


Worker node information on machine and job features

Stefan Roiser
IT/SDC

Description

- A way to provide information about the worker node and job constraints to the VO payload
- Provided via two environment variables pointing to directories containing files with info
  - $MACHINEFEATURES
    - Info on: machine power, # log/phys cores, shutdown time
  - $JOBFEATURES
    - Info/job on: cpu / wall time limit, mem limit, # cpus alloc, ...

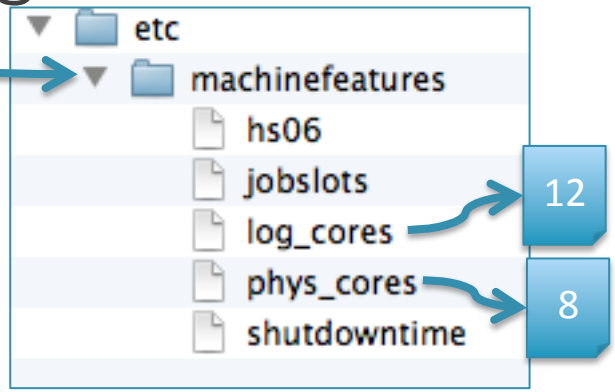

- Described in https://twiki.cern.ch/twiki/bin/view/LCG/WMTEGEnvironmentVariables

Status

- Implementation on sites
  - CERN (LSF)
    - Implemented and working in production
  - NIKHEF (Torque/Maui)
    - Implementation done
  - Other batch systems?
    - PBS, CONDOR, SLURM, SGE
      - (~ 45 % of all installations)
- LHCb used it in the context of VM testing
  - Using fake info on the VM
  - Other uses / tests?

Experiments

- LHCb interested in using it
  - "MC factory" uses info about end of queue to gracefully shutdown the MC job before end
  - Use for virtual machine "pre-notice for shutdown"
- Other use cases
  - Retrieve information about cores on the node in case of multicore submission

What would be needed,

- ... to make a concrete proposal
  - Principal interest from the experiments
    - LHCb is interested
  - Plan for implementation / batch system
  - Plan for grid wide deployment
  - Monitoring of the deployment status



BACKUP

/etc/machinefeatures

- ATLAS supports this initiative
  - https://twiki.cern.ch/twiki/bin/view/LCG/WMTEGEnvironmentVariables
  - Basically a way to publish important per-node information
- Panda pilot is instrumented to collect the information
  - Does not expose them to the server yet, but will soon
- After validation, ATLAS plans to use them for:
  - Monitoring (e.g. HS06 of the job slot)
  - Improve the knowledge of the pilot runtime environment (e.g. CPU and memory limits)
- ATLAS would like to ask sites to follow the Hepix WG recommendations
  - Followed up at the GDB