CERN

LSF – HTCondor migration

Ben Jones IT-CM-IS

Agenda

- Batch Service
- Why exit LSF?
- What is HTCondor?
- Benefits of HTCondor
- Timescale
- How can IT help?
- Questions

Batch Service

- Service used for both grid and “local” submission, with HPC on the way
- Local means open to all CERN users, kerberos, shared filesystem, managed submission nodes
- ~75K cores in LSF pools
- ~26K cores in HTCondor
  - Till now just grid
- ~650K jobs per day

Why exit LSF?

- Proprietary product
- Limits to number of nodes (>5K not advisable)
- Doesn’t scale very well past 180K jobs
- Slow queries, submission
  - All goes through one master
- Security model limits flexibility of submission hosts
- Product seems to be diverging from our use case
  - Scaling into machines, rather than jobs + nodes

What is

- Open Source batch system developed at the CHTC at the University of Wisconsin
- “High Throughput Computing”
- Long history in HEP and elsewhere (including previously at CERN)
- Used extensively in OSG, and things like the CMS global pool (160K++ cores)
- System of symmetric matching of job requests to resources using ClassAds of job requirements and machine resources

Benefits: scalability

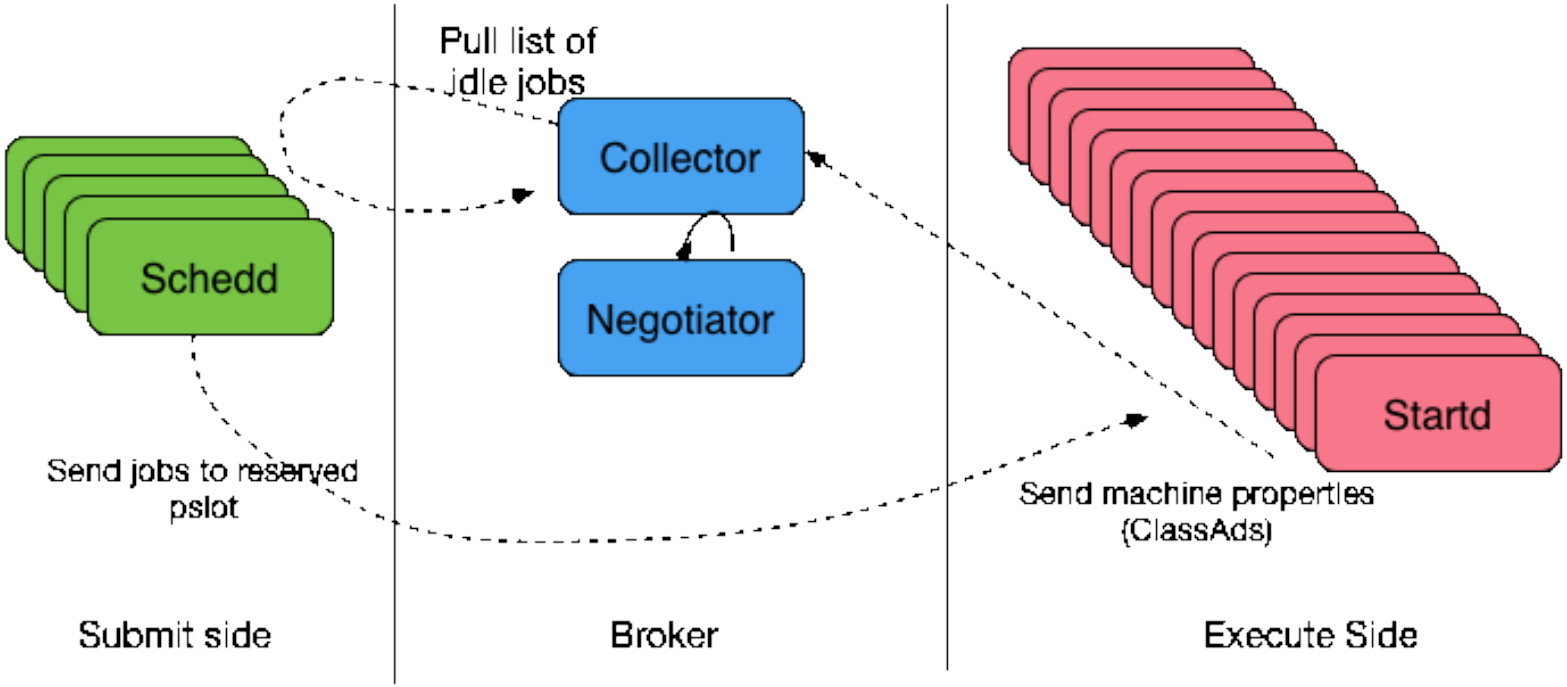

Benefits: Flexibility

- Extra “Universes”
  - Docker, Parallel as well as Vanilla
- DAGs
  - Job dependencies between different submit files
- Condor-G to submit to other systems
  - For example, condor submission to boinc
- Flexible configuration allows routes to clouds, or specific resources, or HPC
- HTCondor can be a single frontend to have jobs run in many different ways on different systems
- Cgroups to ensure jobs can coexist without stepping on each others’ resources

