



Science & Technology Facilities Council

e-Science

CASTOR Database Capacity Planning

Gordon D. Brown

Rutherford Appleton Laboratory

3D Workshop
PIC, Barcelona
20th-21st April 2009





Overview

- Hardware
- Software
- Monitoring
- Production
- Future



Hardware

- What are we all using?
 - 32/64 bit
 - Cores
 - Memory
 - Processor
 - Speed
 - RAID
- Setup
 - RAC
 - Test databases
 - Standby



Space

- Current RAL disk space:
 - $(25 + 11) = 36\text{GB}$
 - Approx 200GB of redo logs per day
- How much space growth should we expect?
 - Exponential
 - Different schemas (stager, DLF)
 - Space for analysis (log miner)
 - Backup policy



Software

- Oracle
 - 10g/11g
 - CPU
 - Test databases
- CASTOR version
 - 2.1.7



Metrics

- What should we record?
 - Growth
 - Transactions per Second
 - Application stats
 - Rows per tale



Interaction with Service Challenges

- **Schedule**
 - STEP09
- **Production**
 - Full running
 - ASM reports



Monitoring

- DB Load
 - Difficult to know if linked to requests/files
 - Tools of CASTOR “load” would be useful
 - Is application “good” at being on RAC
- Oracle Services
 - Currently one “preferred” node and one “available” node for each schema
 - Stagers failover to SRM for example
 - Is two nodes per Stager better?



Science & Technology Facilities Council

e-Science

Questions and (hopefully) Answers

databaseservices@stfc.ac.uk