Virtual Infrastructure with Database as a Service (VIDaaS) Project

Monday 1st December 2011

James A J Wilson
James.wilson@oucs.ox.ac.uk

VIDaaS objectives

- Deliver costs savings / efficiencies to UK research
- Create production-ready ‘Database as a Service’ (DaaS)
- Create hybrid cloud infrastructure on which to host DaaS
- Enable sharing of infrastructure and services
- Provide appropriate user training and support for DaaS
- Ensure services are sustainable beyond the life-span of the project

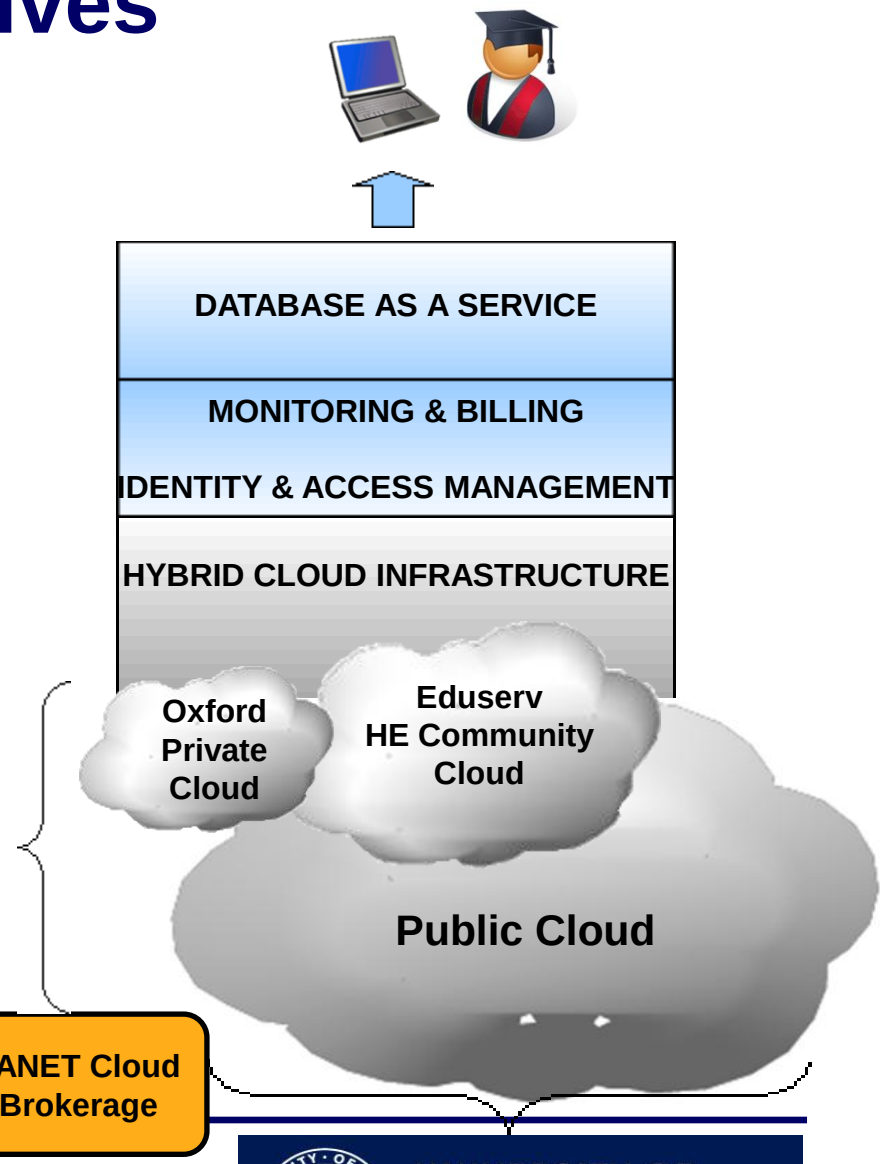

Research data infrastructure at Oxford

- Programme begun in 2008 with an internal scoping study

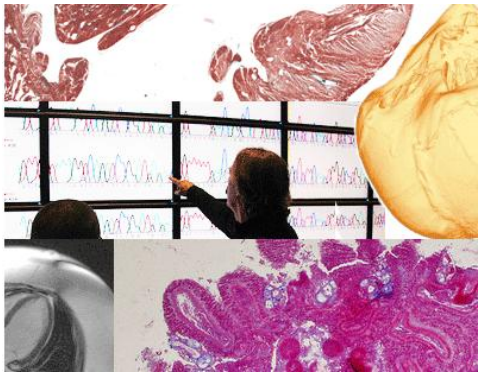

- Eidcsr (JISC funded, 2009-2010)
  - Scoping and piloting institutional data management infrastructure (software, metadata, responsibilities, etc.)
- Sudamih (JISC funded, 2010-2011)
  - Researcher training (organisation, software tools, etc.)
  - Pilot ‘Database as a Service’ (DaaS)

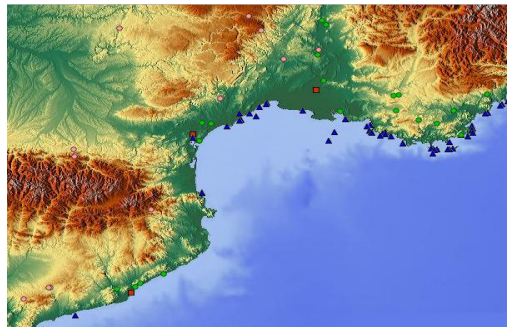

- VIDaaS (JISC & HEFCE funded, 2011-2012)
  - Full production-level DaaS, hosted on cloud infrastructure
- DaMaRO (JISC funded 2011-2013)
  - Integrate existing Oxford data management tools into enhanced infrastructure

