Building A Shared Data Infrastructure: A Narrative

Susheel Varma[1], Stefan Zasada[2], Yves Martelli[3], Debora Testi[4], Steven Wood[5]
Peter Coveney[2], John Fenner[1]

Introduction – Fuelled by their promise to solve the problem of data sharing and distilling valuable information and scientific insight from large data stores and databases in a scalable and user-friendly way, distributed data storage and cloud computing have been one of the priority areas in the bioinformatics field. With a plethora of open source and commercial offerings (iRODS, Lustre, Hadoop, GlusterFS, CouchDB, Cassandra, MongoDB, Redis,…) and a cacophony of technical terms (CAP, Paxos, Merkle trees, gossip, vector clocks, sloppy quorums, MapReduce,…) it is difficult for the scientific community to see the forest for the trees.

Objective – This paper clarifies terms and technologies central to federated data-sharing and describes the trials and tribulations of creating a distributed data-sharing environment. This includes a reflective exercise that attempts to qualify and characterise the performance of a central data-sharing component proposed for a VPH-type environment.

Methods – Distributed data management is a strategic and technological challenge within a complex research environment such as the VPH, where large amounts of data are being created and need to be shared responsibly among the members of the community and beyond. Workflow, reliability, interoperability and security of such a distributed data management solution is of paramount importance to the sustainability and extensibility of the project.

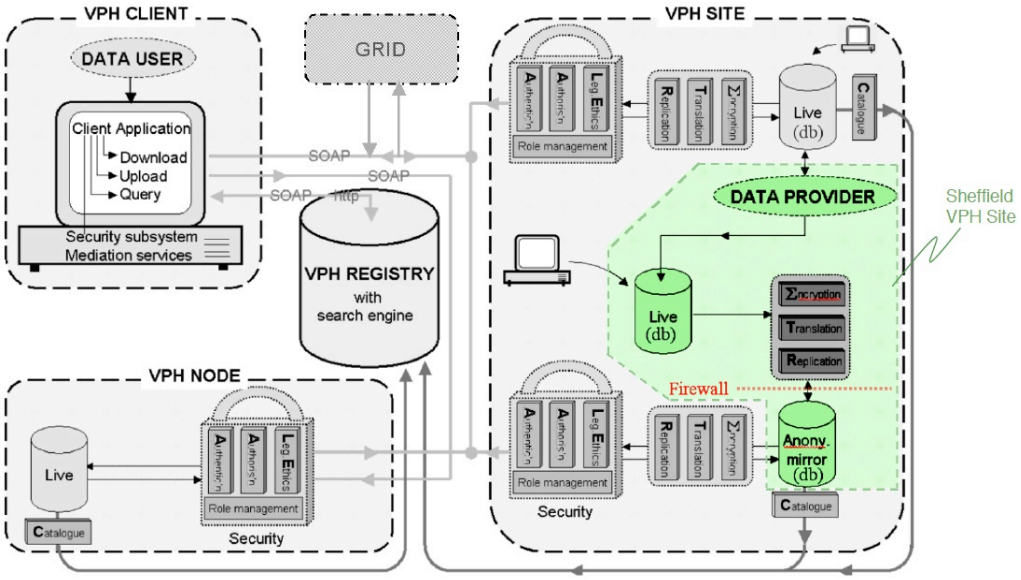

Figure 1: Originally proposed VPH NoE Data Hosting Environment Infrastructure. From [1]

Two necessary features required of the distributed data infrastructure are; *robustness* and *distributed data ownership*. In this context, robustness refers to the capability of the system to persist whilst dealing with heterogeneous environments and distributed data ownership refers to capacity of the system to maintain local control in a distributed environment. Such an idealised IT infrastructure (see Figure 1) might be characterised by three fundamental operations: *Expose* - an operation that reflects the choice of a data provider to expose a part of their data to the community; *Query* – an operation that permits authorised users of the

[1] The University of Sheffield [2] University College London [3] Universitat Pompeu Fabra
[4] Super Computing Services(SCS) [5] Sheffield Teaching Hospitals

data management solution to search and filter content within a VPH repository; *Transfer* – an operation that is performed by an authorised user to move export data to/from a VPH repository. Based on these high-level requirements, an analysis of suitable infrastructures for the VPH/NoE was warranted, classifying key attributes that are relevant and important to those high-level initiatives. The results of this exercise are presented below.

