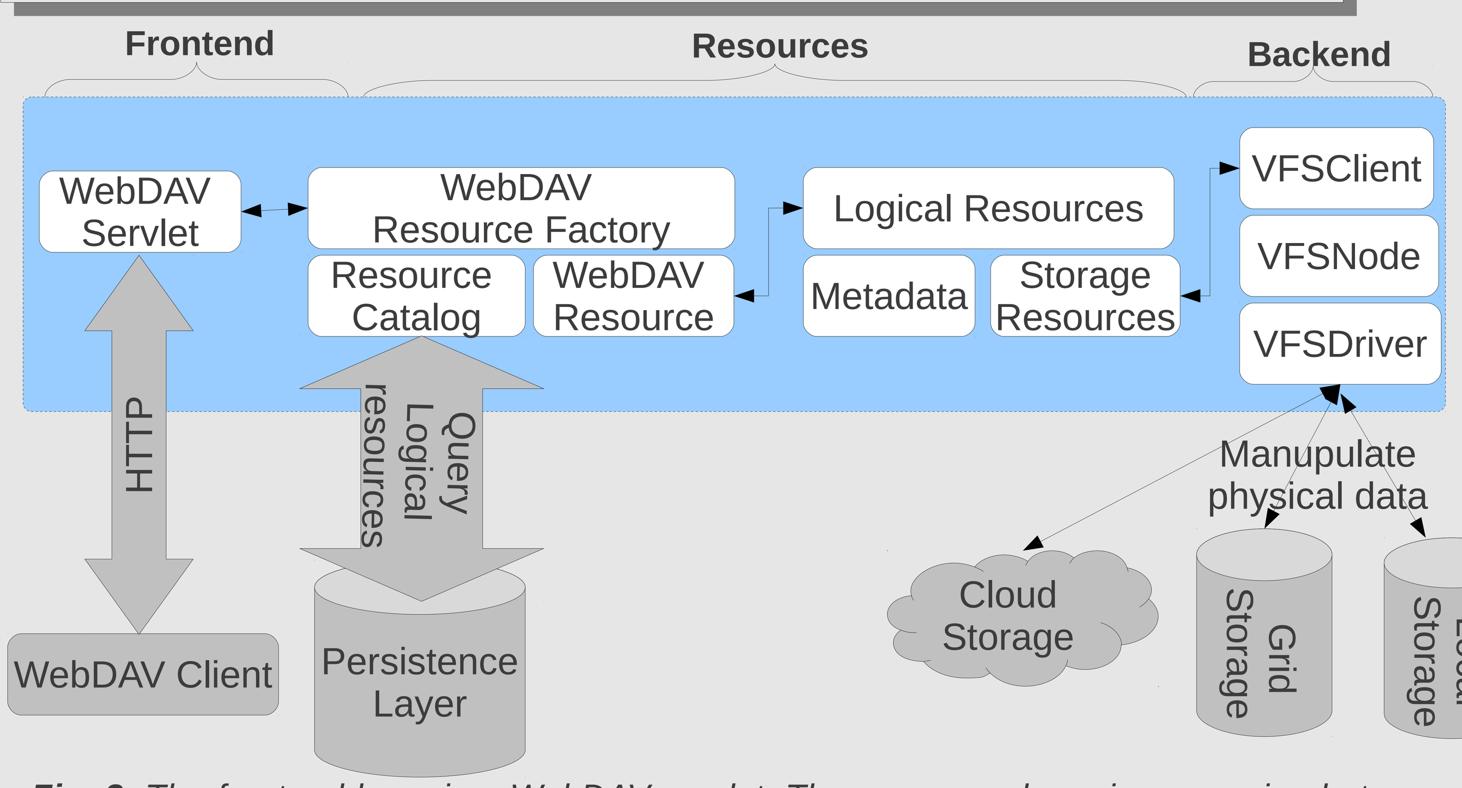

Fig. 3. The frontend layer is a WebDAV servlet. The resources layer is a mapping between the logical resources and the physical data kept in the backend storage. The backend layer is responsible for accessing physical data

The **L**arge **OB**ject **C**loud **D**ata storag**E** fede**R**ation (LOBCDER) is a **storage federation service**. It is a part of the Data and Compute Cloud Platform of the VPH-Share project.