Project team

Prof. Paul Jeffreys – Principal Investigator

Dr. Michael Fraser – Co-investigator

Dr. Stuart Lee – Director of OUCS

Dr. James A J Wilson – Project Manager

John Ireland – Lead Architect (DaaS)

Jon Hutchings – Lead Architect (VI)

Peter Jones – Lead Architect (infrastructure)

Asif Akram – Software Developer

Christian Fernau – Software Developer

Adrian Park – VI Specialist

Dr. Meriel Patrick - Analyst

Elena Blanco – Technical Author

Dr. Miko Flohr – Research Representative

Diane West – PA to the Director of IT

What is the DaaS?

- A web-based system that will enable researchers to quickly and intuitively
  - build a relational database from scratch, or
  - Import an existing database in common formats (such as Access)
- Generic data addition, editing, and querying interfaces
  - Research groups may, if desired, develop their own Web front-end interfaces to databases hosted by DaaS
- Databases centrally hosted, maintained, and routinely backed up
- Access controls to determine who can view or edit each database
  - Easy to share data with colleagues or even the general public
- Metadata capture to improve data rediscovery

DaaS Components

1.

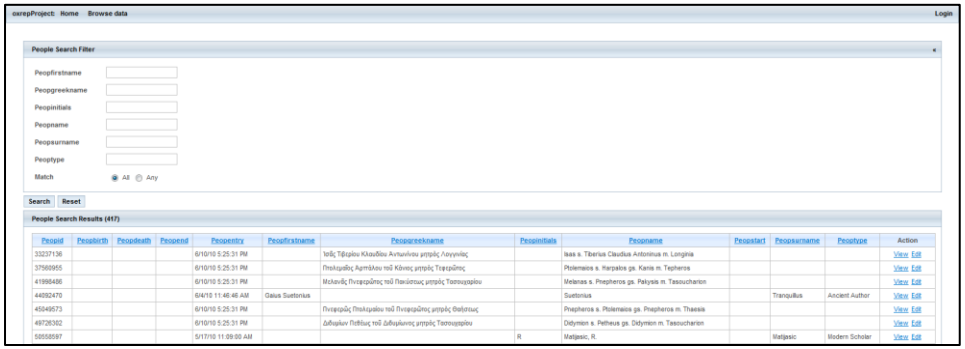

2.

