National e-Science Centre

Campus grids: e-Infrastructure within a University

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

GOSC NGS National Grid Service JISC PPARC e-Science eGee

Thanks to:

- Mark Calleja, Cambridge
- David McBride, Imperial College
- David Wallom, Oxford

for descriptions of their campus initiatives.

National e-Science Centre

Overview

- Goals
- Methods & Examples
- Some opportunities and implications

National e-Science Centre

Goals of campus grids

- Resource utilisation
  - Computers: in each university x,000 PCs are little used
  - Research data
- Collaboration
  - Access Grids – meetings across sites
  - Sharing visualisation, data, programs
- Infrastructure for research
  - From resource sharing in a university to international collaboration with one middleware stack

National e-Science Centre

Overview

- Goals
- Methods & Examples
  - Harvesting CPU cycles
  - Crossing administrative domains
    - Globus
    - Storage Resource Broker
- Some opportunities and implications

National e-Science Centre

Harvesting CPU time

Often-idle processors!!

Teaching labs. +
Researchers

Analyses constrained by CPU time!

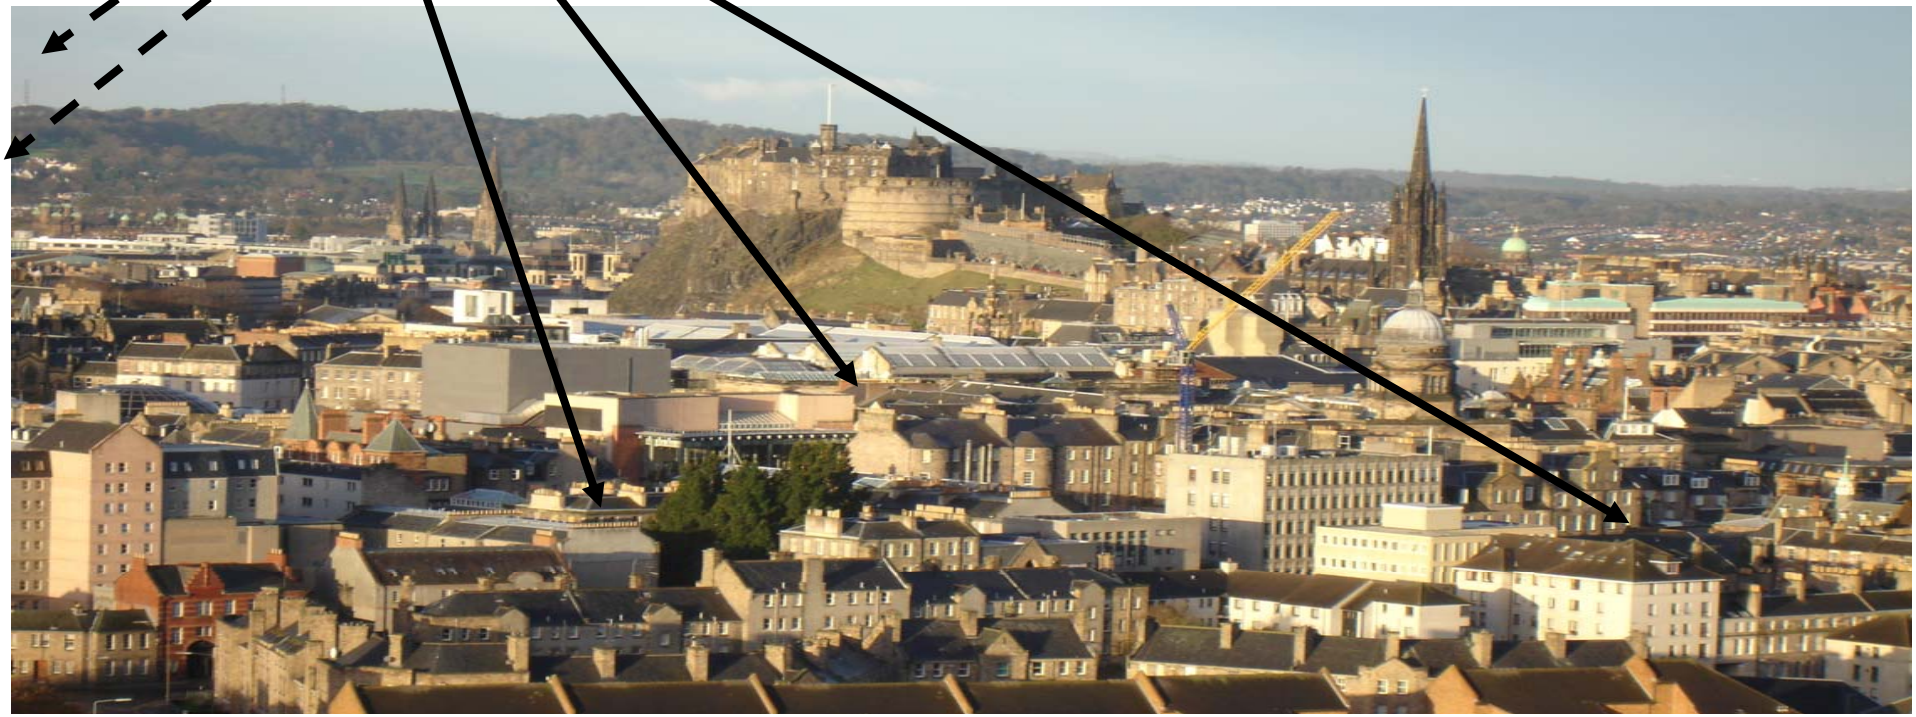

Harvesting CPU time

- Teaching lab machines lie idle for most of the time
- Harvest spare compute cycles to create a low-cost “high throughput computing” (HTC) platform
  - Goal: run many tasks in a week, month, …
  - Typically: many similar tasks invoked from workflow or a script
    - Monte-Carlo
    - Simulation – parameter sweeps
- Pool 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/

Example: viewshed analyses

Viewsheds: what can be seen from point at "+"

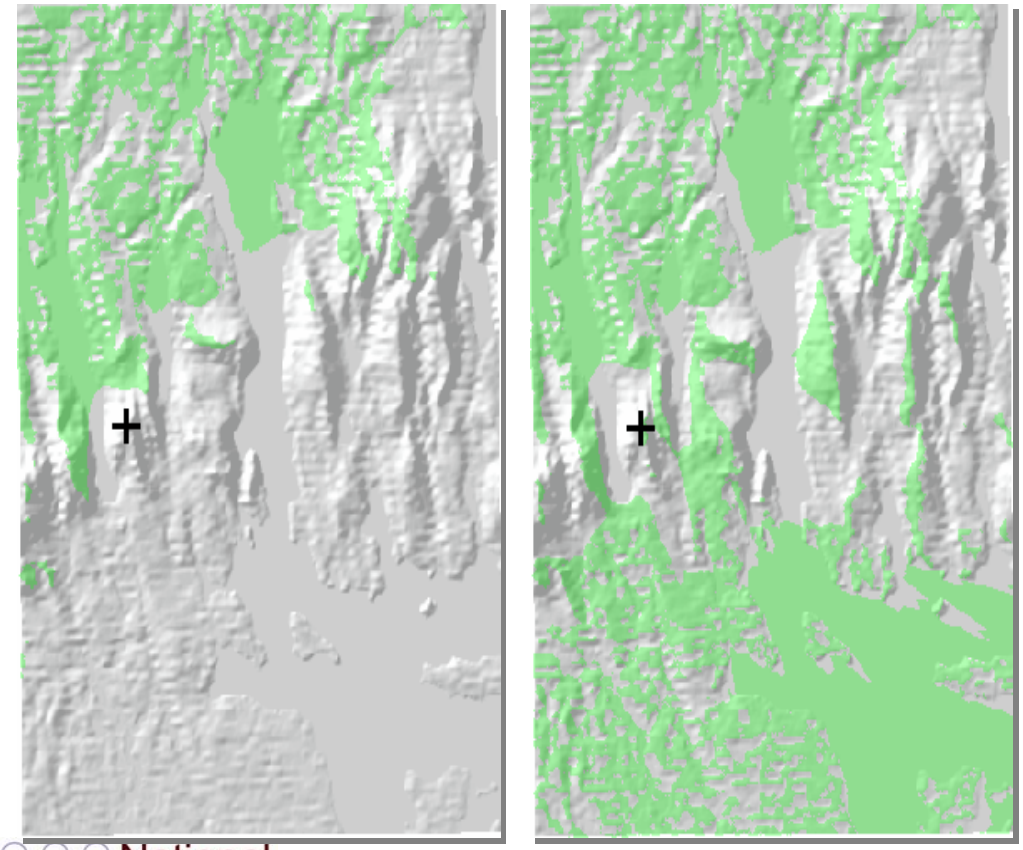

- Derive viewsheds for all points in "digital elevation model" (DEM)
- Build a database to allow
  - Derivation of indices to characterise viewsheds
  - Applications to access pre-calculated viewsheds

Mineter, Dowers, Caldwell

National e-Science Centre

Example: viewshed analyses

Typical run:

39 PCs (Windows)

330 CPU days in

9.94 days elapsed

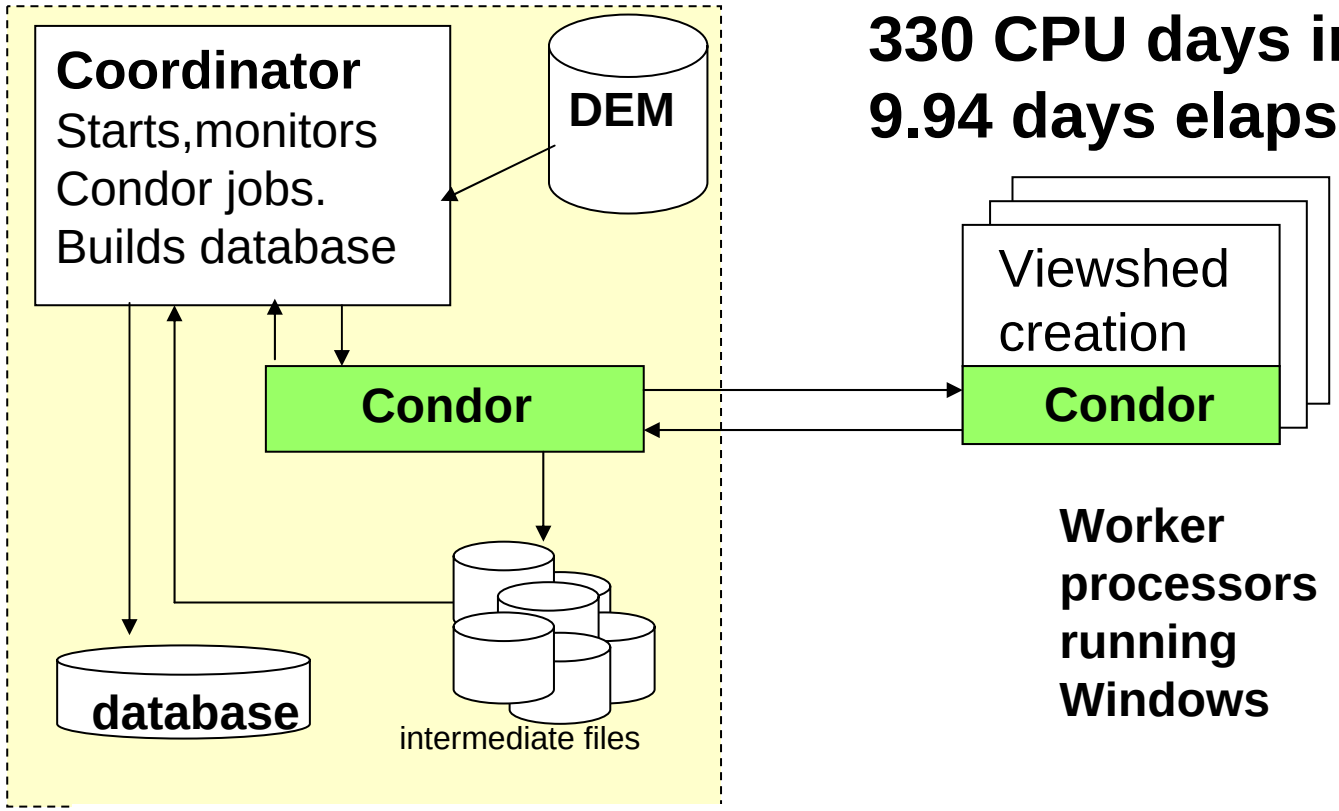

Mineter, Dowers, Caldwell

National e-Science Centre

The UCL Condor Pool

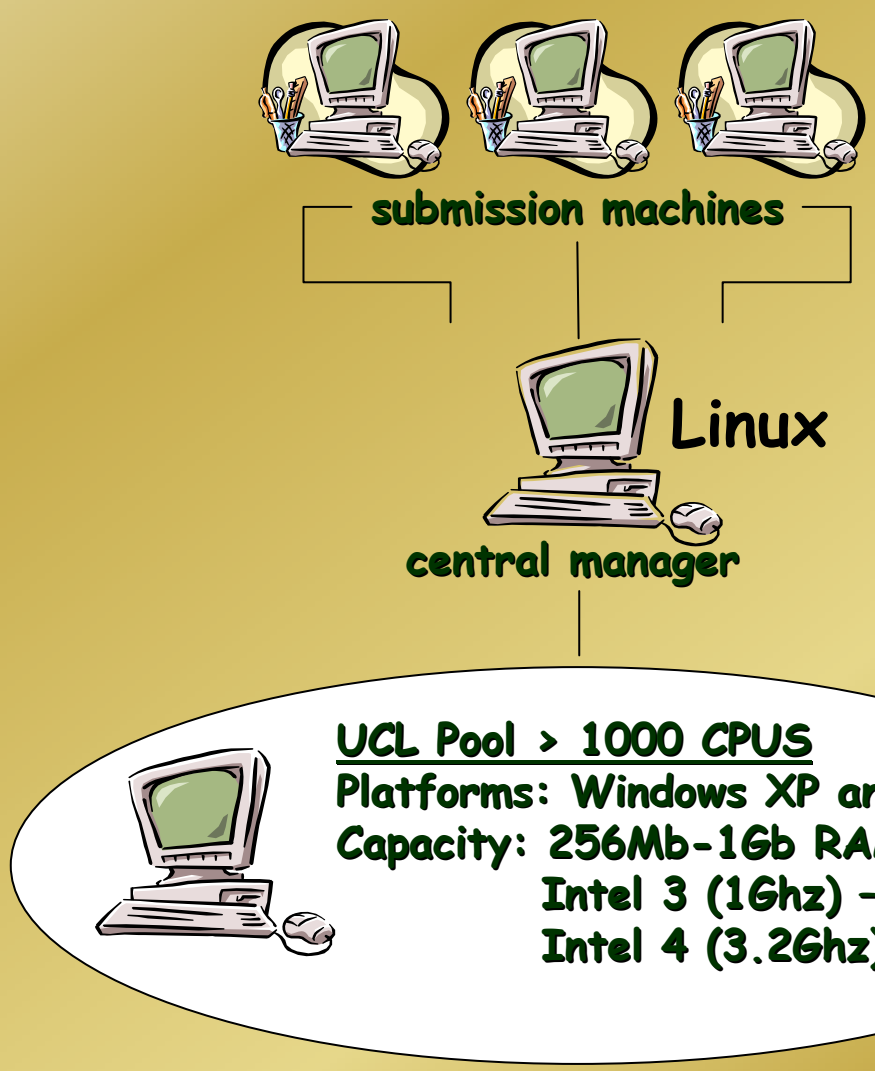

Different projects can have their own submission machine

Note: need Windows executable (.exe file)

Mineral Surfaces

Calculations:
- investigate 10-20 surfaces
- 2 to 5 surface terminations
- 4 to 16 impurity positions
- > 4 concentrations