The datasets exposed will be used by four **workflows:** 1)management of cerebral aneurysms**,** 2)integrated cardiac care, 3) HIV decision support 4) modeling of osteoporosis.

The LOBCDER service is divided into **three main layers**:

•The **frontend** layer presents a logical file system as a WebDAV

•The **resource** layer is the connection point between physically and logical data. Internally this layer maps logical data with their WebDAV representations

•The **backend** offers methods in order to manipulate physical data

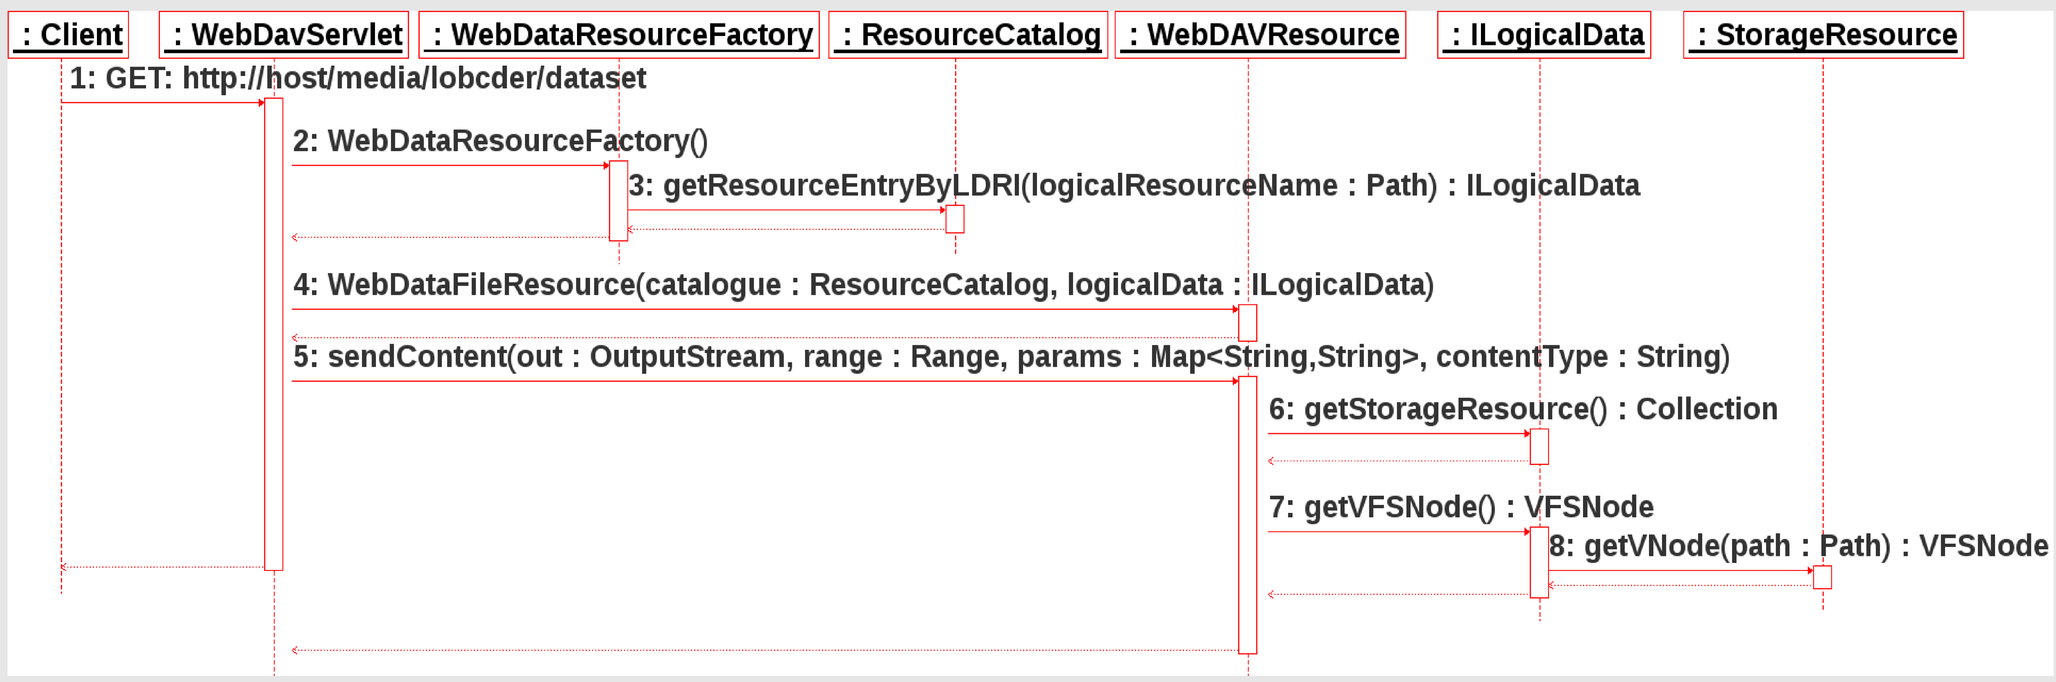

Fig. 4. LOBCDER's sequence diagram of GET request

Results

LOBCDER's **performance was compared with a WebDAV** repository. Both systems are using the local file system as their storage backend.

