Summary of pre-GDB on Batch Systems

Michel Jouvin
LAL, Orsay
jouvin@lal.in2p3.fr

March 2014

Meeting Facts

- Agenda: https://indico.cern.ch/event/272785/
- Well attended
  - 25 person in the room
  - ~15 remote
  - Not only regular (pre)-GDB attendees
  - Only site representatives
- Most "well known experts" of batch systems present
  - CESGA (Grid Engine) apologized not being able to join
  - Covering Torque/MAUI, Grid Engine, LSF, HTCondor, SLURM
  - EGI (Peter) present

Meeting Goals

- Share experience about the different batch systems
  - First part of the meeting was a batch system review by sites with a concrete experience
- Identify strengths and weaknesses
  - Base features of a batch system
  - Multi-core job support
  - Handling of dynamic WNs
  - Proposal: build a summary table (Twiki page)
- Review missing bits for MW integration
  - Job submission and management
  - Accounting
  - Monitoring

Torque/MAUI

- Used by most sites, including T1s
  - Torque reasonably maintained but we are still running very old (unmaintained) versions
    - Still used for Moab, the commercial replacement for MAUI
    - No known showstopper for migration to recent versions but some validation/configuration work to be done (e.g. munge)
  - MAUI is a requirement and has been unmaintained for years
    - MAUI is feature rich when Torque has very basic scheduling capabilities
    - Running unmaintained SW is a potential concern, even though every security vulnerability has been fixed by the community
- PIC and NIKHEF reported a successful experience with Torque/MAUI at the 3K job slot scale
  - Not yet convinced of the benefit of moving to something else
  - No major problem so far with MAUI, take in charge its development remains an option…

Univa Grid Engine

- All the features of major batch systems
  - Fair share, back filling, multi-core job support…
  - Several fair share strategies
- Several big sites (T1s + large T2s) migrated to Grid Engine
  - UNIVA seems the only alive variant
    - Commercial variant with very good support: sites happy
    - Son of GE (open-source) still alive but not used as far as we know
  - Good feedback: presentations given by KIT and CCIN2P3
    - No scalability issues at the 15-20K job slot scale
  - Well integrated with the MW
    - CCIN2P3 using its site specific integration
- Multi-core job support without dedicated resources successfully experimented at KIT
  - Using dynamic reservations: 0.5 to 1% of CPU usage loss

LSF

- Robust, feature rich, commercial batch system
- Used successfully at CNAF and at several INFN sites
  - National license for INFN
  - CNAF: 1400 WNs, 18K job slots, 100K jobs/day
  - Also used at CERN but no report during the meeting
- Lots of tools developed by CNAF to help with LSF monitoring and to integrate it with the dynamic WN infrastructure (WNoDeS)
  - Local development to control packing of jobs on nodes
  - Development in progress for helping with multi-core job placement optimization
- No plan to move to something else
  - But technical feasibility of moving has been assessed recently

HTCondor

- RAL adopted it 6 months ago for its production cluster as a replacement for Torque/MAUI
  - Already used at most OSG sites
  - No major issue migrating: simple configuration, simple to administer, reliable
  - Scalability tests done at a very large scale
    - During test reached 30K simultaneous jobs without problems, 10K in prod
  - Dynamic cluster membership: no predefined list of WN
  - cgroups support may help to prevent resource exhaustion by jobs
- Integrated both with ARC CE and CREAM CE
  - RAL running 3 ARC and 3 CREAM
- Multi-core job support enabled: several features helping with it
  - See detailed presentation at the Multi-core job TF
- Already a couple of other sites in UK, with ARC CE

SLURM

- Modern, highly scalable, open source batch system
  - Easy to configure
  - Good multi-core job support
  - Good community support + commercial support
  - Successfully tested at the scale of 10K jobs, limit probably higher
- Widely adopted in Nordic countries
  - All Finnish scientific computing centers, Sweden moving towards
  - Also adopted by Swiss CSCS: an HPC center and a WLCG T2
- Working with both ARC CE and CREAM CE
  - EMI-3 required for APEL accounting
- Some weak points also…
  - Release quality, preference for a share file system, identical configuration file on every node at any time…

MW Support and Accounting

- MW support (almost) available for all 5 batch systems in EMI
  - Job submission and management for CREAM: BLAH
  - BDII publication: recent fixes released to fix all known issues
- CREAM Accounting: solutions available for the 5 batch systems
  - No problem with ARC accounting (JURA): no parser involved
  - HTCondor: currently based on a script converting to Torque format, need to be enhanced as a real parser.
    - No objection/difficulty to do it but no interest expressed when EMI-3 parsers where written
  - Need to create a small team of experts to assess what is currently published: call for volunteers (contact J. Gordon)
    - Experiments complain about some discrepancies with their own measurement: need to check consistent processing across batch sys
    - Need to check failed jobs handling

Multi-Core Jobs

- Most of the work happening in the Ops Coord TF dedicated to multi-core job deployment
- Pragmatic work to evaluate technical possibilities of each implementation and find appropriate solutions
  - First stage is dedicated workshops on each implementation
  - HTCondor and UGE already reviewed, others to come soon
- Main issue is entropy (mix of job types)
  - Multi-core jobs increase the need for an efficient back filling strategy to avoid wasting resources
  - But back filling requires short single core jobs advertised as such: not currently the case in WLCG
    - Despite many short jobs, e.g. in Atlas
  - Need to discuss more with VOs this need for a mix of job type

Conclusions

- Very good discussions based on actual experiences
  - A lot of valuable information
- A summary table will be produced in Twiki to help sites wanted to review their batch system choice
  - Weaknesses, not only strengths/features…
  - Scale at which problems where observed
  - Contact of reference sites
- A lot of work in progress, in particular for multi-core job support
  - Follow-up in Ops Coordination
  - Summary at a future GDB, probably end of Spring
- Some topics not discussed due to lack of time
  - Dynamic WN handling
  - ARC CE vs. CREAM CE comparison based on recent experiences