Total number of calculations **per impurity: 120-2440**

M. Alfredsson, J.P. Brodholt and G.D. Price et al, Submitted to Nature Materials

CaO-termimated

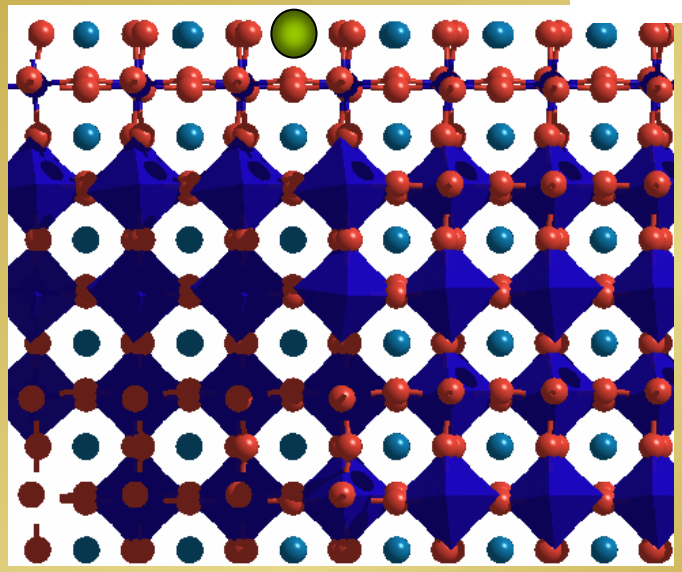

{001} surfaces of $CaTiO_3$

TiO_2-termimated

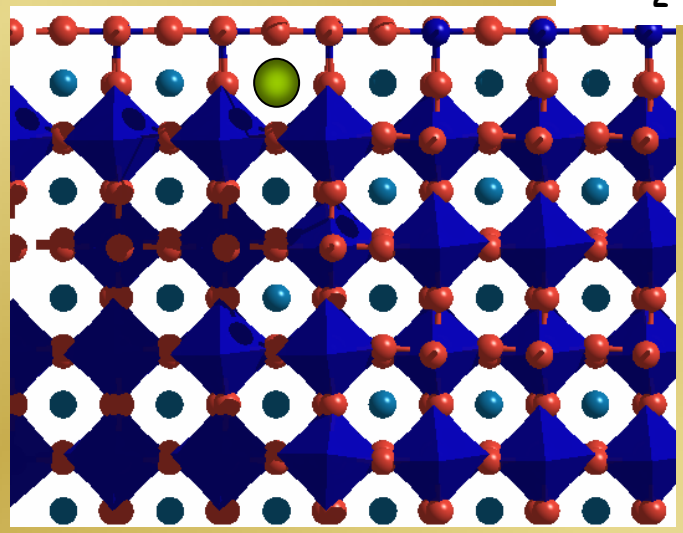

Surface Energy for doped surfaces

X_M:bulk + M_M:surface → M_M:bulk + X_M:surface,

M_M=un-doped material and
X_M=doped material

(Ca,M)TiO_3:*bulk* + $CaTiO_3$:*surface* → $CaTiO_3$:*bulk* + (Ca,M)TiO_3:*surface*

