

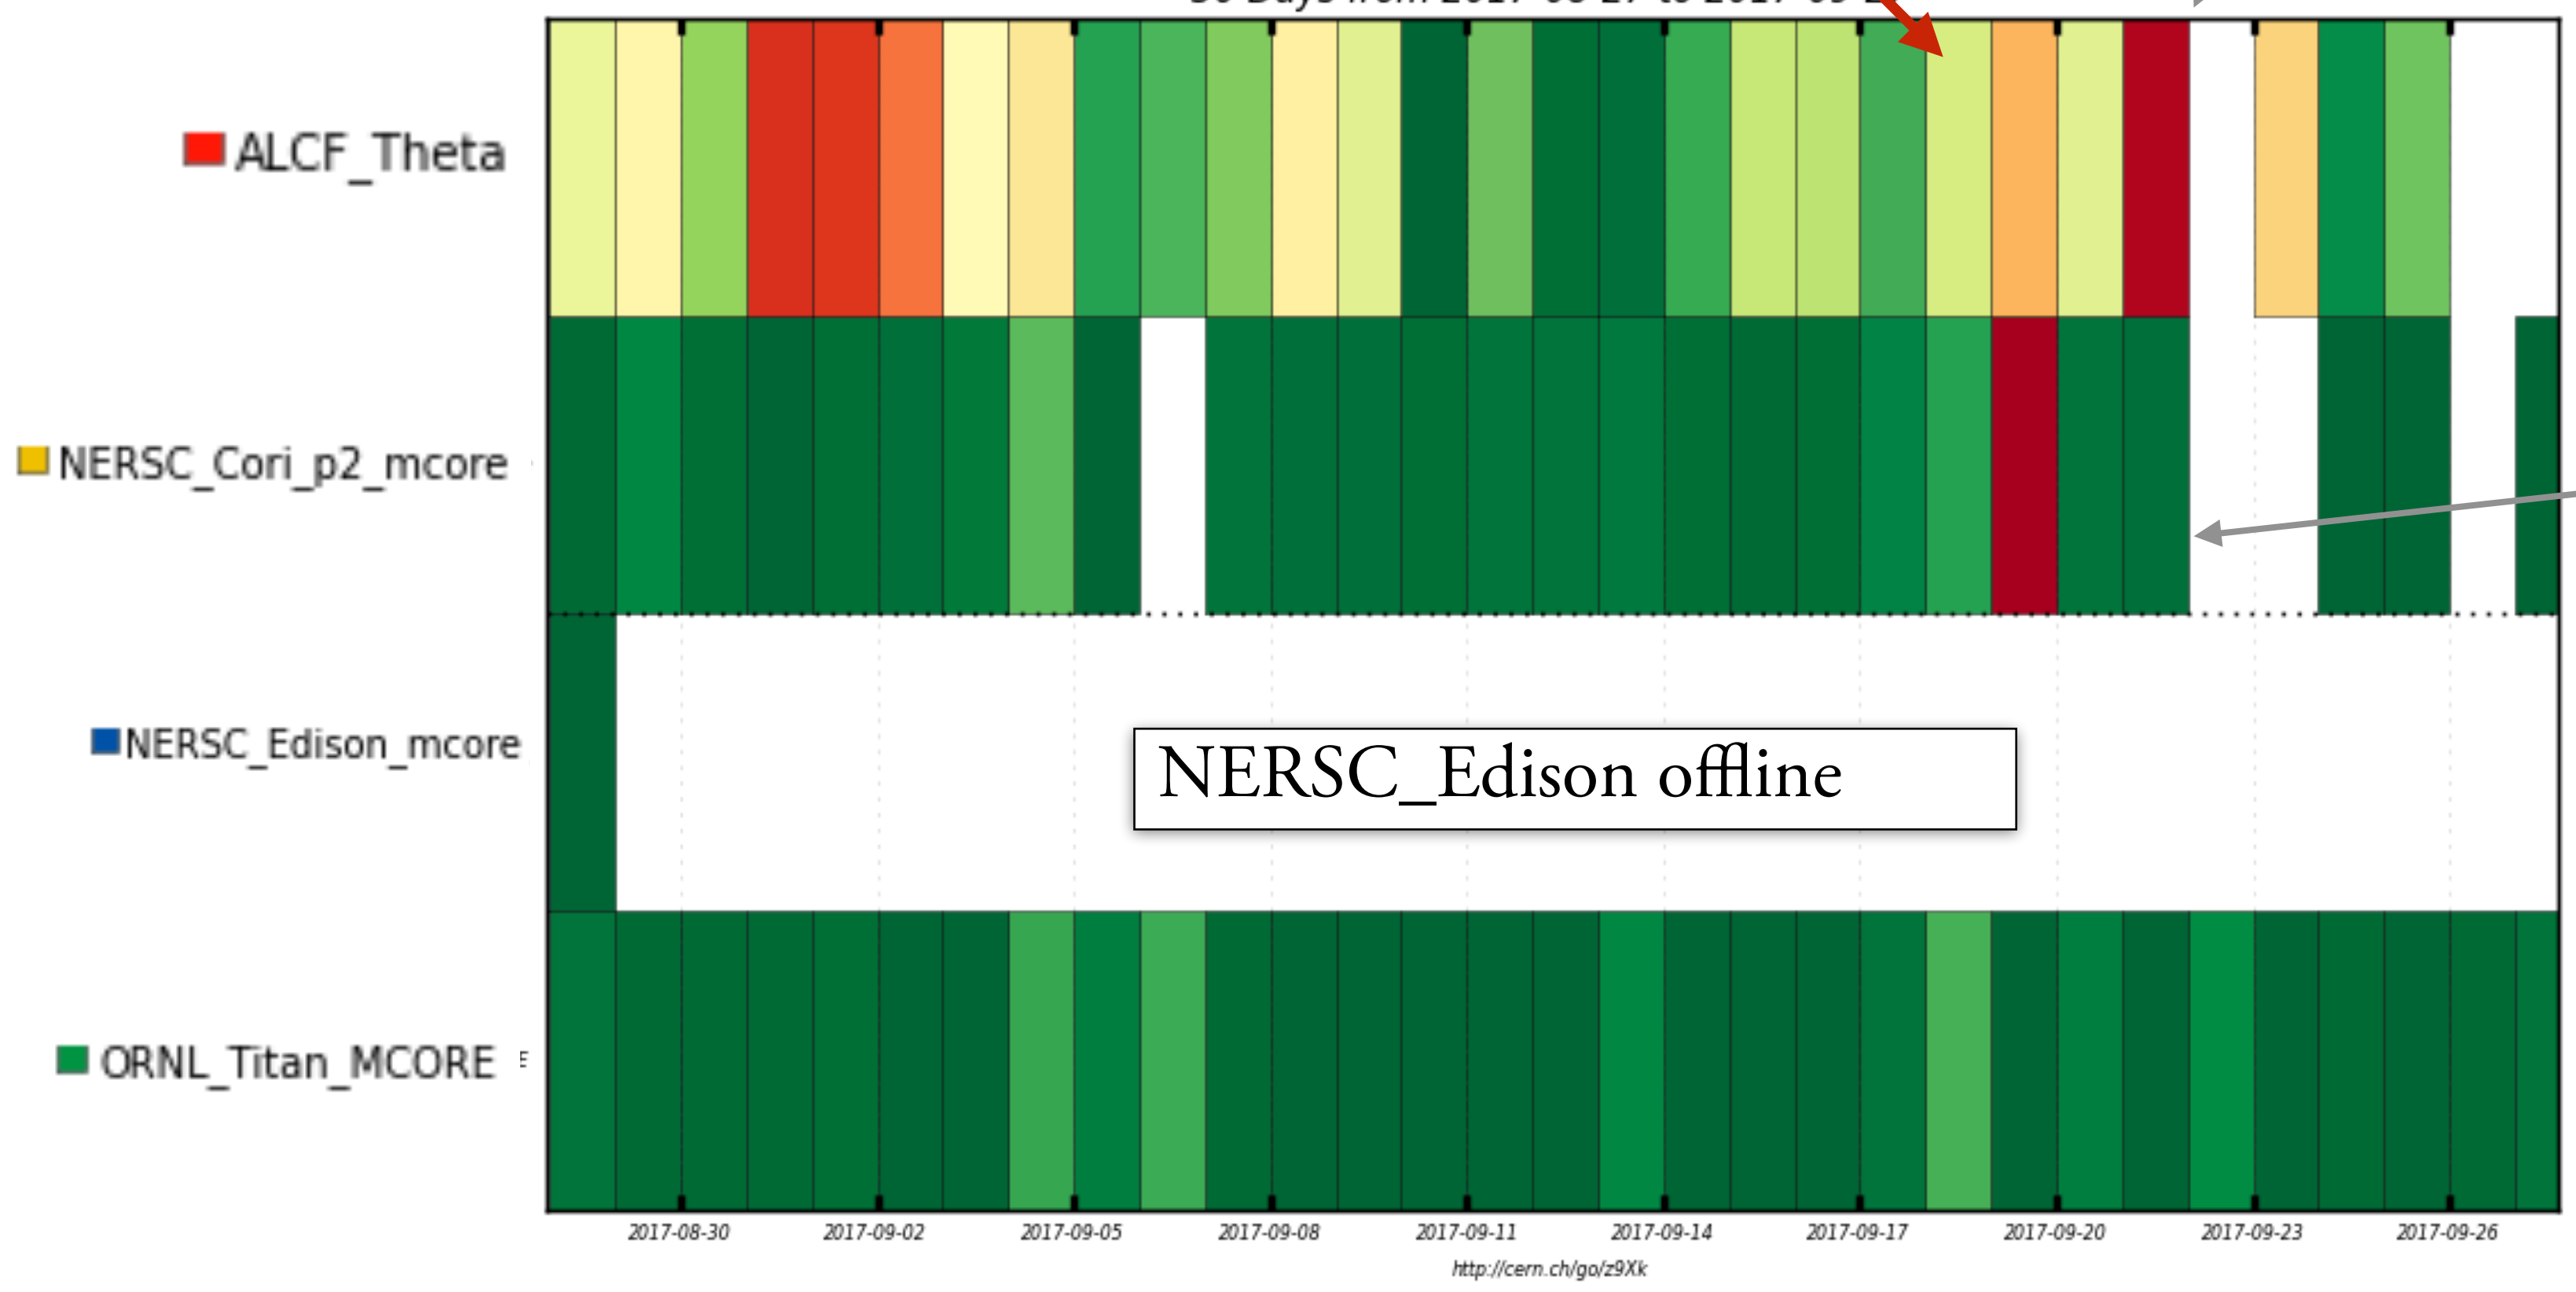
US DOE HPC Usage Summary

ALCF_Theta converted to Globus Online

ALCF_Theta understanding GO and debugging



Efficiency based on success/all accomplished jobs
30 Days from 2017-08-27 to 2017-09-27




ALCF_Theta offline

NERSC_Cori_p2_mcore ran out of allocation

NERSC_Edison offline


ORNL_Titan chugging along.



