T2 discussion

Some points for discussion

1) Resource sharing

2) Manpower availability

3) monthly operation meetings

4) Workflows on Pitz Daint

5) WLCG/CSCS monitoring comparison (discussion if time allows)

6) anything else ?

Resource Sharing

- Starting point: 40:40:20 sharing integrated of the pledged year
  (fractions would need to be reviewed to reflect the present VOs involvement)

Because of the ATLAS computing model (nucleus+satellites) that is not optimal. On top of that ATLAS needs flat efficiency, not the case for CMS/LHCb.

- Improve the present SLURM configuration ? from previous discussions it seems complex to achieve

- Possible alternative: dynamically allocated multicore nodes as implemented by CSCS for the T0 R&D project:
  - Slurm partition with adequate cap on number of nodes per LHC user (paid for + X% ?)
  - Higher priority of incoming jobs (for LHC users: ATLAS, CMS, LHCb) over other users
  - Once using a node, node is exclusive to the LHC user (single and multicore jobs). User segregation guaranteed.

Resource Sharing

- Equalise pledge to capacity: the 10% margin is probably not needed anymore

- Opportunistic usage of Pitz Daint ? (max fraction XX%)
  HPCs (and standard clusters) are used opportunistically by all (ATLAS/CMS/LHCb)
  collaborations to get to >> 100%

Items miscellanea

- ATLAS:
  - ATLAS own job micro-priority (- - nice) => top priority now
  - VO relative share (latest ticket closed, metrics not settled)
  - ATLAS ~flat delivery (+/- 20% from due core count) => now seldom a nucleus site
  - ARC metrics (monitoring and alarms) - since the dismissal of the ganglia monitoring which was available to us (few years)
  - ATLAS HammerCloud status (monitoring and alarms)

- General:
  - Timely dCache maintenance and upgrades to avoid disruptive upgrades. Inform VOs for plans and progress
  - Storage accounting implementation (WLCG / EGI)

People availability

- 2.5 FTEs effective availability ?

- Long delays in replying to operation issues: Is there any way to improve/help the situation ?
  - rotating a contact person ?
  (was less of an issue with Phoenix because several people had access to the machine)

- RT-tickets are sometimes closed without asking feedback from the VO-representative. Having feedback on the implemented changes can prevent mis-understandings/delays

Monthly operation meeting

- some VOs would prefer to have a CHIPP member to chair the meeting
  (probably Mauro TBD)

- chair to compile minutes and list of actions, publish on the agenda and request approval/objections by email, and actively follow up on the actions on a regular basis

- Proposed standard agenda:
  - VO reports
  - T2 site reports (CSCS, UNIBE)
  - T3 site reports (PSI, UNIGE)
  - follow up on actions from f2f/previous month

Workflows

- ATLAS is considering to have the HPC only for some workflows (as already done with the other ATLAS HPC sites)
- CMS wants to have HPC for all workflows
- LHCb is already using HPC only for MC production

- ATLAS is moving to a federated use of resources (CSCS + Bern) in Switzerland Storage will transition first (going in the direction of reducing the pressure on the dCache storage (or reducing the size of dCache)

Other topics (if time allows):

- WLCG/CSCS monitoring comparison (discussion if time allows)

- Anything else ?