



DUNE Production Infrastructure

Ken Herner, Ivan Furic

BNL Meeting

August 2019

What is in this talk

- We'll discuss the current job submission, tracking, and management system that the production group is using
- Expansion / new site commissioning
- New features
- Future requirements

Current setup

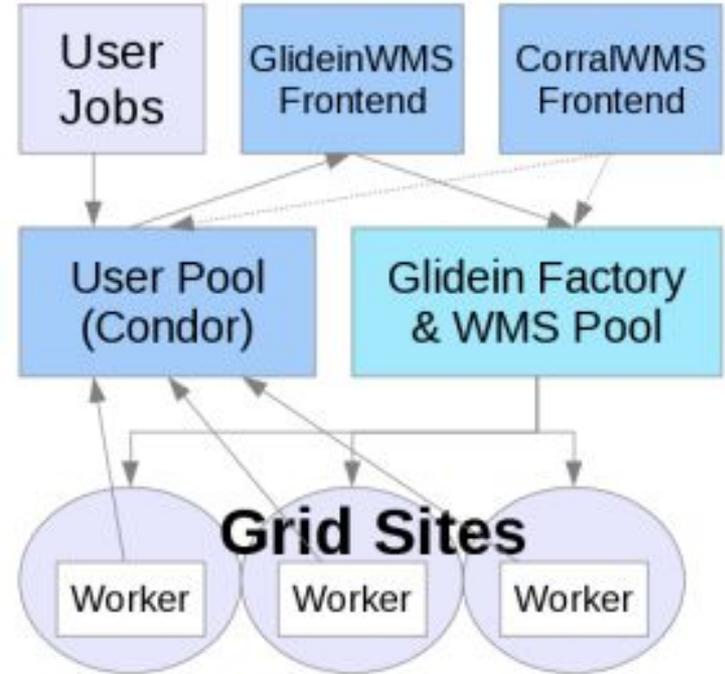
- Job submission is via POMS (see POMS talk for details)
- Resource provisioning is with GlideinWMS
- Submit everything onsite and offsite to both SL6 and SL7 (DP DC3 was SL7-only)
- Copyback is generally to FNAL dCache, but have copied back to other places as well (IN2P3)

Aside: GlideinWMS

GlideinWMS is a workload management system sitting on top of the local batch system
Extensively used by CMS, IF expts, Open Science Grid projects, IceCube, LIGO, etc.

Adding DUNE support is *trivial* if you support any of these other groups

Has long-term support



Resource expansion/site commissioning

Well-established procedures for adding new sites to the ecosystem

Local CE-agnostic nature of gwms helps (currently supports HTCondor, CREAM, ARC)

Typical integration time is <1 week (in fact as short as 2 hours!)

