Computing Update

N. Baltzell - CLAS Collaboration Meeting - March 2, 2021

CLAS12 Chef Documentation

- Weekly meetings with run group chefs and software/computing coordinators
- Documentation cleaned up, expanded, more guidelines and examples, FAQs, common issues and concerns
- Linked from the reconstruction tab of the software wiki
- Useful for new chefs and future run groups coming online in the collaboration, soliciting feedback on what's missing

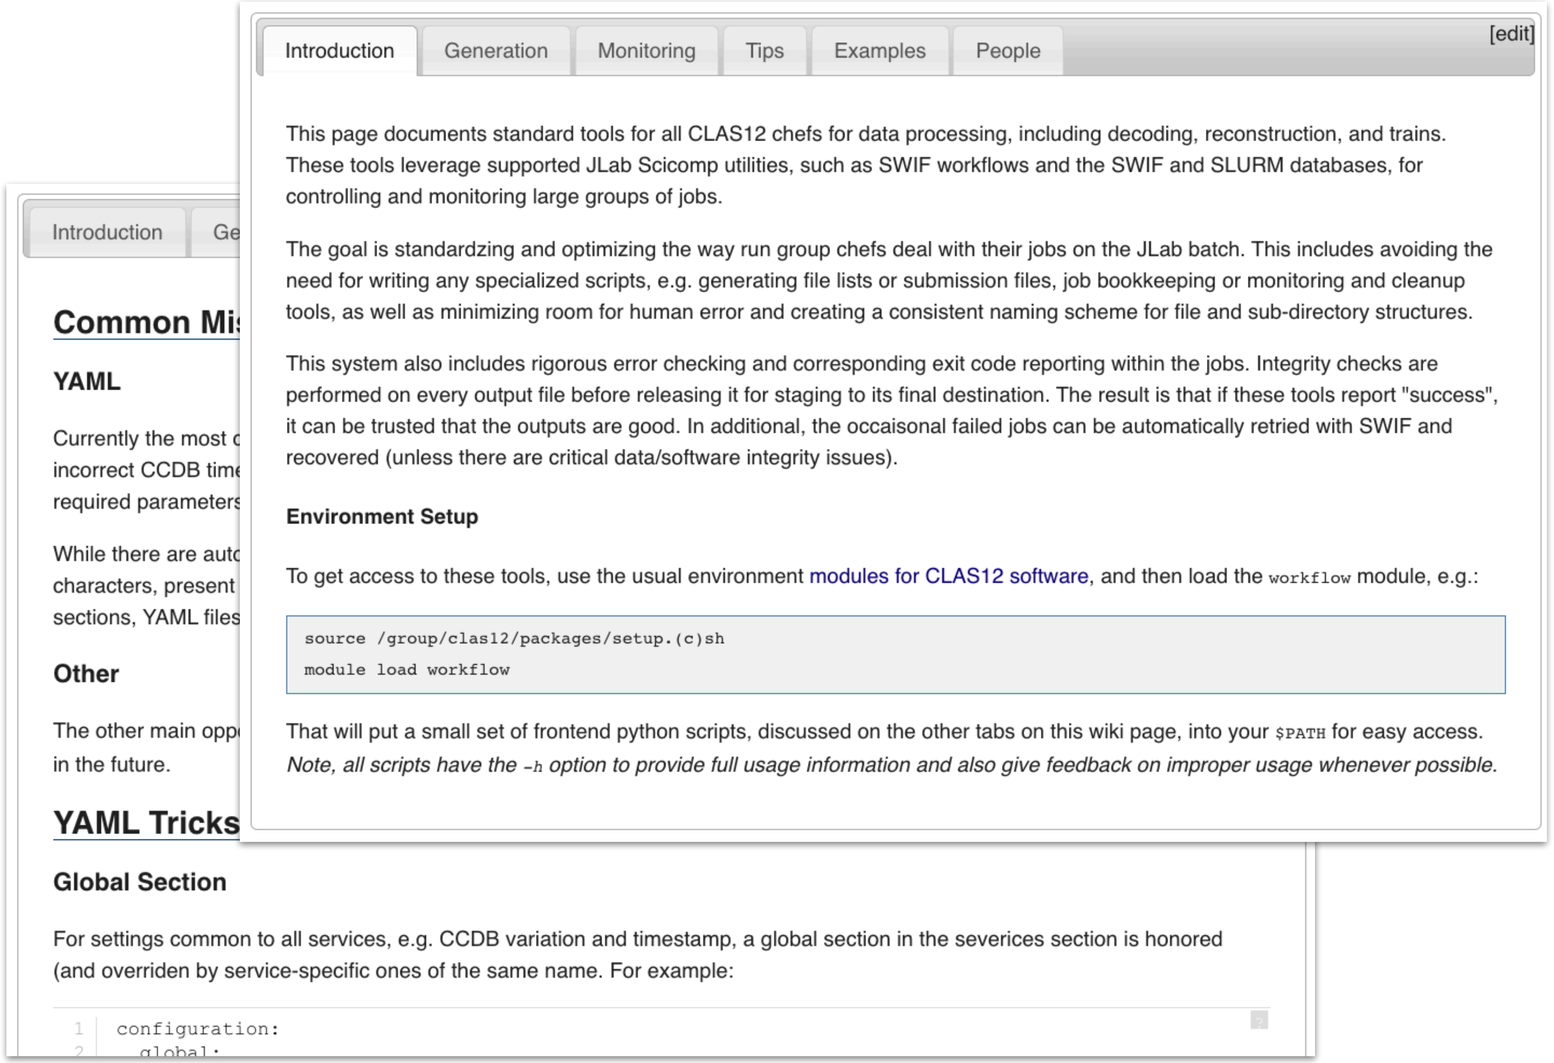

Batch Farm - Memory Requests

- Over the past year, we've gotten much better regarding CLAS/CLAS12 user jobs, but still occasionally some outliers.
- The plot below was yesterday and abnormal. It shows a 50% idle farm, even though plenty of jobs in the queue, purely due to unnecessarily-bloated memory requests!
- Before running many jobs on the batch farm, make sure your requests are appropriate for the jobs in question. Run some test jobs, check memory usage at scicomp.jlab.org, and set requests accordingly.

