



UNIVERSITY OF
OXFORD

The Oxford e-Research Centre

Neil Caithness, Mark Duller, Marina Jirokta,
Jon Lockley, John Pybus, David Spence,
Kang Tang, Matteo Turilli, **David Wallom**,
Steven Young, Xin Xiong



oerc



Overview

- Research Computing Facilities/Activities
 - Campus Grid
 - OSC & Microsoft HPCS08
 - National Grid Service
 - Volunteer Computing
 - Cloud
 - Visualization
- Video Conferencing Technologies & Facilities
- Low Carbon ICT
- ICT Project Requirements Gathering




Oxford Campus Grid

- Single point of entry to research computing resources
- Commonly used for high throughput computing tasks, lots of similar but independent jobs
- Utilisation of spare capacity and investment to improve ROI
- Access to types of resources unavailable locally such as visualisation, national HPC etc.

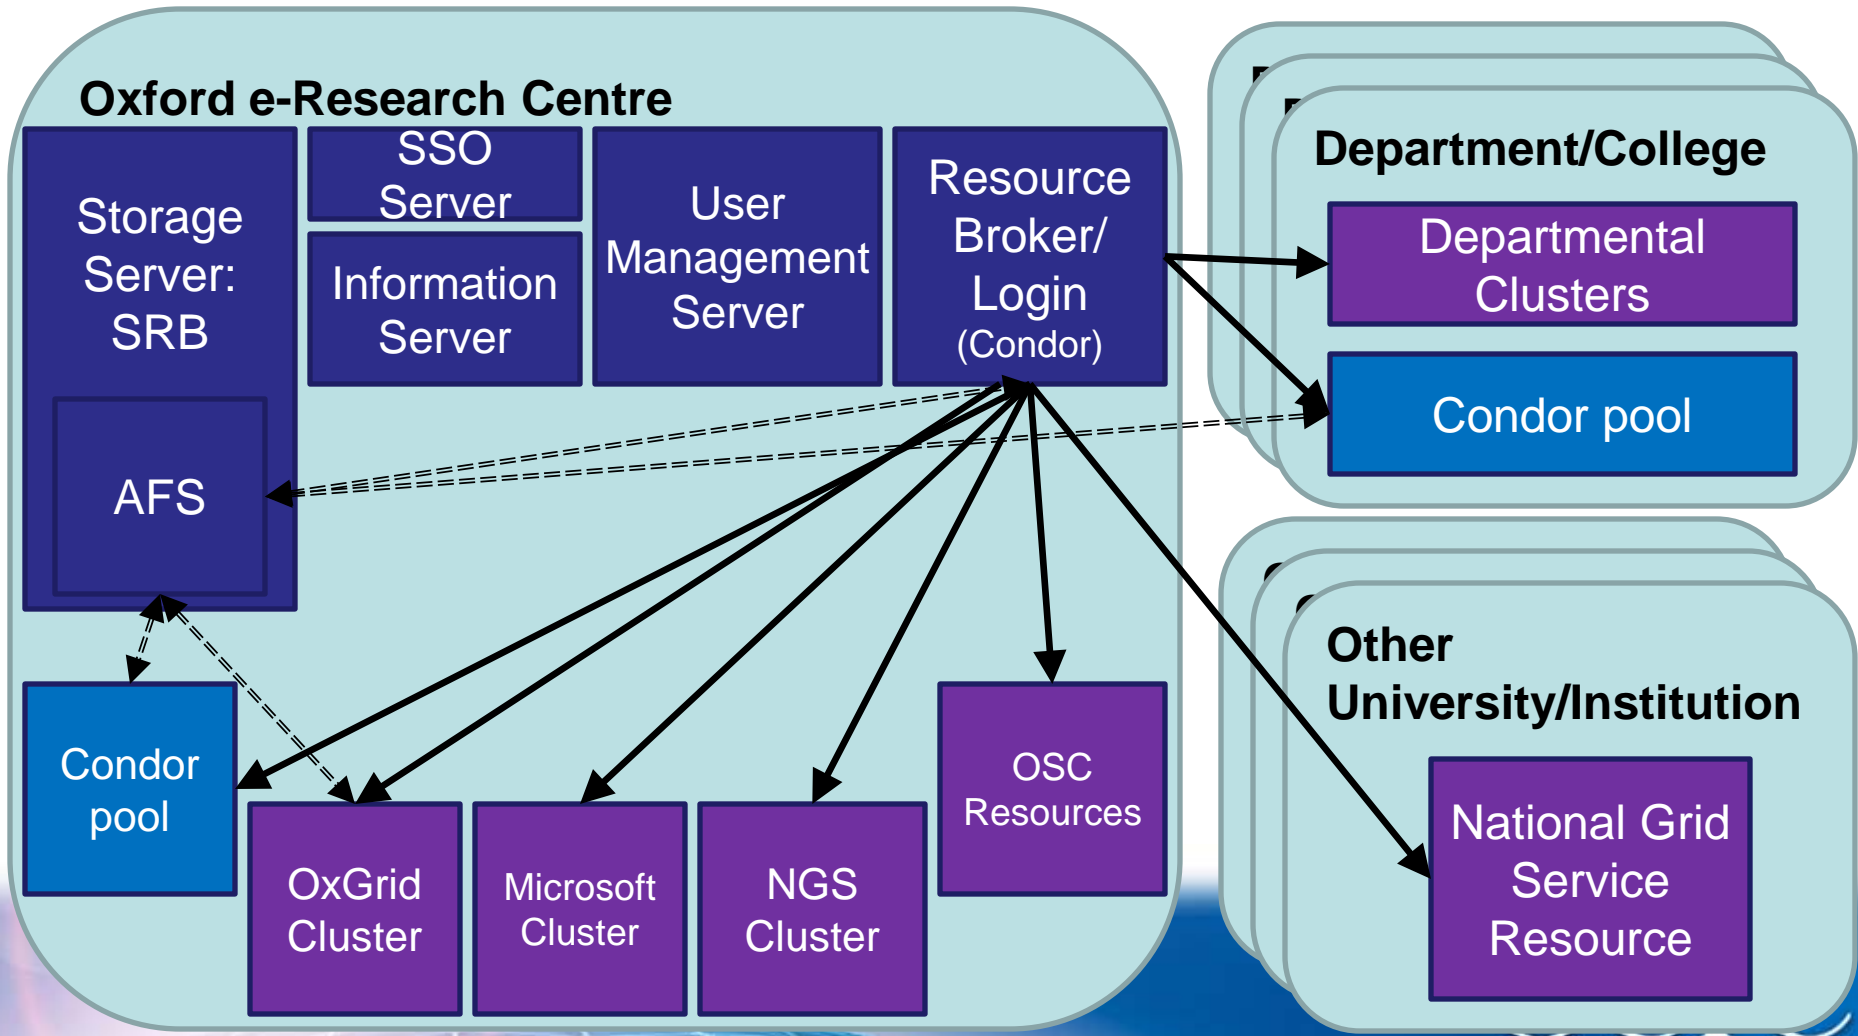
User community

- Malaria Map
- Phylogenetic Analysis of DNA
- Analysis of Clinical ECG data
- Biochemistry
- Astronomy
- Chemistry
- Satellite image analysis
- Heart modelling



<http://www.maps.ox.ac.uk/>

OxGrid Architecture





Joining OxGrid

- Departments/Colleges have under-used desktop machines
 - Set up a Condor pool to make use of this capacity
 - Can still integrate with Wake-on-LAN work see OUCS/OeRC/OUCE LowCarbon ICT project
- Process
 - OeRC provides a departmental Pool headnode
 - This means firewall exceptions can be kept to a single machine
 - ITSS in department install one of our MSIs
 - One-click installation and configuration of CoLinux
 - Can be used with all Windows remote management systems

Windows Client

- Simple MSI
 - Silent install
 - Self configuring
 - Finds pool master
 - Runs with "Low" priority
 - Network shares host's IP
 - Two setups:
 - Background service
