

Computing Session

W T Hewitt

26th October 2005

ASW on Social Simulation

Manchester

Existing Resources

- Desktop/Laptop
- Research Group/Institutional Facilities
- National Facilities
 - The National Grid Service
 - CSAR
 - HPCX
- You will use more than one of these resources
 - Even very big simulations need small, and medium simulations to establish the parameters of the very big simulation



- For the production use of computational and data grid resources, resulting from the UK's e-Science programme.

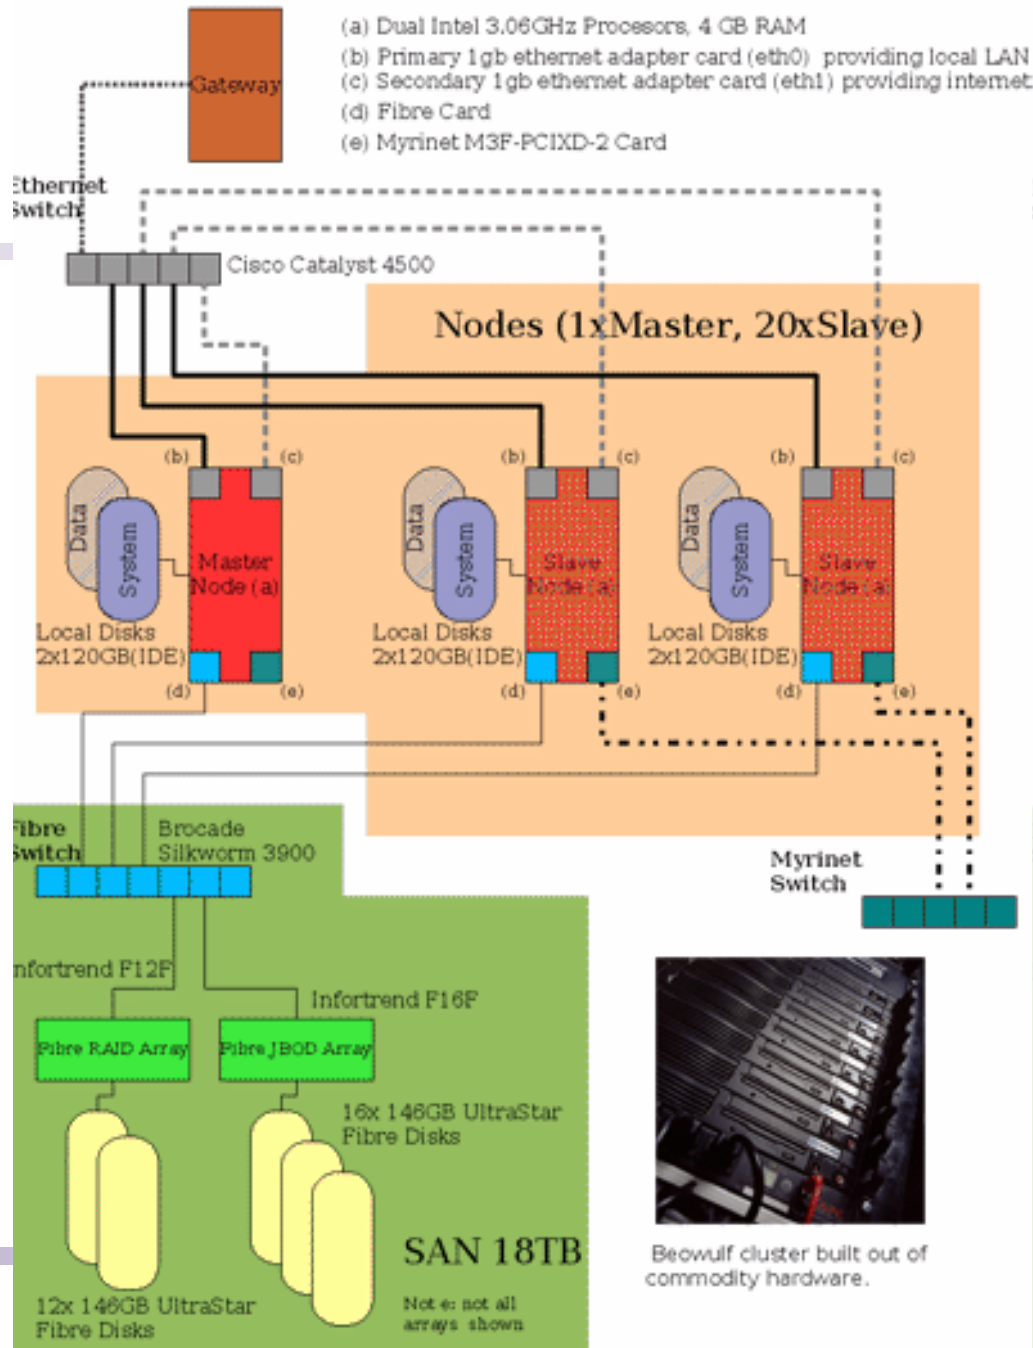
- Types of Nodes
 - Those where you need resources already allocated to use
 - (CSAR, HPCX) through Peer Review by one of the Research Councils
 - Those where you apply to the NGS for resource (or you are a local user)
 - Lightweight peer review at the moment
 - Funded by JISC
 - Give some of your resource to join the NGS