$$E_S(n) = E^0_S + (n/S)\cdot E_{seg}$$

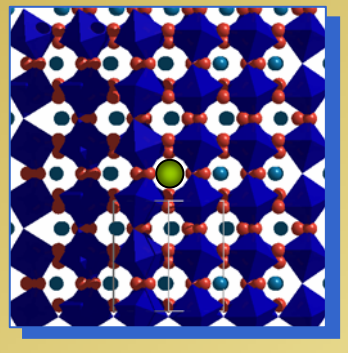

E^0_S = un-doped E_{surf}
E_{seg}= Segregation energy
n = dopants on surface
S = Surface Area

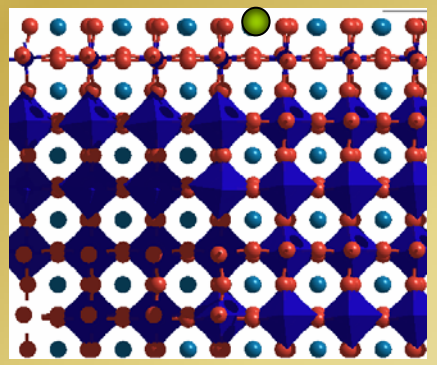

Determine Equilibrium Situation ⇒ Beyond Langmuir Isotherm

Surface Energy for doped surfaces

$$E_S(n) = E^0_S + (n/S)\cdot E_{seg}$$

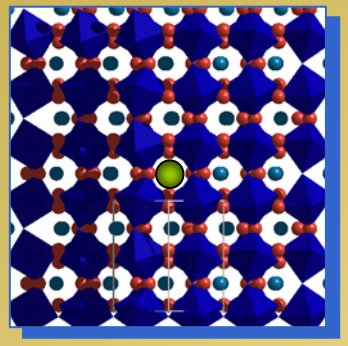

E^0_S = un-doped E_{surf}
E_{seg} = Segregation energy
n = dopants on surface
S = Surface Area

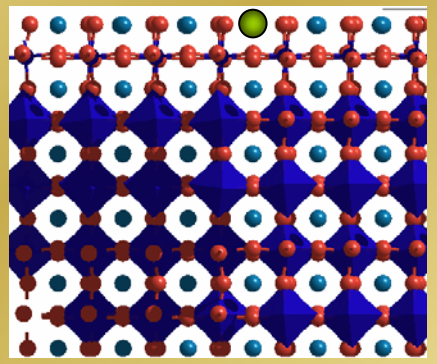

Determine Equilibrium Situation ⇒ Beyond Langmuir Isotherm

Use these surface energies to predict crystal forms as a function of type dopant AND concentration:

