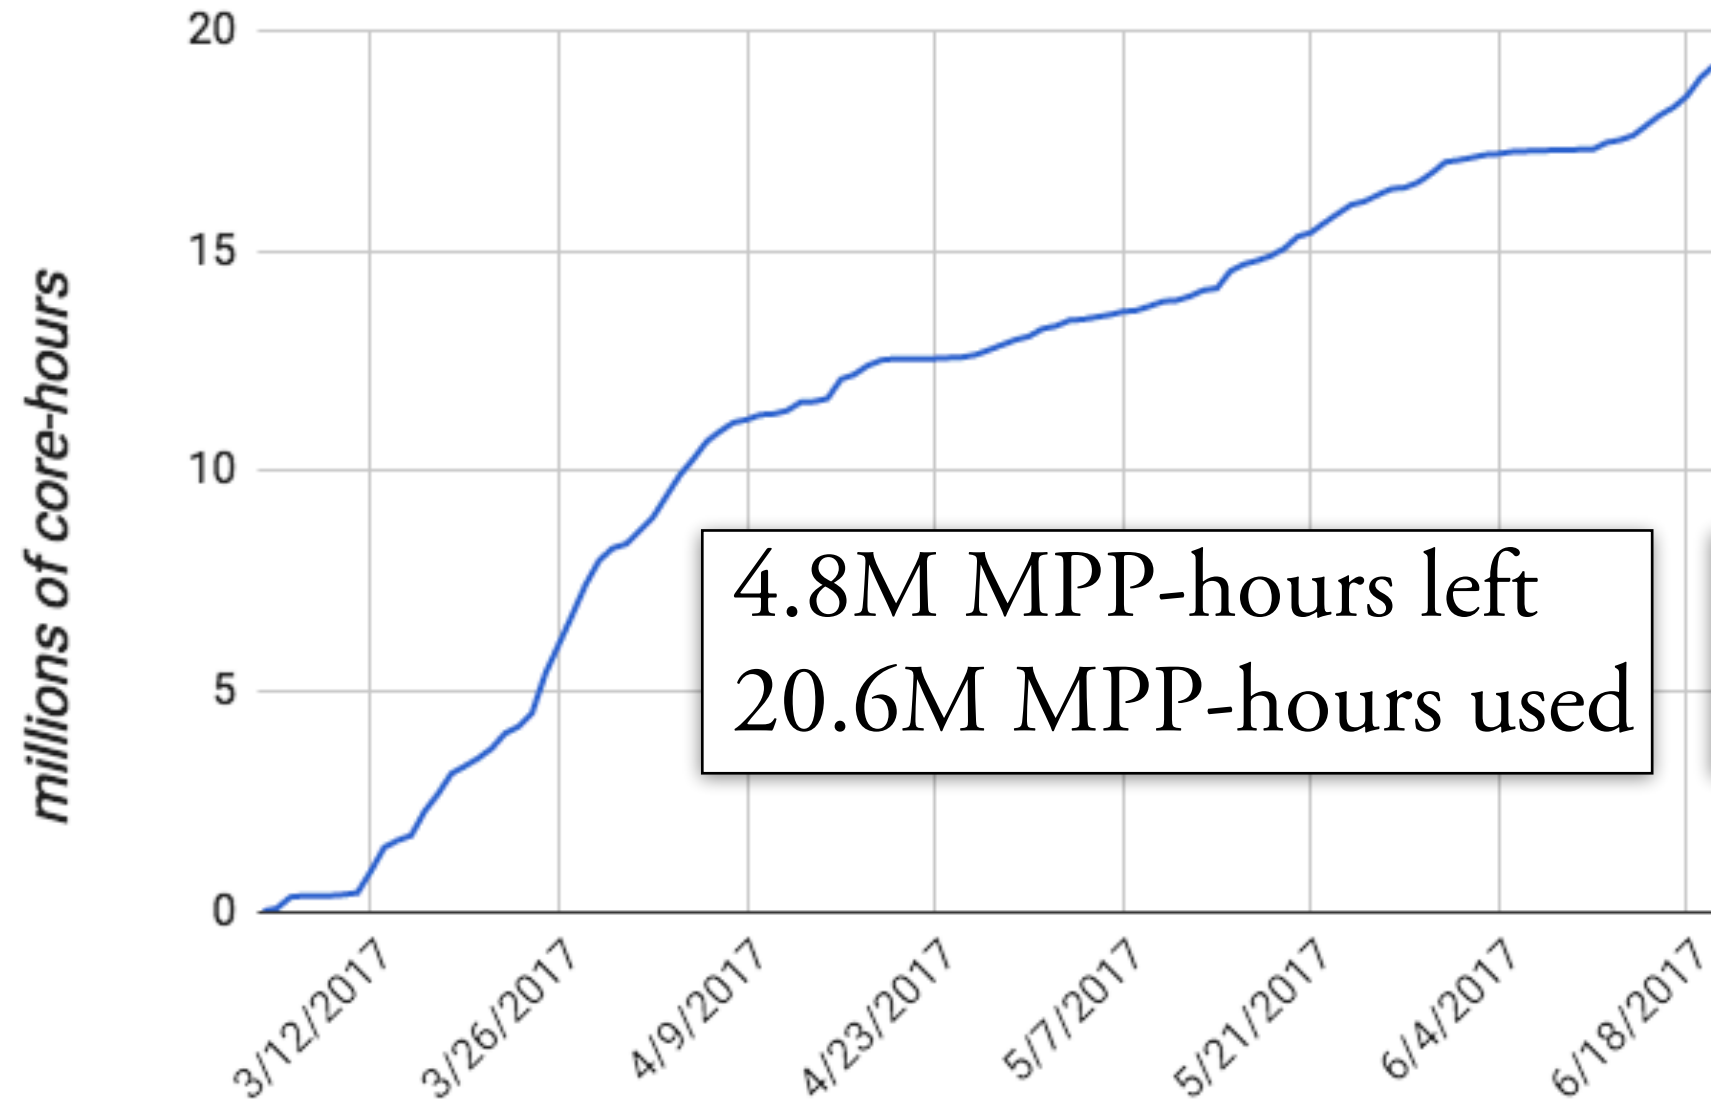


Titan/NERSC Running Summary

NERSC Allocation Usage

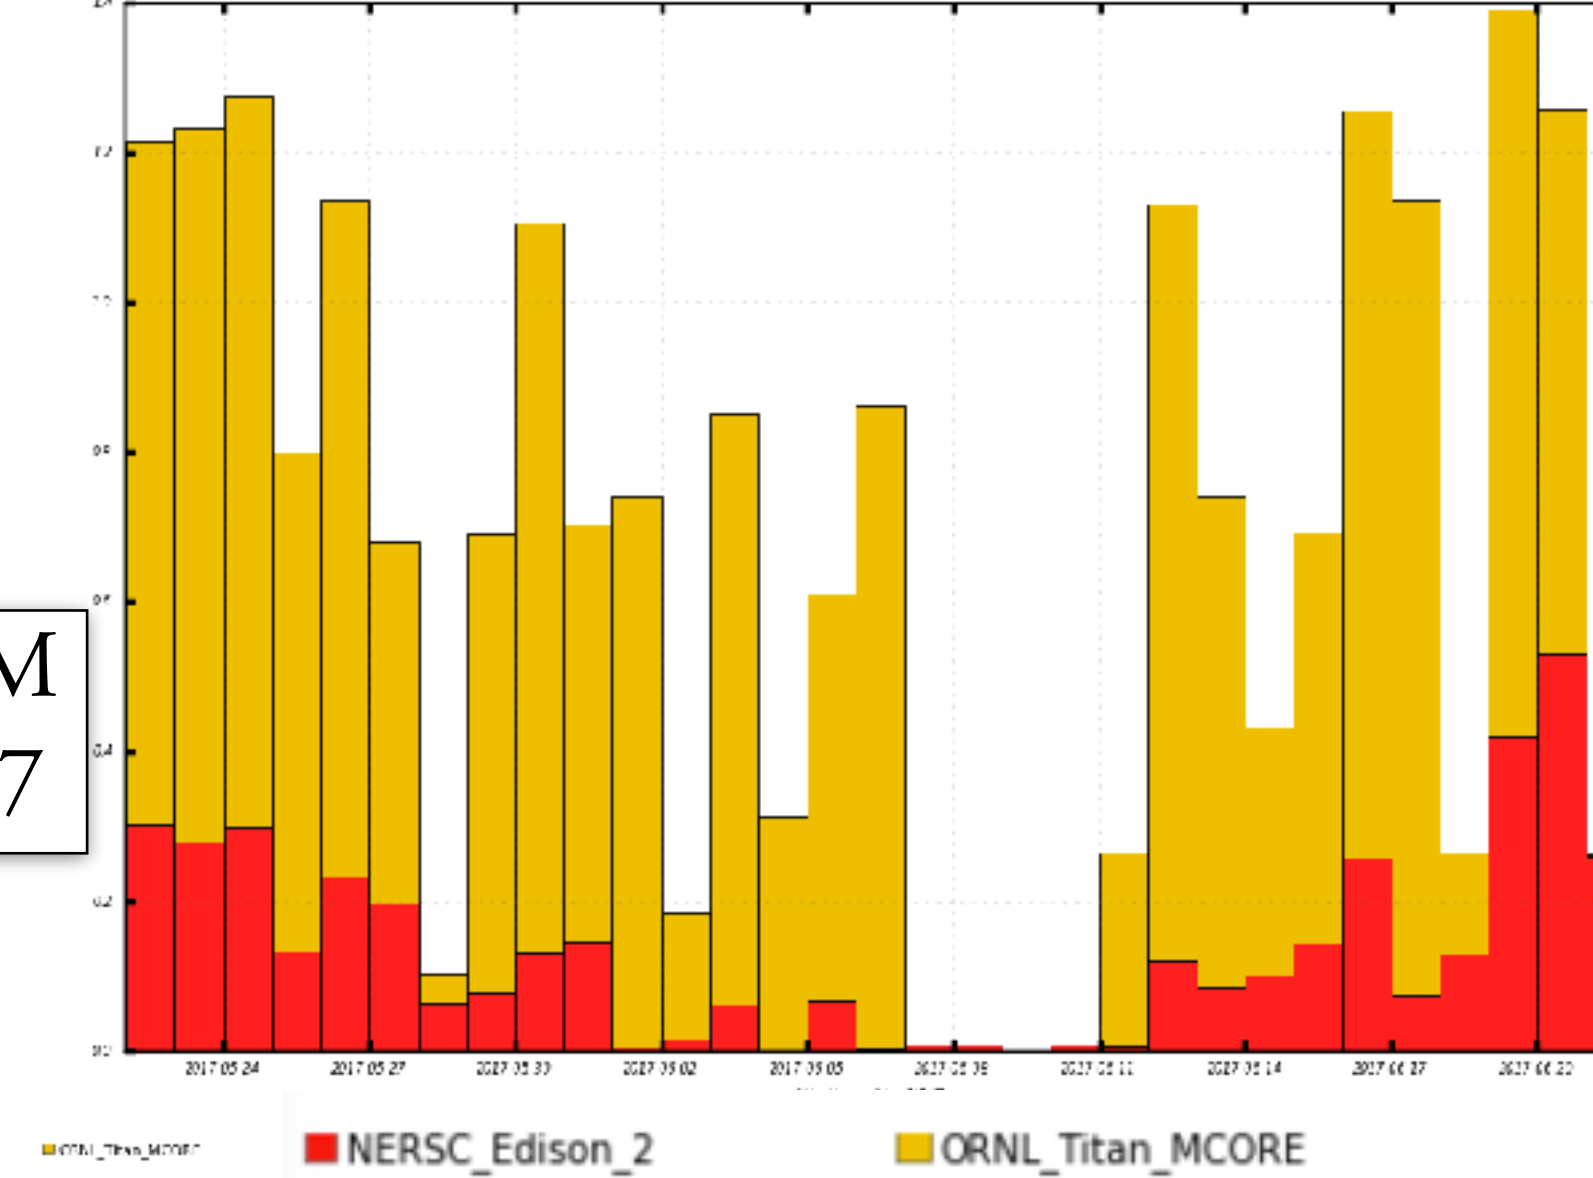


4.8M MPP-hours left
20.6M MPP-hours used

Titan: 72.5 M
hours used in 2017


dashboard

CPU consumption Good Jobs in seconds




dashboard

Efficiency based on success/all accomplished jobs



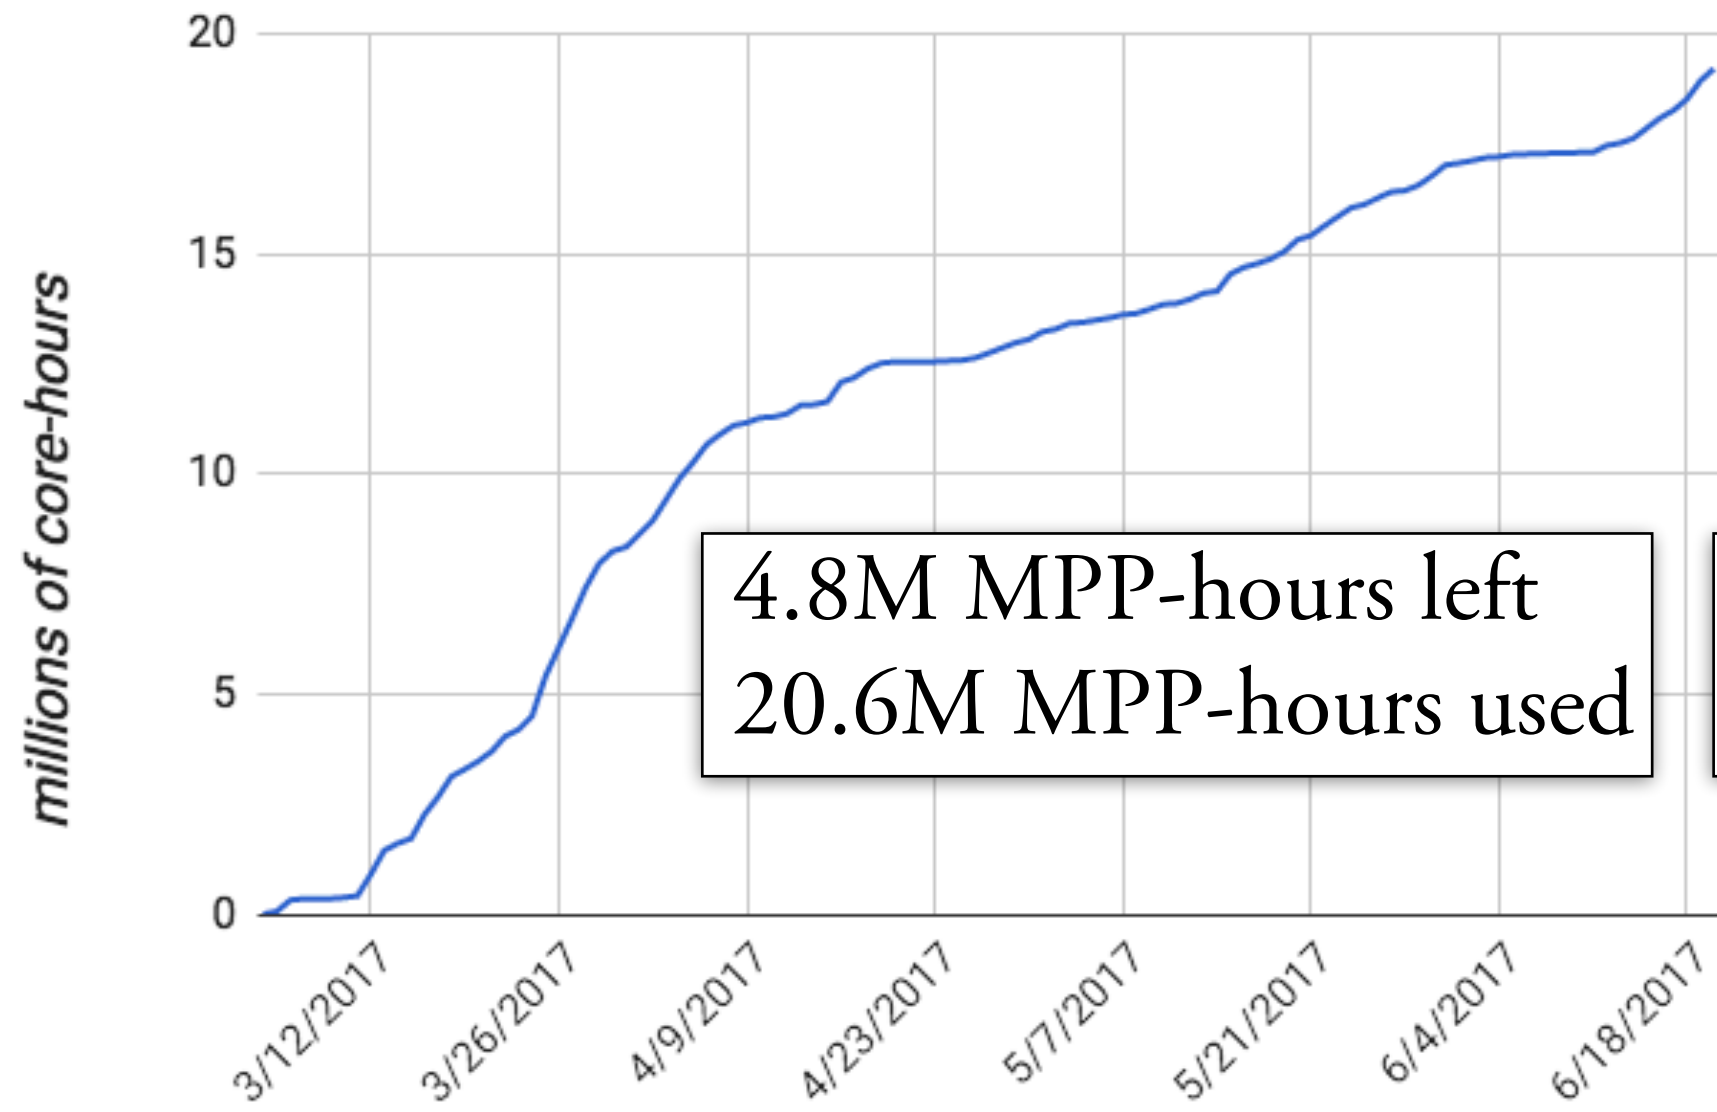
dashboard

Slots of Running Jobs



Titan/NERSC Running Summary

NERSC Allocation Usage

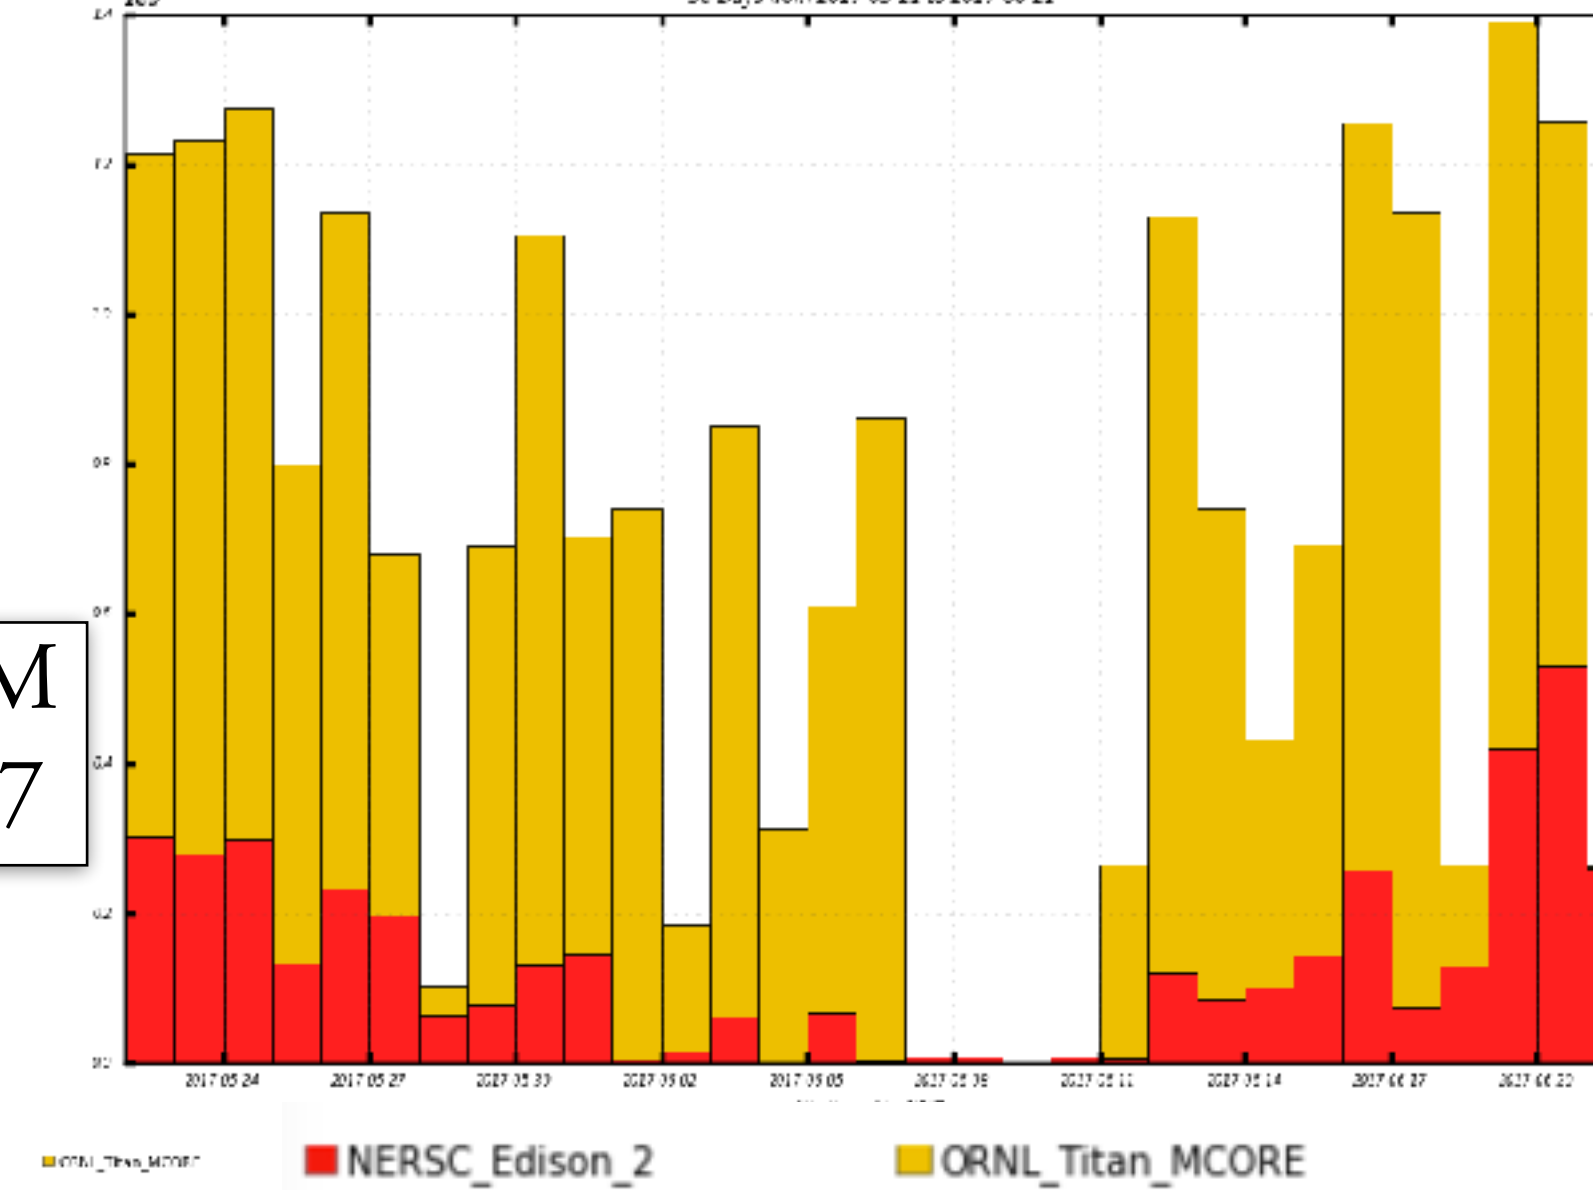


