Karlsruhe Institute of Technology

GridKa

CCRC'08 @ GridKa

Andreas Heiss

Steinbuch Centre for Computing

Forschungszentrum Karlsruhe
in der Helmholtz-Gemeinschaft

Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825

CCRC'08 @ GridKa

In general everything went smooth, business as usual.

Not much to report!

A few minor problems, though...

CCRC'08 @ GridKa

- **SRM timeout problem (begin of May)**
  - **High load on dCache PNFS server observed**
    - **Bulk deletions?**
    - **Atlas T1↔T1 replication tests / 'agressive' FTS parameters?**
  - **Caused many network connections on SRM machine in 'closewait' state.**

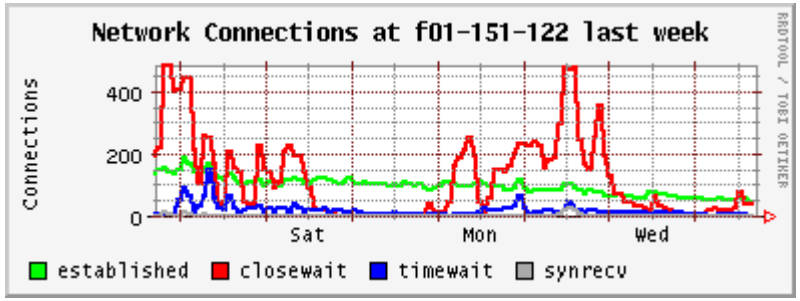

  - **potential bottleneck PNFS 'under observation'.**

CCRC'08 @ GridKa

- **FTS down for some hours between May 14th and 15th**
  - **accidential incomplete FTM installation implemented trigger in Oracle DB but function to be called by trigger was missing.**
  - **Shouldn't be a problem for Oracle (according to Oracle DBA) but FTA/FTS stopped working. (non-zero return code of SQL queuries?)**
  - **GridKa's mistake. Sorry!**

- **FTS setup (3 servers each running FTAs + FTS) works very well under high load (high # of transfers)**

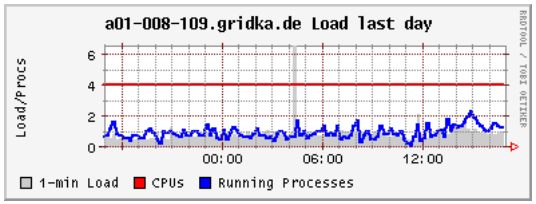

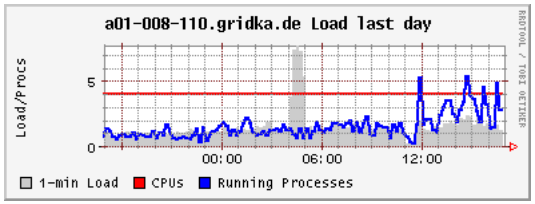

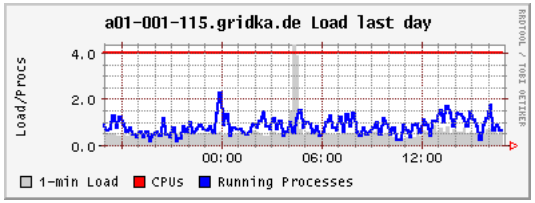

moderate system load and CPU utilisation on FTS/FTA machines.

CCRC'08 @ GridKa

- **Atlas observed LFC timeouts**
  - **Changed # of threads from 20 to 60**
    → problem gone.

- **Nothing to report about other gLite services (BDII, CEs, RBs)**
  → "business as usual" :-)

CCRC'08 @ GridKa – CMS transfers

- CMS admin responsible for GridKa decided to only import production data (no fake), i.e. all imports to tape.
- Period of good transfer rates interrupted
  - VO box credential problem
  - CERN export problem

