

IT Batch Service Outlook

Ben Jones

HTCondor Migration

- Local submissions now possible to HTCondor
 - Docs at cern.ch/batchdocs
 - Tutorials and Training being developed
 - Contact us for help migrating
- HTCondor pool is now ~25% of capacity
 - This will increase as local users are moved over
 - All of Grid submissions will be moved this year
- Committed to supporting LSF till end of run
 - Advantages to moving sooner (proper multicore support, new tools, newer hardware)

HTCondor Service

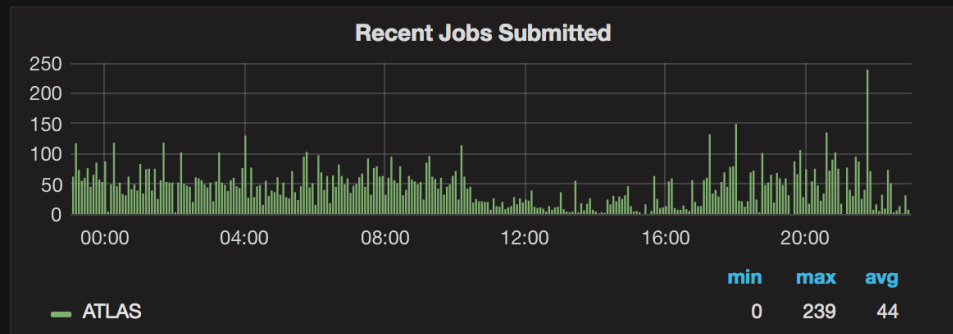
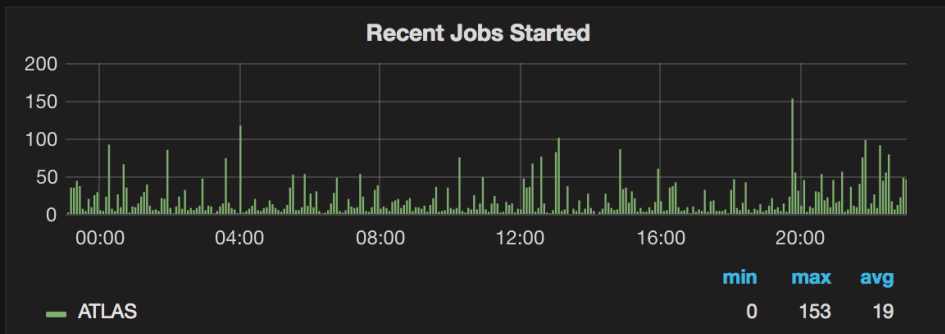
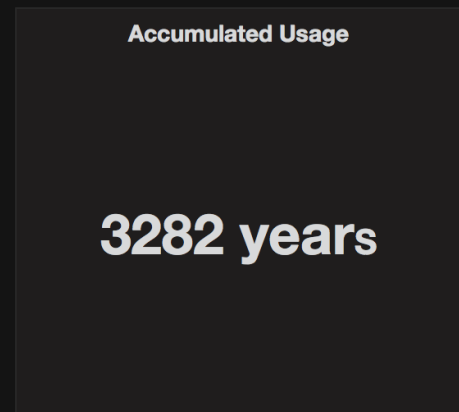
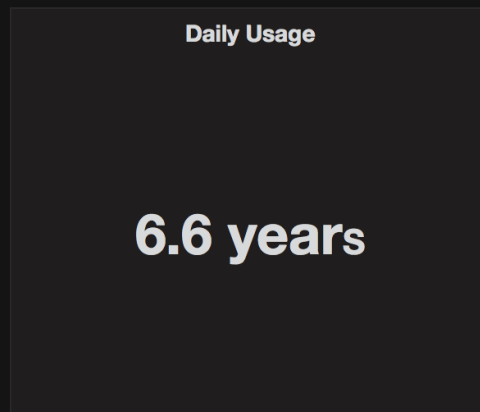

- New grafana dashboards
 - <https://batch-carbon.cern.ch/grafana/>
 - Views for experiments, individual users, overall cluster
- Draining for mcore
 - Will increase mcore pool as needed
- CC7 support
 - Local submissions
 - Grid coming (and some container / cgroup tech)
- CGroups for all jobs
 - Soft limits for memory (standard 2gb / core)

Experiment dashboard

Experiment Batch Details ☆ ↻ 📄 ⚙️ ◀ Zoom Out ▶ 🕒 Last 24 hours Refresh ever... 🔄

cluster: cernprod Experiment: ATLAS

USAGE & PRIORITY



Job Summary

min	max	avg
0	153	19

CPU Efficiency


min	max	avg
0	239	44



User dashboard

Navigation bar with icons for home, User Batch Jobs, favorites, refresh, save, and settings. Right side includes navigation arrows, "Zoom Out", "Last 24 hours", and "Refresh ever..."

cluster: cernprod | user: atlprd21





Daily Usage

43.5 week

Experiment User Priority (Lower is Better)

Metric	Min	Avg	Max	Current
ATLAS	45.15 Mil	80.30 Mil	151.10 Mil	45.50 Mil



CPU Efficiency

00%

Memory Request Efficiency

110%

Wastetime (avg)

1.353

Hardware & Flavours

- Finally some uniformity:
 - LSF: 16 HT cores ~2gb/core
 - HTCondor 8 HT cores ~2gb / core
 - ...ATLAS T0 notwithstanding (no HT, and some 12 core)
- New hardware now with 40 HT cores & 128gb
 - 10 core VMs (8 core not possible due to NUMA)
 - Obviously slightly more memory

Questions?