4.8M MPP-hours left
20.6M MPP-hours used

Titan: 72.5 M
hours used in 2017

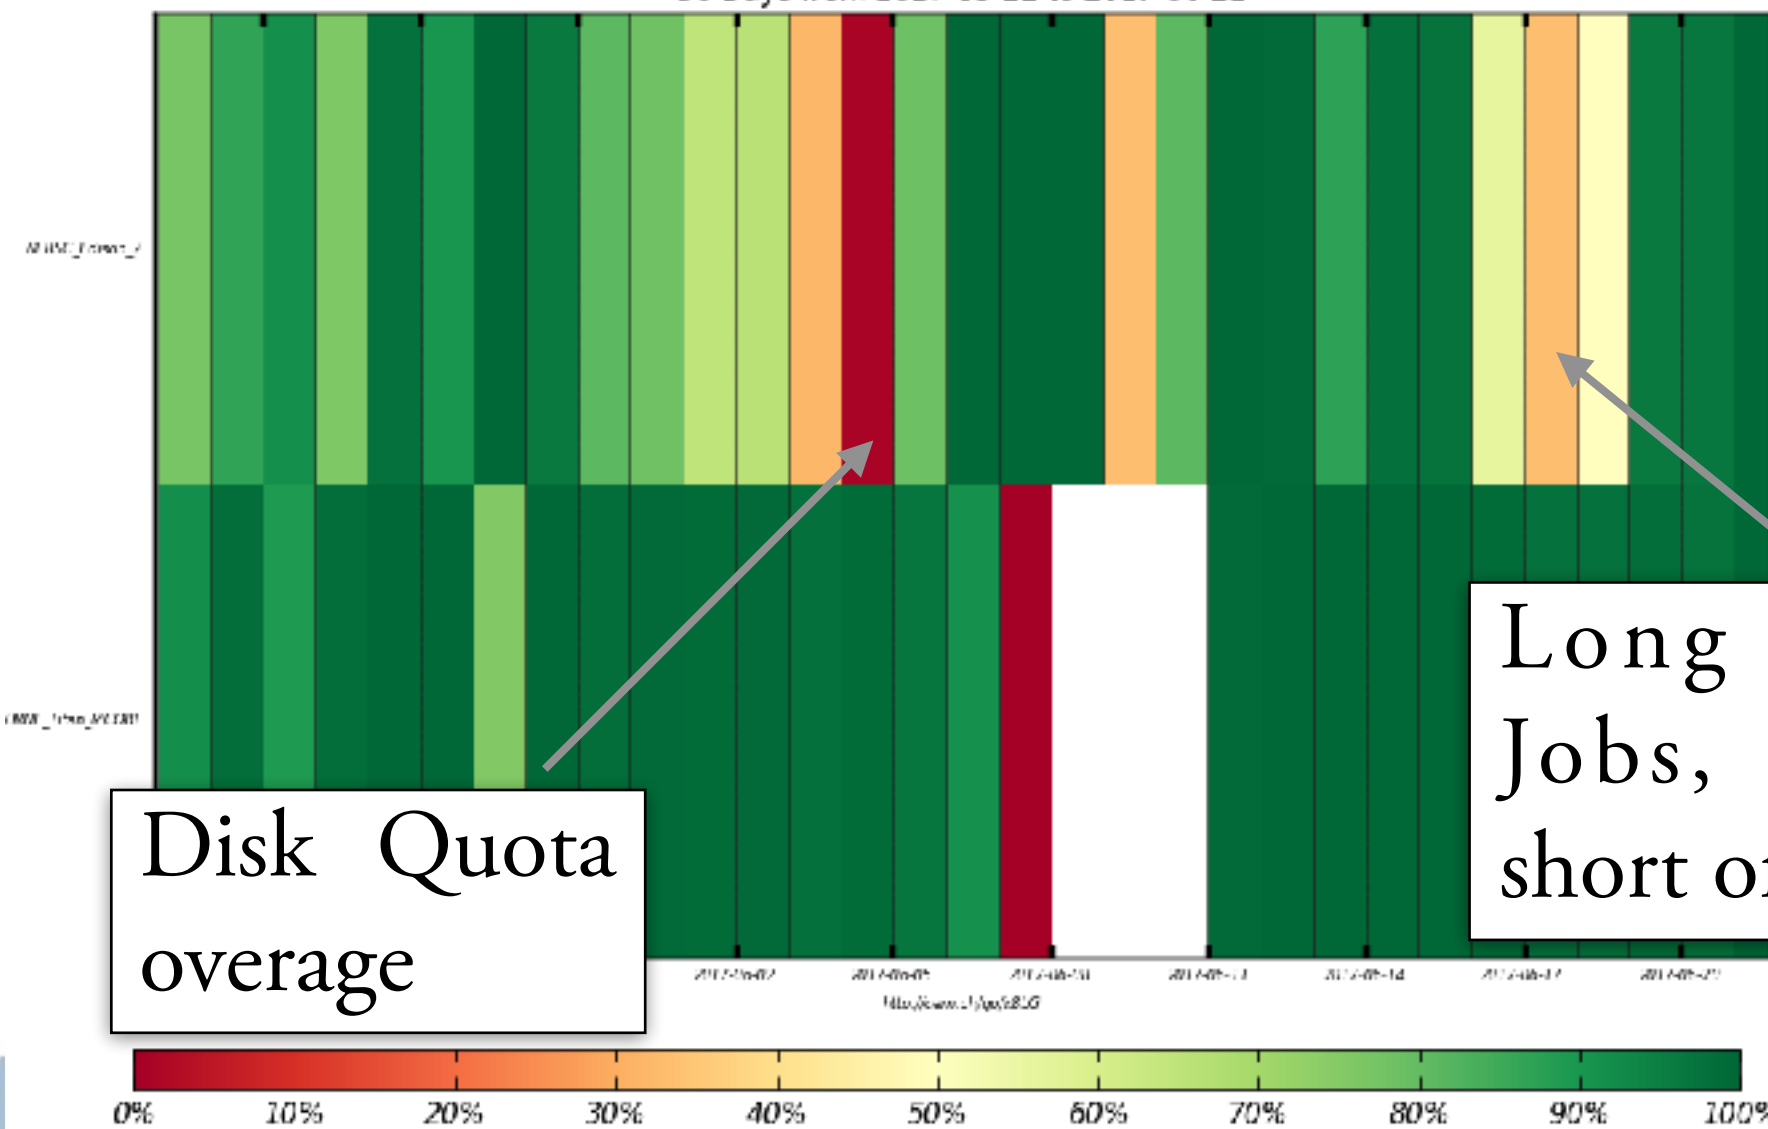
dashboard

CPU consumption Good Jobs in seconds



dashboard

Efficiency based on success/all accomplished jobs
30 Days from 2017-05-21 to 2017-06-21




Disk Quota
overage

Long Running
Jobs, with too
short of wall-time


dashboard

Slots of Running Jobs



Cori Phase-2: KNL Validation

- ▶ Validation jobs for KNL is done
- ▶ Waiting on the validation analysis to be done.



