

Perlmutter

What was, what is, what will be

Sergiu Weisz
sergiu.weisz@upb.ro
18/09/2024

Perlmutter

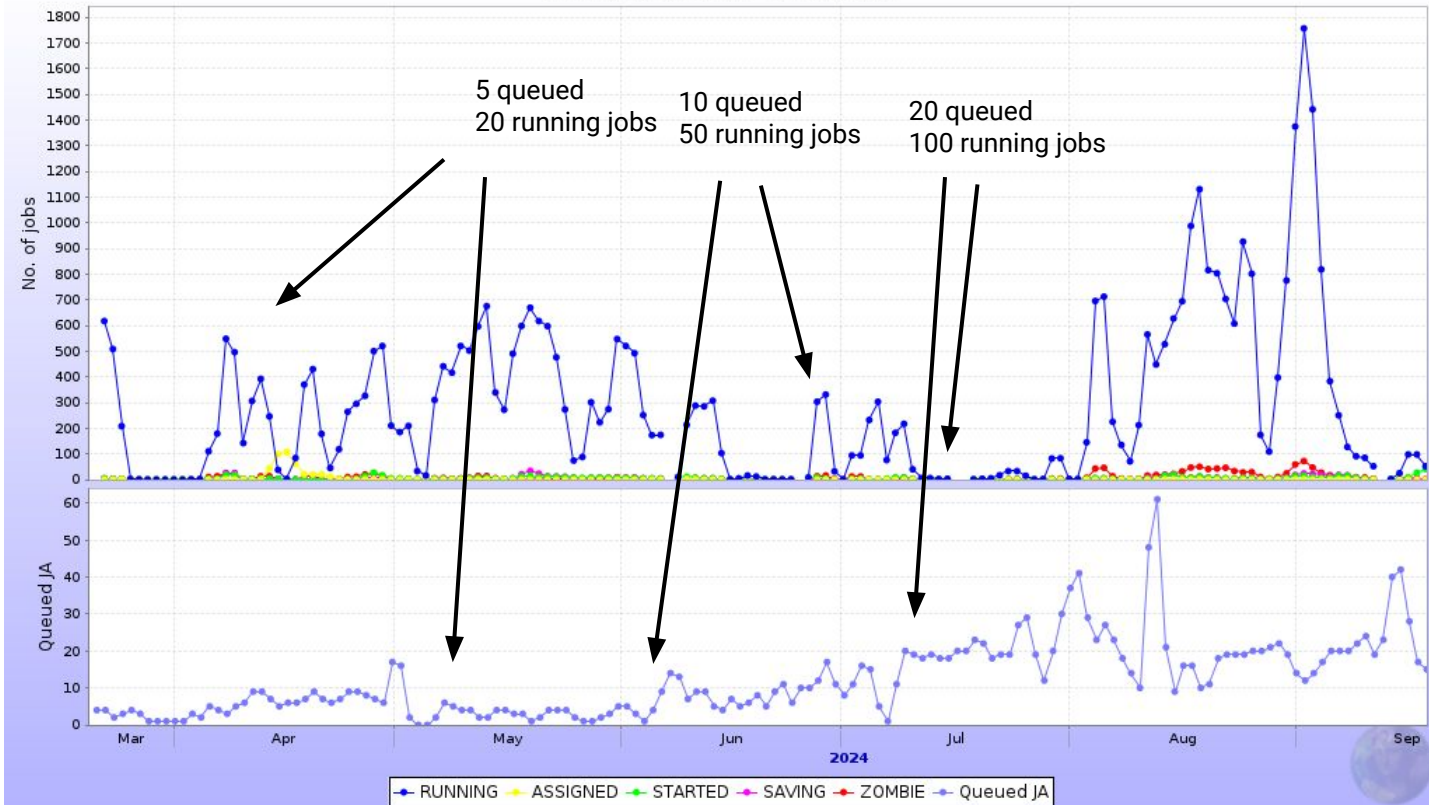
- Runs HPE Cray OS
- Only whole-node scheduling
- 3072 nodes running 2x AMD EPYC 7763, 64 core, 512 GB memory
- Jobs run in Shifter containers
- Outside connection from nodes
- Full CVMFS support
- SFAPI submission

Issues when last speaking

- VO running under my user
 - Moved to aliproduct user on VM
- VO running “by hand”
 - Moved to a systemd user unit
- CE running from JAR, agents running from CVMFS
- Key generated by hand
 - Still a limitation for automation
 - Irakli can setup the keys now