US DOE HPC Usage Summary



Wall Clock consumption Good Jobs in seconds
30 Days from 2017-08-27 to 2017-09-27




NERSC_Cori_p2_mcore

NERSC_Edison_mcore



NEvents Processed in MEvents (Million Events)
30 Days from 2017-08-27 to 2017-09-27




ORNL_Titan_MCORE

ALCF_Theta

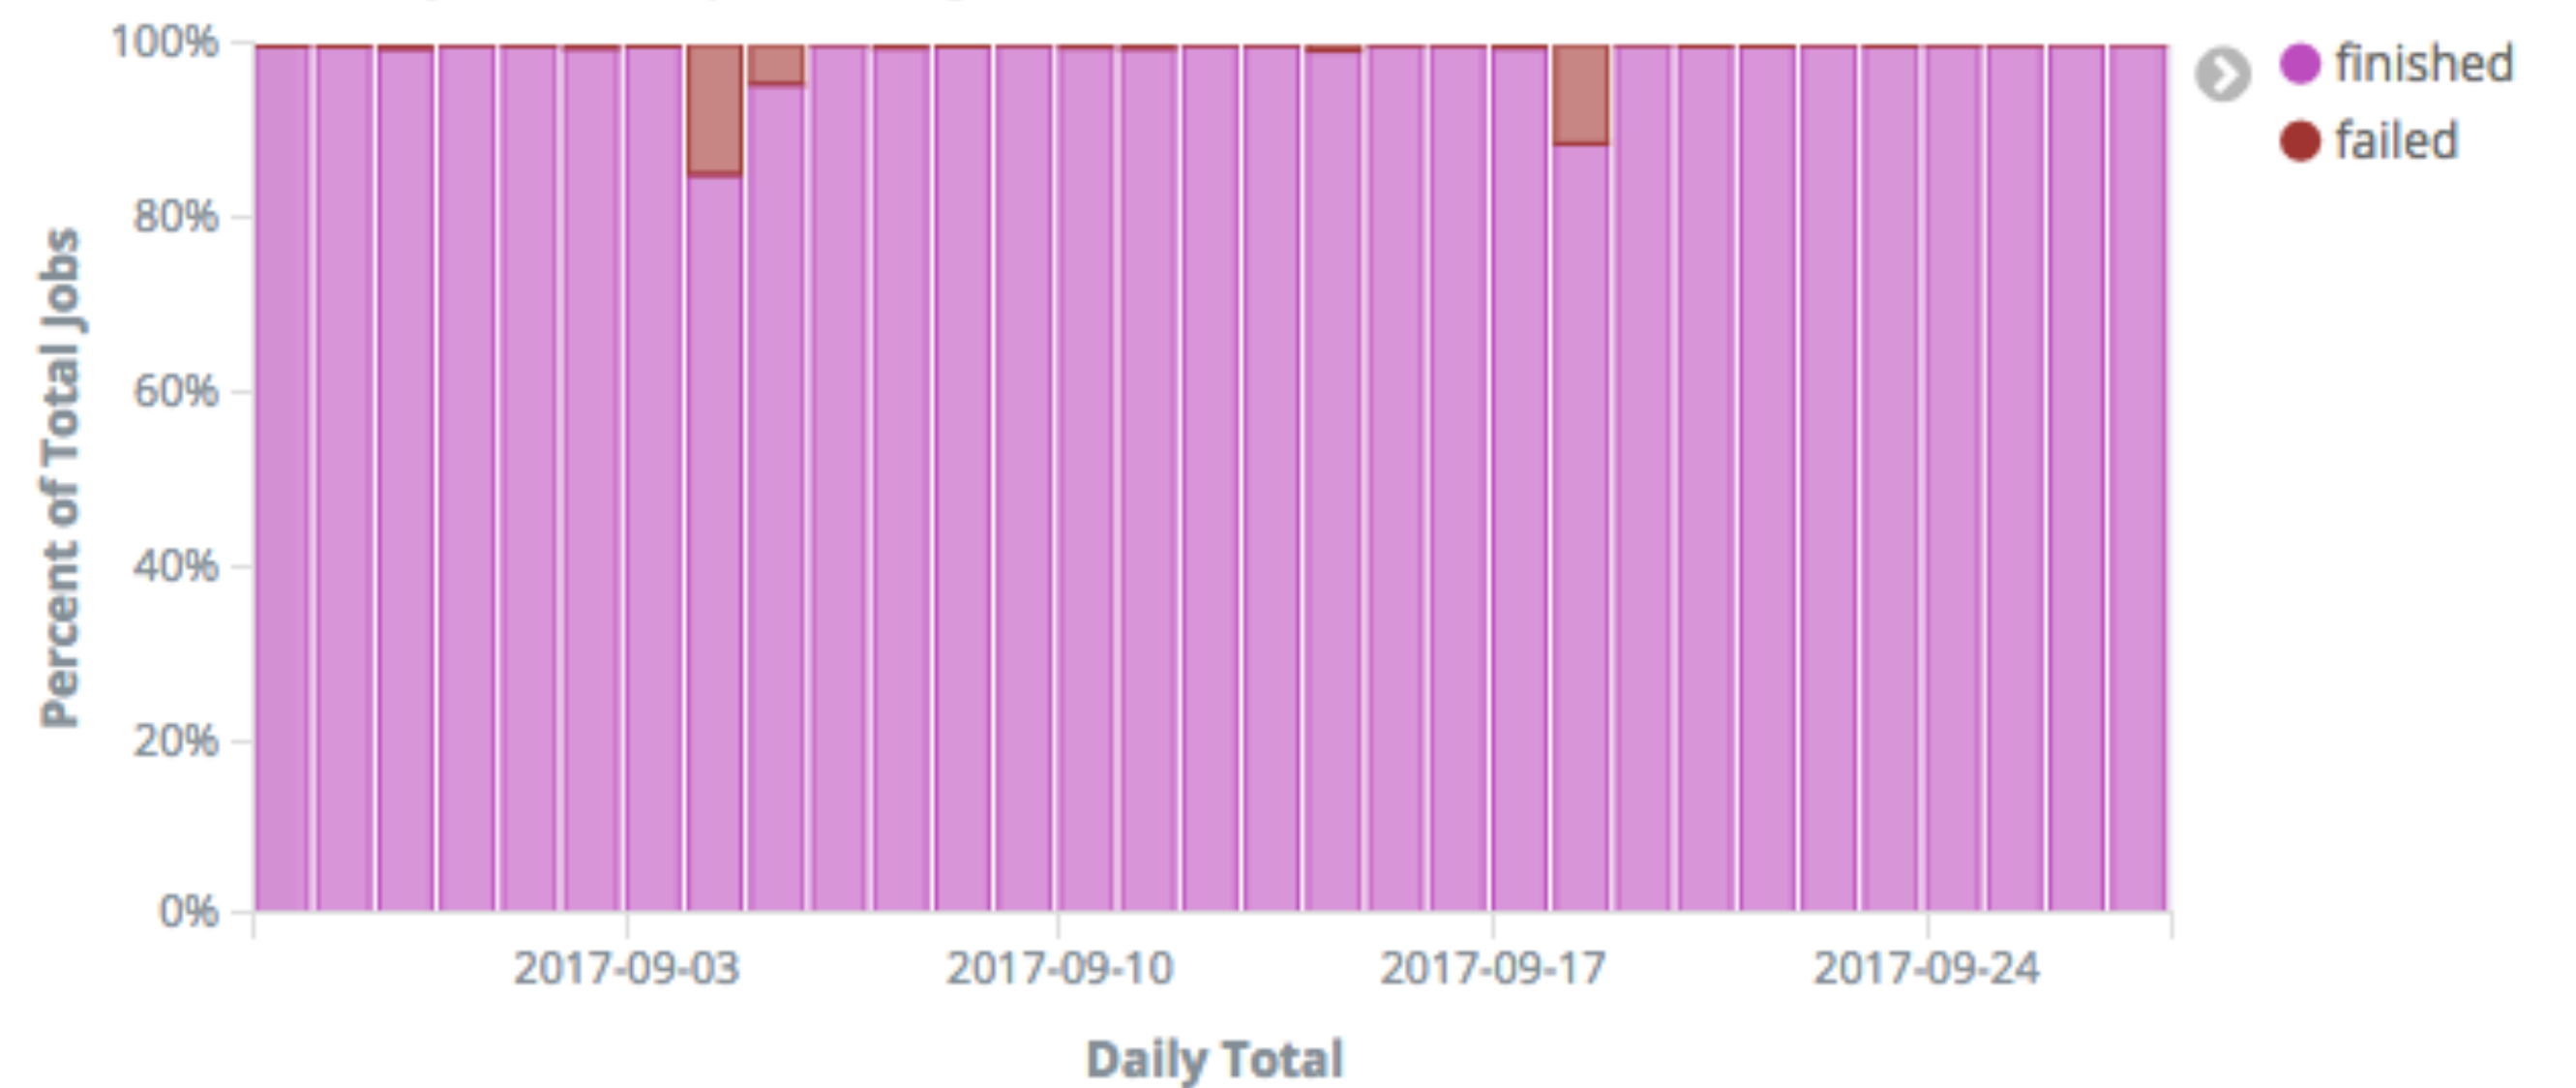