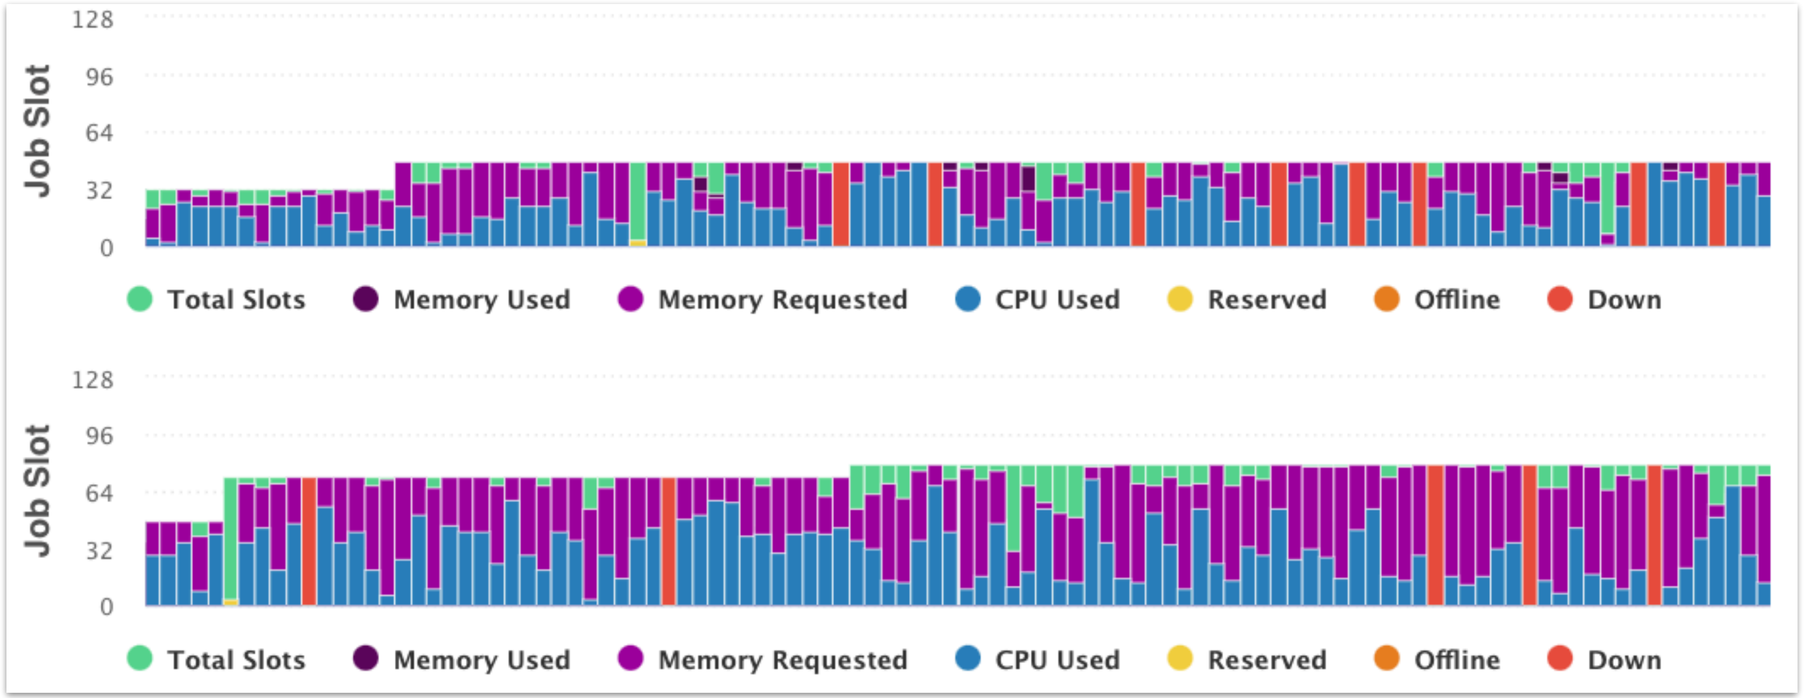

Future COATJAVA Releases - Java 11

- 10-20% reconstruction speedup by compiling and running in Java 11
  - wasn't effective in multi-threaded jobs until CLARA also compiled in 11, recently
- Dependencies also starting to get benefits from going to 11
  - e.g. recent GROOT improvements
  - In order to accommodate that, production code will not be compilable in 8 anymore (without manually reverting to older dependency versions, but we can keep a branch available for that if necessary)
- Java 11 is standard on all/most modern operating systems and is the current LTS release
  - and available in our CUE environment modules
- It's (past) time to make the switch; **expect future releases to not support Java 8 at compile nor runtime …**