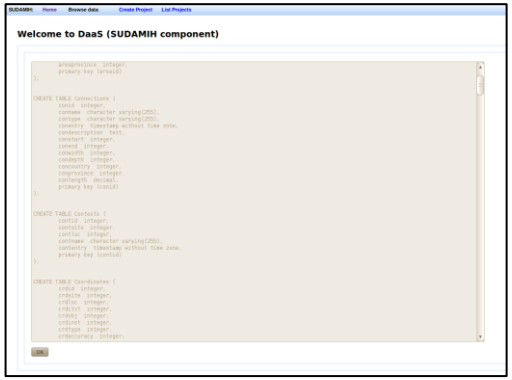

3.

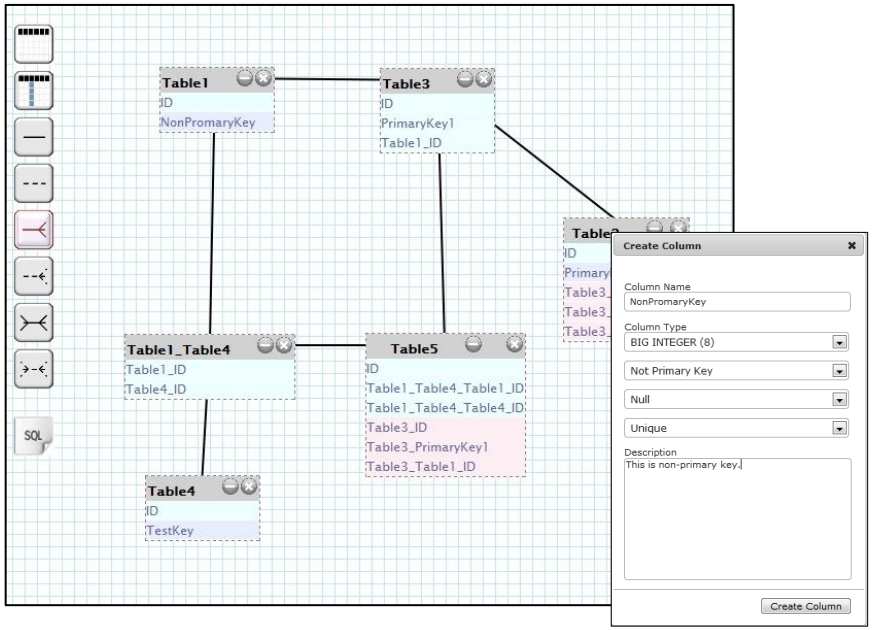

4.

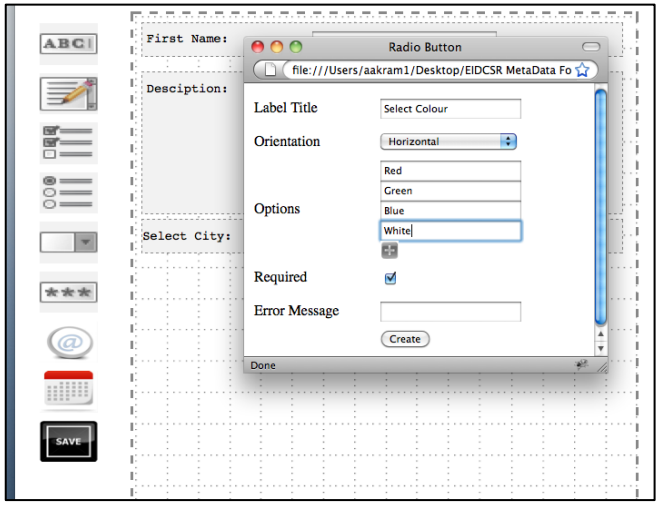

& 5...

