# Consolidated batch queue for LHC Experiments

CERN - WLCG

# Consolidated multi-cloud

- HTCondor-CE as single entry point for submission to multiple clouds
- Experiment workloads submitted with some of the technical details managed
- Differences in clouds: public IP in Exoscale vs NAT at TSY
- Workloads for WLCG experiments managed on behalf of several WLCG sites
- Compute intensive (as opposed to data intensive) still most common workload

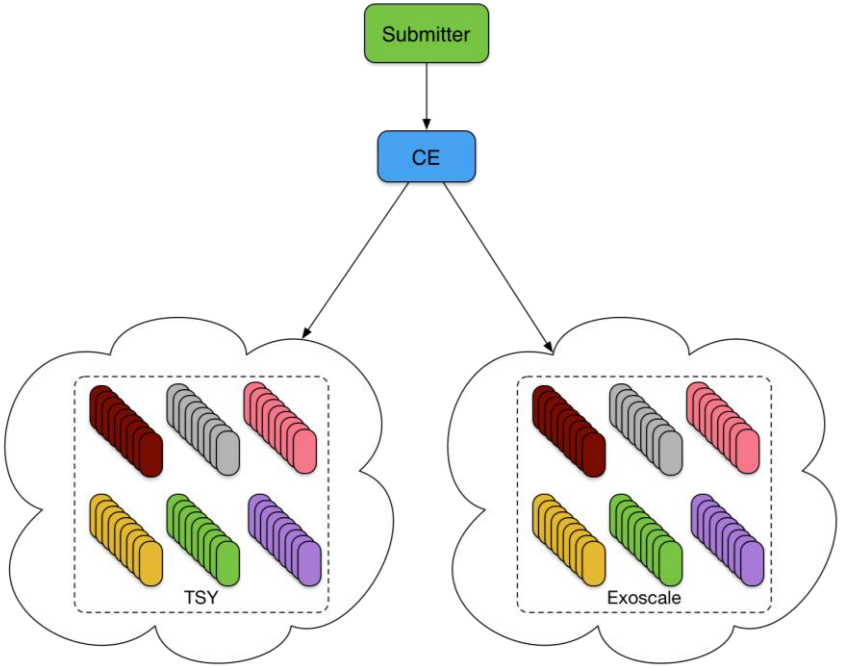

# Feedback from CMS for the last 6 months

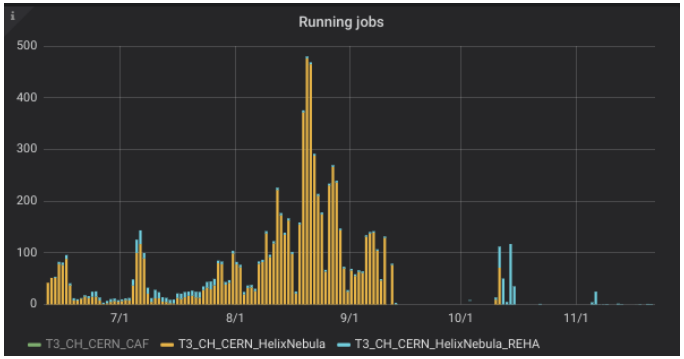


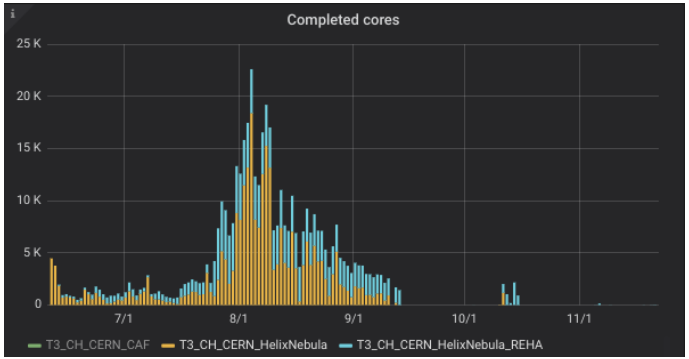


- The HNSciCloud resources integration and validation process was smooth and successful
  - Initial validation done with T-Systems and then included RHEA.
  - Resource integration as CERN overflow, via Unified
- CMS ran more on RHEA than T-Systems
  - This was due to fairshare being applied across both clouds, whereas demand especially for multicore slots, was not even across clouds
  - This is reflected in the plots… most of the resources used by CMS are from T-Systems
- Did not experience inefficiencies and/or failures
- Utilization affected by close of resources by T-Systems then loss of pressure by CMS