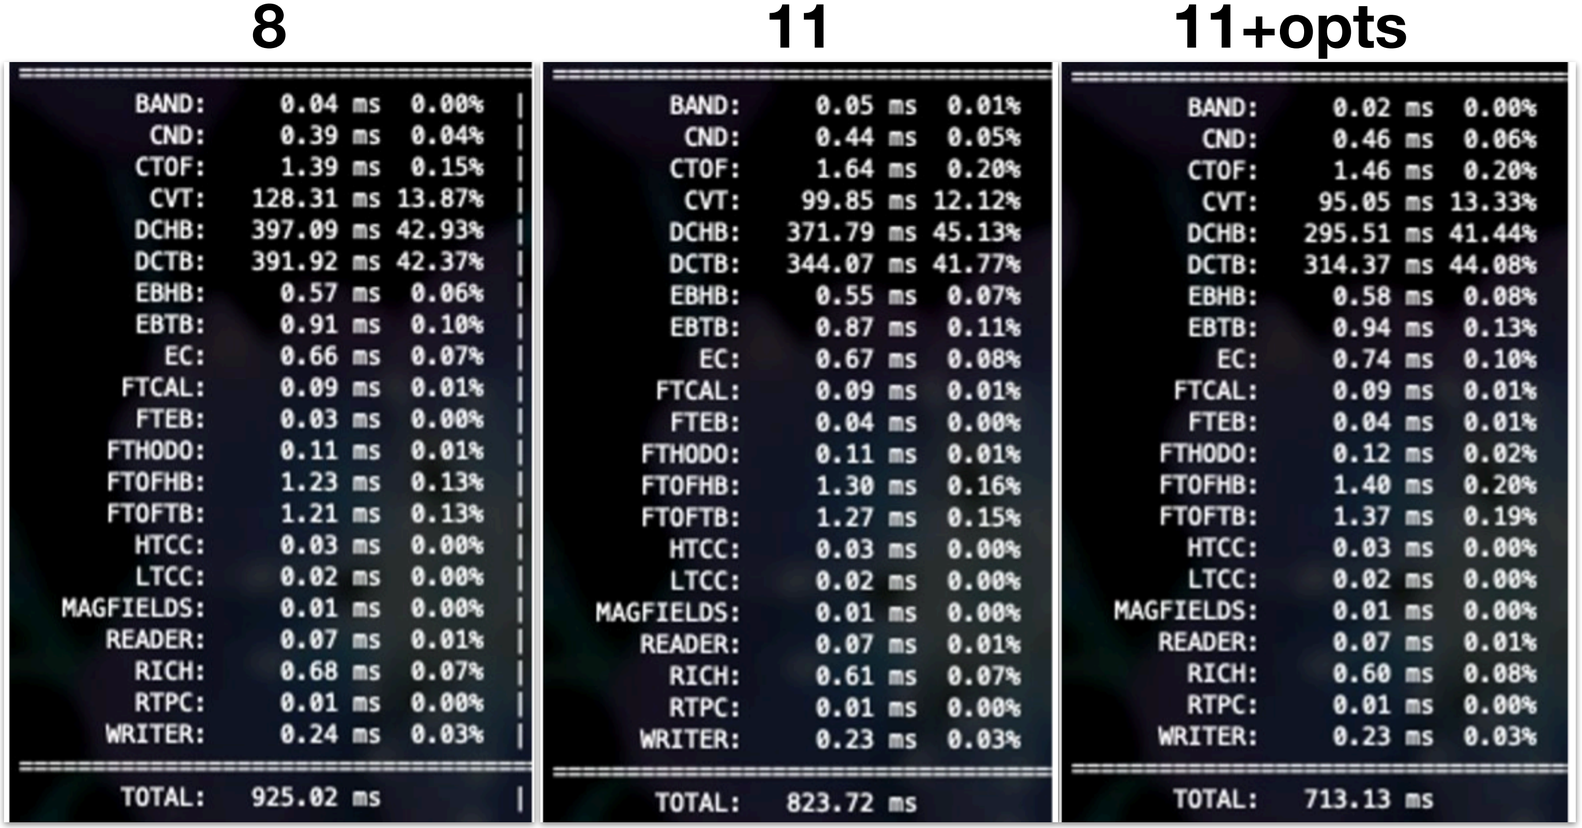

There's also potential speedup from 14 in garbage collection, which we have a lot of, and we may test and run production chef jobs in that soon. But since 14's not a LTS, and less easily available in many distributions, standard releases will not be migrated yet.

Scicomp (1)