Results – After suitable deliberation, key requirements/components for building a data sharing service include: *Data Discovery, Data Management, Policy Management, Security, Metadata Catalog, Standards, Data Registry, Query Service, Co-Processing, Push-Pull API, Workflow Manager, Resource Manager, Federation, Performance, Community, Documentation, Usage* and *Licensing*. These were examined in the context of candidate technologies and their suitability in respect of an infrastructure capable of supporting federated VPH/NoE data-sharing.

	iRODS	HDFS	Lustre	GlusterFS	CouchDB	MongoDB	Cassandra	Redis
Data Discovery	No	No	No	No	No	No	No	No
Data Management	Yes	Yes	No	Yes	Yes	Yes	Yes(key-value)	Yes(in-memory)
Policy Management	Yes	No	No	No	No	No	No	No
Security	Good	No	No	Limited	Limited	Limited	No	No
Metadata Catalog	key-value	No	No	No	JSON	BSON	key-map	key-value
Standards	Yes	No	Yes	Yes	Yes	Yes	No	No
Data Registry	Yes	No	No	No	Yes	Yes	No	No
Query Service	Yes	No	No	No	Yes(API)	Yes(API)	Yes(API)	Yes(API)
Co-Processing	Yes	Yes	No	No	Yes	Yes	No	Yes
Push-Pull API	Yes	Yes	No	No	Yes	Yes	Yes	Yes
Workflow Manager	Yes	Yes	No	No	Yes	Yes	No	Yes
Resource Manager	Yes	Yes	Yes	Yes	No	No	No	No
Federation	Yes	Limited	No	Limited	Yes (P2P)	No	No	No
Performance	Average	Average	Excellent	Excellent	Good	Good	Poor	Excellent
Community	Average	Good	Average	Good	Good	Good	Average	Average
Documentation	Poor	Average	Poor	Average	Good	Good	Poor	Good
Usage	Excellent	Excellent	Poor	Average	Excellent	Excellent	Poor	Average
Licensing	BSD	Apache	GPL	GPL	Apache	AGPL	Apache	BSD

Table 1: Qualitative Survey of Distributed File systems and NoSQL databases. From [2]

