



101010001000000100100
101010001000000100100
?
@
1010100010000001

The UK National Grid Service & MPI

Matthew Viljoen, STFC RAL



The Mission of the NGS

The Mission of the National Grid Service

- To provide 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



National Grid Service

- Largest National Grid Initiative outside US
- approx. 500 (real!) Users + Projects
- Has the largest grid PKI infrastructure
- Integration of Grid and HPC
- *Heterogeneous* hardware and software
- Full, partner and affiliate membership status



NGS & Partners, 2008





Middleware

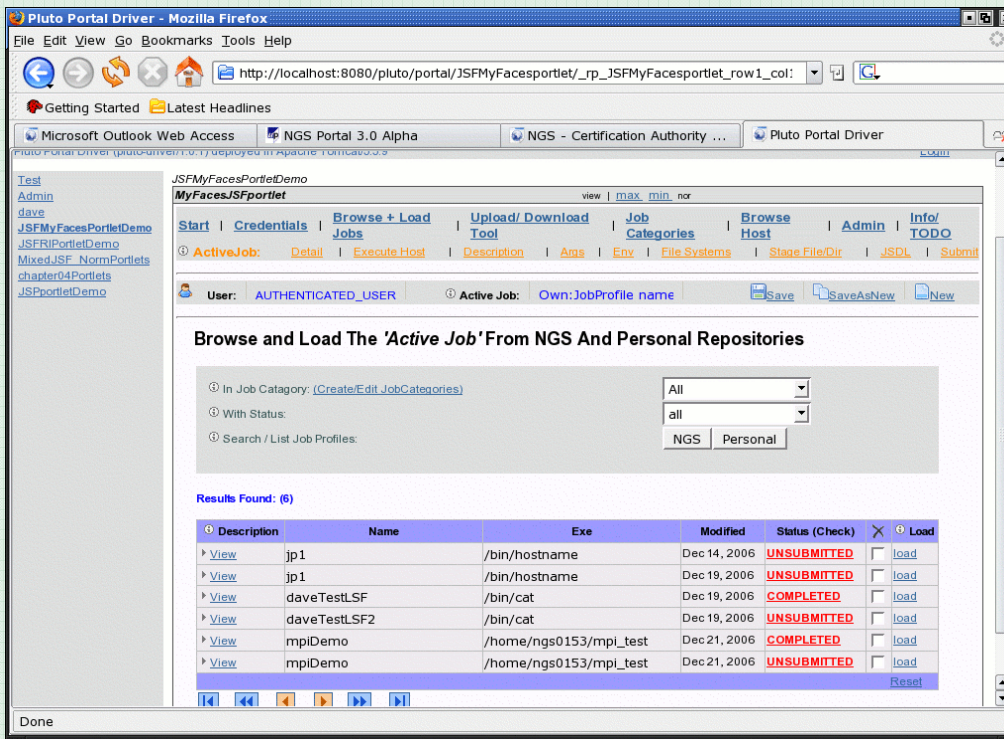
- Resource Broker (gLite WMS-LB)
- CE's based on vanilla VDT (*not* EGEE-type CE) No YAIM style configuration
- VO Support / VOMS available
- Migrating to GT4
 - Retain GT4 pre W/S for GT2 compatibility (this WILL be retained for the foreseeable future!!)
- GridSam Interfaces
- GRIMOIRES Registry
- Portals
 - JSDL Application Repository
 - (previously known as NGS Portal)
 - Updated P-Grade portal (University of Westminster)
- User Accounting System
- Support Centre
 - Fully integrated with EGEE GGUS Support Centre
 - Defined SLD – defined Service provision
- ...



