Production Issues

Kaushik De

University of Texas at Arlington

Tier 2 Workshop, UTA

December 8, 2006

Production Status

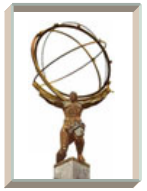

- Great progress since last meeting:
  - CSC production running well with Panda in the U.S.
  - 12 sites busily churning out events for physicists
  - U.S. production accounts for 28% of ATLAS wide production
  - Thanks to BNL and all the Tier 2's we heard from already
- But we are discovering scaling issues:
  - November produced record data rates – also many problems
  - Eowyn – stops sending jobs if too many status updates
  - PandaDB – frequent lock wait timeout's
  - DDM – transfer backlogs, lost files
  - All systems and services already stressed – at <10% of required scale in 2008: fixing this will be critical focus in 2007

CPU Consumption by Successful Jobs

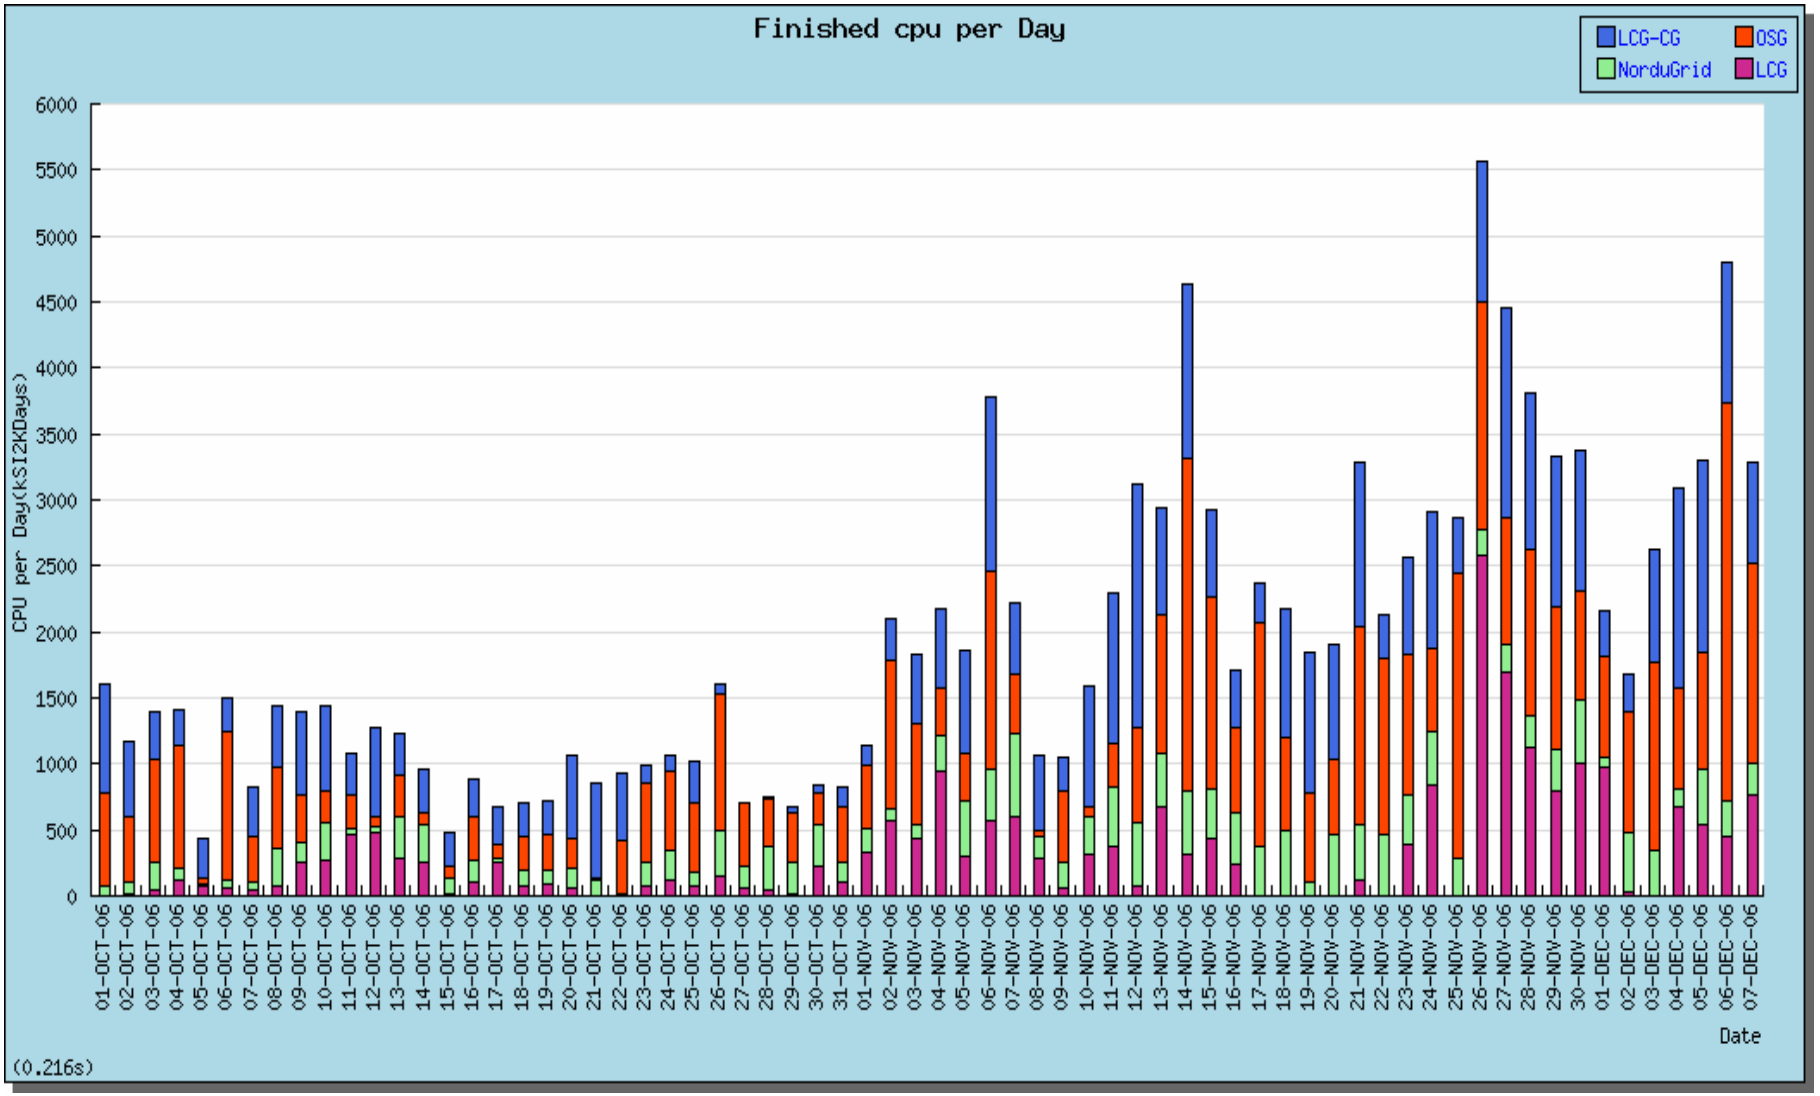

