

WLCG workshop  
Database BOF Session  
1 September 2007

# LHCb Scalability Tests at CNAF

Marco Clemencic  
[marco.clemencic@cern.ch](mailto:marco.clemencic@cern.ch)

- Procedure:
  - Retrieval of a full snapshot of the CondDB content
  - Many concurrent jobs
    - Start the retrieval at the same time
- Goal:
  - Discover bottlenecks

- Expecting to be able to run ~1000 jobs
- Actually:
  - Connections: **OK** (random problems)  
(1 connection with 2 user sessions per job)
  - Network: doesn't seem to be limiting
  - CPU: **bad** (saturated 2 nodes)
    - problem with non-optimized COOL query (~3%/job)
    - optimized query is better, but CPU saturated with system activity (context switching?)

# Comments

- COOL team is aware that the query I was using is not optimized and it is now high priority
- The situation is not extremely bad
  - We expect ~200 jobs/site
  - On average 1 job starting every 2-3 min.  
(actually starting in small bunches)
- We would like to have some safety margin  
(server is shared: LFC and CondDB)