



# Monitoring and Accounting on the NGS

Guy Warner  
NeSC TOE Team

# Policy for re-use

- This presentation can be re-used for academic purposes.
- However if you do so then please let [training-support@nesc.ac.uk](mailto:training-support@nesc.ac.uk) know. We need to gather statistics of re-use: no. of events, number of people trained. Thank you!!



# Acknowledgements

- The slides in this presentation are taken from presentations by:
- S. Pickering at the e-Science All Hands Meeting 2005.
  - <http://www.ngs.ac.uk/guide/NGS-Partners-AHM05-Pickering.ppt>
- J. Schopf
  - [http://www-unix.mcs.anl.gov/~schopf/Talks/mds4Inca\\_lcg\\_nov2004.ppt](http://www-unix.mcs.anl.gov/~schopf/Talks/mds4Inca_lcg_nov2004.ppt)
- S. Smallen and K. Ericson at Super Computing 05.
  - [http://inca.sdsc.edu/downloads/inca\\_sc05.pdf](http://inca.sdsc.edu/downloads/inca_sc05.pdf)
- K. Weeks – presentation at All Hands Meeting 2007



# NGS Grid Monitoring

- Service Reliability
- Performance Monitoring
- Benchmarking
- Site Interoperability Certification
- Software Stack Validation
  
- Customisations
- Archiving
- Integration – PBS, GITS, Ganglia, INCA



# Collecting Information

- System Administration
  - Operating System
  - Disk
  - Network
  - Problem detection
- User Information
  - Software/Modules
  - Queues
  - Resources



# What is monitoring?

- Discovery and expression of data
- Discovery:
  - Registry service
  - Contains descriptions of data that is available
  - Sometimes also where last value of data is kept (caching)
- Expression of data
  - Access to sensors, archives, etc.
  - Producer (in consumer producer model)



# What is Grid monitoring?

- Grid level monitoring concerns data that is:
  - Shared between administrative domains
  - For use by multiple people
  - Often summarized
  - (think scalability)
- Different levels of monitoring needed:
  - Application specific
  - Node level
  - Cluster/site Level
  - Grid level
- Grid monitoring may contain summaries of lower level monitoring



# Grid Monitoring Does Not Include...

- All the data about every node of every site
- Years of utilization logs to use for planning next hardware purchase
- Low-level application progress details for a single user
- Application debugging data
- Point-to-point sharing of all data over all sites





# INCA & Ganglia

- INCA
  - a framework for the automated testing, benchmarking and monitoring of Grid resources
  - INCA on the NGS - <http://inca2.ngs.ac.uk/>
- Ganglia
  - Each node broadcasts information (UDP Multicast)
  - One node listens
  - Good for current CPU/Memory usage
  - Ganglia on the NGS - <http://ganglia.ngs.rl.ac.uk/>
    - Only the front page is available to users. You will get "Page not found" or equivalent errors if you try and drill down into ganglia.

# Grid Accounting

- Accounting for any production grid is an important part of the monitoring process
  - Pricing policies may be introduced to grids in the future
  - To uphold policies relating to grid use and allocated hours
  - To monitor systems – particularly important for funding and future planning
  - To have an overview of the system – how much are we allocating? How much is being used? How much spare capacity do we have? How much are our biggest users using?
- It's an issue many grids now face

# Grid policing

- Users are allocated *limited* resources
- Important to know how much of those resources have been consumed
- Users tend to go over quota even when monitored
- Need to 'lock-out' users who go over quota
- There is an important distinction between accounting and policing
- Retain integrity of application and peer-review process



# Policing the NGS

- User Accounting System (UAS) queries the RUS every day for total CPU and disk space for every user
- A warning email is sent out when you reach 90% of your CPU allocation
- The account is automatically locked and an email sent when you reach 100% of your CPU allocation



## Policing the NGS (2)

- When an account is locked, you can apply for more resources
  - Via application form
  - Via your account details
- When your application is successful, your account is automatically updated with your new allocation and account is 'active' again
- An email is sent to you letting you know you're back within your limits
- Your account will be active within the hour.



# Accessing your details

- Users wanted to know how much of their allocation they had used
- <https://www.ngs.ac.uk/useraccountinfo.php>
- Certificate access to account details
  - Not supported by Oracle Apex
  - Needed a workaround to take certificate details from browser
- Also provides ability to change contact details
- Renewals can be done through their own account