 - + Windows Condor
 - Monitors keyboard & mouse
 - Can run Windows tasks
- CoLinux based client
 - Runs Linux kernel as a Windows driver
 - Appears as service in Windows
 - 98% speed of native
 - 3 Gb Debian 4.0 image
 - Data: read-only access to AFS applications and scratch space
 - Software: Matlab Compiler Runtime


Oxford Supercomputing Centre



Four state of the art high performance computers



Five full time support staff providing the OSC services



Assistance & advice on scientific computing from application to programming support



Software licenses & support e.g. Gaussian, ADF, Matlab, etc




Help with research proposal for HPC time




Free training courses


OSC Shared Memory Systems



Shared memory –
lots of processors
and memory in a
box




All sub processes
in one box



Orac

- SGI Altix 4700
- 256 processors
- 1TB memory



OSC Cluster Systems



Clusters

- lots of separate machines “wired” together
- cooperate on large problems



Redqueen, Hal, Queeg

- 512 processors
- 0.5/1TB memory
- Infiniband network(s)





UNIVERSITY OF
OXFORD

Microsoft HPC Server



- Campus wide service, available to all researchers
- Mix of HPC Server 2008 and Cluster Compute Server
- Hardware
 - 20 x dual Xeon CPU, dual core, 2 GB Memory/core, 'thin-node'
 - 2 x quad Xeon CPU, dual core, 4 GB Memory/core, 'thick-node'
 - Gig-E interconnect
 - 10TB storage in MS fileserver
 - 510 Gflops (peak)
- Connected to UK-Light lambda network
- Available primarily as a Distributed MATLAB service

oerc



NGS

NGS Mission and Goal

To enable coherent electronic access for UK researchers to all computational and data based resources and facilities required to carry out their research, independent of resource or researcher location.