# LHCb on RHEA and T-Systems – normalized CPU

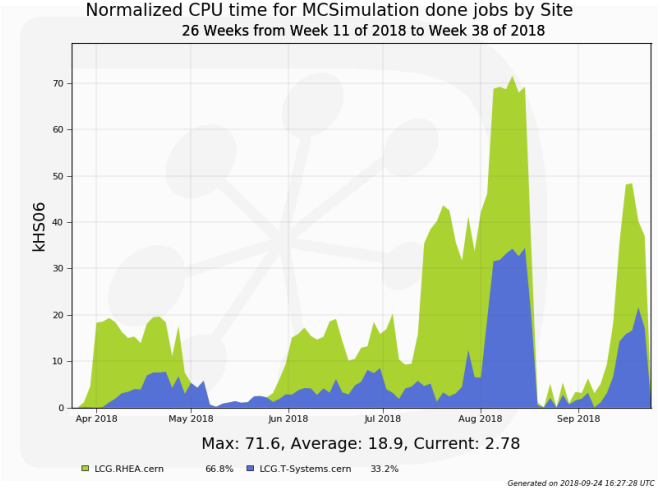


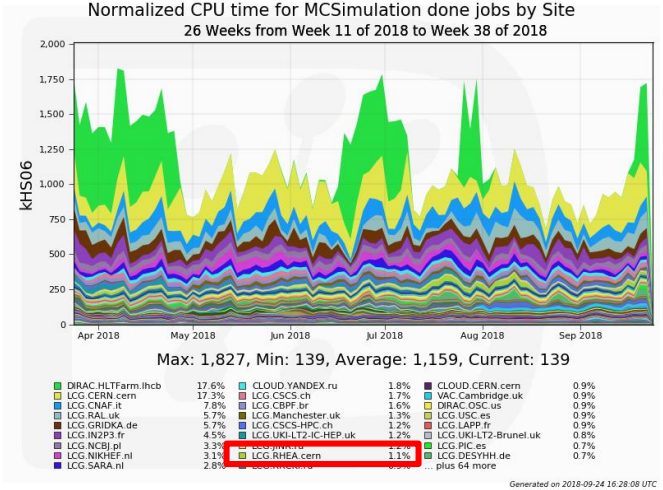


- Normalized CPU: ~6kHS06 on TSYS, ~12kHS06 on RHEA for 6 months
- Shares of MCSimulation on the Grid: ~0.5% on TSYS, ~1.1% on RHEA
  - **In total, an addition of ~1.6% of CPU resources for MC over 6 months**