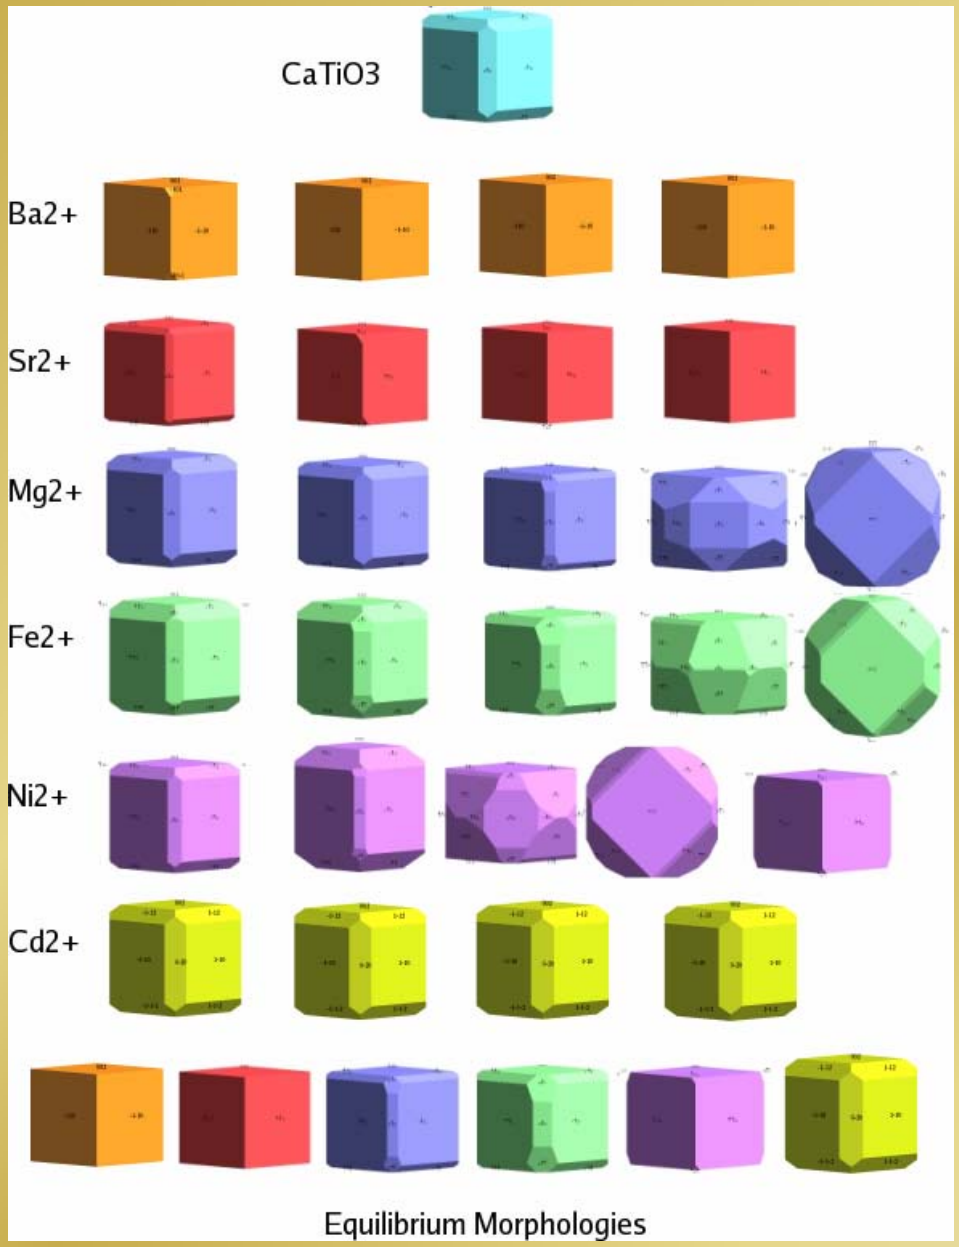

Maria Alfredsson, UCL

Crystal Morphologies of $CaTiO_3$

Predicted Morphology of $CaTiO_3$

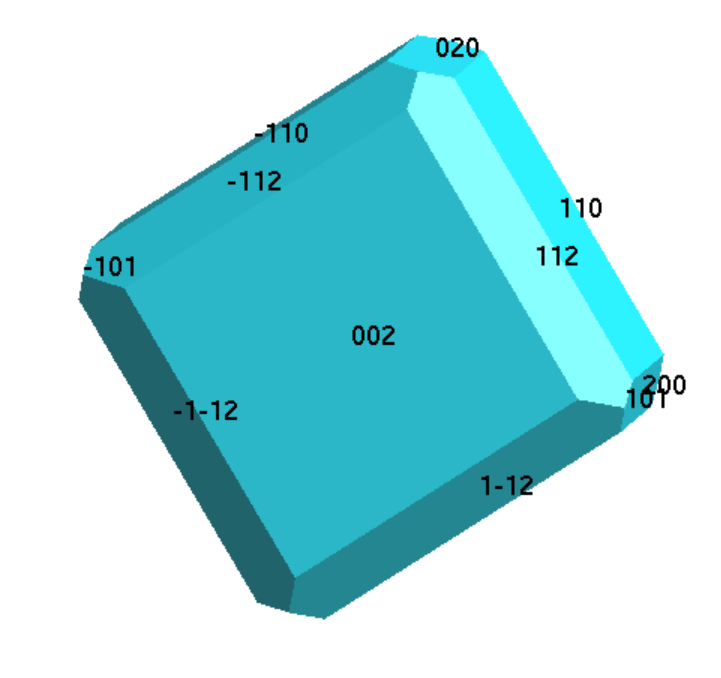

SEM-picture of pure $CaTiO_3$

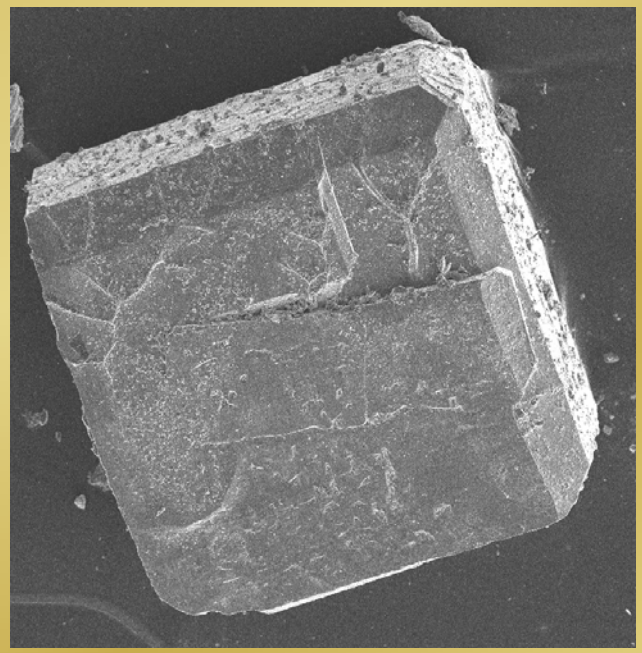

Campus *grids*

- 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

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

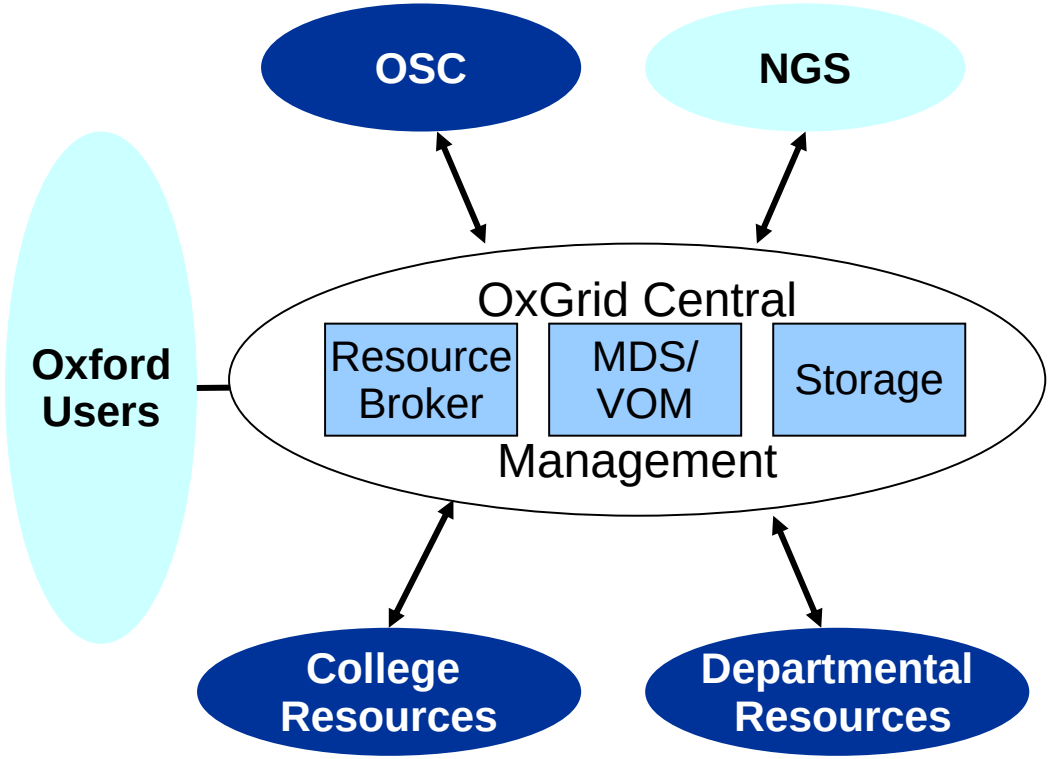

David Wallom

Oxford Interdisciplinary e-Research Centre

Middleware for campus grids

- Globus toolkit
  - Tools built on Grid Security Infrastructure and include:
    - Job submission (GRAM) : run a job on a remote computer
    - Information services: So I know which computer to use
- Storage Resource Broker
  - Virtual filesystem: for files held in multiple locations
  - NIEES offers a testbed to give SRB experience
- SRB and Globus Toolkit 2 are part of the NGS stack

National e-Science Centre

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/

National e-Science Centre

Globus is a Toolkit

- To submit a job to run on a remote computer

```
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
```

- NEED
  - Brokers to allow jobs to be submitted to “a grid”
  - Portals… to empower those interested in their research rather than UNIX and scripting languages!

