LHCb on RHEA and T-Systems

Andrea Valassi (IT-DI-LCG)

HNSciCloud pilot phase open session, 14th June 2018

LHCb use cases for HNSciCloud tests

- During HNSciCloud tests, **LHCb only uses CPU resources**
  - not the storage, not the GPUs
  - this is likely to be the case for any future production usage of clouds
- During HNSciCloud tests, **LHCb only runs Monte Carlo Simulation jobs**
  - event generation (no input data) + detector simulation (no reconstruction)
  - CPU intensive, output data uploaded back at the end
- *MC Simulation is the largest fraction of CPU used by LHCb on the Grid*
  - will be even more so in LHC Run3 (data reconstruction done online)
- *MC Simulation requires no data management operations*
  - user analysis would instead require a large operational overhead

LHCb interaction with HNSciCloud resources

- Requirement: **manage RHEA and TSYS as two independent sites**
  - separate monitoring for the two sites (see plots on following slides)
  - be able to start/stop/vary job submission number/type separately
  - we did not request to see RHEA sub-sites independently
    - compromise between enough flexibility and not too much overhead

- Job submission via **HTCondor CEs**
  - two separate CEs to submit batch jobs to RHEA and TSYS
  - LHCb often operates cloud resources using VAC and Vcycle instead

- **HTCondor CEs are managed for LHCb by CERN IT-CM**
  - decouple LHCb infrastructure from low-level cloud technicalities
  - decouple LHCb human operations from interaction with cloud providers