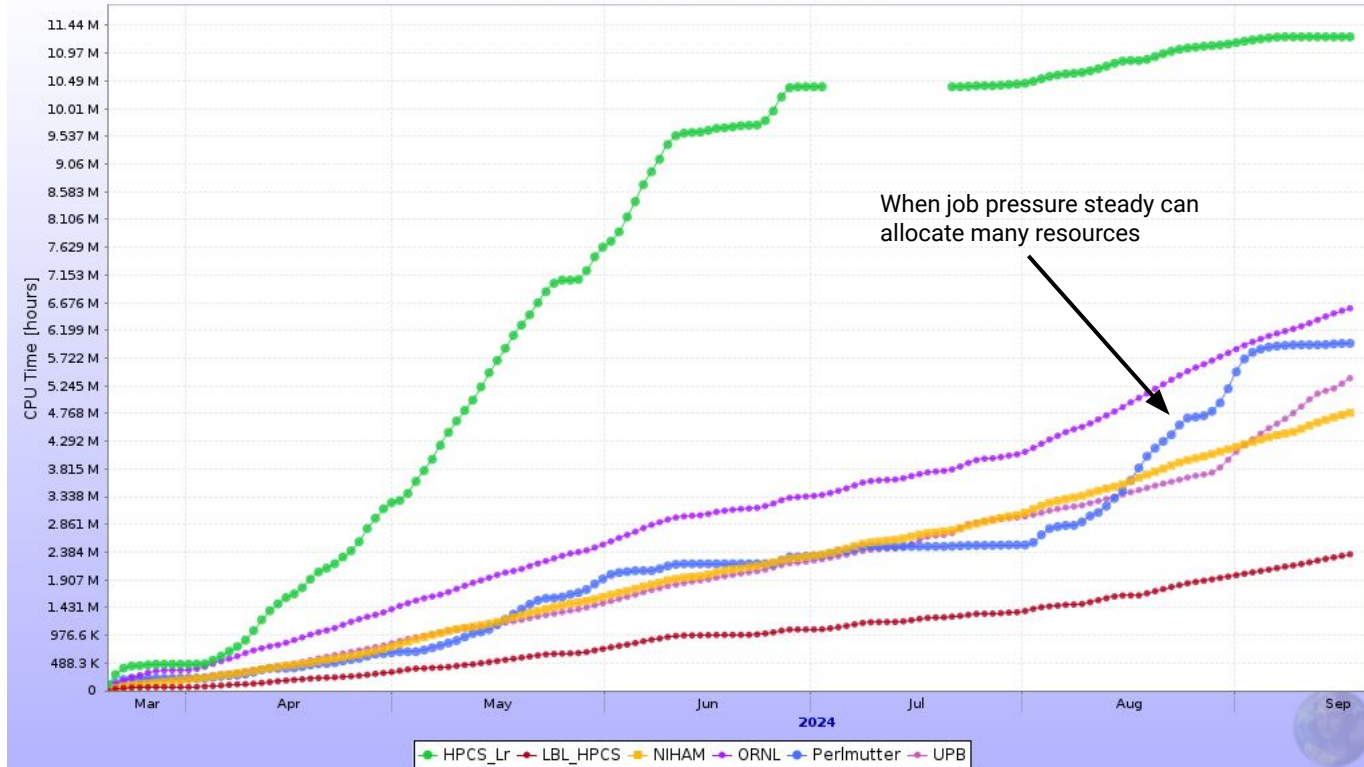
Story so far

Active jobs in Perlmutter



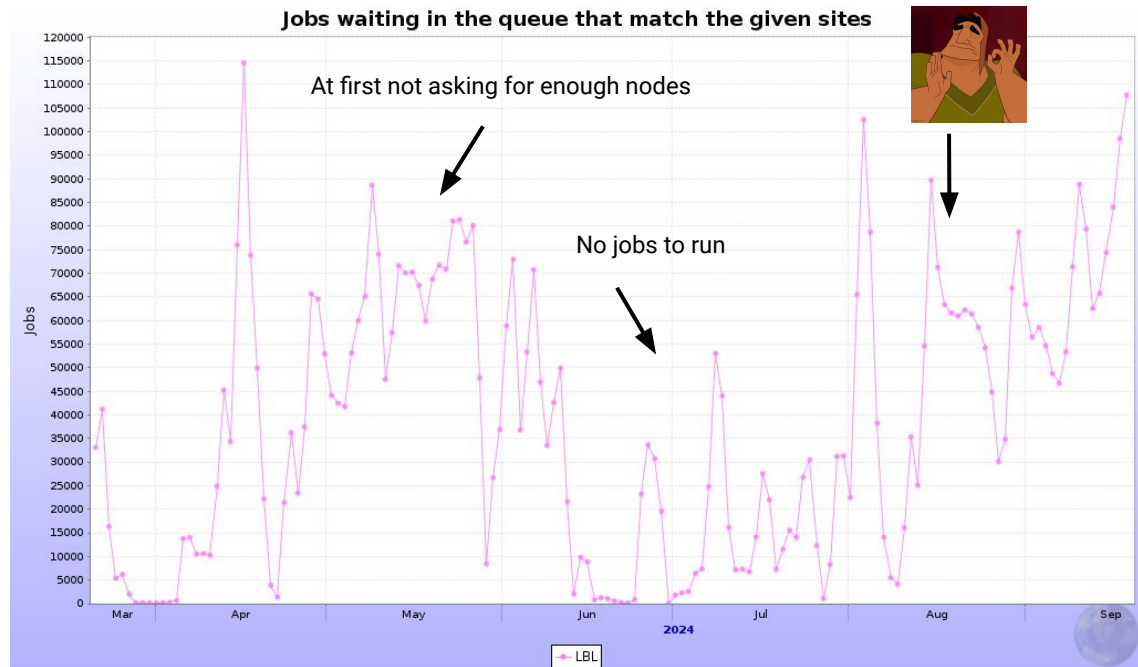
CPU hours offered

Total CPU time for ALICE jobs



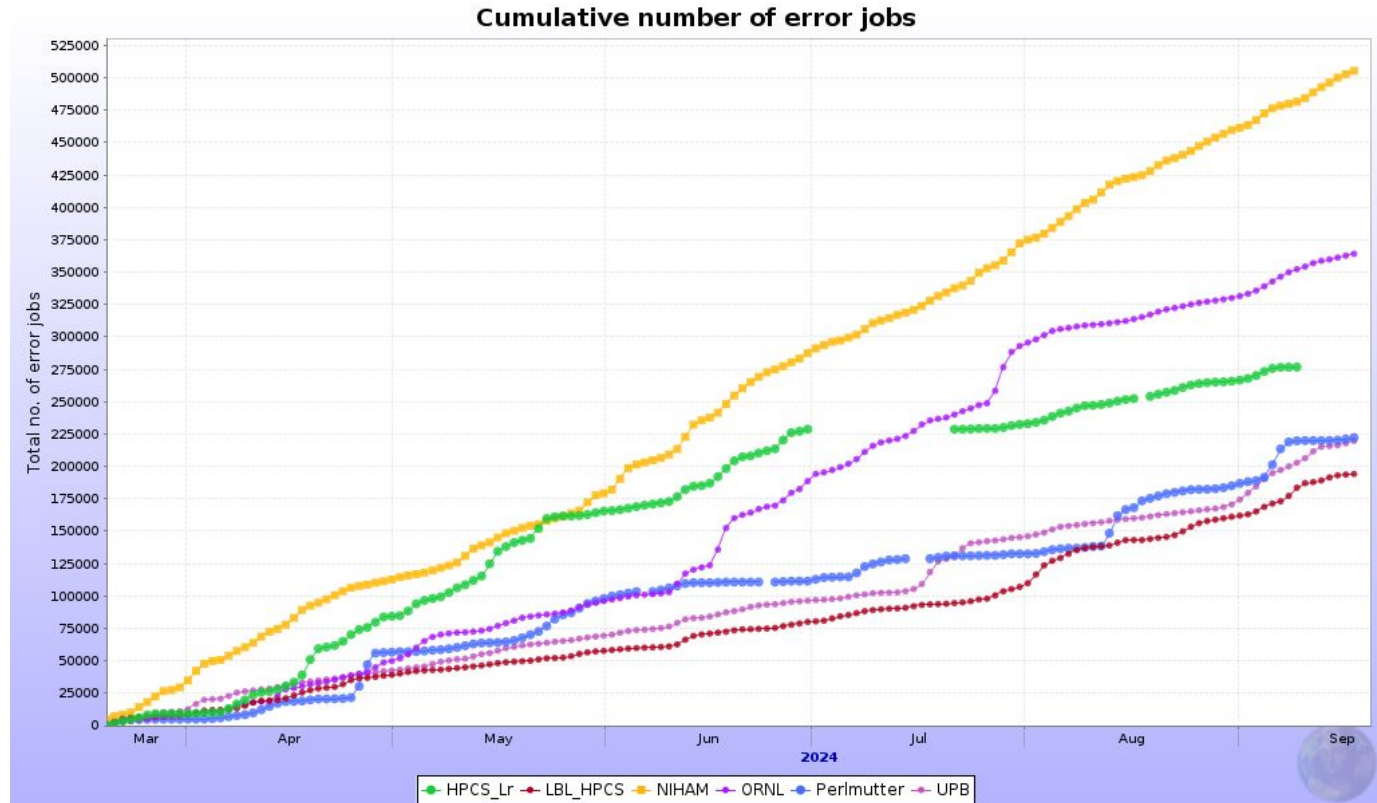
Why is job allocation an issue

- Only running MC and DAQ jobs