Goal:

- To enable a production quality e-infrastructure
 - Expand towards all Higher Education Institutes.
 - Continue to support research
- To deliver core services and support
 - Collaborative research
 - Computation and Data.
- Integrated with international infrastructures
- Support inter-institutional and international research collaboration



NGS

What does the NGS offer?

- Compute services
 - Access to more and different resources
 - Different ways of running jobs e.g. multi site MPI
- Data services
 - Access to data storage, relational databases, metadata capable data stores, user interface tools and technologies
 - Support and advice
 - New ways of accessing data
- Access Services
 - User facing services providing methods to use available resources
- Central Support services
 - Individual services needed to support user access, control and management



UNIVERSITY OF
OXFORD

Volunteer Computing

- A specialized form of “distributed computing” which is really an “old idea” in computer science -- using remote computers to perform a same or similar tasks
- Along the lines of Condor, which now supports BOINC VC projects as “backfill” – so we’re all grid now! 😊
- Was around before '99 but took off with SETI@home, which was the impetus behind development of BOINC (large, open source project).
- Offers high CPU power at low cost (need a few developers/sysadmins to run the “supercomputer”); e.g. CPDN can offer up to 60TF.



oerc



OeRC's Volunteer Computing Group

- Currently running the Climateprediction.net (BBC Climate change Experiment) project, with 40-60k active hosts and nearly 44 million model years of output (24TB of data).
- The infrastructure is easy to set up, requiring only a database/web server and sufficient storage space for the anticipated data.
- Well suited to computationally intensive jobs that can be broken into small units for transmission to clients.



oerc



UNIVERSITY OF
OXFORD

OeRC Cloud Infrastructure

- Intel Cluster, 128 CPUs, 64 computing nodes, 1 Head Node, 1 data node with 3TB ScSI HD array
- Rocks 5.1 (CentOS)
- Eucalyptus 1.4.2
- Managed network mode allows for:
 - User defined ‘named network’ = EC2 ‘security groups’
 - User allocation of public IPs = ‘elastic IP’
- Default Eucalyptus Web Interface for user registration and image management
- Compatible with third party cloud/image management solutions as RightScale



oerc



Development Roadmap

1. Migration to Eucalyptus 1.5.1 due to:
 - stability problems
 - Improvements of the web interface
2. Development of three customised images for end-users and sysadmins training
3. Integration of Eucalyptus with the UK e-Science CA and X509 based AAA
4. Customisation of the web-interface



Users

- Examples
 - NGS training team:
 - (automated) deployment of a on-demand grid infrastructure for end-user, developers and sysadmin training courses
 - OxGrid users:
 - Service development and testing
 - Legacy application hosting
- Open to all research users throughout the university
- Please contact us if you have user communities that are interested.



Visualization

- Research datasets are increasingly complicated
- Visualization techniques are now able to assist interpretation of these
- OeRC has physical facilities as well as expertise to work with researchers

Visualization Wall



- 12 x 24" wide screen displays, arranged in 4 x 4 array
- Active viewing area of 7600 x 6400, 50M pixel
- Controlled from Dell Precision 309 workstations with nVidia Graphics cards
- Can be seen on BBC video presentation by Dr Peter Kohl on Heart Simulations

<http://news.bbc.co.uk/1/hi/health/774016.stm>

Video Conferencing Technologies





Video Conferencing Technologies: Physical Rooms

- Two Access Grid (AG) rooms
 - The e-Science Laboratory, Keble Road
 - IAT building, Begbroke Science Park
- Multiple cameras and projectors, echo cancellation hardware, Table top, ceiling and presenter radio mikes
- A number of software/technology/protocol solutions supported (see next slide)
- We charge for Operator time. We will happily train relevant staff so units can use the room without being charged.
- **Streaming Media, Visualisation, Digital whiteboard technologies**

Video Conferencing Technologies: Software Solutions

- Access Grid
 - Video/Audio over multicast
 - Multi-site video conferencing, phone bridging possible at some AG sites
- H.323 (Polycom PVX)
 - Video/Audio Point to point (though with an external MCU can be multisite)
- EVO evo.caltech.edu (JAVA based)
 - Multi-site video or teleconference. UK phone bridge supported by JANET
- Skype
 - Skype video conferences have been used successfully for interviews and remote presentations.