Access Mechanisms

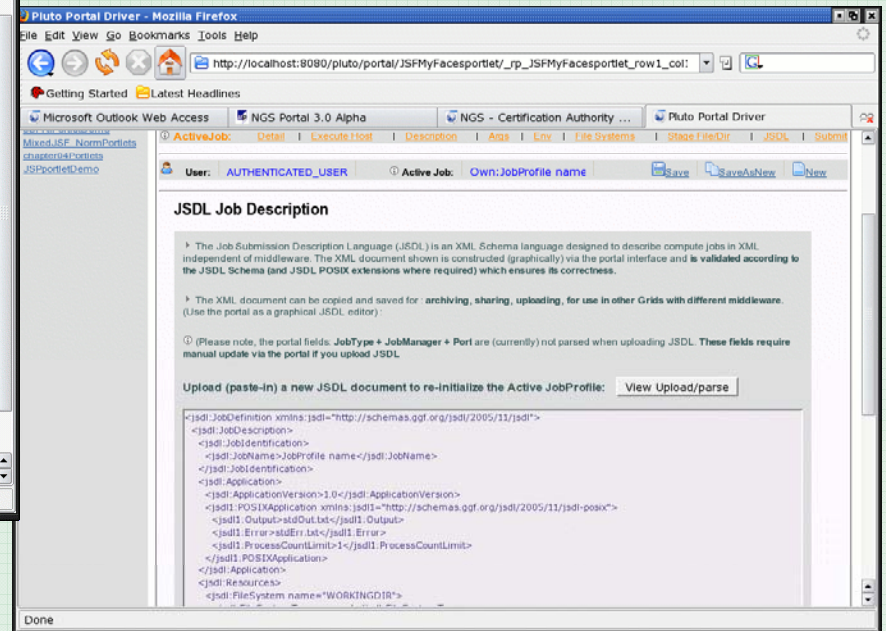
- Direct gsissh login
- Globus tools
- Portals
- Resource Brokers / UI machines
- AHE (Application Hosting Environment)
- GridSAM
- *.....your own application...*

NGS Portal (https://portal.ngs.ac.uk)



The screenshot shows the 'JFSMyFacesPortletDemo' interface. The main heading is 'Browse and Load The 'Active Job' From NGS And Personal Repositories'. Below this, there are search filters for 'In Job Category', 'With Status', and 'Search / List Job Profiles'. The search results show 6 jobs found, listed in a table:

Description	Name	Exe	Modified	Status (Check)	Load
View	jp1	/bin/hostname	Dec 14, 2006	UNSUBMITTED	<input type="checkbox"/> load
View	jp1	/bin/hostname	Dec 19, 2006	UNSUBMITTED	<input type="checkbox"/> load
View	daveTestLSF	/bin/cat	Dec 19, 2006	COMPLETED	<input type="checkbox"/> load
View	daveTestLSF2	/bin/cat	Dec 19, 2006	UNSUBMITTED	<input type="checkbox"/> load
View	mpiDemo	/home/ngs0153/mpi_test	Dec 21, 2006	COMPLETED	<input type="checkbox"/> load
View	mpiDemo	/home/ngs0153/mpi_test	Dec 21, 2006	UNSUBMITTED	<input type="checkbox"/> load




The screenshot shows the 'JSDL Job Description' page. It includes a description of JSDL (Job Submission Description Language) and instructions on how to use the portal. Below the text, there is a section for 'Upload (paste-in) a new JSDL document to re-initialize the Active JobProfile:' with a 'View Upload/parse' button. The JSDL XML content is displayed in a text area:

```
<jSDL:JobDefinition xmlns:jSDL="http://schemas.ggf.org/jSDL/2005/11/jSDL">
  <jSDL:JobDescription>
    <jSDL:JobIdentification>
      <jSDL:JobName>JobProfile name</jSDL:JobName>
    </jSDL:JobIdentification>
    <jSDL:Application>
      <jSDL:ApplicationVersion>1.0</jSDL:ApplicationVersion>
      <jSDL:POSIXApplication xmlns:jSDL="http://schemas.ggf.org/jSDL/2005/11/jSDL-posix">
        <jSDL:Output>StdOut.txt</jSDL:Output>
        <jSDL:Err>StdErr.txt</jSDL:Err>
        <jSDL:ProcessCountLimit>1</jSDL:ProcessCountLimit>
      </jSDL:POSIXApplication>
    </jSDL:Application>
    <jSDL:Resources>
      <jSDL:FileSystem name="WORKINGDIR">
```