User requirements

- Intuitive interface
- Better collaboration functions (data editing and sharing; multiple permission levels)
- Straightforward means of publishing datasets
- Automated back-up
- Ability to import and export data in various formats
- Training on principles of database design
- Cheap
- Great visualization tools [mostly out of scope]

VIDaaS enhancements

- Extended functionality beyond the humanities
- Frequent testing with ‘early adopters’ group
  - In exchange for 3 years hosting without charge
- Improved user interface, documentation, and support
- Enabling ‘publication’ of data, or sub-sets of data
  - Both dynamic and static data publishing options
- XML databases as well as relational databases

How the DaaS works

When a user registers a project they are given a unique URL where they can access their database(s)

Each project = 1 WM

VMs can be moved between cloud infrastructures to cheapest / most appropriate hosting environment

Users need not be kept informed about where their application is being stored behind the scenes

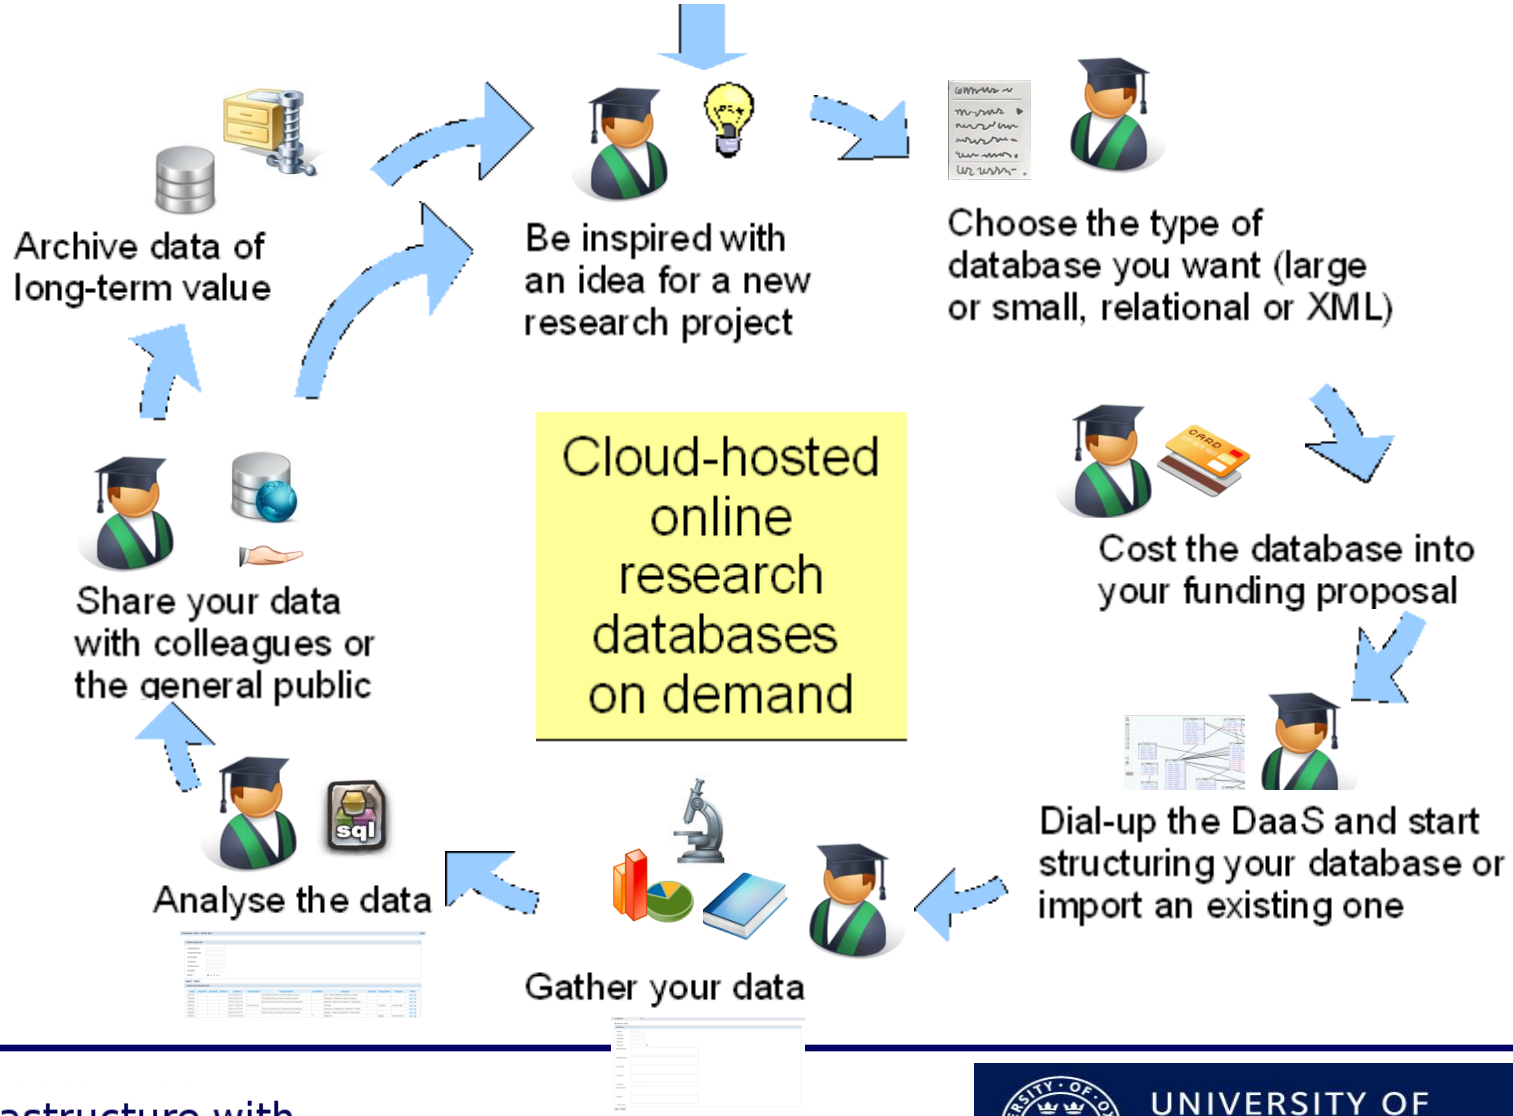

Delivery dates and dependencies

Deliverable	Due Date	Dependency
VMWare vSphere hybrid cloud in place	August 2011	VMware
Non-humanities user requirements identified	August 2011	
Pilot DaaS on Vmware delivered to Eduserv	December 2011	Eduserv
Launch of DaaS to early adopters	December 2011	
Launch of DaaS as production service within Oxford Oxford Research Database Service (ORDS)	January 2011	