US DOE HPC Usage Summary

HPC DOE sites events per hour per core




Playing with Kibana

HPC DOE sites job status percentage

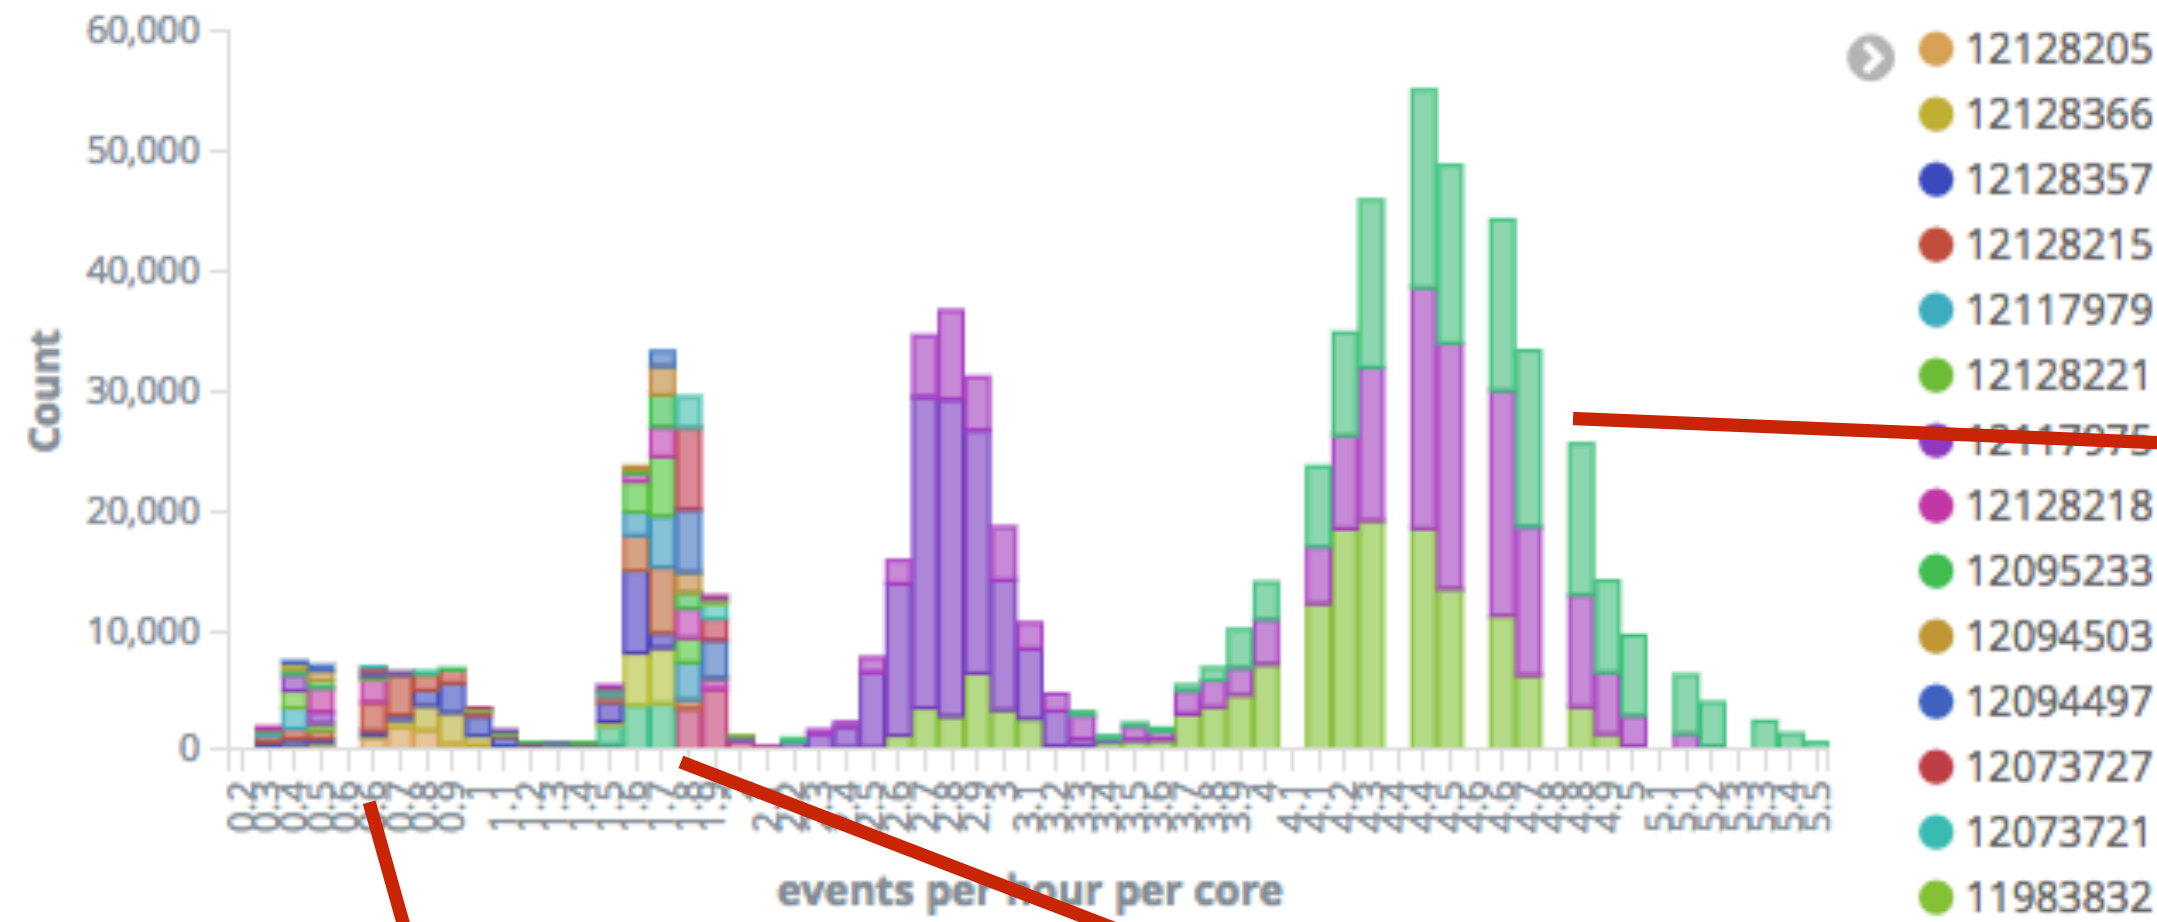