Portals: 2

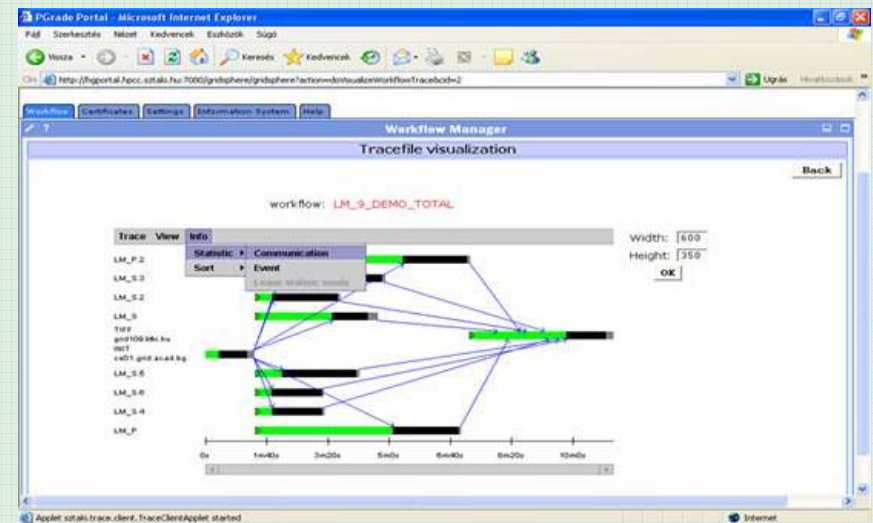
P-GRADE Portal (<http://ngs-portal.cpc.wmin.ac.uk/>)



The screenshot shows the NGS P-Grade GEMLCA portal interface. At the top, there are logos for NGS, P-Grade, and the University of Westminster. A navigation menu includes 'Workflow', 'Certificates', 'Settings', 'Information System', 'GEMLCA Administration Tool', and 'Help'. Below this is a 'Workflow Manager' section with a 'Workflow list' table. The table contains two entries: 'mttc_fromgt2' and 'wflow', both with a status of 'finished'. A message at the bottom says 'Message: Press a button.' and the date 'June 14, 2006' is displayed.

Workflow	Status	Size	Quota (10 Mb)	[Output]	[View]	[Action]
mttc_fromgt2	finished	779 B	0.0%	<input checked="" type="checkbox"/>	Details	Submit Attach Delete
wflow	finished	737 B	0.0%	<input checked="" type="checkbox"/>	Details	Submit Attach Delete