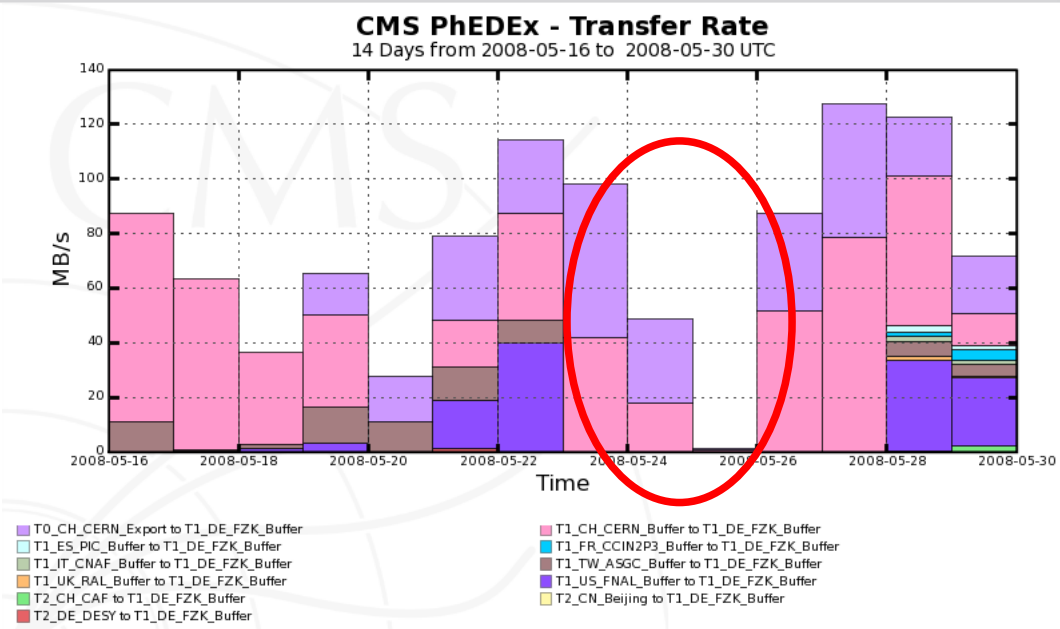

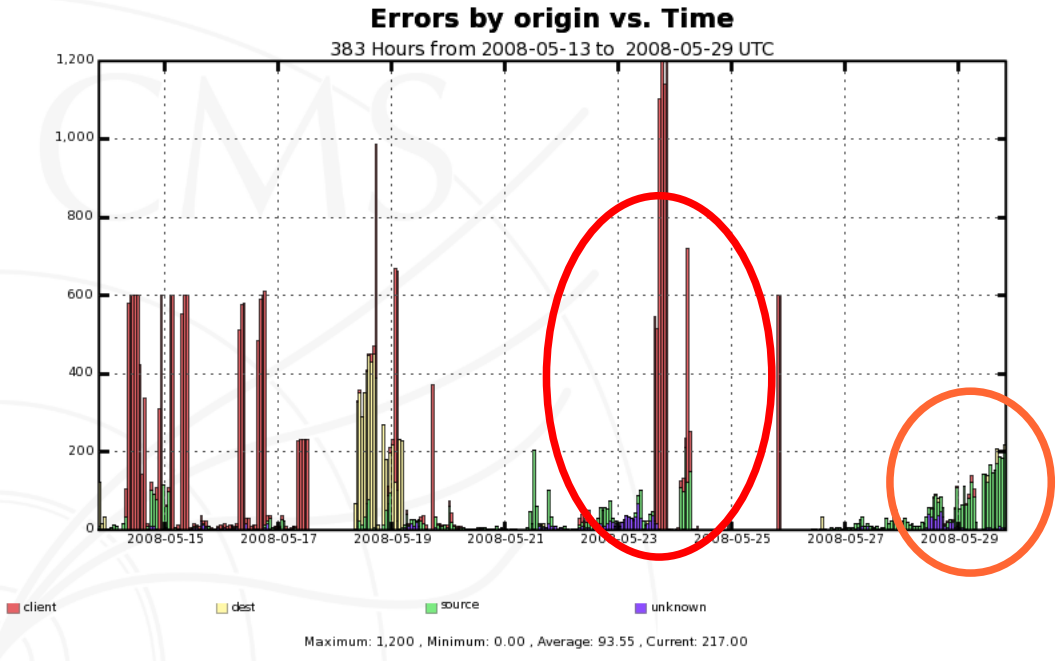

CCRC'08 @ GridKa - Skimming jobs 1/2

- CMS skimming jobs
  - saw low cputime/walltime ratio between 0.05 and 0.2
  - dCache delivered close to 1.5 GB/s sustained over several days. (some 2 GB/s theoretical limit for CMS with current setup.)
  - CMS share at GridKa is currently approx. 850 job slots in average.
  - Sometimes >2000 jobs were running.
  - read-ahead configuration (1MB buffer)