National e-Science Centre

Storage Resource Broker

User sees a virtual filesytem:

- Command line (S-Commands)
- MS Windows (InQ)
- Web based (MySRB).
- Java (JARGON)
- Web Services (MATRIX)

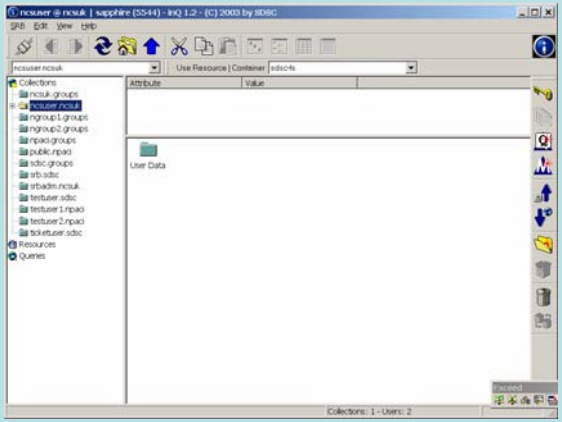

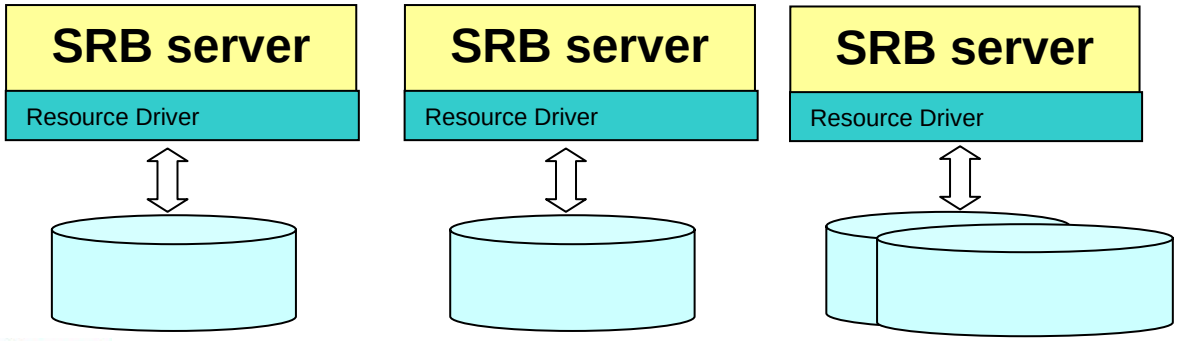

Filesystems in different admin. domains

National e-Science Centre

How SRB Works

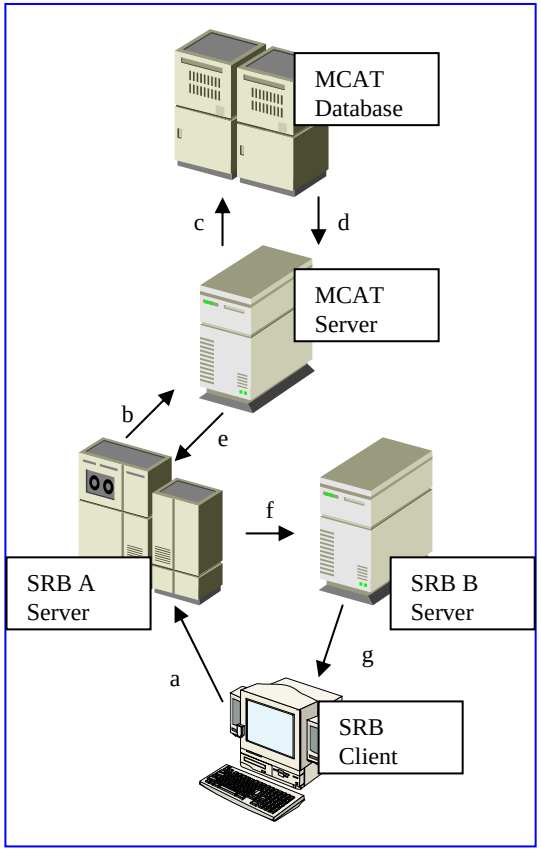

- 4 major components:
  - The Metadata Catalogue (MCAT)
  - The MCAT-Enabled SRB Server
  - The SRB Storage Server
  - The SRB Client

National e-Science Centre

Overview

- Harvesting CPU time
  - Computers within an “administrative domain”
- Grid enabling
  - Computers
  - Data held by collaborating researchers
    - Files, in this talk, databases deferred to the next talk!
- **Some implications**

National e-Science Centre

Before you start!

- 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!

- Be alert to National Grid Service deployment of middleware, and potential for interoperability

Summary

- By using established middleware ….
  - Condor: high-throughput computing
  - Globus Toolkit: basis for resource sharing across admin domains… i.e. a grid
  - Storage Resource Broker: virtual filesystems
- ….can build admin and researcher experience in distributed, collaborative computing
- …gaining the advantages of using the same stack as the National Grid Service
  - Expertise and support available
  - Smaller leaps to partnership in and use of the NGS
- ….and once the infrastructure exists
  - Enhanced potential for interdisciplinary research

National e-Science Centre