Finished vs Failed Walltime Usage

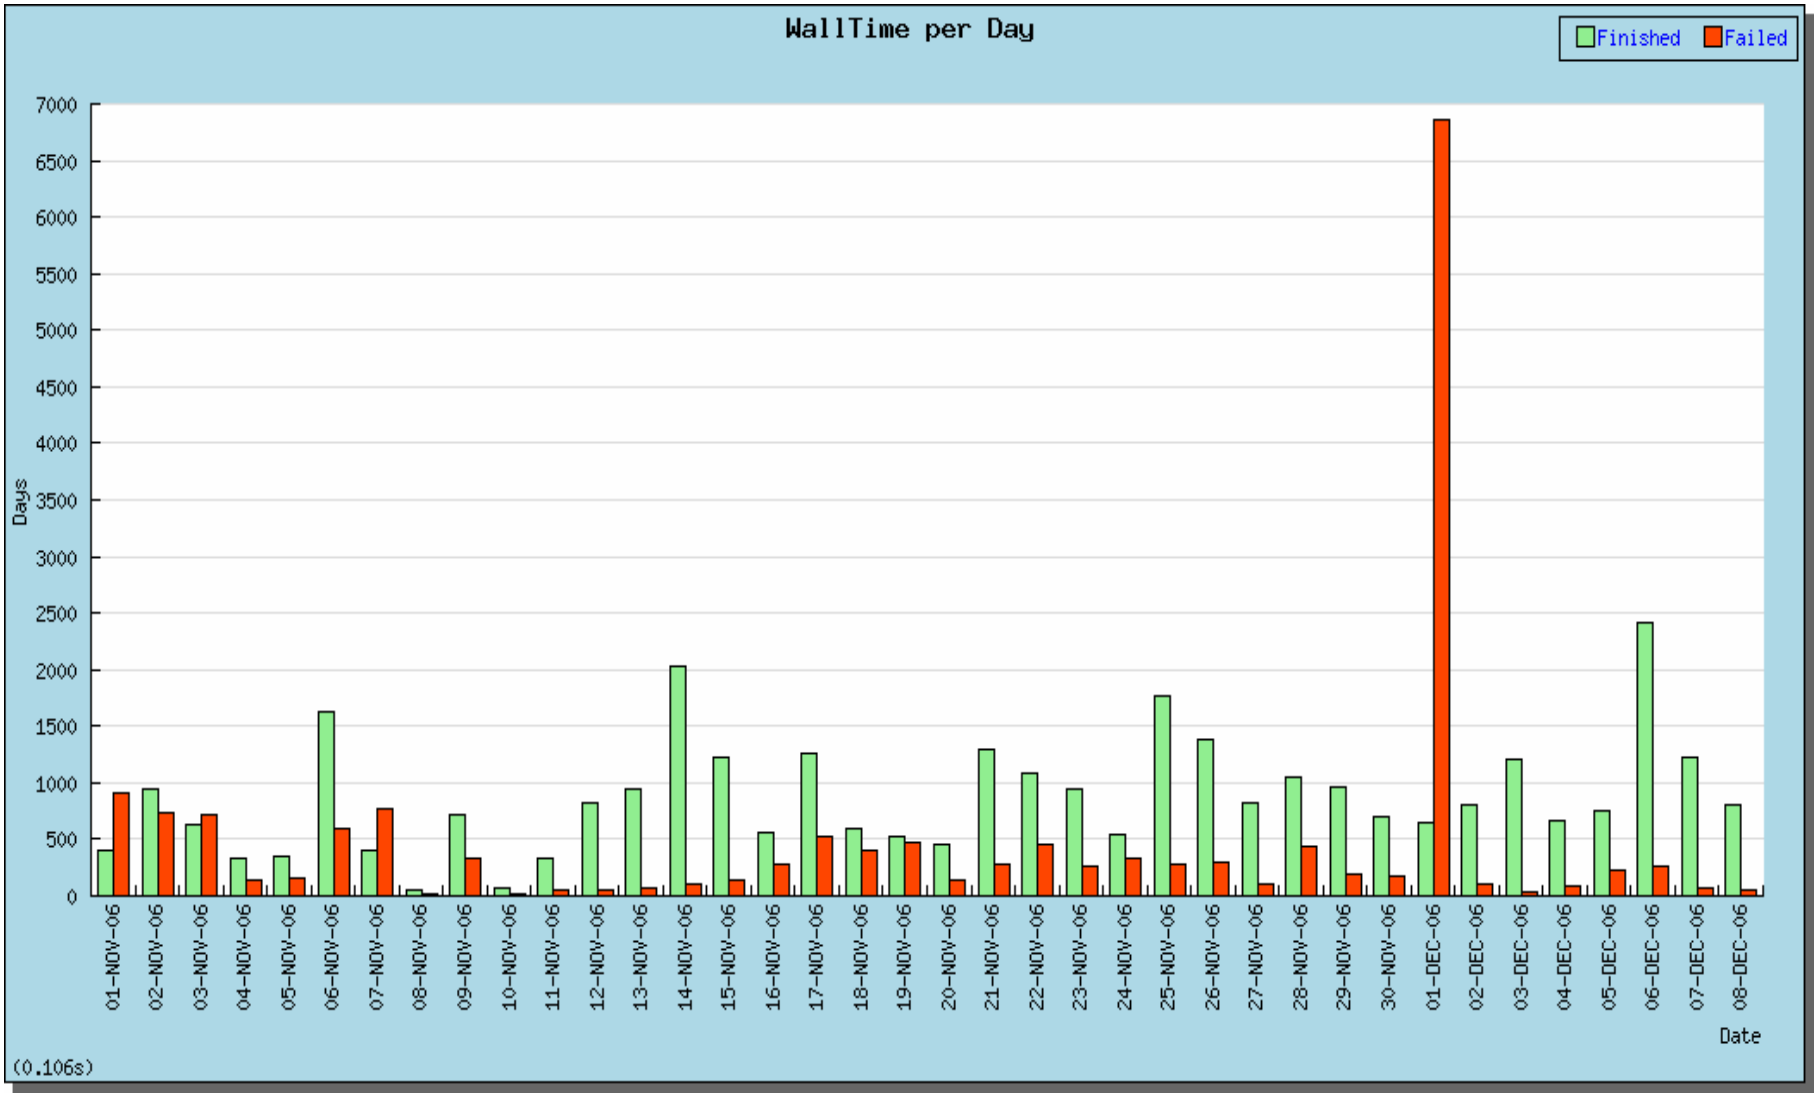

CPU Usage/day For Successful Jobs

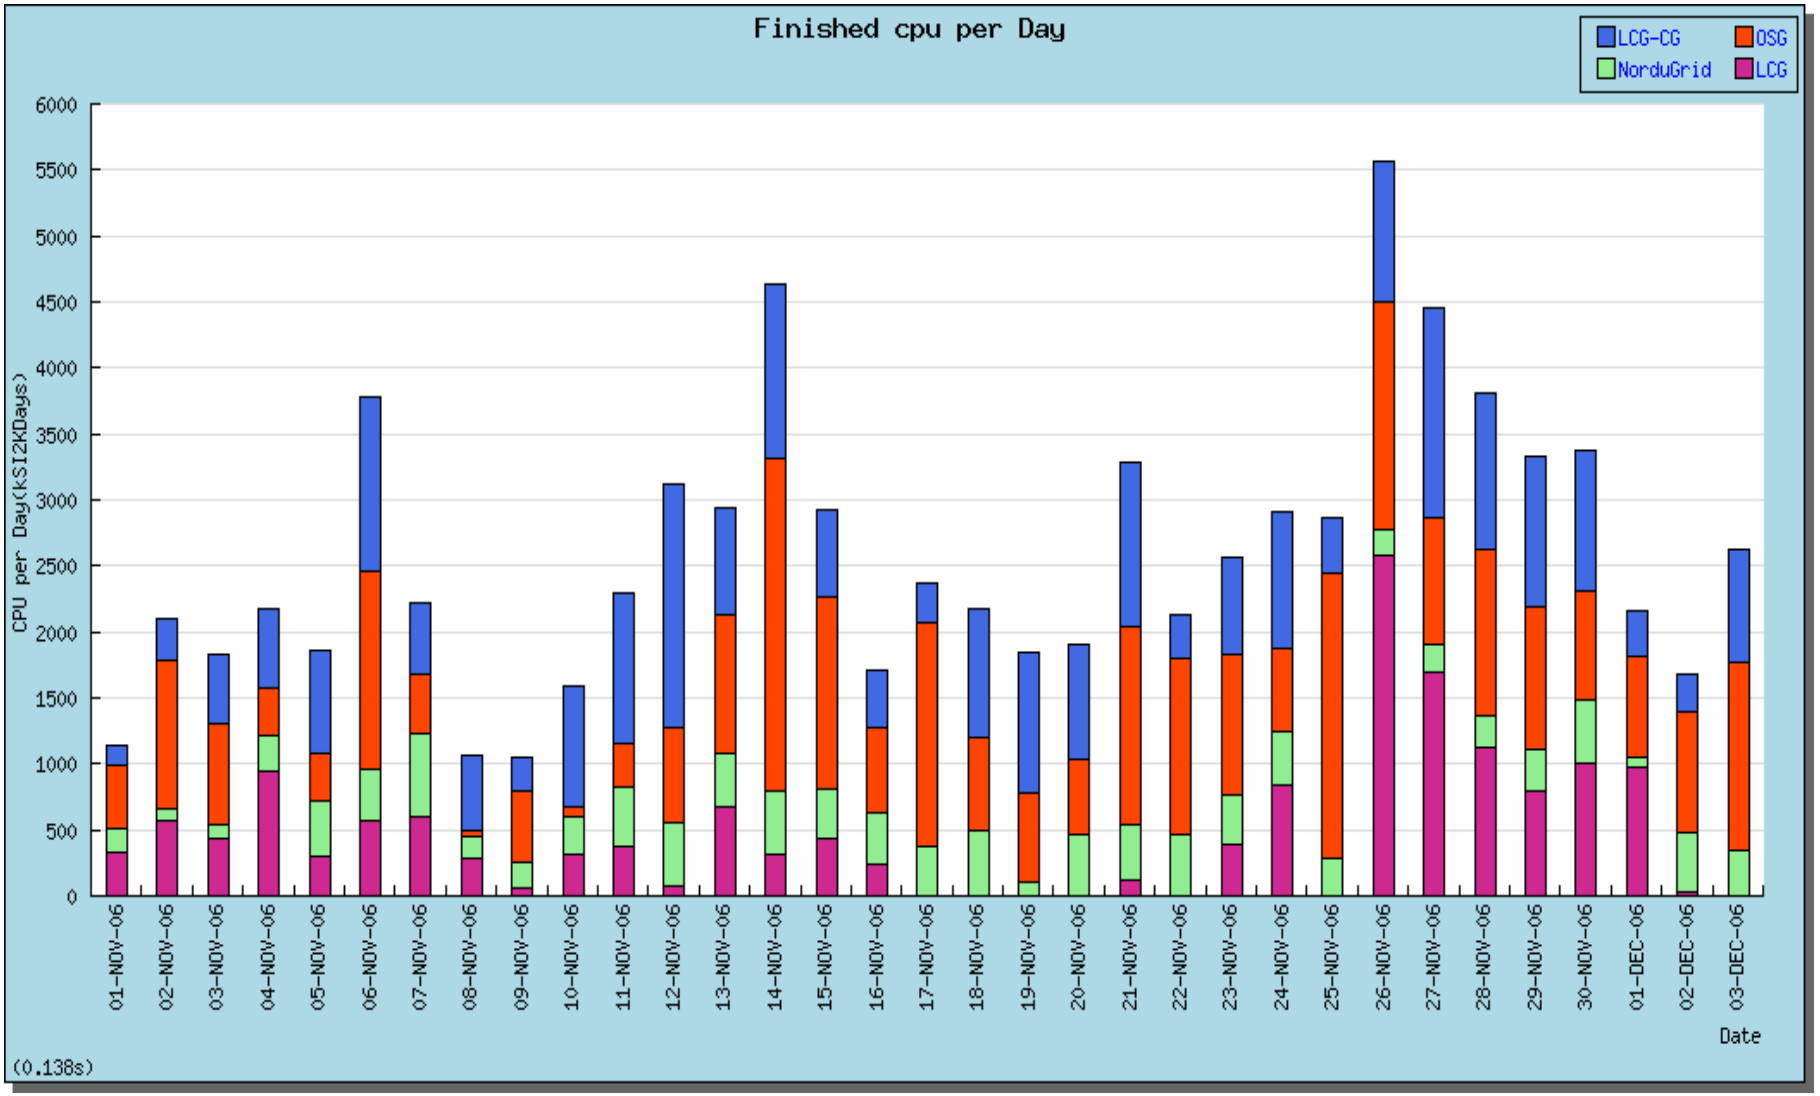

CSC Production – Task Breakdown

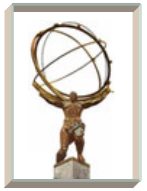

Number of Tasks