Out with the old…

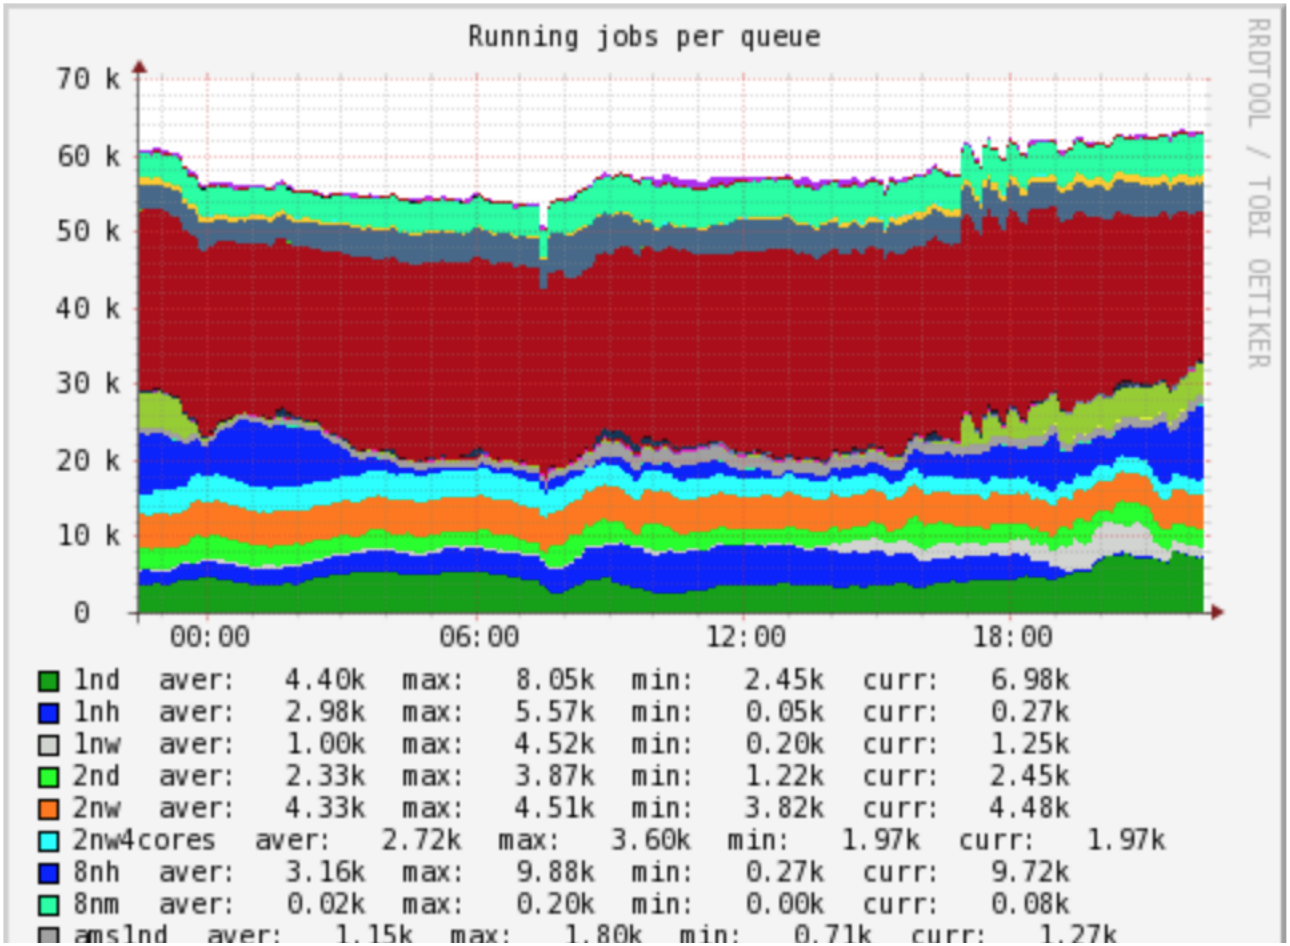

Benefits: community

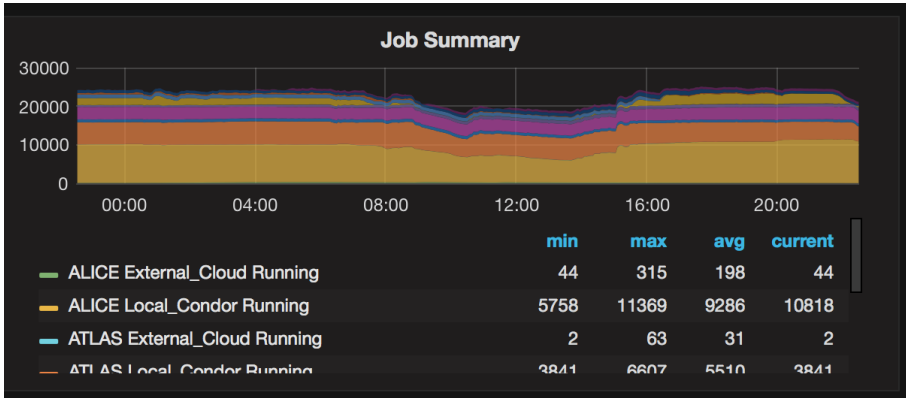

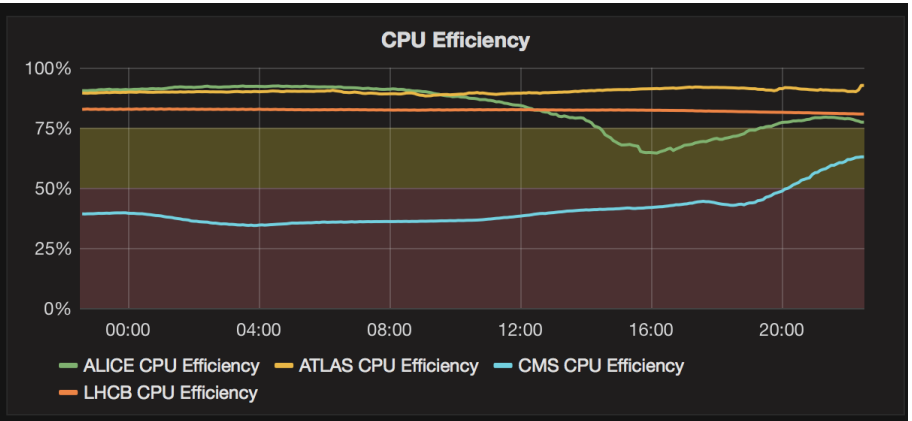

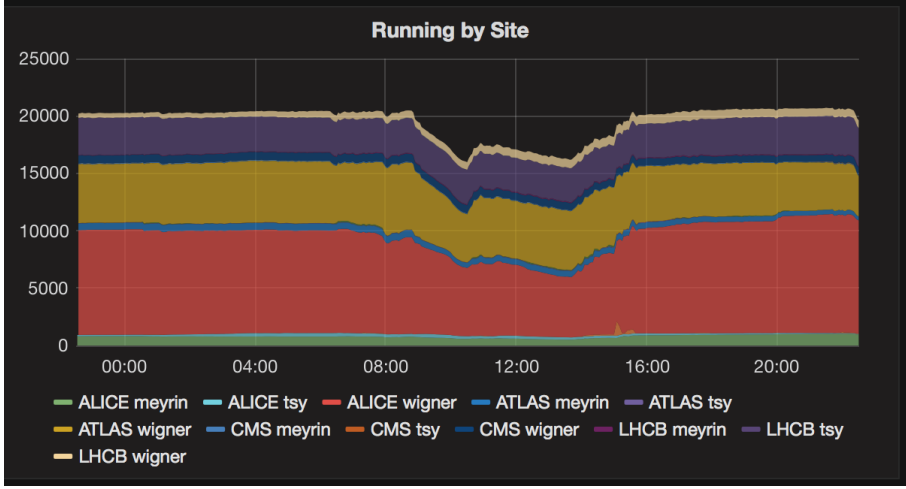

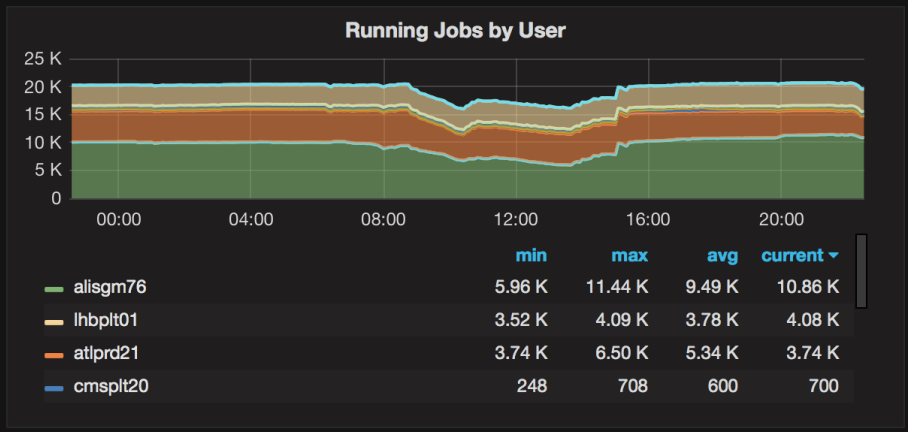

Timescale

- Grid is prod since November
- Local required work with upstream for kerberos renewal, now no technical issues
- IBM support till end of 2017
- IT support for LSF till end of Run 2

How can IT help?

- Some help available with migration
  - We can help advise on submission scripts etc
- Migration can be easy for most use cases
- User training classes foreseen
- Documentation for local use iterating with early pilot users
- batch-operations@cern.ch / SNOW to batch team / contact us directly

Questions?