LHCb on RHEA and TSYS – 2½ month overview

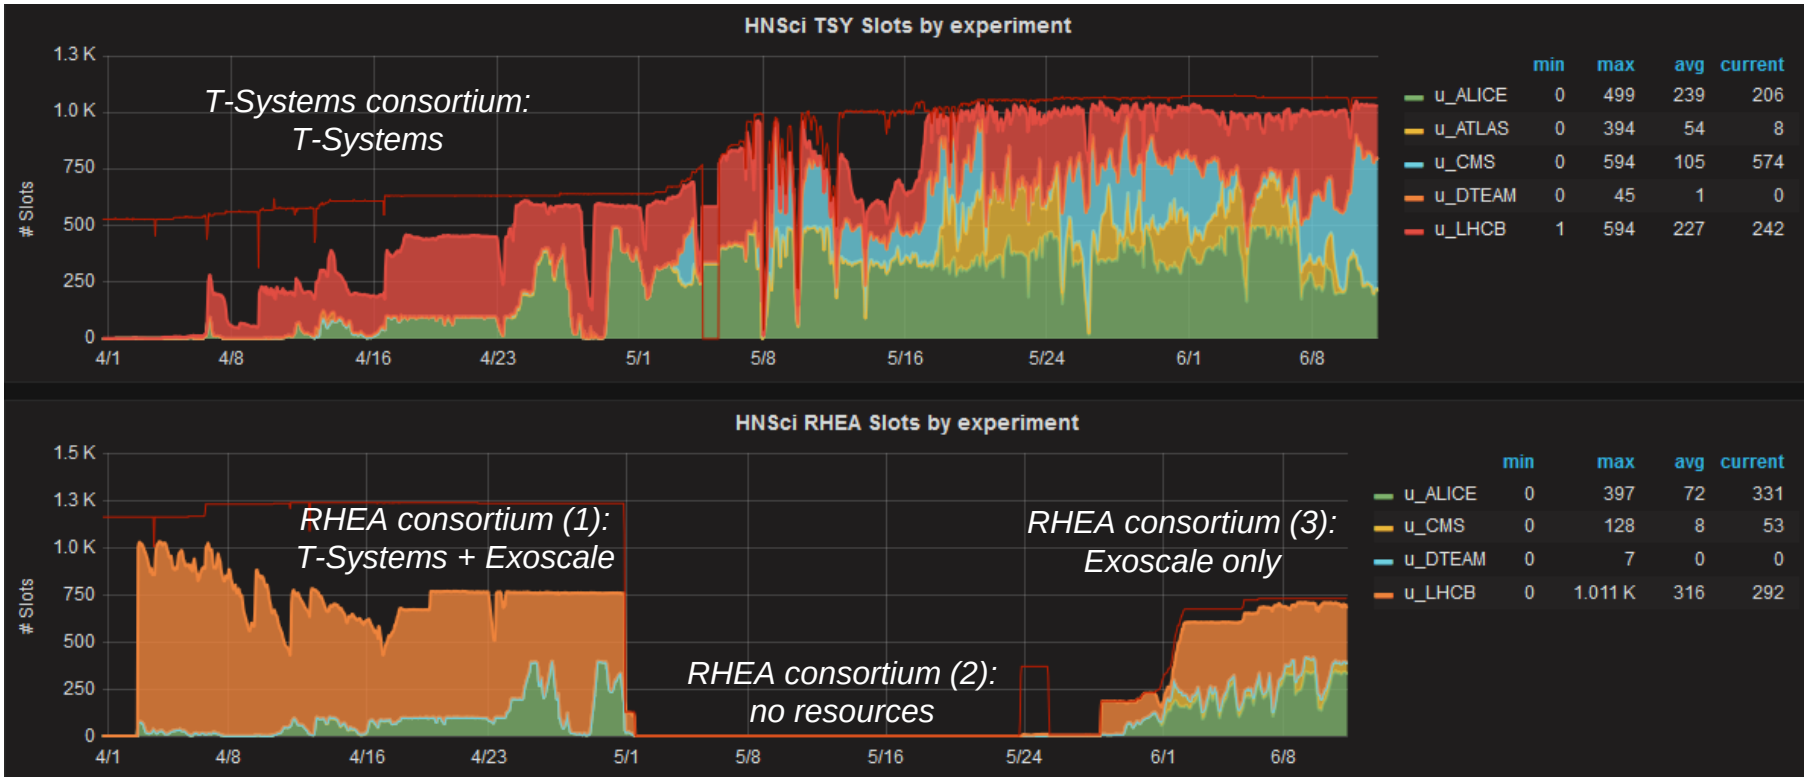

- TSYS: shared with ALICE, ATLAS, CMS – slower start, 200-400 LHCb jobs running
- RHEA: shared almost exclusively with ALICE – three distinct phases
  - 1st month: 600-1000 LHCb jobs running on T-Systems and Exoscale
  - 2nd ~month: no resources
  - more recently: ~400 LHCb jobs running on Exoscale only

LHCb on RHEA and TSYS – more details

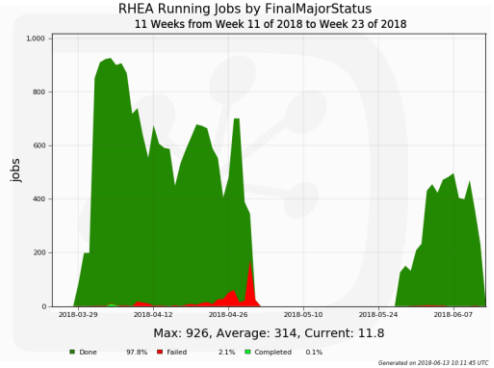

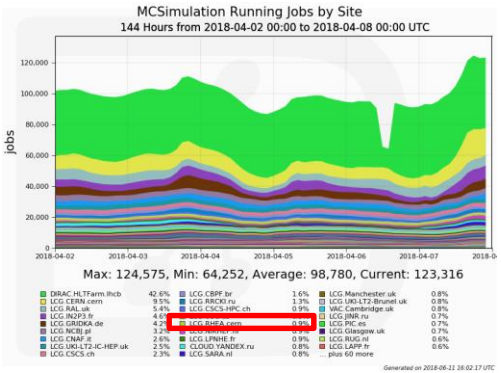

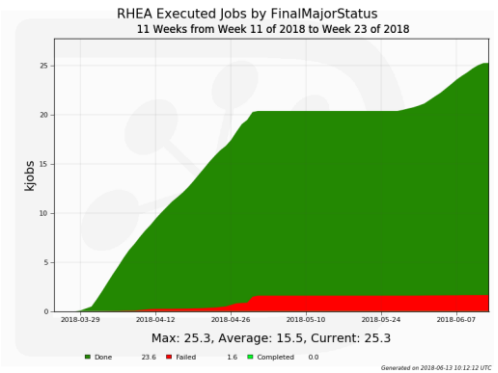

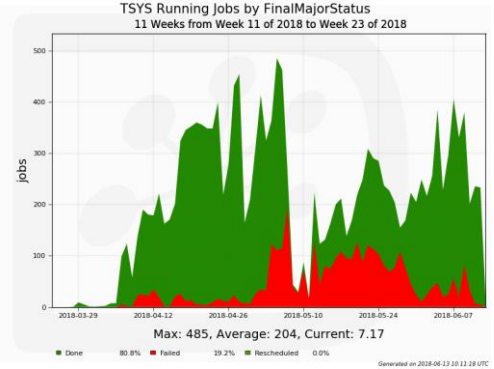

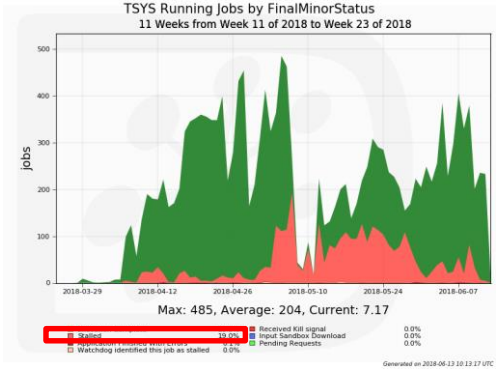

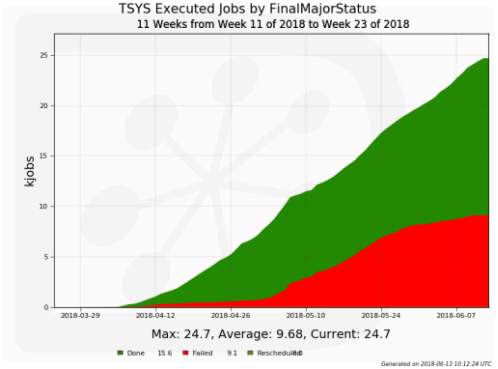

- **RHEA consortium: excellent success rate (network with public IP)**
  - 20.6k done jobs, 1.6k failed jobs (mostly killed when Exoscale was removed)
  - *Job rate of Exoscale alone (~400) still lower than for Exoscale + T-Systems (~600-1000)*
    - Was one of top 15 sites for MC production (0.9% of total running MC jobs in LHCb)
- **TSYS consortium: very high rate of jobs failed due to lost heartbeat (network with SNAT)**
  - 12.9k done jobs, 8.5k failed jobs (stalled, i.e. lost heartbeat)
  - job rate (~300-400) generally smaller than RHEA, as TSYS also used by ATLAS and CMS