# HNSciCloud usage by ALICE

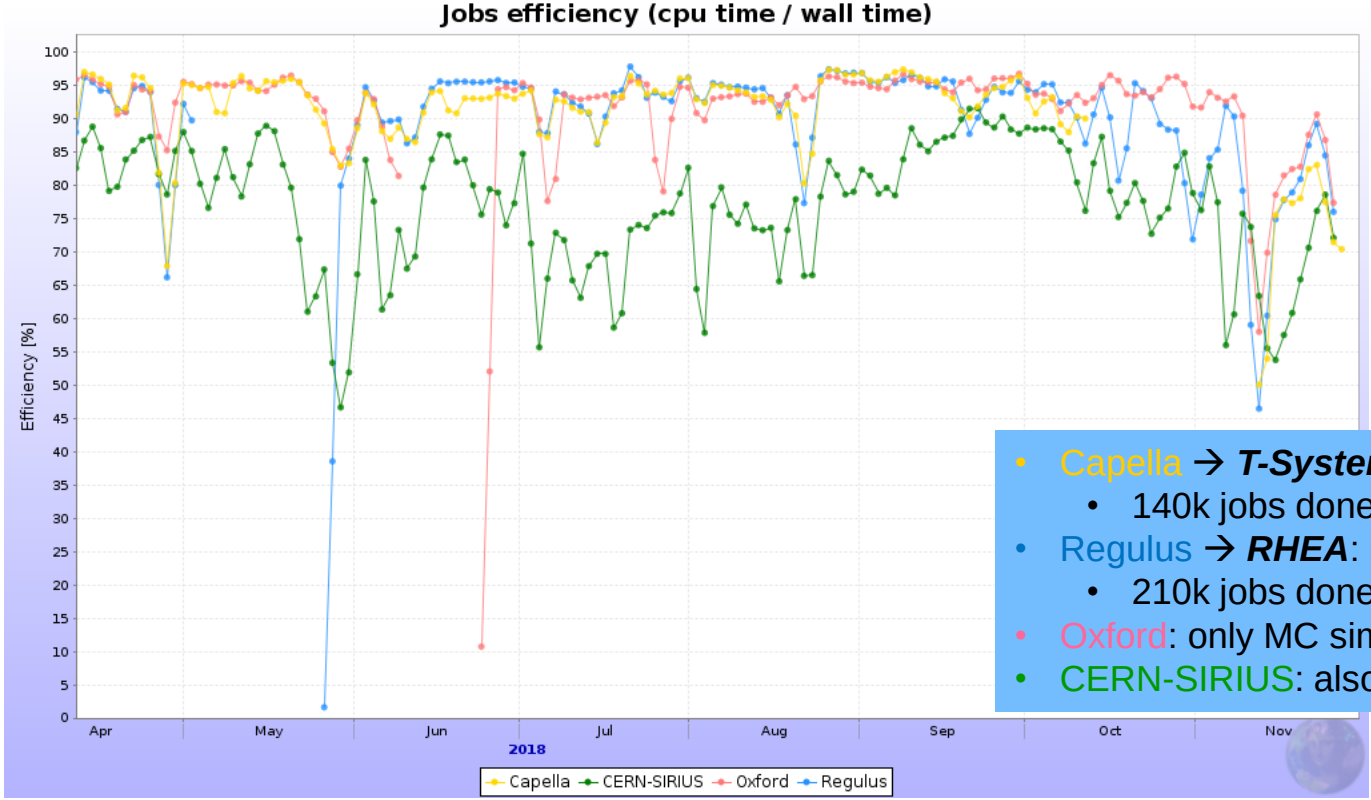


- Capella → ***T-Systems***: mostly MC simulation
  - 140k jobs done so far
- Regulus → ***RHEA***: mostly MC simulation
  - 210k jobs done so far
- Oxford: only MC simulation
- CERN-SIRIUS: also reco and lots of analysis

# HNSciCloud usage by ATLAS

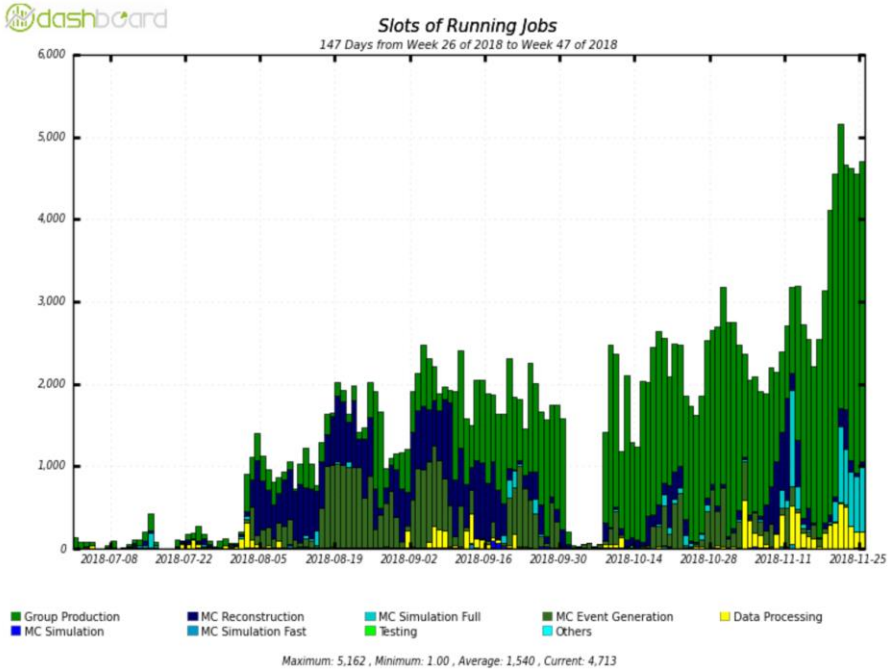


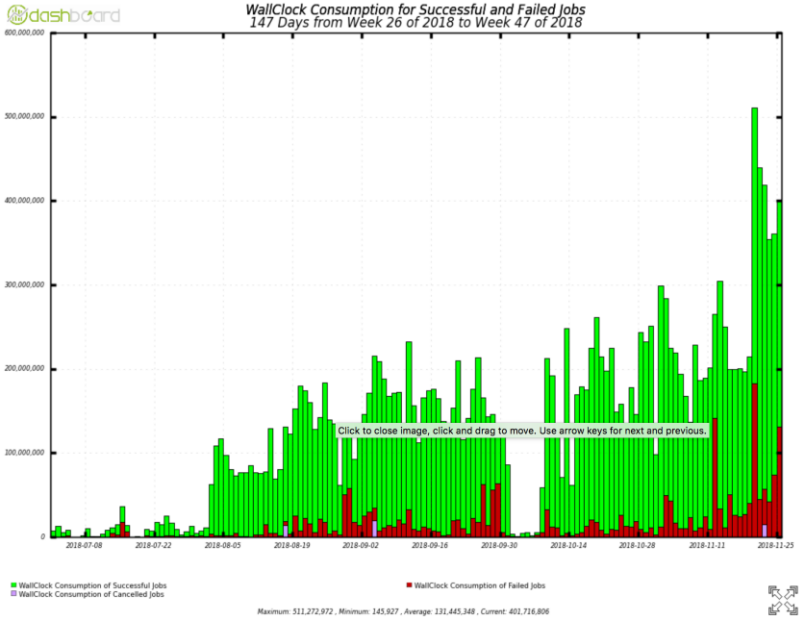

# Conclusions

- Consolidating cloud resources behind single entry point [using HTCondor] is an easy model with potential operational overhead savings
- Demand continues to be variable from experiments; fairshare helps utilization of reserved instances
- Would prefer not to manage NAT
- Abstraction layer (terraform now, k8s later) crucial for multicloud