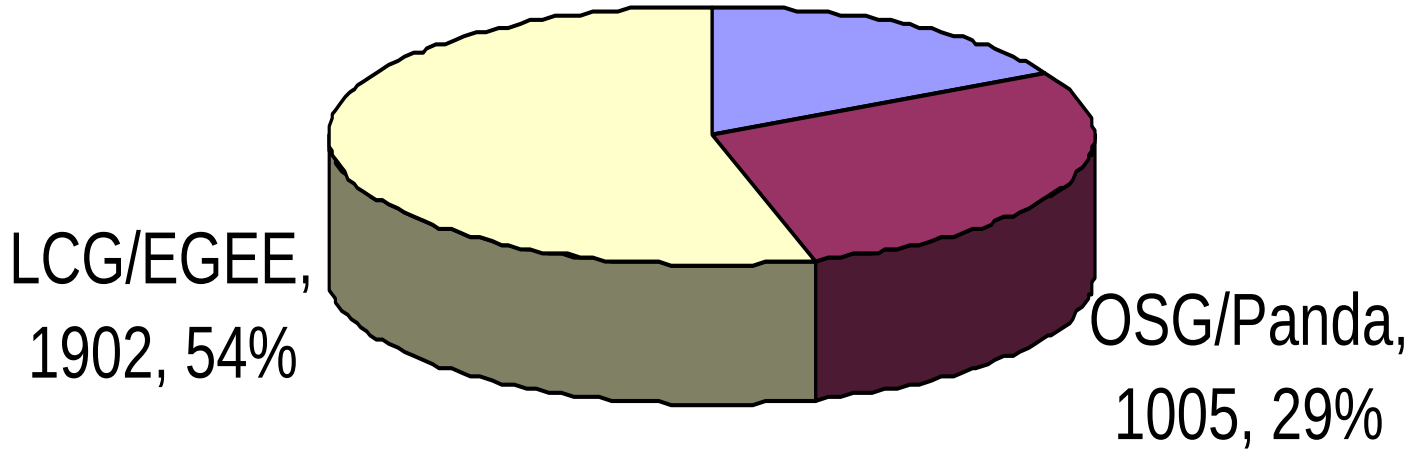

CSC Production – Job Breakdown

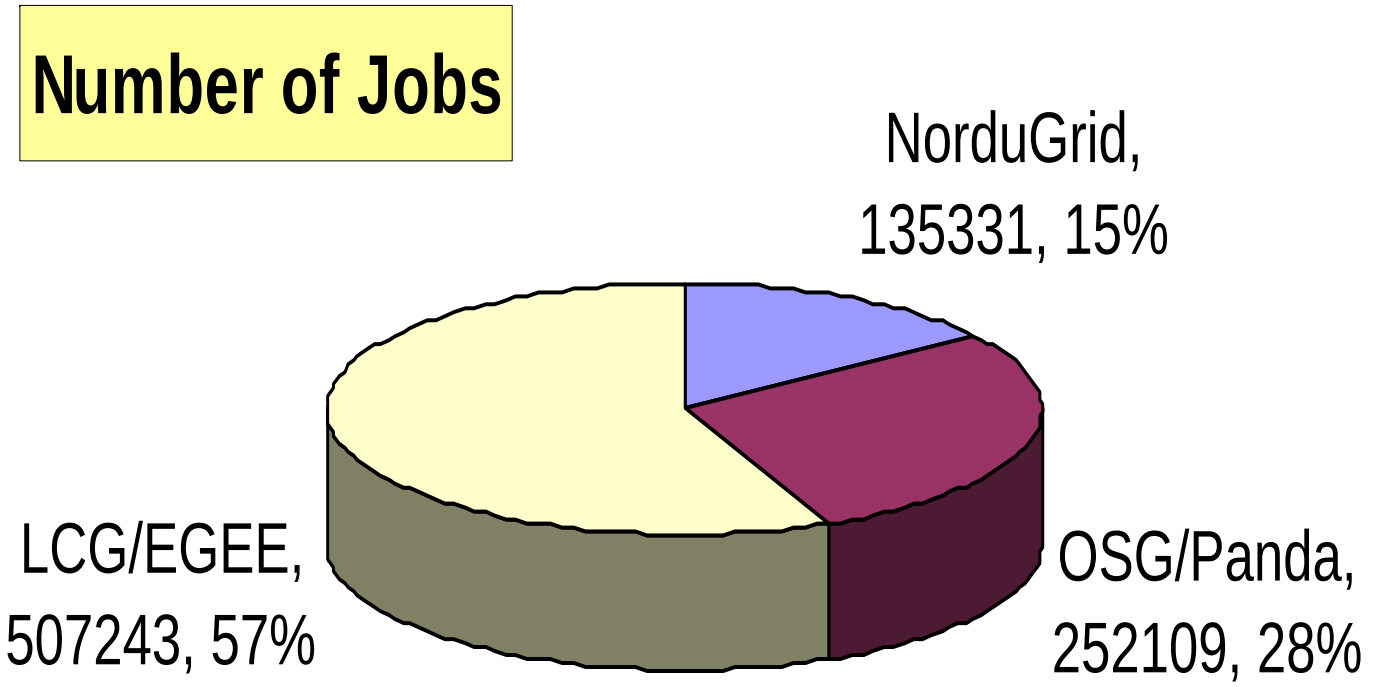

Panda production – 50% of the jobs done on Tier 1 facility at BNL
50% done at U.S. ATLAS Tier 2 sites