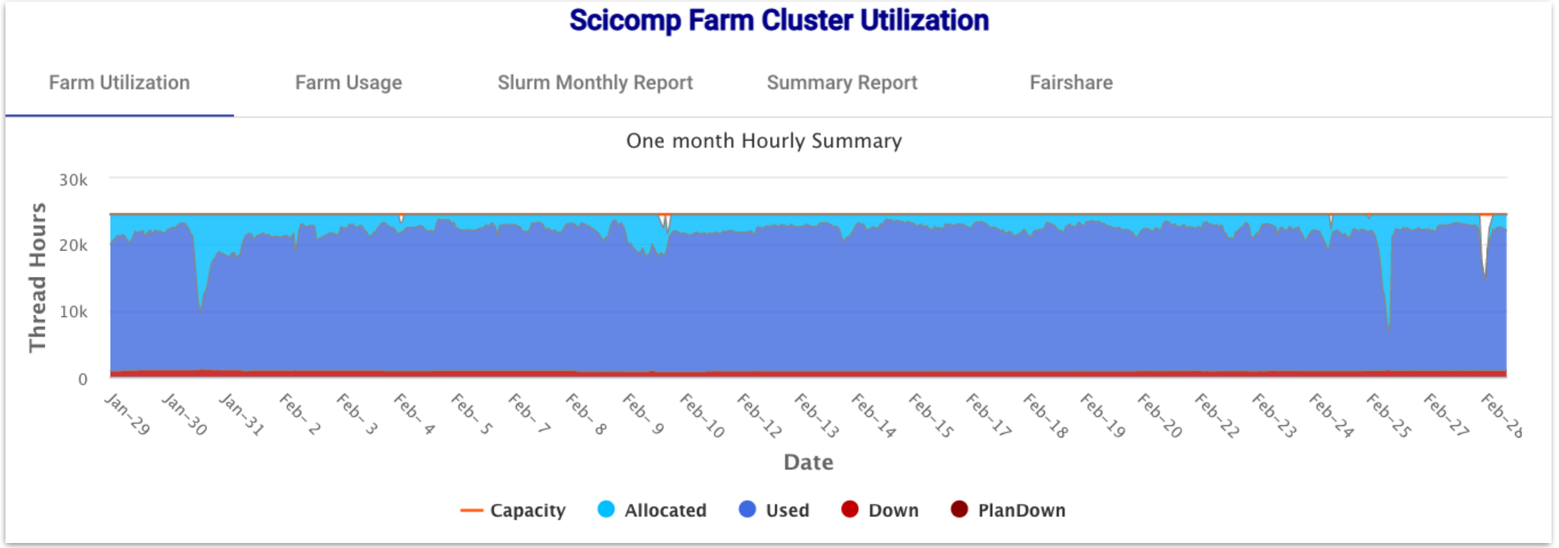

- New, more I/O-performant /work fileserver in procurement stages
  - should be available later this year
- New scicomp-2020.jlab.org website developed and available
  - Much faster load times, more user-friendly, more batch farm bookkeeping, e.g. quarterly, yearly
  - Does not provide support for Auger jobs, and Auger will eventually be removed and replaced by SWIF
    - *If you're still submitting jobs with Auger, you should consider updating to SLURM, or SWIF if you need file staging*
- swif2 is available for testing
  - see /site/bin/swif*, overhauled for offsite and grid jobs
  - user interface is very similar to the original swif (now called swif1)
  - see scicomp's website for documentation

Scicomp (2)

- New jupyterhub.jlab.org notebooks requested by CLAS collaboration members and in development
  - one with python kernel and pyROOT
  - another with C++/ROOT kernel, and clas12root support coming soon
- Full access to JLab filesystems
- Must be in a SLURM account for access, which is already true if you've run a batch job before
  - otherwise we submit a request at cc.jlab.org

