

Whole Node Scheduling

Pete Eby
ebypi@ornl.gov

Steve Moulton
moultonsa@ornl.gov

Slides:
Dec 7th 2022 @ ORNL

ORNL is managed by UT-Battelle LLC for the US Department of Energy

Process

- Modified Slurm configuration (server: slurm)
 - Simplified partition list, since number of cores to be used is no longer an issue
 - `PartitionName=compute default=yes MaxTime=25:00:00`
`Nodes=alice-[001,003-056,060,061,063-070,079-084,086-102,104-105,107-112],alice-cray[011-018]`
 - Cleanup: removed `SlurmdTimeout` – not needed
 - Modified `maxjobs` from 5000 to 500. Probably no net effect.



- Modified CE environment

```
< export SUBMIT_ARGS="-o /tmp/jobagent_%j.out -e /tmp/jobagent_%j.err  
--gres=MaxJobs:1"  
> export SUBMIT_ARGS="-o /tmp/jobagent_%j.out -e /tmp/jobagent_%j.err  
--nodes 1 --exclusive"
```

Process (2)

- Changes to LDAP
 - Handled elsewhere ☺
- While system load was light, updated compute nodes, rebooted (some nodes up over 180 days)
 - In passing: Some service hosts up > 490 days, still running happily
- Modified startup/shutdown, etc scripts. Much simpler now.

```
[aliproduct@voboxslr ~]$ cat alien-services-start.sh
#!/bin/bash
/cvmfs/alice.cern.ch/scripts/vobox/jalien-vobox.sh start monalisa
/cvmfs/alice.cern.ch/scripts/vobox/jalien-vobox.sh start ce
```

Outcomes

- Initial job instability
 - Jobs requiring large amount of memory, OOMs
 - Instability was in O2 MC jobs; coincidental to whole node scheduling change
 - Looks good now.
- Spot problem
 - Some nodes will go a long time without getting a job. How to debug?

Per Node Efficiencies (Purple is Good)