Community Supported Portals
e.g. University of Westminster
NGS P-Grade Portal






Applications: 1

substrate complex

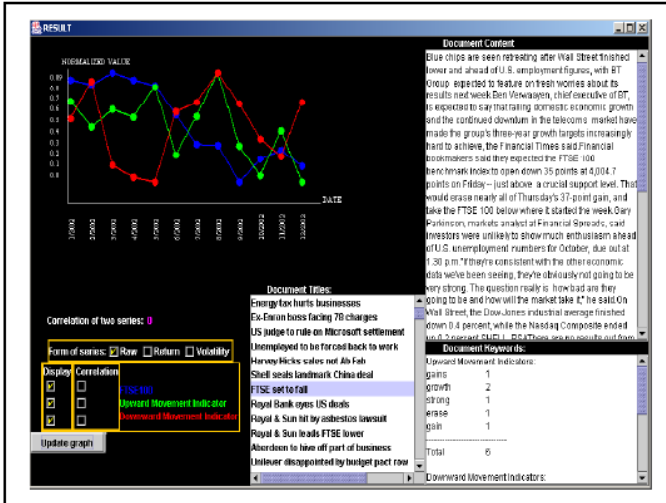
product complex




Molecular Dynamics




Lattice Boltzmann
Text mining

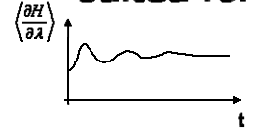


Calculating drug affinities 

Thermodynamic integration is ideally suited for a HPC Grid

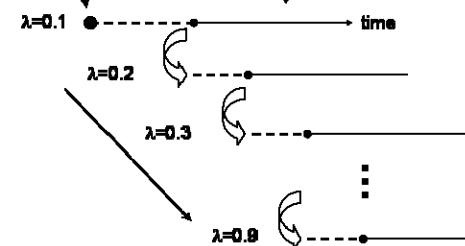


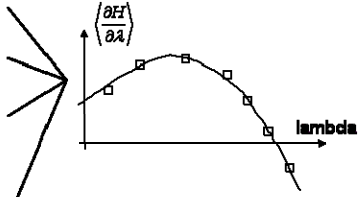
Starting conformation



Check for convergence


Use steering to launch, spawn and terminate λ -jobs





Combine and calculate integral

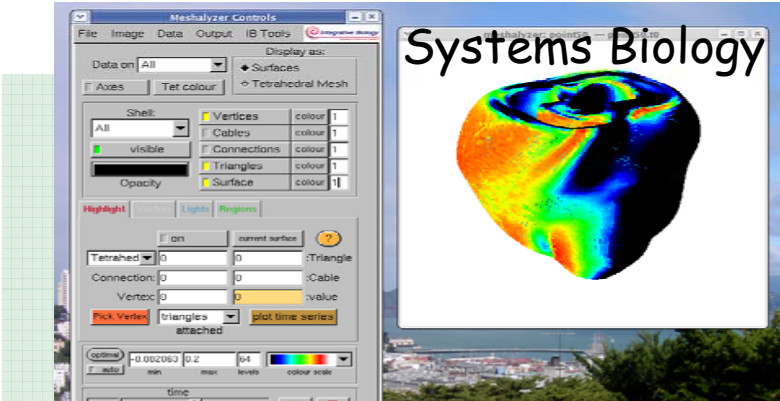
Run each independent job on the Grid



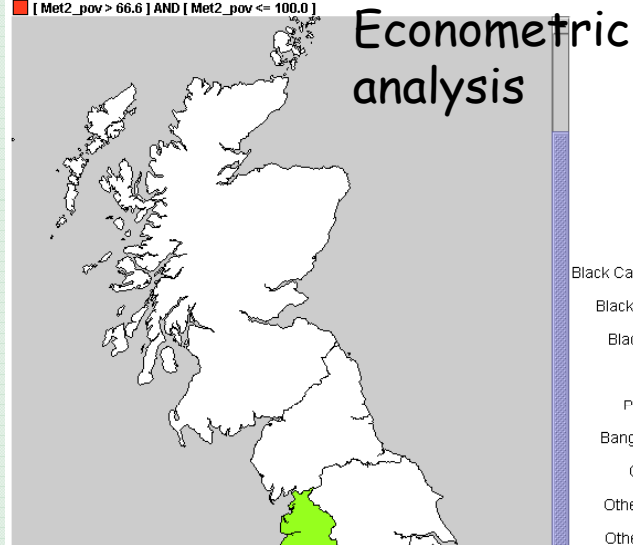
MPI Workshop, 0



Applications: 2

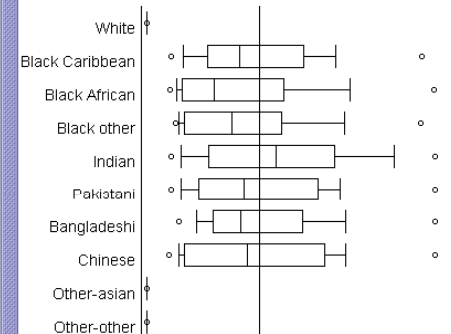


- [Met2_pov = 0]
- [Met2_pov > 0.0] AND [Met2_pov <= 10.0]
- [Met2_pov > 10.0] AND [Met2_pov <= 20.0]
- [Met2_pov > 20.0] AND [Met2_pov <= 25.0]
- [Met2_pov > 25.0] AND [Met2_pov <= 33.3]
- [Met2_pov > 33.3] AND [Met2_pov <= 50.0]
- [Met2_pov > 50.0] AND [Met2_pov <= 66.6]
- [Met2_pov > 66.6] AND [Met2_pov <= 100.0]

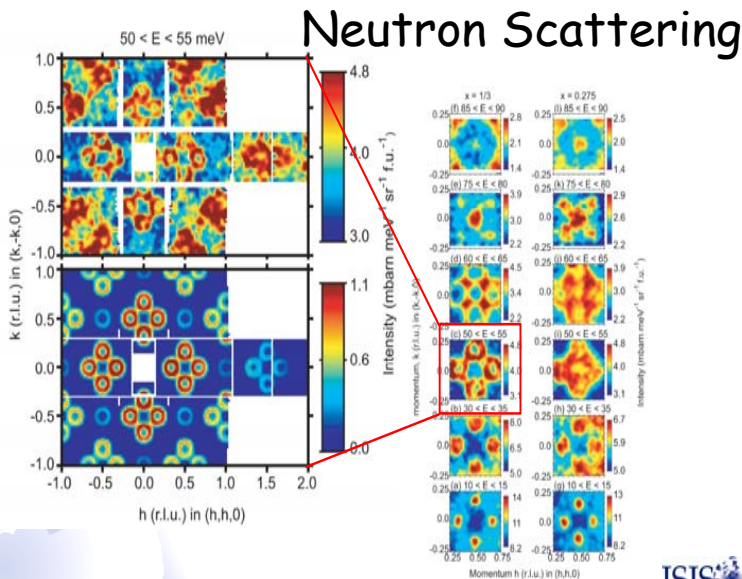


- White
- Black Caribbean
- Black African
- Black other
- Indian
- Pakistani
- Bangladeshi
- Chinese
- Other-asian
- Other-other
- Male
- Female
- All