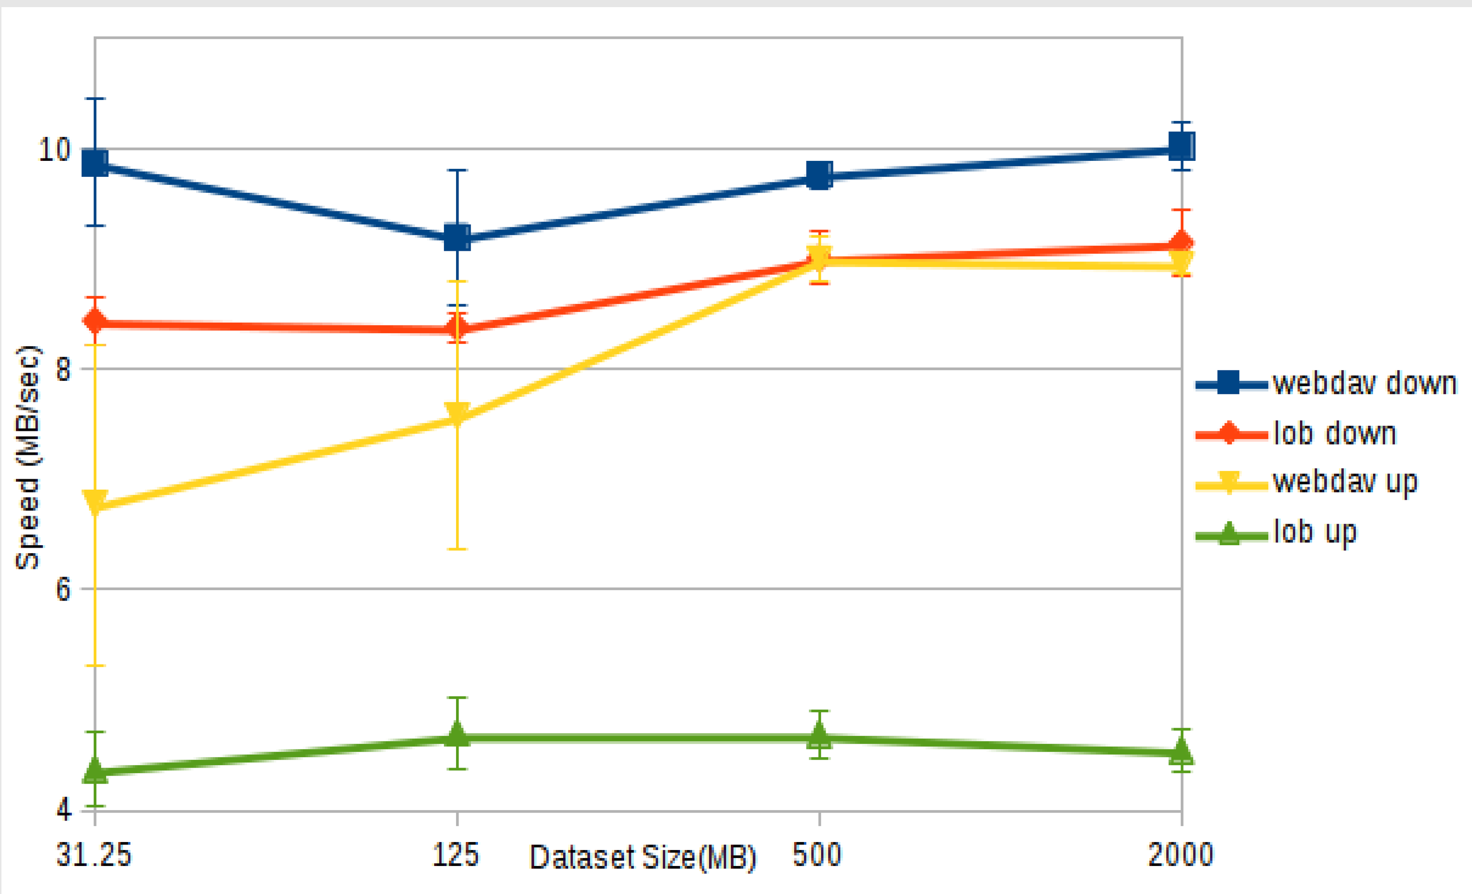

Fig.5. The results show a comparison of LOBCDER against WebDAV with both systems using the local file system as their backend. The graph shows the speed for uploading and downloading dataset D (see Table 1).

Dataset Name	Num. Of Files	Total Size (MB)
A (100kB each file)	10	0.97
	40	3.9
	160	15.62
	640	62.5
Sum	850	82.99
B (400kB each file)	10	3.9
	40	15.62
	160	62.5
	640	250
Sum	850	332.02
C (1600kB each file)	10	11.71
	40	46.87
	160	187.5
	640	750
Sum	850	996
D (3200kB each file)	10	31.25
	40	125
	160	500
	640	2000
Sum	850	2656.25

Tab.1. Datasets used to measure the upload and download performance of LOBCDER. Each of the datasets contains four subsets of 10, 40, 160, and 640 files.

Figure 6 shows **which component of LOBCDER takes the stress** for downloading dataset A, B, C, and D. For dataset A **much of the downloading time is spent on the backend**. As file sizes get larger **more than 90% of the downloading time, is used to copy data from the resource layer to the client.**