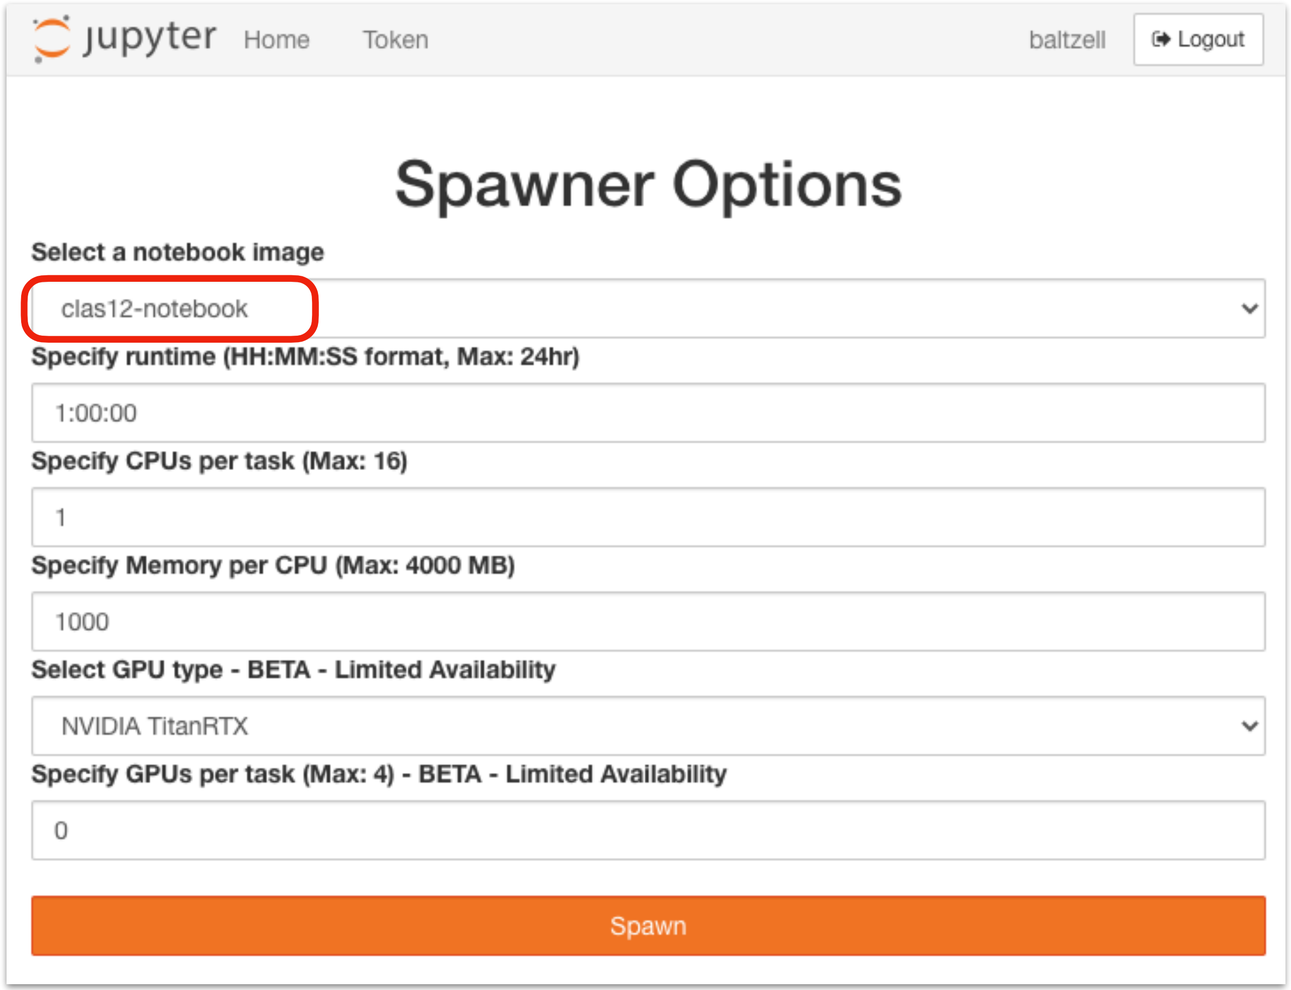

Summary

- CLAS12 chef documentation improved and expanded, soliciting feedback from chefs
- Batch farm job memory requests still important (and not always good, sometimes severely!)
- Moving COATJAVA and some dependencies (at least GROOT) to Java 11 for future releases
  - with a major version bump
- Scicomp
  - New and improved scicomp website
    - Doesn’t support Auger, which will be getting phased out eventually, move your job submissions to SLURM/SWIF
  - New notebooks at jupyterhub.jlab.org for CLAS in progress
  - New /work fileserver in procurement, anticipated ready for use this year