UK Male Imputed Income

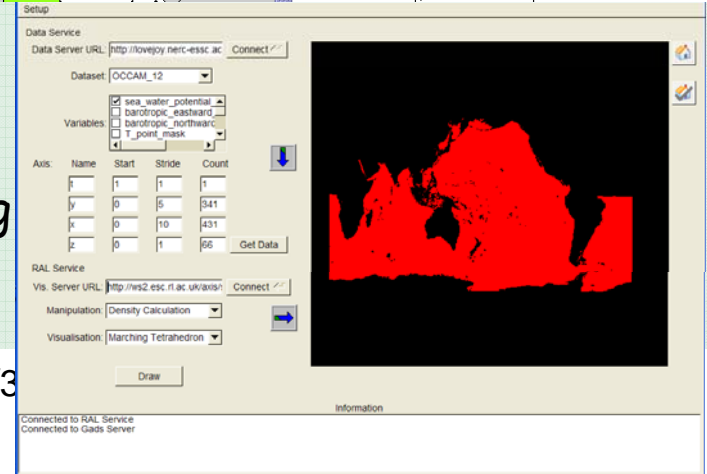


Example: $\text{La}_{2-x}\text{Sr}_x\text{NiO}_4$



H. Woo et al, Phys Rev B 72 064437 (2005)

Climate modelling



Workshop, CNAF, 19/3



Resource Broker

- EGEE gLite WMS-LB (+ UI) from 2007/04
 - The NGS Resource Broker and User Interface is based on the gLite 3.0 WMS-LB and UI nodes from the EGEE Project. The UI is accessible to all members of the NGS by GSISSH on port 2222 at the following node:
`ngsui01.ngs.rl.ac.uk`
- Gridway
 - <http://asds.dacya.ucm.es/GridWay/>
 - Metascheduling technologies for the grid

Jobs Submission

- Mostly submitted via portal & gsissh logon with globus-job-submit
- Central RB submission still new concept
- Most Jobs invoking locally installed applications (including MPI)

(GAUSSIAN, LAMMPS, SIESTA, DL_POLY, BLAST, EMBOSS, AMBER etc.)



Central Job Submission: UEE

- Applications installed at different sites in *different ways*, with different environments.
- Uniform Execution Environment (UEE) guarantees the same wrapper script name for any application.

Site A: /usr/ngs/GROMACS_3 -> /apps/Gromacs/runscript/GROMACS_3

Site B: /usr/ngs/GROMACS_3 -> /opt/Gromacs/bin/gromacs

- UEE scripts published to BDII via
`GlueHostApplicationSoftwareRunTimeEnivonment`

MPI Status

- MPICH-GM, MPICH-MX and Scali-MPI implementations
- both mpirun & mpiexec
- Now rolling out MPI_START from I2G to enable job submission via RB
- Coscheduling already being done by GENIUS project using HARC at UCL

Issues with WMS-LB

- Running older version of gLite-3.0 WMS-LB. No out-of-box MPI support
- Hardcoding of paths in JobWrapper, assuming gLite CE *not* vanilla VDT
- Needed to tweak LSF and PBS template files to work with MPI-START:

```
/opt/glite/etc/templates/template.mpi.lsf.sh  
/opt/glite/etc/templates/template.mpi.pbs.sh
```

Issues with MPI support

- MPI seems to be added to WMS-LB as an afterthought – not primary usecase
- Can't MPI-START be part of WMS?
- Problems with error messages
 - Failed: Status Reason: hit job retry count (3) with NodeNumber > 1
- LRMS configuration at WMS instead of nodes. Practical for different setups?

Future plans

- Upgrade to latest WMS
- Move WMS into production
- Role out MPI-START to all supporting sites
- Integrate MPI support into portals



Acknowledgements

- Stephen Childs
- Maarten Litmaath
- the I2G MPI-START team
for help with MPI

- Slides from Andrew Richards



Thank you

Questions?

support@grid-support.ac.uk