



東京大学
THE UNIVERSITY OF TOKYO

Analysis in STEP09 at TOKYO

Hiroyuki Matsunaga
University of Tokyo

WLCG STEP'09 Post-Mortem Workshop

Site Configuration during STEP09


- Tier-2 dedicated to ATLAS

- SE:

- DPM 1.7.0
- 13 disk servers + 1 head node

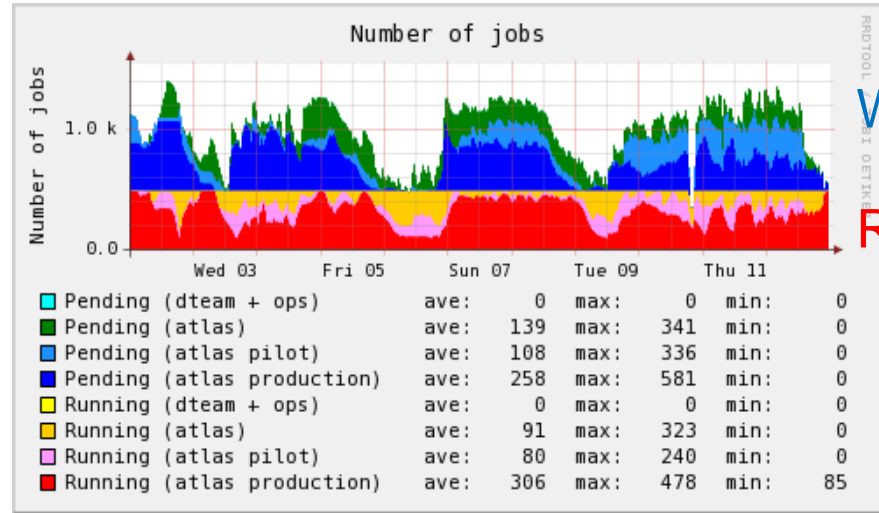
- WN:

- 4 cores/node
- 120 nodes (480 cores)




Job efficiencies

- Job submission rates were not stable
 - Affected fairshare history
- Efficiency (CPU/Wall time) ~ 0.5 for both WMS (direct RFIO) and Panda (staging)

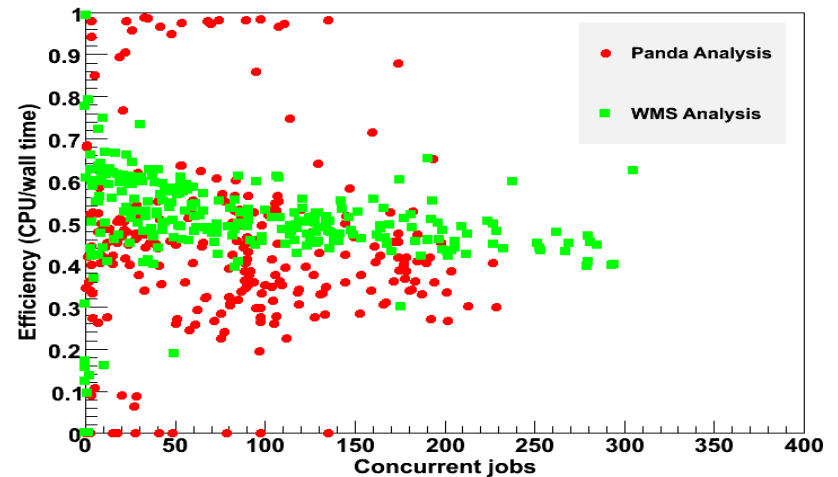


Waiting
Running

Job Efficiencies

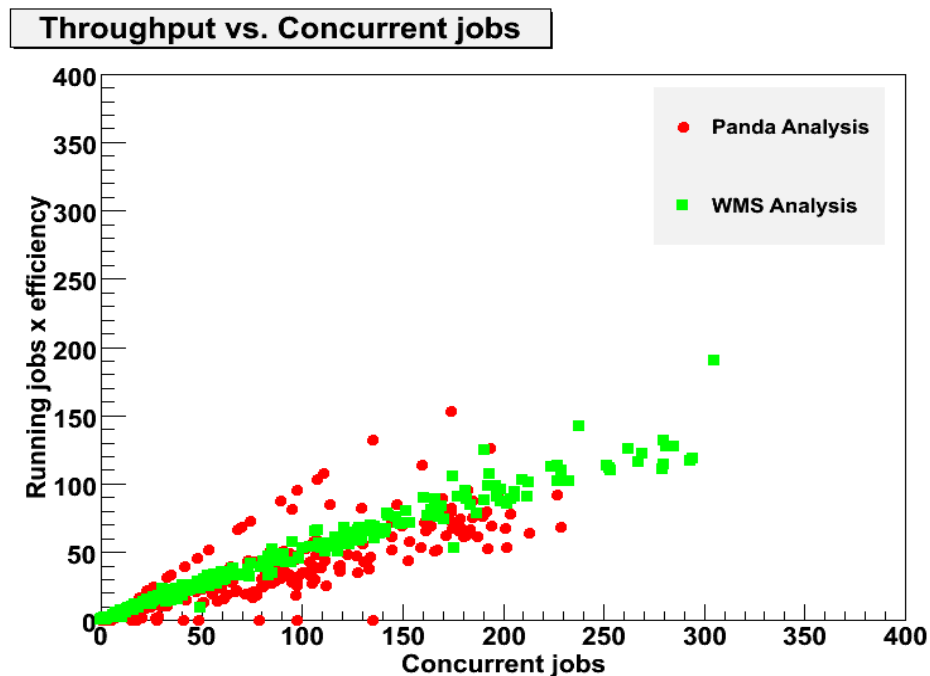


Job Efficiency vs. Concurrent jobs



“Throughputs”


- No clear saturation seen in the current scale
- More scattered for Panda jobs
 - Efficiency was time-dependent




SE-WN network bandwidth

Overall network traffic

6GB/s



Network throughput vs. Concurrent WMS jobs




● Good correlation between network traffic and number of WMS jobs

Storage Performance


- Data transfer not affected by user analysis
- Data delivery from Lyon Tier-1 went well

SE inbound network traffic (T1->T2)


500MB/s



DPM headnode CPU usage



DPM diskserver CPU usage




Conclusions

- No admin's intervention during STEP09
- Throughput scales up to ~300 jobs
- Storage worked:
 - 450MB/s reading from a disk server
 - No network saturation
- Some concerns
 - WN local disk (size and speed)
 - 8 cores (or more)/node in the near future
 - Interconnection between core switches


Backup slides

CPU utilization of WNs



Network traffic vs. all analysis jobs

Network throughput vs. Concurrent jobs



Network traffic for a Panda/WMS job

Network throughput per job vs. $N(\text{Panda-WMS})/N(\text{Panda+WMS})$