- EPSRC funds the Grid Operations Support Centre (NO application

support)

- Nodes (to date)
 - CSAR – RCs funded
 - HPCX – RCs funded
 - Manchester – JISC Funded data node
 - Rutherford Lab – CCLRC Funded data node
 - Leeds – JISC Funded computational node
 - Oxford – JISC Funded computational node
 - Bristol – self funded
 - Cardiff – self funded
- Use these systems through Grid software, typically
 - Globus, OGS-DAI, OGSA-DQP, SRB

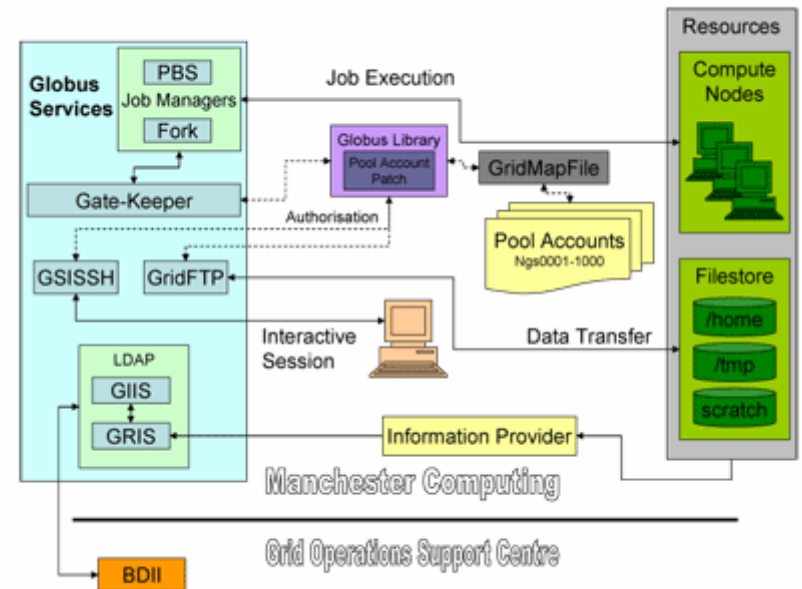
MC NGS Data-Node Hardware Overview



Manchester Node

- Globus
- Storage Resource Broker (SRB)
- OGSA-DAI

- Usual Tools
 - Compilers, Libraries, MPI



- <http://www.kato.mvc.mcc.ac.uk/blog/NGS/>

- Induction to Grid Computing and the National Grid Service
 - 13:00 02 November, 2005 to 16:00 3 November, 2005
 - Rutherford Appleton Laboratory

- <http://www.nesc.ac.uk/esi/events/633/>
- Bring your own laptop (or share with someone)

- Registration should be made via the NeSC website

UK National HPC services

- CSAR
 - 2 tightly coupled systems – Origin and Altix
 - provided by CfS consortium
 - University of Manchester (MC)
 - Computer Sciences Corporation (CSC)
 - SGI

- HPCx
 - based around a single large IBM cluster
 - provided by UoE HPCx consortium
 - EPCC
 - DL
 - IBM

