HK T2 site status

By Roger Wong

11 Apr 2018

Agenda

- HK-LCG2 site overview
- Cluster configurations
- Cluster status
- Future plan

HK-LCG2 site overview

- Background
  - An ATLAS Tier 2 under FR-cloud
  - Endorsed in ICB meeting at CERN in Dec 2017
  - Still processing the MoU and the pledge
- Cluster
  - Initially 1,008 cores (HS06: 11,894.4) and 436TB disk storage
  - The hardware of another 654TB disk storage has been ready and will be deployed to SE soon
- OS and Software
  - CentOS 7 for ARC CE + HTCondor, worker nodes and Frontier squid
  - Scientific Linux 6 for DPM
- Support
  - Six IT staff in central IT department manages the ATLAS cluster, a Central HPC platform and some smaller computing clusters

Cluster hardware configurations

- 18x worker nodes
  - CPU: Dual E5-2680v4
  - Memory: 176GB
  - 10Gbase-T NIC
  - 3x 2TB RAID 0 HDD
  - Will adjust after the pledge is confirmed
- 3x DPM disk nodes (436TB)
  - 12x 8TB HDD per disk chassis
  - 2x disk chassis per disk node
  - 10Gbase-T NIC
- 4x DPM disk nodes (654TB) – planned to be deployed in summer
  - 12x 10TB disk chassis + 12x 8TB disk chassis per disk node
  - 10Gbase-T NIC
- 2x Frontier Squid nodes

Cluster configurations

- CE
  - HK-LCG2_ARC: Single core production queue
  - HK-LCG2_MCORE_ARC: Multicore production queue (8 cores/job)
  - ANALY_HK-LCG2_ARC: Analysis queue
- SE
  - 380TB for data disk
  - 20TB for scratch disk

Cluster status

- Still learning how to monitor the healthiness of the cluster and fixing the errors in the cluster
- Below are some charts showing the healthiness of the cluster

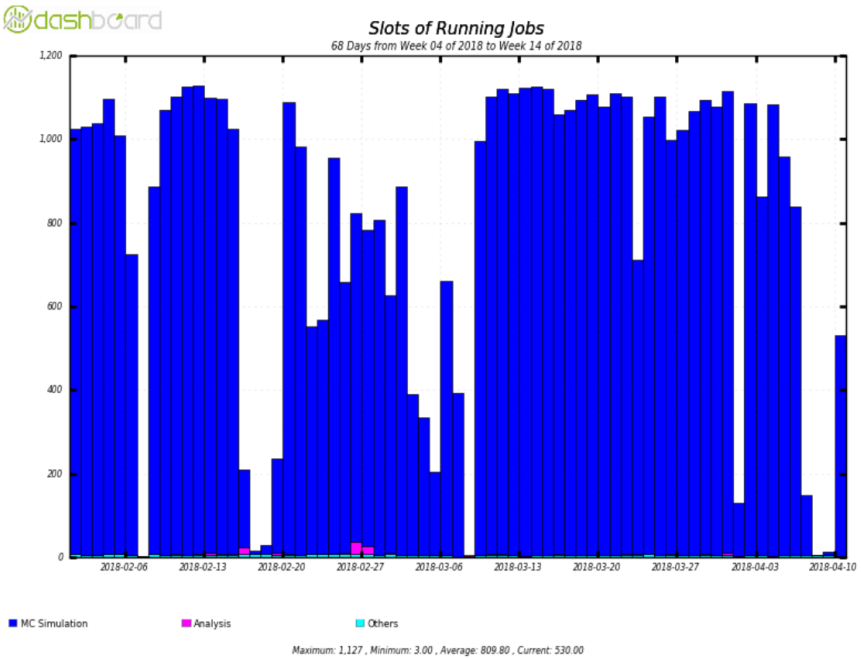

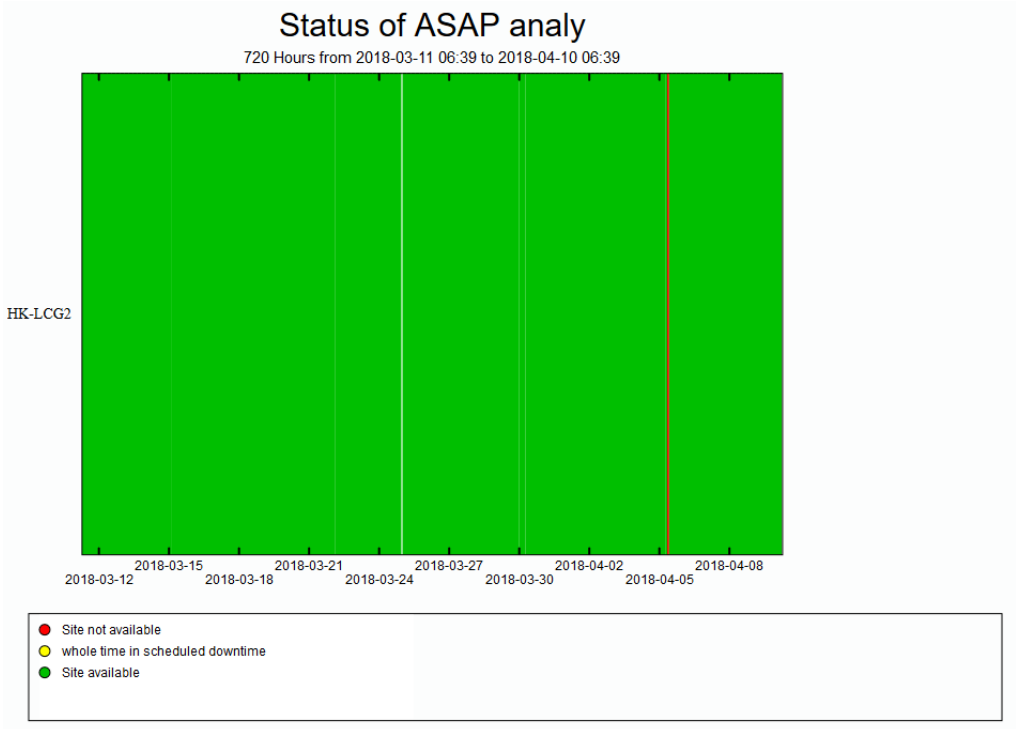

Site monitoring

- At this moment, we normally use the following tools for checking the healthiness of the ATLAS cluster
  - Slots of running jobs under ATLAS Dashboard
  - Metric, Transfer plots and Details under ATLAS DDM Dashboard
  - ASAP metric
  - SAM Visualization
  - Status of functional tests under HammerCloud
  - Searching for job errors in BigPanda and try to fix them
  - EGI ARGO

- Any other suggestion?

Future plan

- Upgrade SE from 400TB+ to 1PB+
  - In 2018 summer
- Optimize the network by LHCONE peering
  - Move the ATLAS cluster to a subnet for LHCONE routing
    - Now in a subnet mixed with other HPC
  - Configure PerfSONAR node to facilitate the tuning
    - http://ps1.itsc.cuhk.edu.hk/toolkit/
- Upgrade to IPv6
  - By 2018

Thank you!