US DOE HPC Usage Summary




US DOE HPC Usage Summary

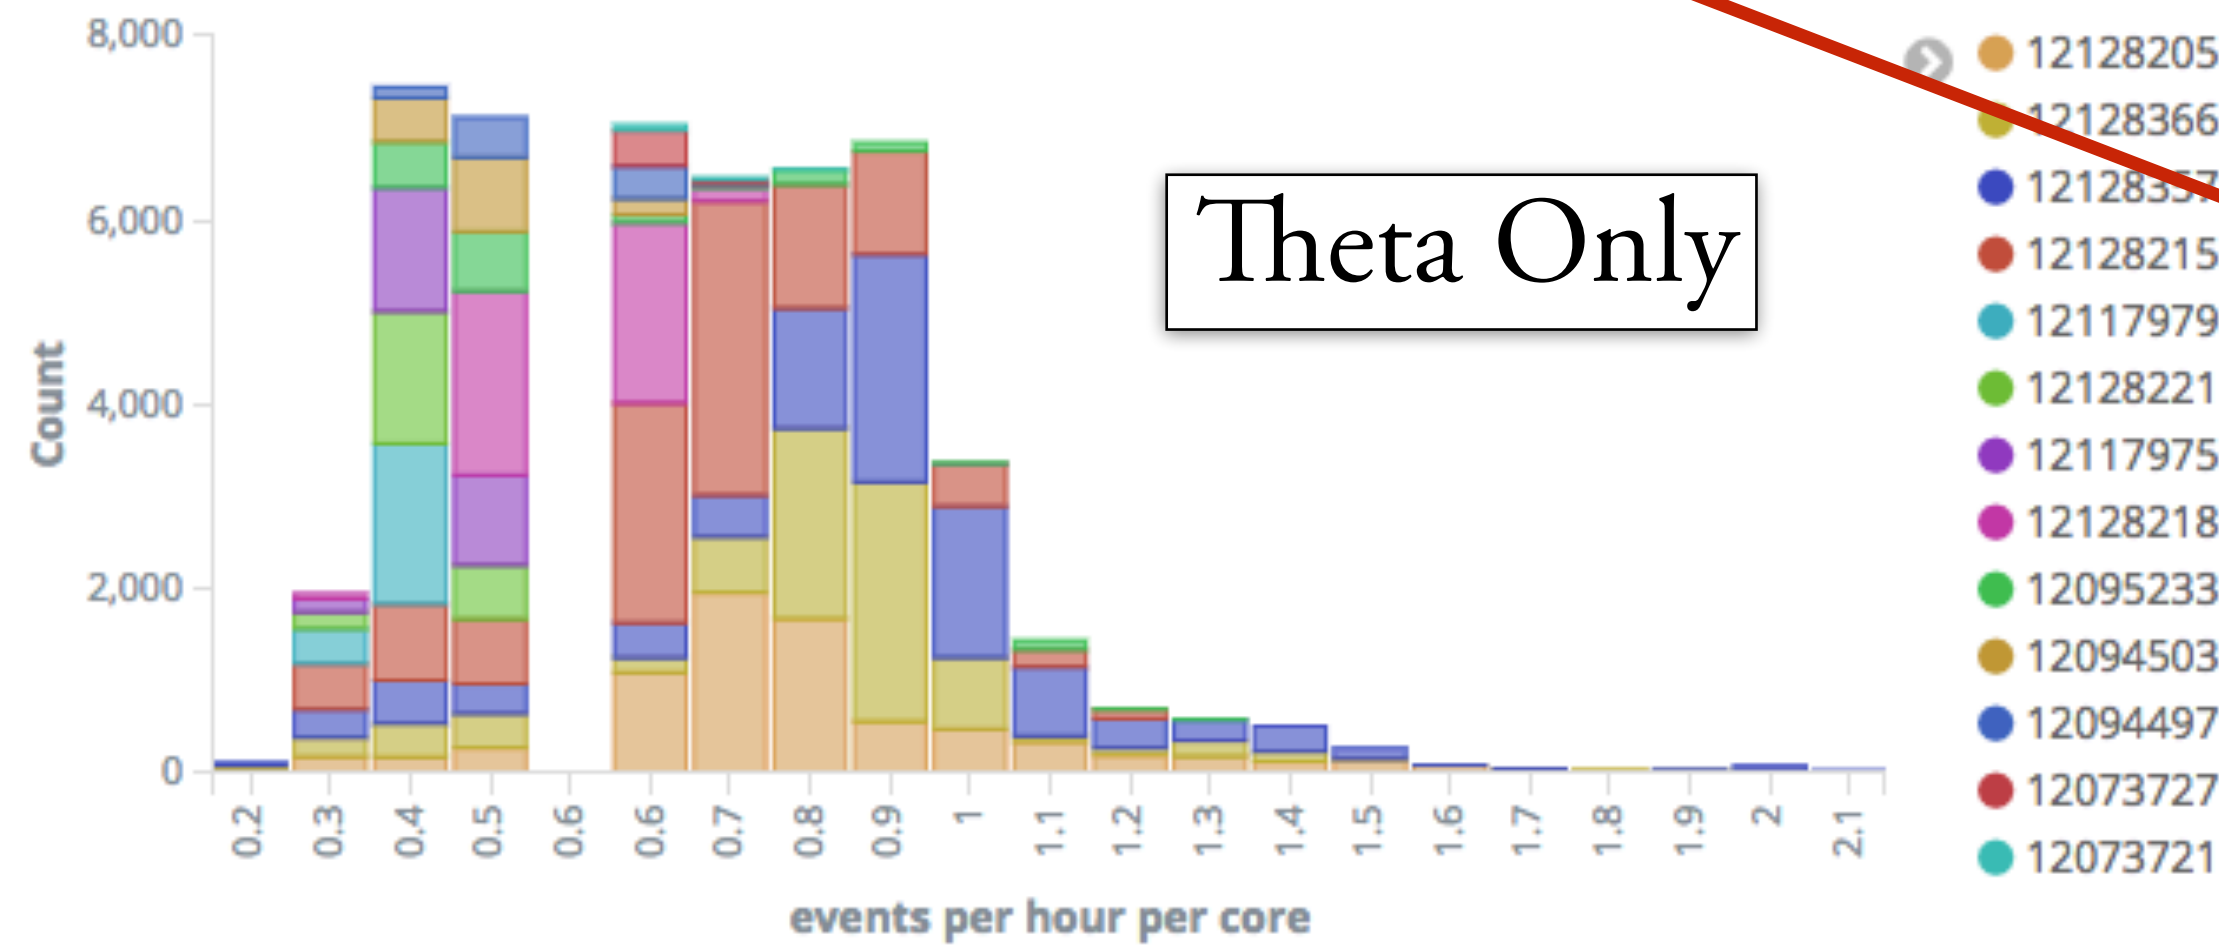
HPC DOE sites events per hour per core per taskid



