HK T2 Technical session

By Roger Wong

11 Apr 2018

Agenda

- IPv6 upgrade
- Site monitoring
- DPM permission issue
- DPM OS upgrade
- Calculation of HS06
- Jobs requested for extremely large memory

IPv6 upgrade

- The deadline for upgrading all T2 with IPv6 is end of 2018
- Will start to plan for this upgrade shortly
- Any experience to share on adding IPv6 to your cluster?

DPM permission issues

- Files created by atlas/Role=pilot cannot be deleted by DDM
- Temporary fixed by running the following commands regularly
  - dpns-chgrp -R 105 /dpm/atlas.cuhk.edu.hk/home/atlas/atlasscratchdisk/rucio/panda

  - while
  - dpns-listgrpmap | grep 105
  - 105 atlas/Role=production

- What is the correct permission setting for various folders in DPM?

DPM OS upgrade

- The DPM in HK-LCG2 is running on SLC6
- Scientific Linux 6 would be end-of-life by Nov 2020
- Does any site could share the experience on upgrading the OS of DPM?

Calculation of HS06 (1)

- SPEC2006 was bought and HS06 benchmark has been executed on the worker nodes (dual E5-2680v4)
- HS06 is 10.60161 per core for 32-bit code
- HS06 is 11.82321 per core for 64-bit code (this is set in the APEL configuration)

- From a HS06 benchmark web site
  - https://w3.hepix.org/benchmarking/sl6-x86_64-gcc44.html
  - A dual E5-2680v4 should result in HS06 of 678 on a SLC6 with 64-bit code, i.e. HS06 is 12.10714 per core for 64-bit code

Calculation of HS06 (2)

- There are 1,008 cores in HK-LCG2 and each with HS06 of 11.8
- The HS06 for the whole cluster should be 11,894
- However, the CPU HS06 shows in the statistics sometime even over 50,000

- Is there anything wrong?

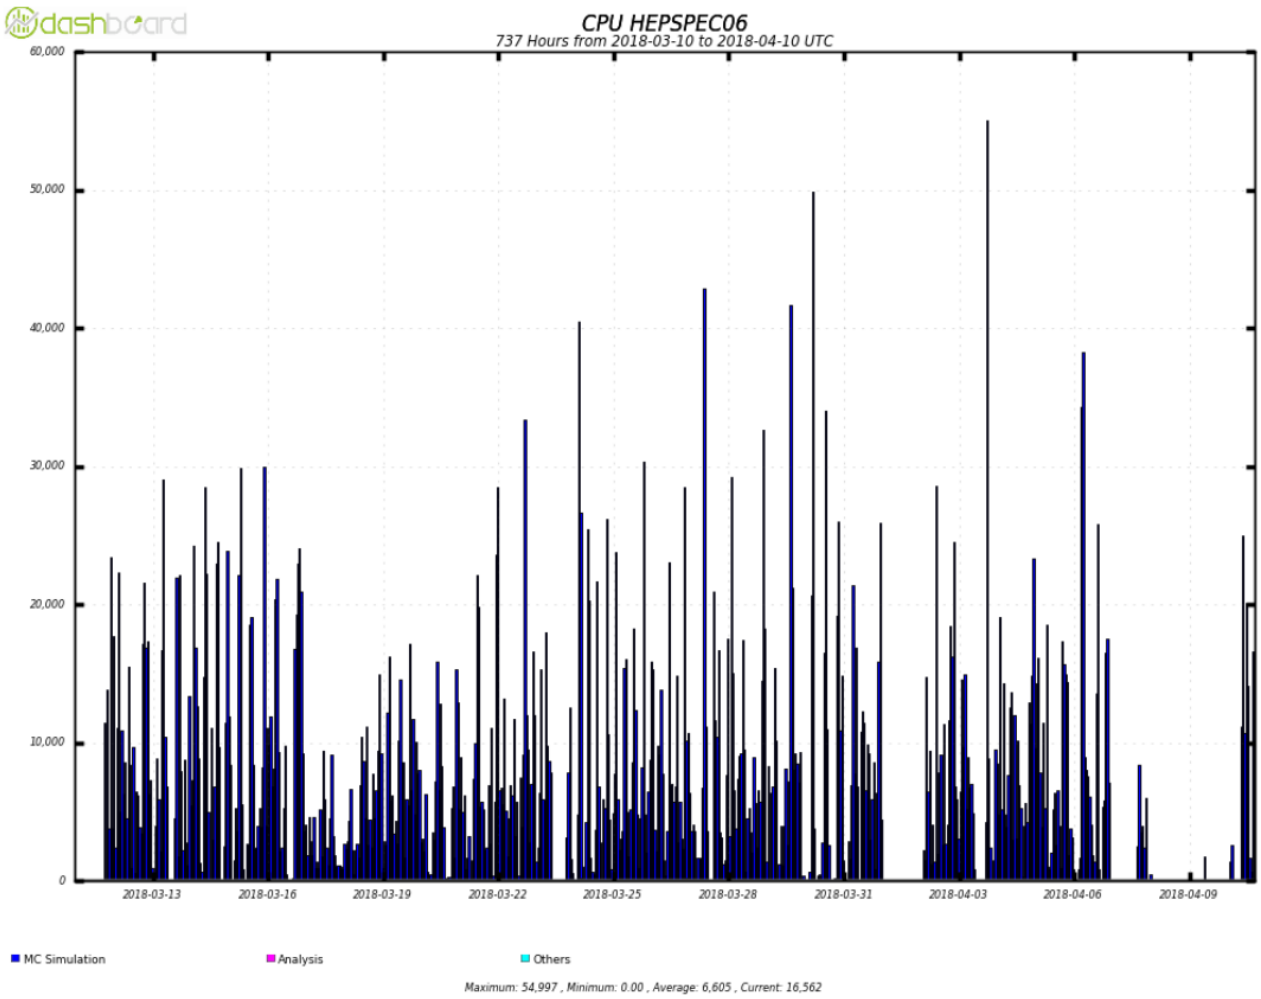

Jobs requested for extremely large memory

- Some grid jobs are requesting extremely large memory

  Job 405054.000 defines the following attributes:

  DiskUsage = 25
  RequestDisk = DiskUsage
  RequestMemory = 536870912

- It seems that they are requesting 500TB+ memory and they are idle on the condor queue forever
- How do such kind of grid jobs be generated? How could we avoid such kind of grid jobs?
- Should we just kill these kind of grid jobs from the condor queue? Or even other grid jobs that have been idle on condor queue for a long time?

Thank you!