Cori Phase-2 is the new Intel Xeon Phi (KNL) machine. Currently, free to use. #5 on Top500

HPC PanDA Integration with Harvester

- ▶ Doug has Harvester running on Theta using Mini-pilots.
 - ▶ Everything working from stage-in to stage-out
- ▶ Taylor has completed Yoda testing on Theta and Cori end-to-end without Harvester.
 - ▶ Next step is to test being launched by Harvester



ATLAS Software Distribution Strategy for non-standard sites

- ▶ Doug and I have been working to reduce the myriad of strategies used by different sites without native CVMFS access on their worker nodes.
- ▶ The idea is to have ideally 1, maybe 2, solutions that work everywhere.
- ▶ Stratum-R
 - Uses Globus Online to copy CVMFS repo contents (file-by-file) from one place with CVMFS access to another place without CVMFS access
 - Used by Blue Waters and others
 - Doug tested this at NERSC and ALCF.
 - NERSC's quota system tracks bytes & inodes, users start with 10M inodes, we quickly filled that up, NERSC gave us 30M, again we filled that up. They refused to give us more. We estimated about 90M would be needed.
 - ALCF does not track inodes and doesn't care, but we'd like a common solution so we stopped Stratum-R testing at ALCF as well.
 - Dave is working on an upgrade that uses images (like .iso) which removes the inode issue.
- ▶ CVMFS-preload + Parrot
 - Doug started running the CVMFS-preload at NERSC to build up the cache