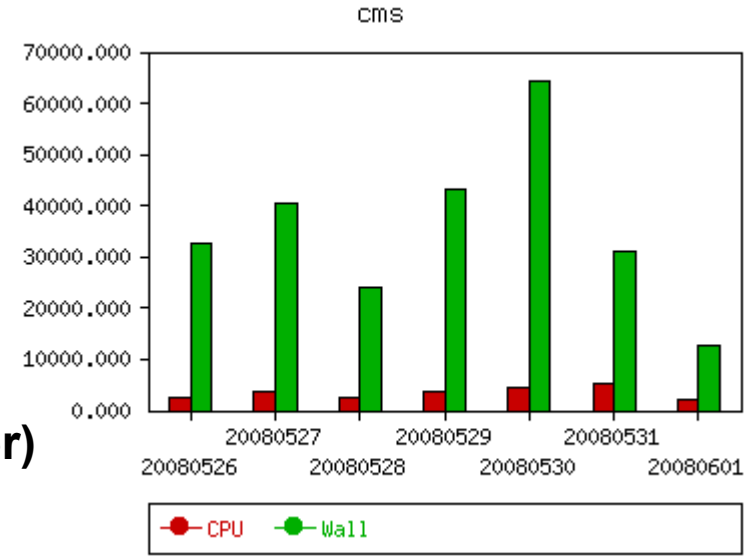

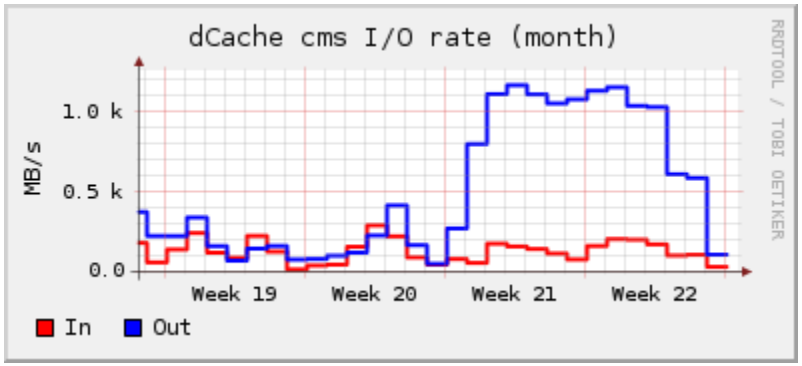

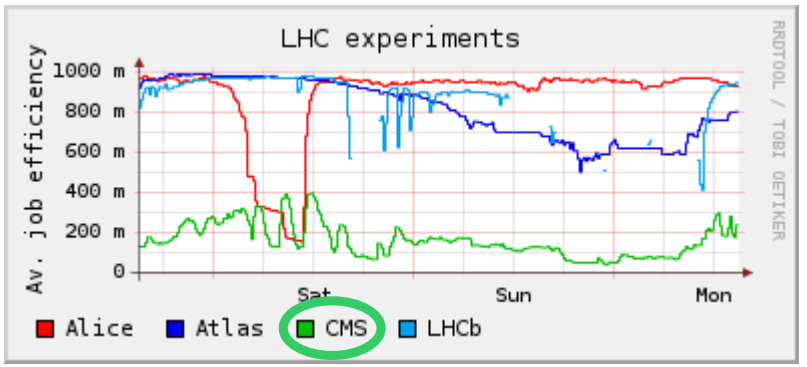

CCRC'08 @ GridKa - Skimming jobs 2/2

- skim jobs will always be I/O bound.
- cputime/walltime ratio will never be close to 1.

- But we can try to minimize the total wall time for a given set of skim jobs.
  - limit the number of concurrent jobs.
  - avoid running more than 1 or 2 skim jobs on same 8-core WN
  - optimize dCache server layout (no. of pools, pool size, distribution of pools on servers etc.).
  - ...
  - *Need to be able as a site to identify skim jobs?*

CCRC'08 @ GridKa

- **CCRC'08 @ GridKa summarized**
  - less problems than expected
  - gLite / 3D services peformed better than expected
  - dCache (1.8.15p2) pushed to the limits mainly by CMS and it was quite stable!
  - Approx. doubled CPUs and storage before CCRC
    - no scaling problems
  - Confident that we are ready for LHC data taking.
  - Still room for improvements, of course.

FNAL and BNL provided some slides last minute.

- BNL:
  - smooth excercise
  - Network cut CERN-BNL. Backup took over with initially only 1Gb/s (fixed within 2 hours)
- FNAL:
  - Some troubles with Phedex (stopped working, stressed PNFS)
  - CRC errors on transfers from CERN.
  - Transfer rate probably limited by Castor at CERN.

BNL and FNAL slides at

http://indico.cern.ch/contributionDisplay.py?contribId=6&confId=23563