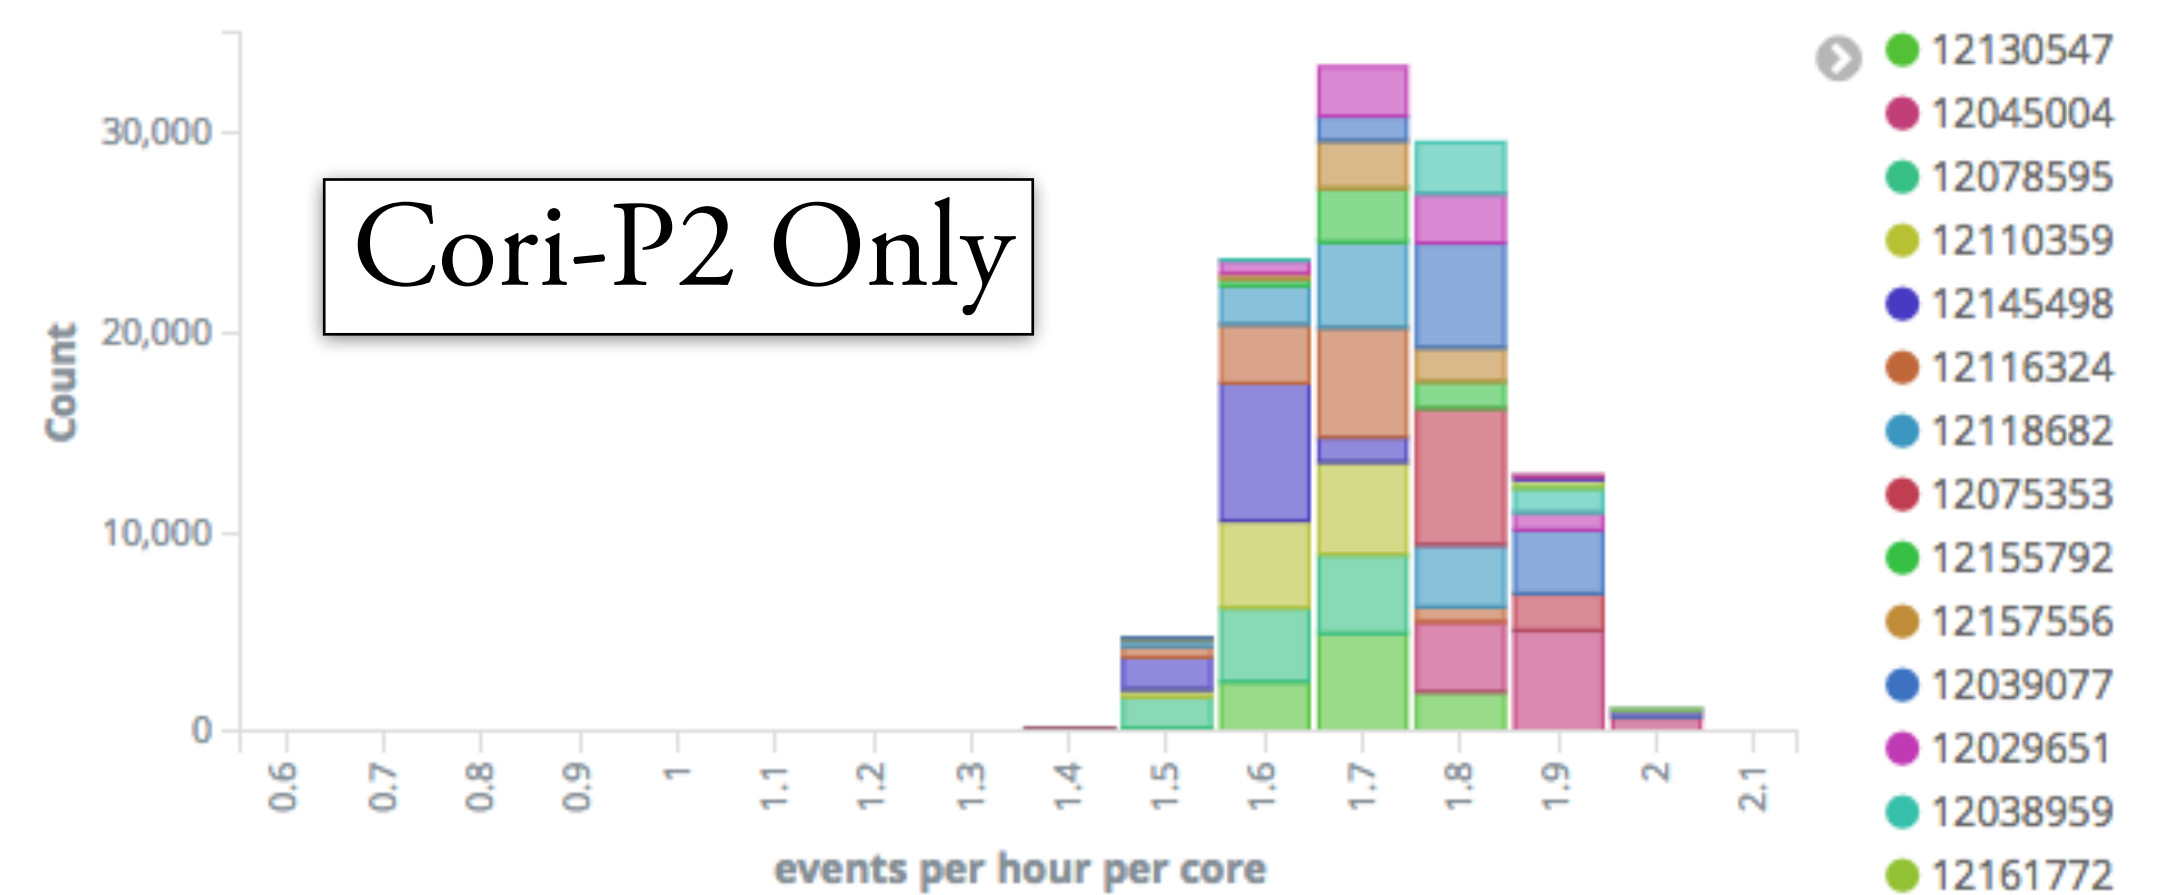
HPC DOE sites events per hour per core per taskid