U.S. ATLAS Production

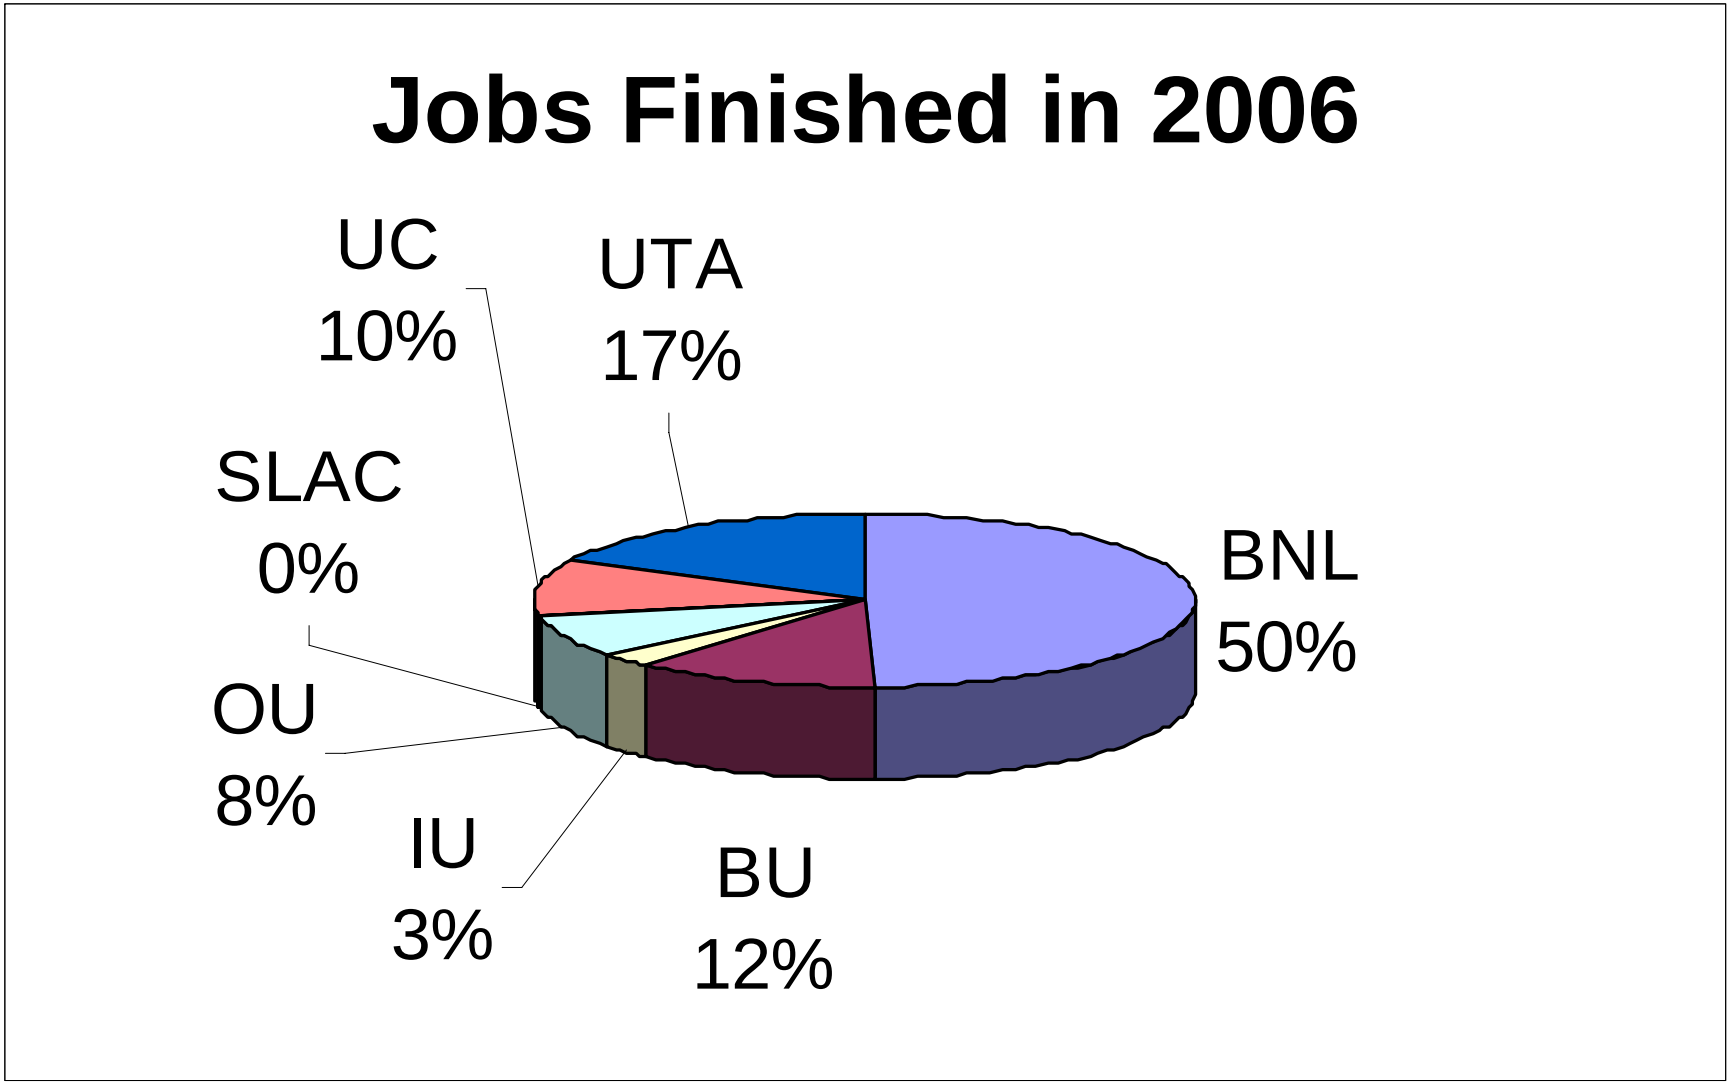

Panda Status

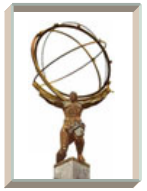

- Overall Status
  - Working very well for distributed MC production on OSG -28% of all ATLAS wide jobs now run in the U.S. (BNL + all ATLAS Tier 2 sites)
  - Distributed analysis through Panda widely used (>75 users)
- Recent improvements
  - Multiple pilot submission mechanisms for robustness
    - Local pilot submission to batch systems at 2 largest sites – BNL, UTA
    - CondorG based jobscheduler at remaining Tier 2 sites
    - New OSG funded pilot factory running at 186 different gatekeepers (250 different queues)
  - Multi-tasking pilot now running – will allow additional DA resources
  - User accounting and quota systems – fully operational
  - Improved error handling (pilot job recovery), robustness against grid failures etc.

Panda Plans

- Goals for next 6 months
  - Increase number of sites served by Panda (from ~20 to ~200 to be used for production and distributed analysis)
  - Support for managed regional or local user production (distinct from user analysis)
  - Support for additional Physics Analysis use cases
  - Scalability improvements – recent bottlenecks seen at 20k simultaneous jobs in Panda (causing timeouts in MySQL PandaDB) and 15k simultaneous dataset subscriptions in DDM
  - Expect Panda to meet all distributed computing needs for U.S. physicists when ATLAS data arrives in 2007

Scaling Issues

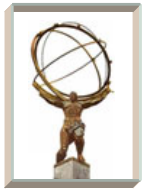

- Eowyn
  - ATLAS prodsys component that sends managed production jobs
  - Karthik (OU) is working on tuning performance for Panda
  - Main issue – single threaded program, which alternates between sending jobs and checking status, which may cause bottlenecks at high rates
- PandaDB
  - In memory MySQL cluster with two nodes
  - 'Lock wait timeout' problems frequent
  - Intense work over past 6 weeks at BNL (and with Panda team) to resolve this
  - Not fully resolved yet
- DDM issues – next talk

Live Demo

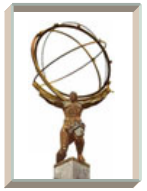

http://gridui02.usatlas.bnl.gov:25880/server/pandamon/query?dash=prod