Video Conferencing Technologies: Desktop solutions

Koala - Steven Young in Marble Beer House

Koala Call Bookings Configuration Search Language

Current Meeting AV Controls Shared Files

Marble Beer House 2

Steven Young EVO User

Video Audio Whiteboard Leave this meeting

Available

- My EVO Buddies
- My EVO Communities
 - Universe
 - janet JANET (UK)
 - GridPP

Universe JANET (UK) GridPP Marble Beer House

[12:55:16] David Wallom joined
[12:56:57] David Wallom Hi
[13:00:15] David Wallom left

Connected to Panda UKERNA_UK HELP! 13:01:33

68 kb/s

Steven Young



Video Conferencing Technologies: Plans

- HD video conferencing
 - *1920pixels wide?*
 - Requires various compression hardware (in camera and streaming unit)
 - Plans for trials for lecture series between Oxford and Melbourne
- We want to develop a Video Conferencing community presence
 - Please talk to us (Mark Duller or Steven Young)
 - Possibility of a site discussed with Groupware Sharepoint requirements team

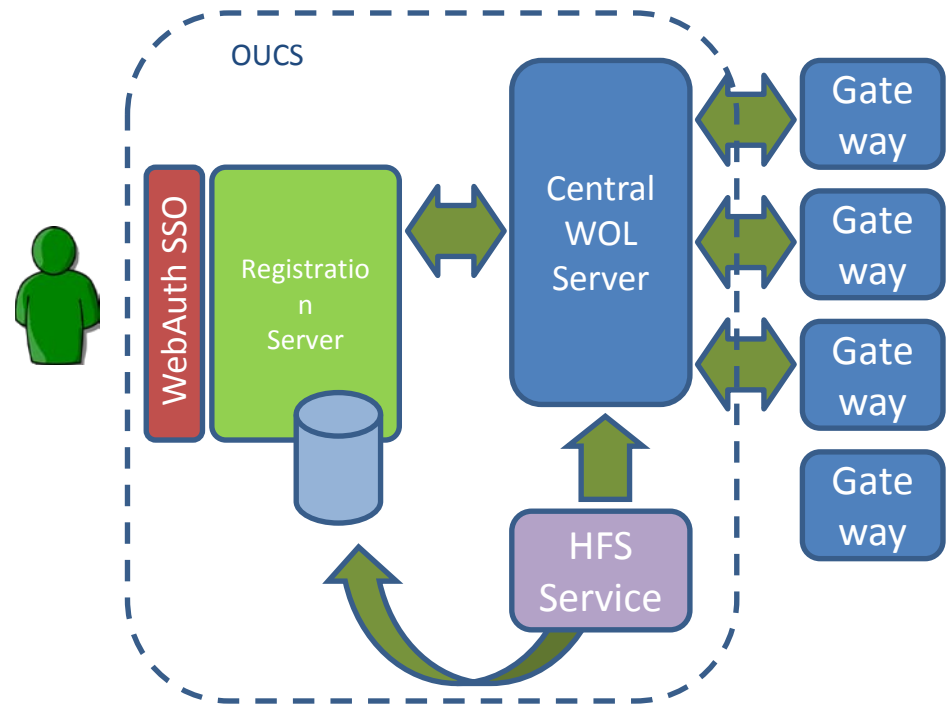


Low Carbon-ICT


- Externally funded institutional exemplar
- Aims to reduce carbon footprint in institutions caused by desktop computers.
- Some reasons desktop computers are left on;
 - So they can be accessed remotely
 - For system update or backup
 - For use with Campus Grid
 - For use as service provider
- Development of services that allow them to be turned on only when necessary

Wake-on-LAN Service overview

- Central server maintained by Sysdev
- Gateway servers deployed in 6 units in Oxford
- Integration with SSO
- Integration with 3rd party services (HFS)
- Central managed database
- Secured communications (SSL, PKI)



Monitoring Service





Requirements Group

- Practice-oriented Design
- Variety of skills and disciplinary backgrounds that provide the following expertise:
 - Methods and approaches for eliciting and analysing requirements
 - methods for engaging users in the design process
 - techniques for analysing interactions in and through physical artefacts in research settings with a view to supporting these interactions digitally
 - methods for providing seamless transition between the physical and digital worlds
 - evaluating prototypes and technologies in research contexts
 - techniques for embedding innovations in research environments
 - techniques for understanding the impact of technological innovations on the scientific record.



Conclusions

- Research computing facilities to assist your staff members and possibly yourselves
- New services are being made available to ensure that we are at the cutting edge including 'cloud'
- Also the host for facilities outside of just supporting research including innovative communication technologies, low carbon
- Other services available to support interdisciplinary research collaboration