HPC DOE sites events per hour per core per taskid





HPC DOE sites events per hour per core per taskid



OLCF Highlights


- ▶ Titan processed 27M events in September
- ▶ Running in backfill
- ▶ Running tests with validation task and allocation setup
- ▶ Eventually move to Harvester with two queues, one for backfill, one for allocation.
- ▶ After FS issues, increasing their batch job sizes again, which can be seen below.




NERSC Highlights

- ▶ September NERSC processed 96M events
- ▶ ERCAP (from my.nersc):
 - Allocated 76.5M core-hours since March
 - Used 94.6M core-hours since March
- ▶ ALCC (from my.nersc):
 - Allocated 70.5M core-hours since January
 - Used 68.2M core-hours since January
- ▶ Totals (from nim.nersc):
 - Used 158M core-hours since January
 - Charged for 133M core-hours since January
 - That's 25M core-hours in backfill time.
- ▶ Ran out of allocation over last week.
- ▶ This reduces us to about 1M core-hours per day.
- ▶ There also seems to be a bug in SLURM such that when we run out, our queued jobs do not get changed to 'scavenger' and they never execute, requiring by hand intervention.
- ▶ Received 10M more core-hours this week.

NERSC Allocation Usage




NERSC Usage



ALCF Highlights



- ▶ Running in ManyToOne mode with mini-pilot
- ▶ Currently running 128-node batch jobs on Theta with up to 8 concurrent jobs, 1024 nodes or 25% of machine.
- ▶ Processed about 15M validation events in September
- ▶ Getting a handle on Harvesters resource utilization
- ▶ Transitioned to Globus Online transfers on Sept 17.
- ▶ Debugging this and understanding features of GO
- ▶ Discussed Harvester readiness at CERN:
 - Need to finalize GO transfers, mainly pooling
 - Cleanup and publish mini-pilot for harvester
 - Define how to handle scheduled downtimes
- ▶ Then Harvester can be released to OLCF for deployment

- ▶ Otherwise, harvester has been doing its job. We've used the 20M core-hours and were just awarded another 45M core-hours to use in FY18.



ALCF Highlights

- ▶ Running in M
- ▶ Currently run
- ▶ 8 concurrent
- ▶ Processed abo
- ▶ Getting a han
- ▶ Transitioned t
- ▶ Debugging th
- ▶ Discussed Ha
 - Need to fin
 - Cleanup an
 - Define how
- ▶ Then Harvest
- ▶ Otherwise, ha
- ▶ 20M core-ho
- ▶ hours to use i



ATLAS Software Distribution via Containers

- ▶ Doug & Wei discussed this last week and in today's meeting
- ▶ Starting to work out the details of how to automate the creation/distribution