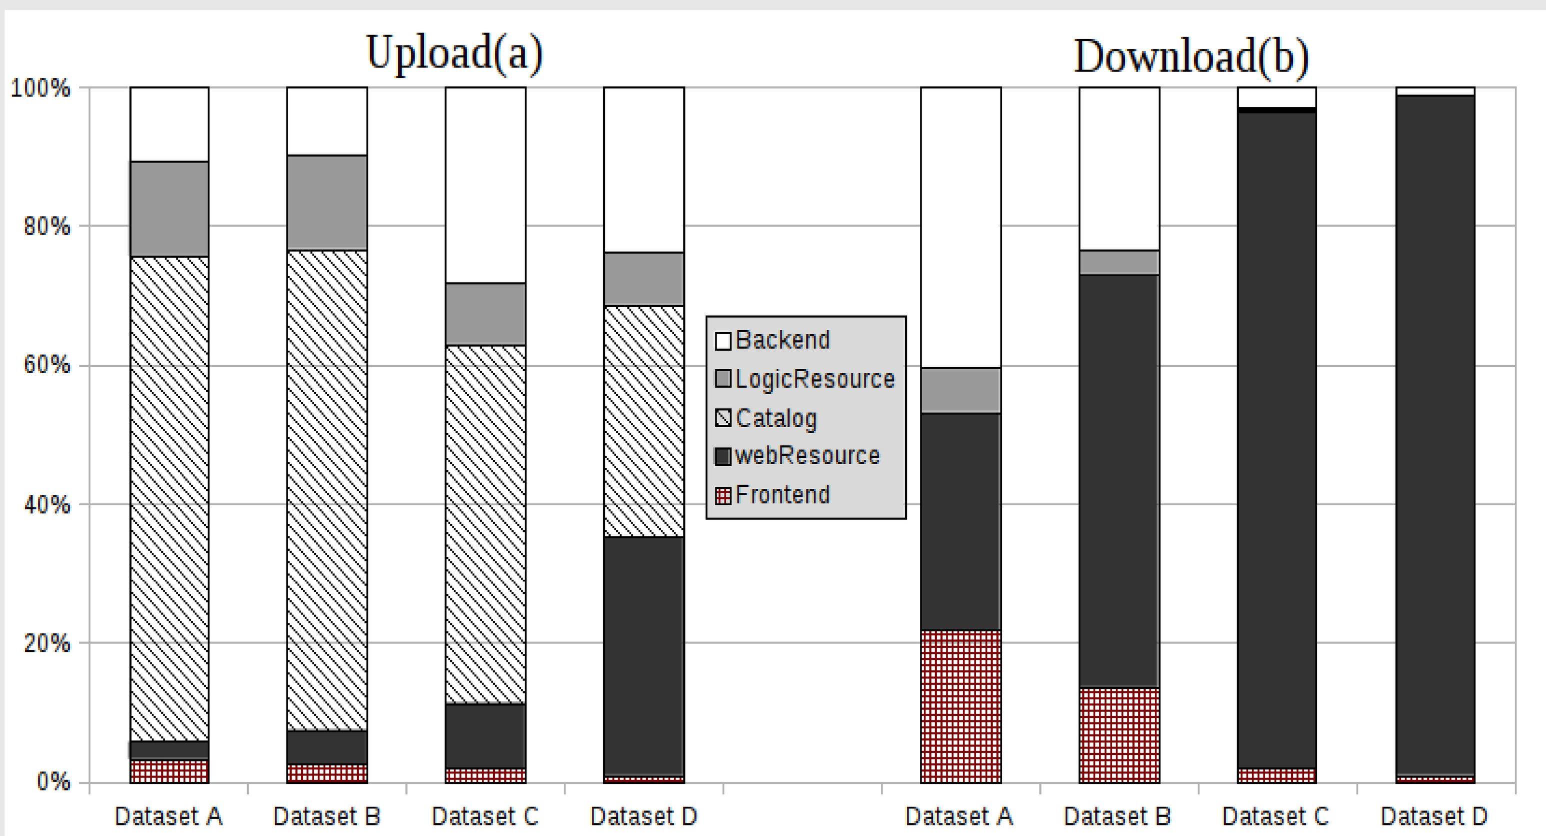

Fig.6. Break down of the time spent for uploading (a) and downloading (b) datasets A,B,C and D as measured in LOBCDER. Each bar shows as a percentage the relative time spent on each component of LOBCDER

Conclusions & Future Work

•LOBCDER's performance is comparable to WebDAV.

•The persistence layer introduces some performance penalties, but they may be easily reduced

•**Observing data access patterns** will help **identify dataset access frequency** and used by the **replication algorithm.**

•LOBCDER will be **deployed on multiple resources** to enable requests for datasets to be **redirected to the most suitable instance.**

References

[1] S. Koulouzis; R.S. Cushing; K.A. Karasavvas; A.S.Z. Belloum and M.T. Bubak: "*Enabling Web Services to Consume and Produce Large Datasets*", IEEE Internet Computing, vol. 16, nr 1 pp. 52-60. 2012.

[2] A. Belloum, M. Inda, D. Vasunin, V. Korkhov, Z. Zhao, H. Rauwerda, T. Breit, M. Bubak, and L. Hertzberger, "*Collaborative e-science experiments and scientific workflows*," Internet Computing, IEEE, vol. 15, no. 4, pp. 39 –47, july-aug. 2011.

[3] S. Koulouzis; E.V. Zudilova-Seinstra and A.S.Z. Belloum: "Data transport between visualization web services for medical image analysis", in ICCS 2010, (Proceedings of the 10th ICCS), vol. 1, nr 1 pp. 1721 - 1730. Elsevier B.V., Amsterdam, May 2010

Contact:
S.Koulouzis@uva.nl
Website:
http://staff.science.uva.nl/~gvlam/wsvlam/

VPH-Share
http://www.vph-share.eu
COMMIT/
http://www.commit-nl.nl