

Campus grids: e-Infrastructure within a University

Mike Mineter
National e-Science Centre
mjm@nesc.ac.uk

14 February 2006



Thanks to:

- Mark Calleja, Cambridge
 - David Wallom, Oxford
 - David McBride, Imperial College
- for descriptions of their campus initiatives.

Current campus grids

- Present focus is on enhanced use of resources
 - e.g. to save purchasing new clusters
- Cambridge, Bristol, UCL, Imperial, Oxford, Cardiff, Glasgow and others

Two classes:

- Within an administration domain: processors on a LAN... (e-infrastructure, but not a grid!!)
- Across administrative domains
 - E.g. Multiple institutes, each with own systems management, authorisation, authentication

Pooling processors on a LAN

- Teaching lab machines lie idle for most of the time
 - Use the processing power!
- Harvest spare compute cycles to create a low-cost “high throughput computing” platform
- Create a pool of processors as a batch processing resource
- Submit jobs that run when a machine is free
- Condor most common approach
 - <http://www.cs.wisc.edu/condor/>

Condor


- Converts collections of distributed workstations and dedicated clusters into a **high-throughput computing** (HTC) facility, a Condor Pool
 - HTC concerned with harvesting maximum compute time for many jobs... e.g. months CPU time in days elapsed
- Manages both resources (machines) and resource requests (jobs)
- Includes :
 - ClassAd Matchmaking
 - Process checkpoint / restart / migration

Terminology

- **Machines**
 - **Pool** – a collection of computers used by Condor
 - **Central manager** – one per pool
 - ▶ **matchmaking**
 - ▶ **pool management**
- **functionalities**
 - **user submit** -> schedd with queue of jobs
 - **job execute** -> startd
 - **matchmaking** ->
 - **collector & negotiator**
 - **control of a node** -> master

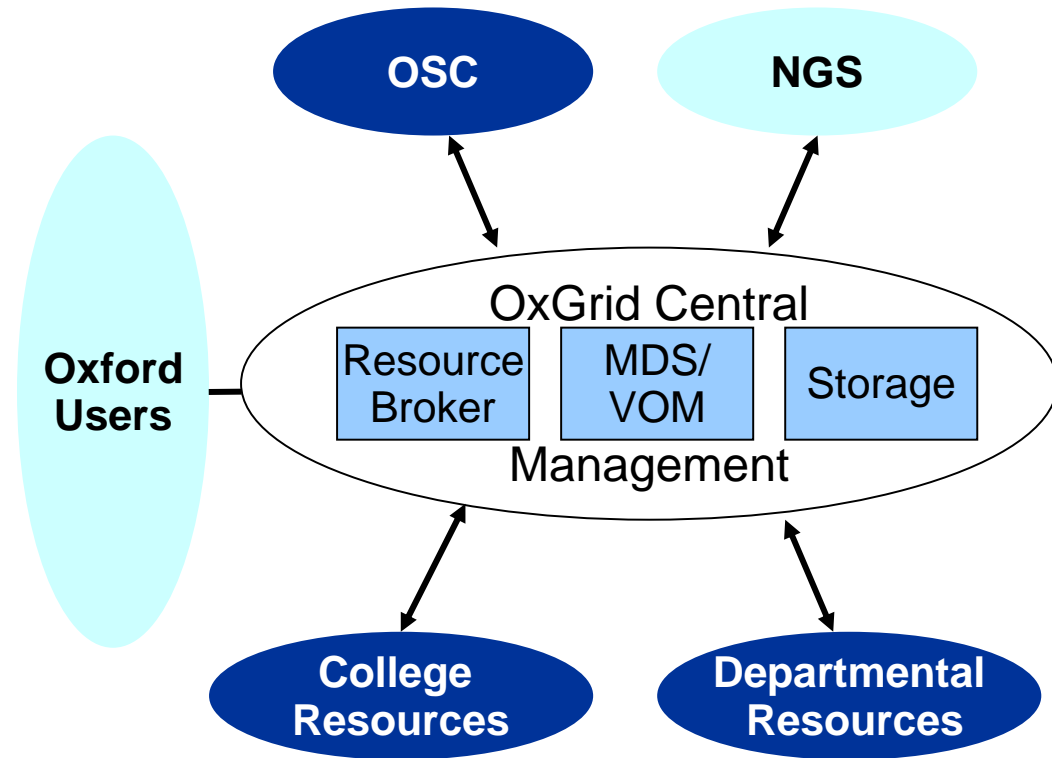
Campus grids

- Single sign-on to virtualised resources in multiple administrative domains
- Need AA mechanisms
 - X 509 certificates
 - For users that only wish to access internal (university) resources, a Kerberos CA (e.g. Oxford, Imperial College)
- Need brokering – where should a job be run?
- Needs information systems – what are the options?!
- Scalability requires each VO or institute contributes its average requirement




Example: OxGrid, a University Campus Grid

- Single entry point for Oxford users to shared and dedicated resources
- Seamless access to National Grid Service and OSC for registered users
- Single sign-on using PKI technology integrated with current methods



David Wallom

Oxford Interdisciplinary e-Research Centre



Middleware for campus grids

- Globus toolkit
 - Tools built on Grid Security Infrastructure to include:
 - **Job submission** (GRAM) : run a job on a remote computer
 - **Information services**: So I know which computer to use
 - File transfer: so large data files can be transferred
 - Replica management: so I can have multiple versions of a file “close” to the computers where I want to run jobs
- Storage Resource Broker
 - **Virtual filesystem**: for files held in multiple locations
- Condor “flocking”
 - Allow access to multiple Condor pools

Globus

- A software toolkit: a modular “bag of technologies”
 - Made available under liberal open source license
- *Not* turnkey solutions, but *building blocks* and *tools* for application developers and system integrators
- International production grids are (currently) based on the Globus Toolkit release 2
- Globus Alliance: <http://www.globus.org/>

Running a job with GT2

- GT2 Toolkit commands: an example
 - Job submission – need to know UNIX and name of a CE to use
 - Higher tools allow jobs to be submitted to “the grid” – to a resource broker

```
globus-job-submit grid-data.rl.ac.uk/jobmanager-pbs /bin/hostname -f
```

```
https://grid-data.rl.ac.uk:64001/1415/1110129853/
```


```
globus-job-status https://grid-data.rl.ac.uk:64001/1415/1110129853/
```

```
DONE
```

```
globus-job-get-output https://grid-data.rl.ac.uk:64001/1415/1110129853/
```

```
grid-data12.rl.ac.uk
```

How SRB Works



- 4 major components:
 - The Metadata Catalogue (MCAT)
 - The MCAT-Enabled SRB Server
 - The SRB Storage Server
 - The SRB Client

Obvious issues

- Need specific initial user communities
>> vague sense this is a good idea!
- Engage with systems managers from first thoughts
- Operations effort must be factored in and justifiable! Researcher enthusiasm is not enough!

Platform for interdisciplinary research??

- Campus grid gives an infrastructure spanning multiple institutes
- Potential for VOs spanning those institutes also
 - Doing more than sharing spare compute cycles
- Enabling easier access / reuse of research data
- Higher level services + ontologies/semantics + workflow