- Depend on MC job pressure and Perlmutter queue



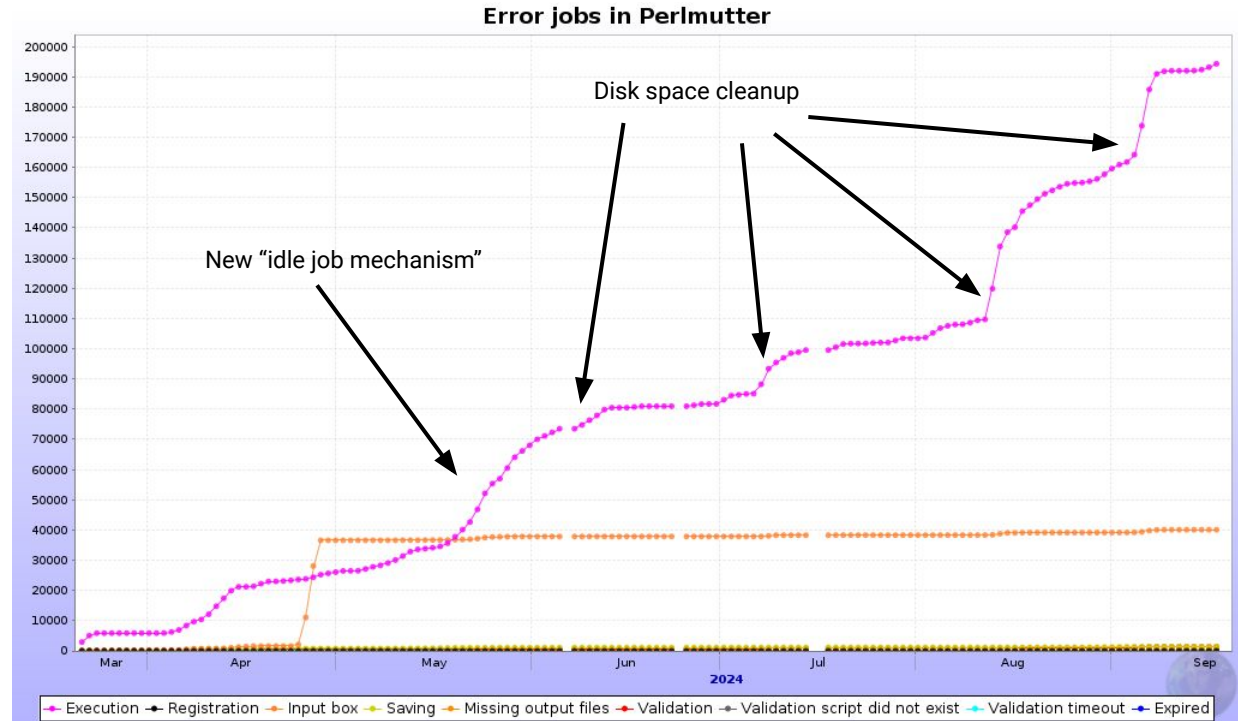
Perlmutter errors

Same as similarly sized sites



Deeper dive into errors

- Sudden fails leave uncleaned workdirs
- Running hyperloop jobs lead to ERROR_E jobs



Current issues

- Patches needed to be able to run CE out of CVMFS
 - Hardcoded paths for SFAPI key
- Need to automate job workdir cleanup
- Perlmutter queue contention may be high
- No OSG reporting
- No automated or longterm key management



Perlmutter queue contention

Active jobs in Perlmutter