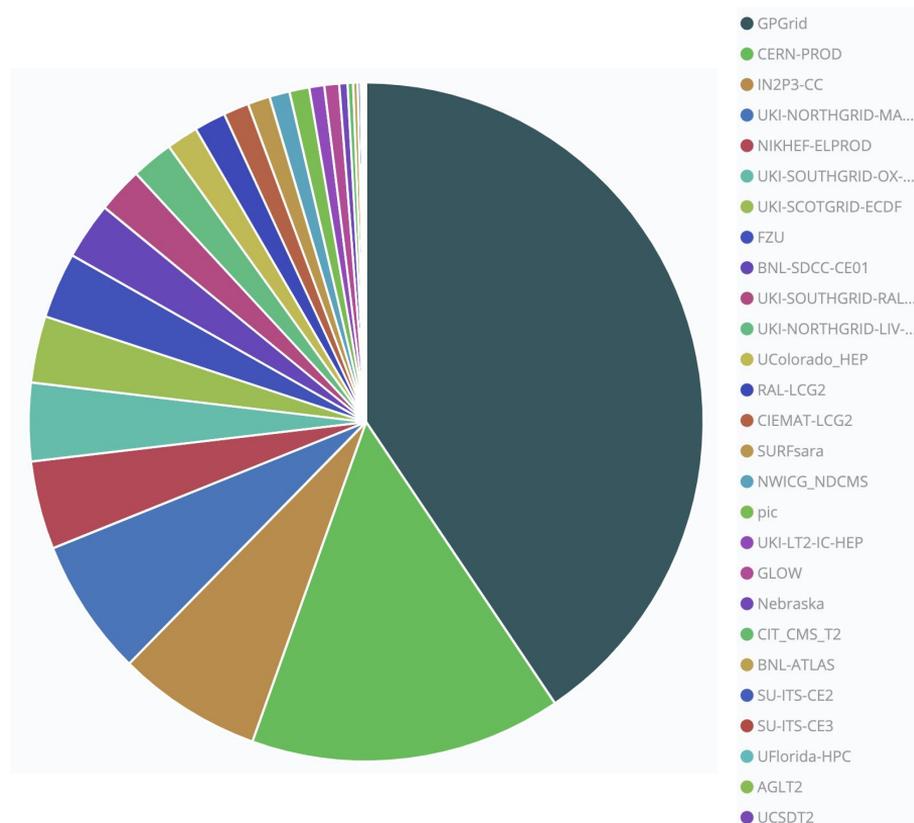
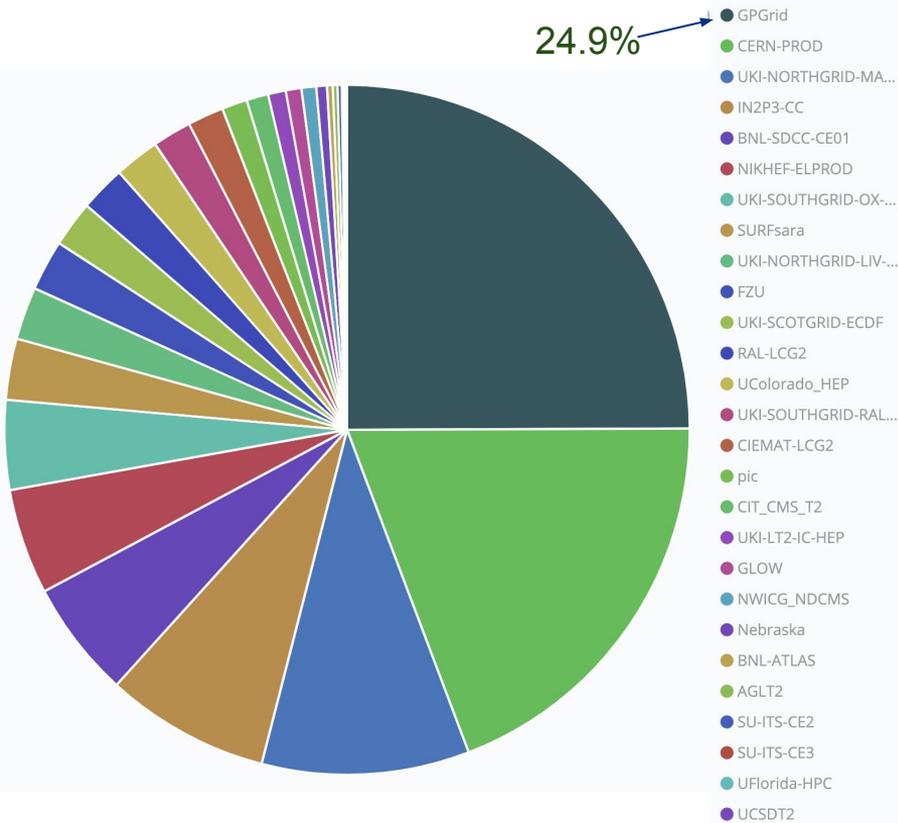
Main sticking point now is dealing with time zones

Max allowed run time, number of cores, memory, etc. can be tuned for each site

Multiple sites added in the past 1.5 years have worked flawlessly starting from the very first test job:



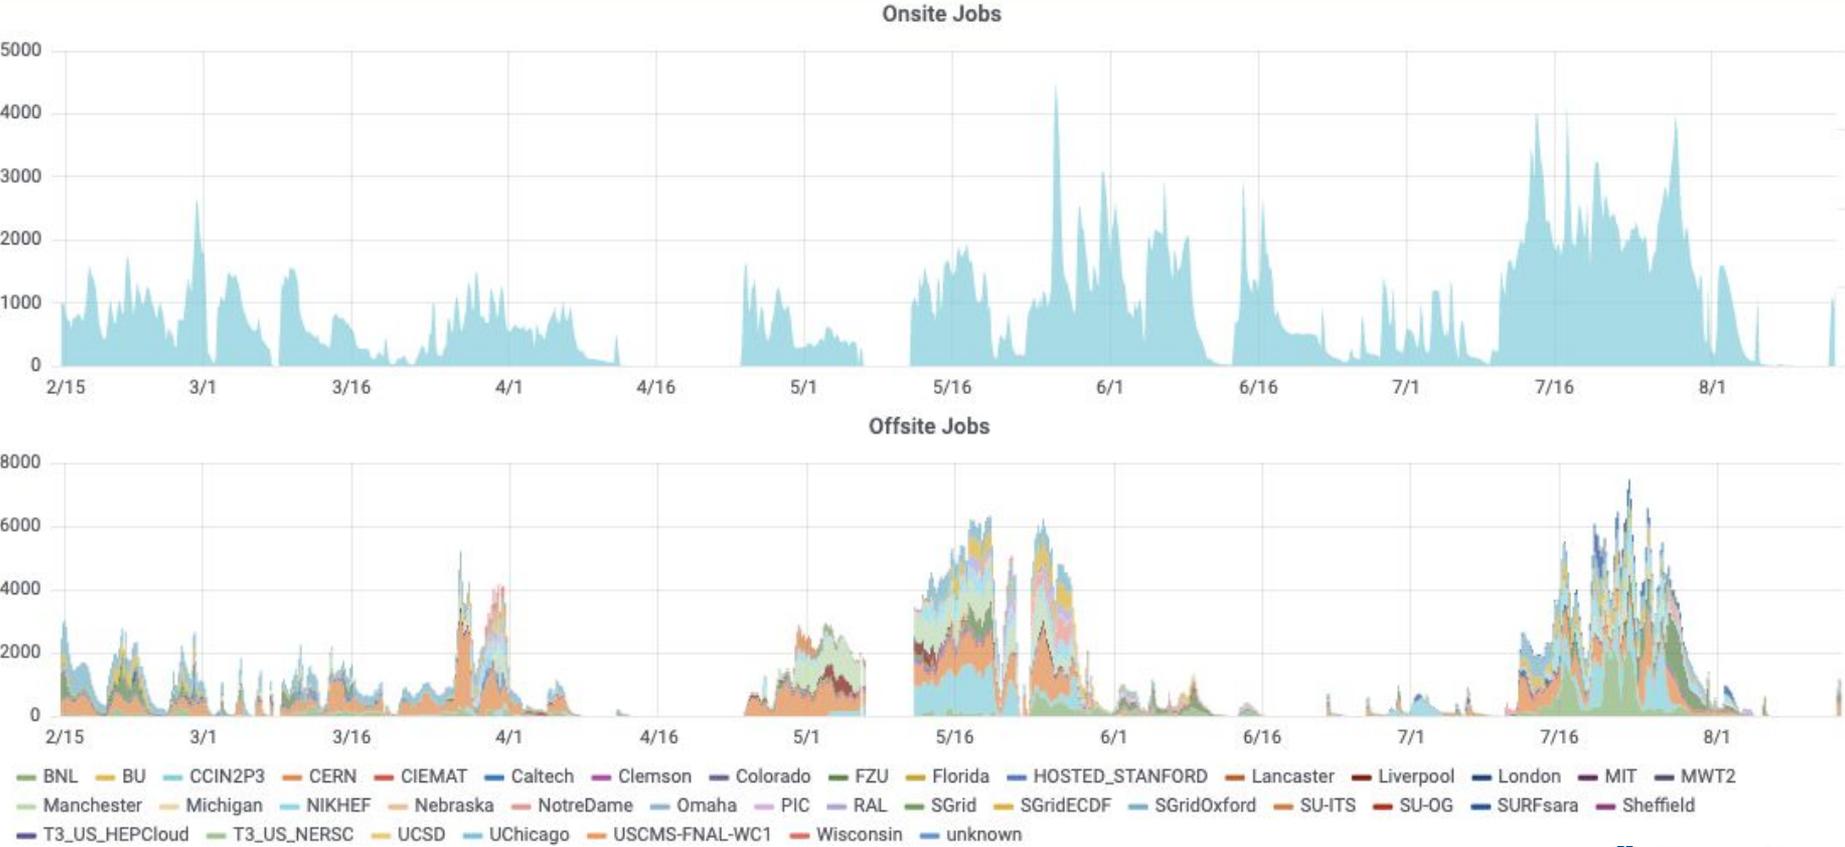
Resource Distribution, past 6 months



All successful production jobs, count only

All successful production jobs, weighted by wall time

Resource distribution, last 6 months



Recent feature additions

Singularity support via GlideinWMS

Glidein wrapper handles singularity invocation; user only supplies an image. Should work at all sites/Singularity

Automatically release jobs that exceed resource requests

Desired long-term features

Support for MPI jobs (not many such workflows now, but they're inevitable). Multi-node pilots?