Launch of DaaS on Eduserv infrastructure	January 2012	Eduserv
Final ROI & business case	February 2012	
Functioning IAM and monitoring & accounting system in place	February 2012	
All documentation and training materials delivered	March 2012	DCC
Final report	March 2012	

Business models & sustainability

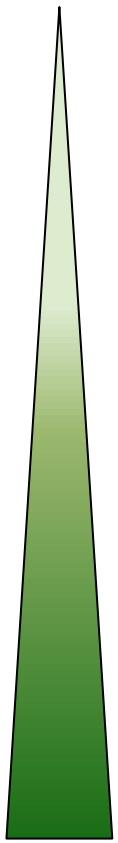

Model
Packaged software Simplest model & necessary starting place. Operational burden on customer
Pre-configured software Simplifies customer installation
Appliance Highly portable solution with reduced burden on customers & greatly increased return on VI investment
Cloud SaaS No technical responsibilities for customers. Requires substantial development of self-service interface and charging mechanisms
National Service Maximises economies of scale. Facilitates researcher mobility. Reduces HEI autonomy. Requires establishment of governance and operational organisations

Some basic sustainability models:

1. Each institution hosts DaaS themselves (whether on their own or national cloud infrastructure).
   Institution provides user support.
   An organisation ‘owns’ the service and coordinates software updates.
   A ‘service board’ collectively develops software, documentation, & training.
2. As 1, but ‘owner’ also offers ‘train the trainer’
3. ‘Owner’ provides full user support via central helpdesk.

Contacts

http://vidaas.oucs.ox.ac.uk/

http://sudamih.oucs.ox.ac.uk/

vidaas@oucs.ox.ac.uk