 - Taylor worked on compiling Parrot for the Cori compute nodes. This exposed quite a few bugs in Parrot so we got a Parrot developer on Cori and they built Parrot.
 - Initial tests show Parrot running and working on Cori-Haswell & Cori-KNL.
 - Parrot is about 2x slower in initialization than when using pacman installed releases and databases.
 - Going to test scalability at NERSC, if this goes well, then test at ALCF is next.



ATLAS Software Distribution Strategy for non-standard sites

- ▶ Doug and I have been working to reduce the myriad of strategies used by different sites without native CVMFS access on their worker nodes.
- ▶ The idea is to have ideally 1, maybe 2, solutions that work everywhere.
- ▶ Containers
 - Wei is producing containers at NERSC that we are also going to test.
 - While ALCF will not support shifter, there is indication that they might support Singularity so this could be a common solution
 - Doug points out, then we need a sustainable way to automatically produce containers
- ▶ Native CVMFS?
 - I met with ALCF operations people to lay out our normal grid site services and show the lengths to which we have to go when they don't exist, then I asked them. Can you install FUSE so we can run CVMFS? They are not deliberating. I'll keep pushing, but if they install FUSE on the Theta compute nodes, it makes life much easier.
 - In addition, I would push for a squid server or two on site. They seem less squeamish about that than FUSE.