Support copying outputs back to local SEs; Rucio handles replication (possible now but a bit of a manual process)

Support for “pipelined” resources

Summary

Production workflows make extensive use of well-tested technologies

Very easy to add new resources as they arrive

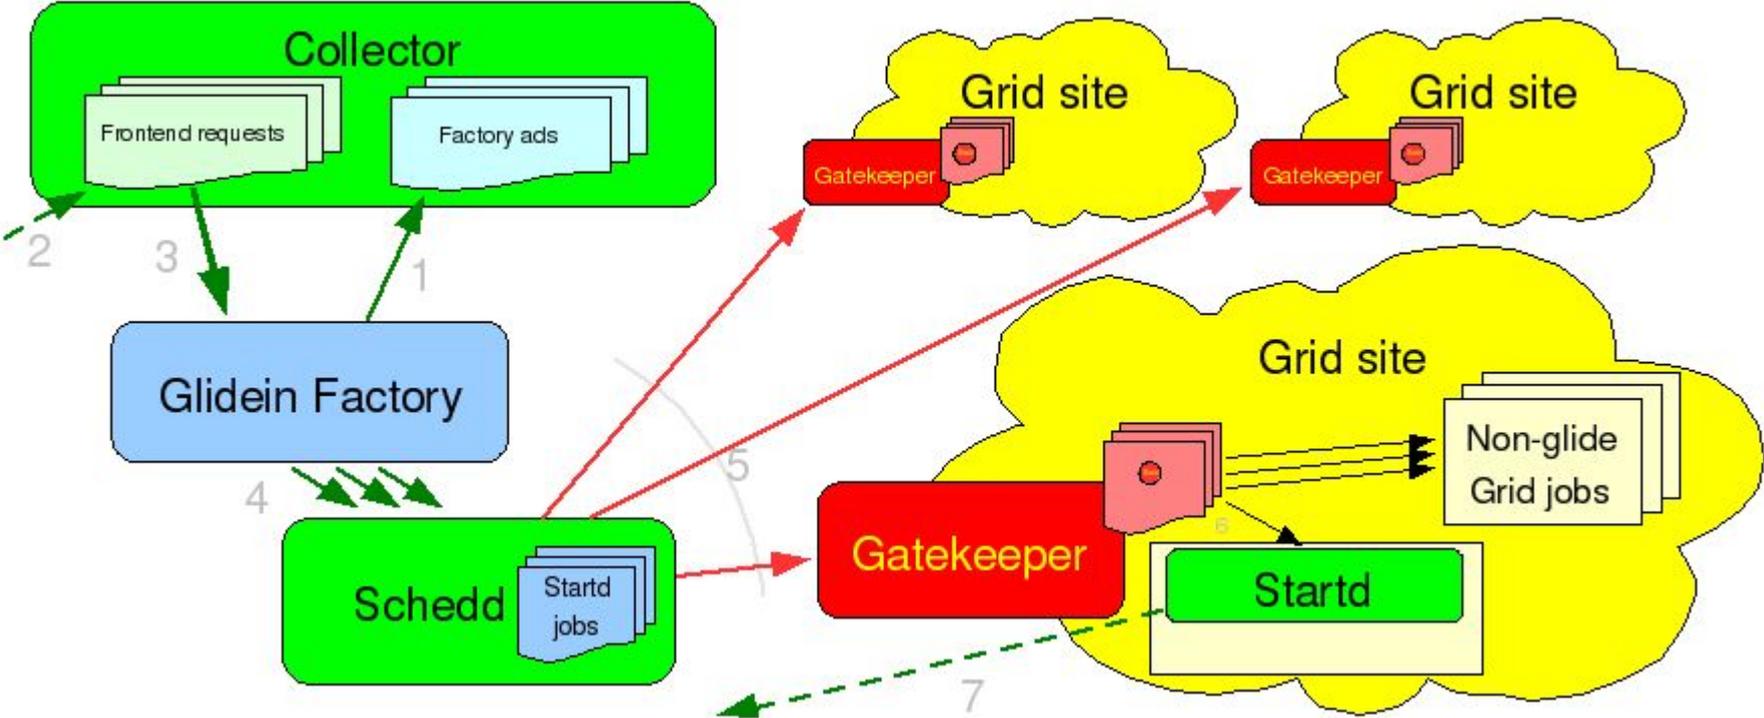
Long-term support of technologies behind the scenes exists

Future workflows will be more heterogeneous than today's.

Future systems must be flexible enough for DUNE's needs!

BACKUP

GWMS Factory



GlideinWMS Frontend

