



Running Jobs Everywhere: an Overview of Batch Services at CERN

Jérôme Belleman (on behalf of the CERN Batch Team)

Batch Services: High Throughput

- HTCondor capacity: now $> 100k$ cores
- Almost $2\times$ LSF capacity ever
- LSF reduced to 20k cores before Spring 2018
- LSF service ends once LHC Experiments move
- End of LHC Run 2 at the latest
- Training
- XBatch with Helix Nebula Science Cloud:
 - Tactical engagement with Oracle Bare Metal Cloud
 - Docker universe

cluster: cernprod -

Batch Services: Monitoring

CURRENT JOB STATUS

Total
124397

Running
73757

Idle
46742

Held
3898

Recent Jobs Started

Recent Jobs Submitted

CPU Utilization

RESOURCE UTILIZATION

Memory Utilization

Disk Utilization

Current CPU Efficiency

80%

Current Memory Efficiency

58.60%

Current Disk Efficiency

91.9%

RUNNING JOBS

Running Jobs by User

Running Jobs by Runtime

BEER, Batch on EOS Extra Resources

- Andrea Sciabà, Andrey Kiryanov's talks (HEPiX Spring 2017)
- Ongoing tests
- Disk server slots as batch slots, no lower-SLA
- Segregating batch functions from disk node:
 - Using cgroups/systemd for the HTCondor service
 - Docker universe (testing CPU set)
 - Dockerised CVMFS

Batch Services: High Performance

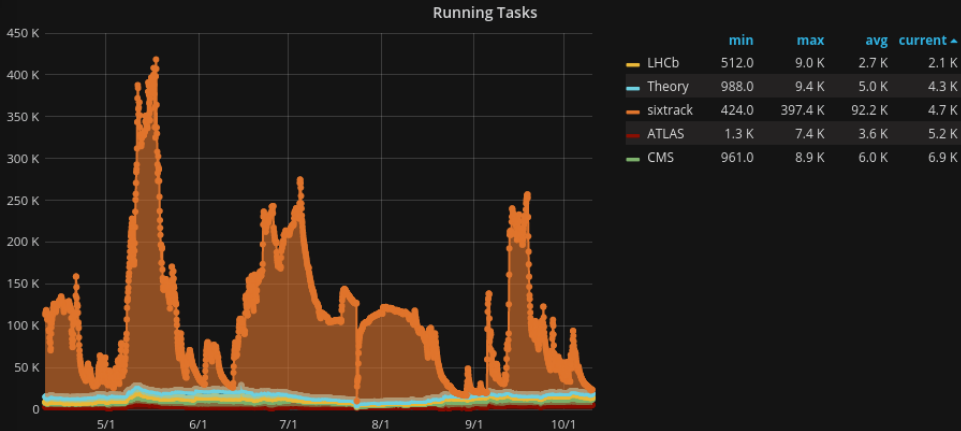
- MPI, shared memory across nodes, InfiniBand
- SLURM batch cluster for HPC
- Backfill via HTCondor/SLURM interface

Advice welcome on this!

Volunteer Computing: LHC@Home

- Accelerator physics simulations (SixTrack): native BOINC app, now also Android
- ATLAS (also Singularity), CMS, LHCb, Theory simulations under CernVM and VirtualBox
- Job management back-end integrated with Condor: wide range of (low-IO/high CPU) apps
- Working with the BOINC community to evolve the BOINC software

Volunteer Computing Consolidation: LHC@Home





home.cern