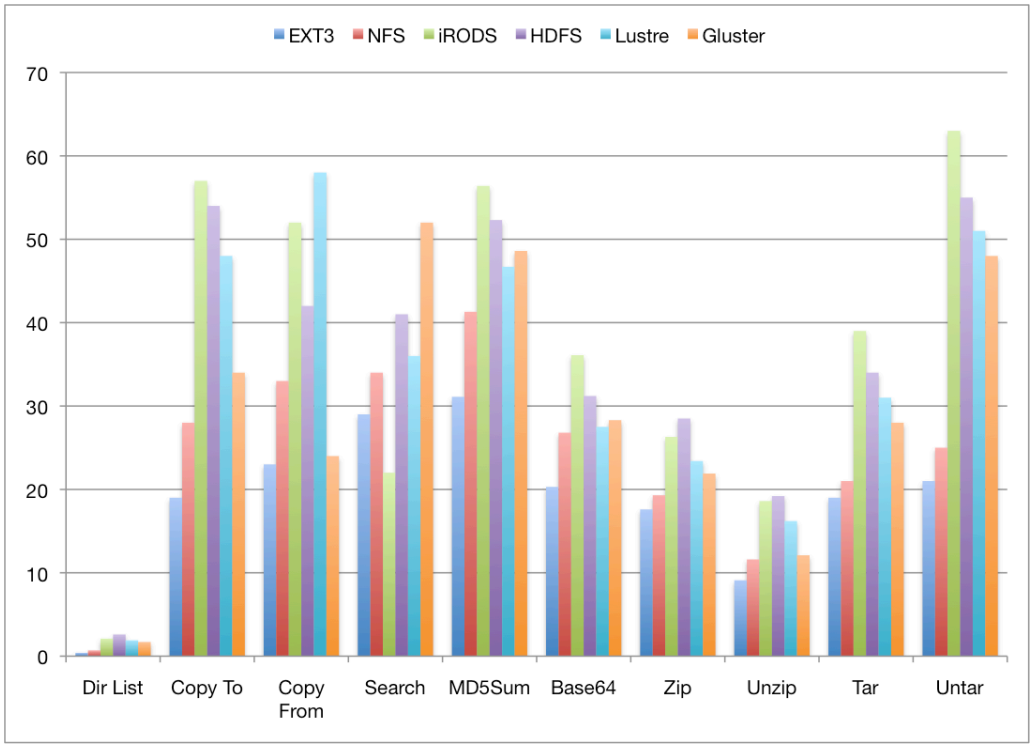

Figure 2: Distributed File-system benchmarks. From [2]

Discussion – The results of qualifying and quantifying the landscape of candidate technologies for the central data-sharing component of the VPH environment is presented in Table 1 and Figure 2. Notably, although iRODS [3] wasn't the fastest technology, it was the only candidate that effectively supported federation. As a model for data sharing, we subsequently deployed iRODS across a small selection of VPH related projects and noted the particular challenges of integrating with a distributed/federated data sharing service. One of the key insights gleaned from our deployment and evaluation exercise was the impact of poorly coherent metadata on federated discoverability (search). Consequently, a case for a minimally semantic data middleware for the VPH environment emerged from this exercise, as one of the over-riding requirements for federated data sharing. This echoes a growing awareness of the importance of coherent metadata annotation and ontological definition in VPH research, which is increasingly being captured by developing resources such as PhysiomeSpace [4] and further by infrastructure deployments like VPH-Share [5], p-Medicine [6] and EUDAT [7].

Providing a uniform middleware abstraction through semantic data across multiple data centres should enable a previously impossible level of coordination in the enforcement of metadata and configuration polices. Optimal metadata and configuration policies can lead to enhanced data provenance tracking, improved search and retrieval of data, increased interoperability, improved storage cost, better security and privacy protection. We envision a new information architecture based on a semantic data middleware as the basis for a new class of applications, where it can provide freedom and flexibility to all applications, and these advantages could be passed on to users.

Conclusion – We propose that the VPH embark on the construction of a minimally semantic data middleware, that provides a uniform, high-level abstraction to collect metadata information from heterogeneous systems into a standard repository, which can later be linked to data management applications, thus exposing data and more importantly metadata as first-class candidates in the data sharing environment.

Acknowledgments – The authors would like to thank the VPH-NoE [8] (FP7-2007-IST-223920) for their funding and support and in particular the exemplar project partners of VPH-NoE: EP8 for their valuable contributions to this paper.

References

[1] J. Fenner, S. Wood, and S. Zasada, "The NODE: an embryonic Data Hosting Environment (DHE) for the VPH Network of Excellence (NoE)," NoE: Exemplar Project DHE Proposal - USFD/UCL, 2009.

[2] S. Varma, S. Zasada, Y. Martelli, S. Wood and J. Fenner, "Research data storage infrastructure survey - part 1," VPH NoE - EP8 Report - Call 6 and the Challenge of the NoE, 2011.

[3] iRODS, i Rule-Oriented Data System. 2012 [Online] http://www.irods.org

[4] PhysiomeSpace. 2012 [Online] http://www.physiomespace.com

[5] VPH-Share, Virtual Physiological Human: Sharing for Healthcare – A Research Environment. 2012 [Online] http://www.vph-share.eu

[6] p-Medicine, Personalised Medicine. 2012 [Online] http://www.p-medicine.eu

[7] EUDAT, European Data Infrastructure. 2012 [Online] http://www.eudat.eu

[8] The Virtual Physiological Human Network of Excellence. 2012 [Online] http://www.vph-noe.eu/