**In-depth
Applications
Support is provided**


Pump-priming resources are free

**Apply for resources through
Research Grant Application to
Your favourite Research Council**

HPCx technology overview

www.hpcx.ac.uk

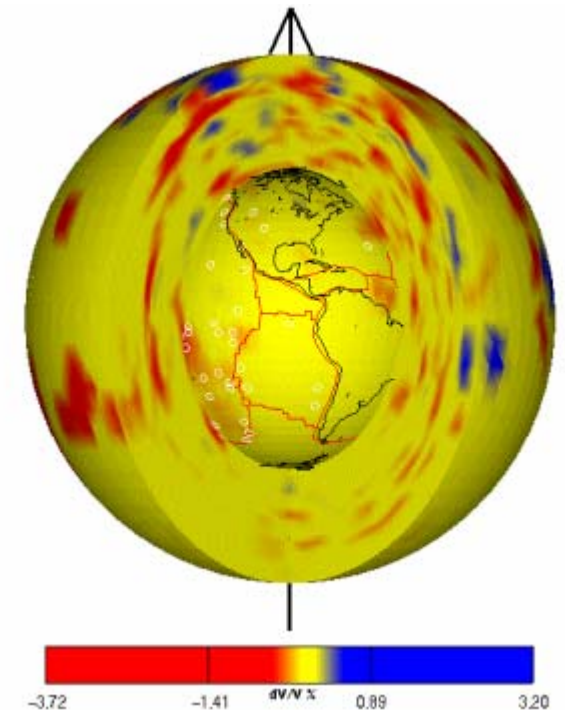
- IBM p690 cluster
 - Phase 2 (2004-now)
 - 6 TFLOPS LINPACK performance
 - 1600 x IBM Power 4+
 - Phase 2a (November)
 - ~1500 x IBM Power 5
 - Phase 3 (2006-08)
 - 12 TFLOPS LINPACK performance
- AIX, Fortran, C, MPI, ...
- £53M



CSAR technology overview

www.csar.cfs.ac.uk

- SGI supercomputers
 - “0.2 TFLOPS NPB”
 - 512 Itanium 2 CPUs SGI Altix 3700 (1TB memory)
 - 512 MIPS CPUs SGI Origin3000
 - £26M
- Linux, IRIX, Fortran, C, MPI, ...



Next national service: HECToR

- HECToR : 2007 – 2013?
 - £200M? (£50M phase 1 hardware?)
- HECToR back of envelope calculations show
 - ~75TFLOP peak phase 1 (doubles each phase 2 & 3)
 - ~5000 processors
 - ~1-2MW power required
 - ~1-4000 sq ft
- HPC budgets
 - “huge” compared to typical academic projects
 - tiny compared to other govt IT projects (health, ATC, ID, etc)

Discussion Questions

Is a Grid Critical or a pre-requisite or a red-herring?

- Is the availability of large scale computational Grids critical to the development of social simulations?
- What is the most important functionality that they provide?
 - Exploring large parameter spaces
 - Performing more detailed (statistical) analyses of simulation results
 - Storing voluminous simulation results
 - Advanced visualization

Who Pays? Grid is NOT free!

- Who is going to provide the computing/storage power?
- Should we rely on specialised scientific grids or set up a volunteers'-powered grids?
- Advantages of general purpose scientific grids (e.g., UK National Grid)
 - Immediate availability.
 - More homogeneous computing environment.
 - Greater reliability (particularly for long-term storage).
- Advantages of volunteers'-powered grids
 - (e.g. SETI@Home or World Community Grid)
 - Potentially, much greater computational power (NGS has 2K CPUs, the largest supercomputer around 30-50K, SETI@Home 400K, WCG 150K)
- Involvement of the general public in social research

- On the base of what criteria should Grid resources be allocated?
 - Scientific merit (evaluated by peer review).
 - Credit system based on the amount of resources provided to the Grid.

What model of distributed computation?

- Ensemble Modelling: multiple copies of the same simulation running on individual processors but with different parameters
- Distributed worlds: a single simulation running on multiple processors, each taking care of a part of the simulation space
- Services networks: simulations exploiting specialised distributed services (e.g. geographical databases, data archives, neural network processors, advanced visualisation facilities; use cases: simulations with sophisticated agent models, simulation validation, micro-simulations).

How to Contact Me

Mr W T Hewitt
Manchester Computing
University of Manchester
Manchester M13 9PL
United Kingdom

tel +44 161 275 6095

fax +44 161 275 6800

w.t.hewitt@manchester.ac.uk

www.manchester.ac.uk

www.mc.manchester.ac.uk

www.mimas.ac.uk

www.csar.cfs.ac.uk

The University of Manchester

World Leading Research, Teaching & Services

www.manchester.ac.uk