





FNAL Perspective

Dirk Hufnagel (Fermilab) for Fermilab/HEPCloud
HPC Strategy Meeting Jan 30 - 31, 2025



- FNAL is not just CMS in terms of Computing, CMS is about half of the deployed resources right now (but this is expected to increase with HL-LHC)
- There is DUNE (strategically arguably more important for FNAL, even if smaller in terms of Computing resource impact)
- There are other experiments (SBND, Nova, Mu2e, g-2, etc)

- Integration of external resources should be useful/usable by all experiments
- CMS potentially will be able to get the biggest absolute benefits of such integration (depending on our ability to secure appropriate HPC allocations)
- Other experiments might benefit more in relative terms (DUNE for instance due to stronger reliance on GPU resources)

- HEPCloud (vision): a scientific gateway to resources beyond local and grid worker nodes, expanding into high performance computing (HPC) centers and the cloud.
- Started about 10 years ago
- Originally also envisioned automated decision making about what type of resource to use (seamless extension of FNAL into external resources). This is part of the code base but was never really implemented by the experiments, they instead rather control directly where they want to sent work.
- Current status is functional cloud integration (for special use cases) and mature HPC integration (more or less continuous use by CMS for the 'easy' US HPC like [ACCESS](#), [TACC Frontera](#) and [NERSC](#), and less continuous use by others).
- Integration of DOE [Leadership Class Facilities](#) (LCFs) has been an R&D area for years (CMS, DUNE and others), but is difficult due to their restrictions (no outbound network from batch nodes) that block a lot of things we assume to 'just work'.

- Functional Cloud integration is strategic, but we don't expect to use it much due to costs.
 - But useful to get access to specialized hardware (FPGA, Tensor chips, latest GPUs, ARM)
- Integration of 'easy' HPCs is stable, we developed a toolbox that so far seems to provide what we need. Not much growth potential here (only when HPC systems are replaced). Caveat: New NERSC-10 should be 'easy' to integrate and be roughly at Exascale (but most of that coming from the GPU partition like with Perlmutter), will be a very interesting machine.
- To grow beyond this, we need to be able to target (Exascale and post-Exascale) LCFs
 - Difficult integration (usual HEPCloud toolbox only partially applicable)
 - Need GPU workflows or use them for 'just' ML/AI (problem for experiments)
- Something new that **might** make LCF integration easier in the future
 - DOE IRI (Integration Research Infrastructure) initiative (next slide)

- Vision is an integrated research infrastructure across all DOE resources (current focus definitely on NERSC and LCF, long term might expand)
 - Allowing user to easily move between facilities and also run multi-facility workflows
 - Mostly about software and middleware to make that happen
 - Hopefully common solutions and not everyone doing their own thing
 - Also about facility policies to allow that middleware to work
- Currently mostly focused on experiments with realtime processing needs (DUNE has a compelling use case here), but the (HL-)LHC type of science is also (supposed to be) represented (“long term campaign”)
- Has triggered a lot of conversations
 - Want to make sure our type of science and use cases are well represented
- Top down approach means there is a vision, but little yet at the technical/policy level to support that vision (and seemingly also not much dedicated funds for now)
 - My personal take: Don’t expect a lot of progress fast here

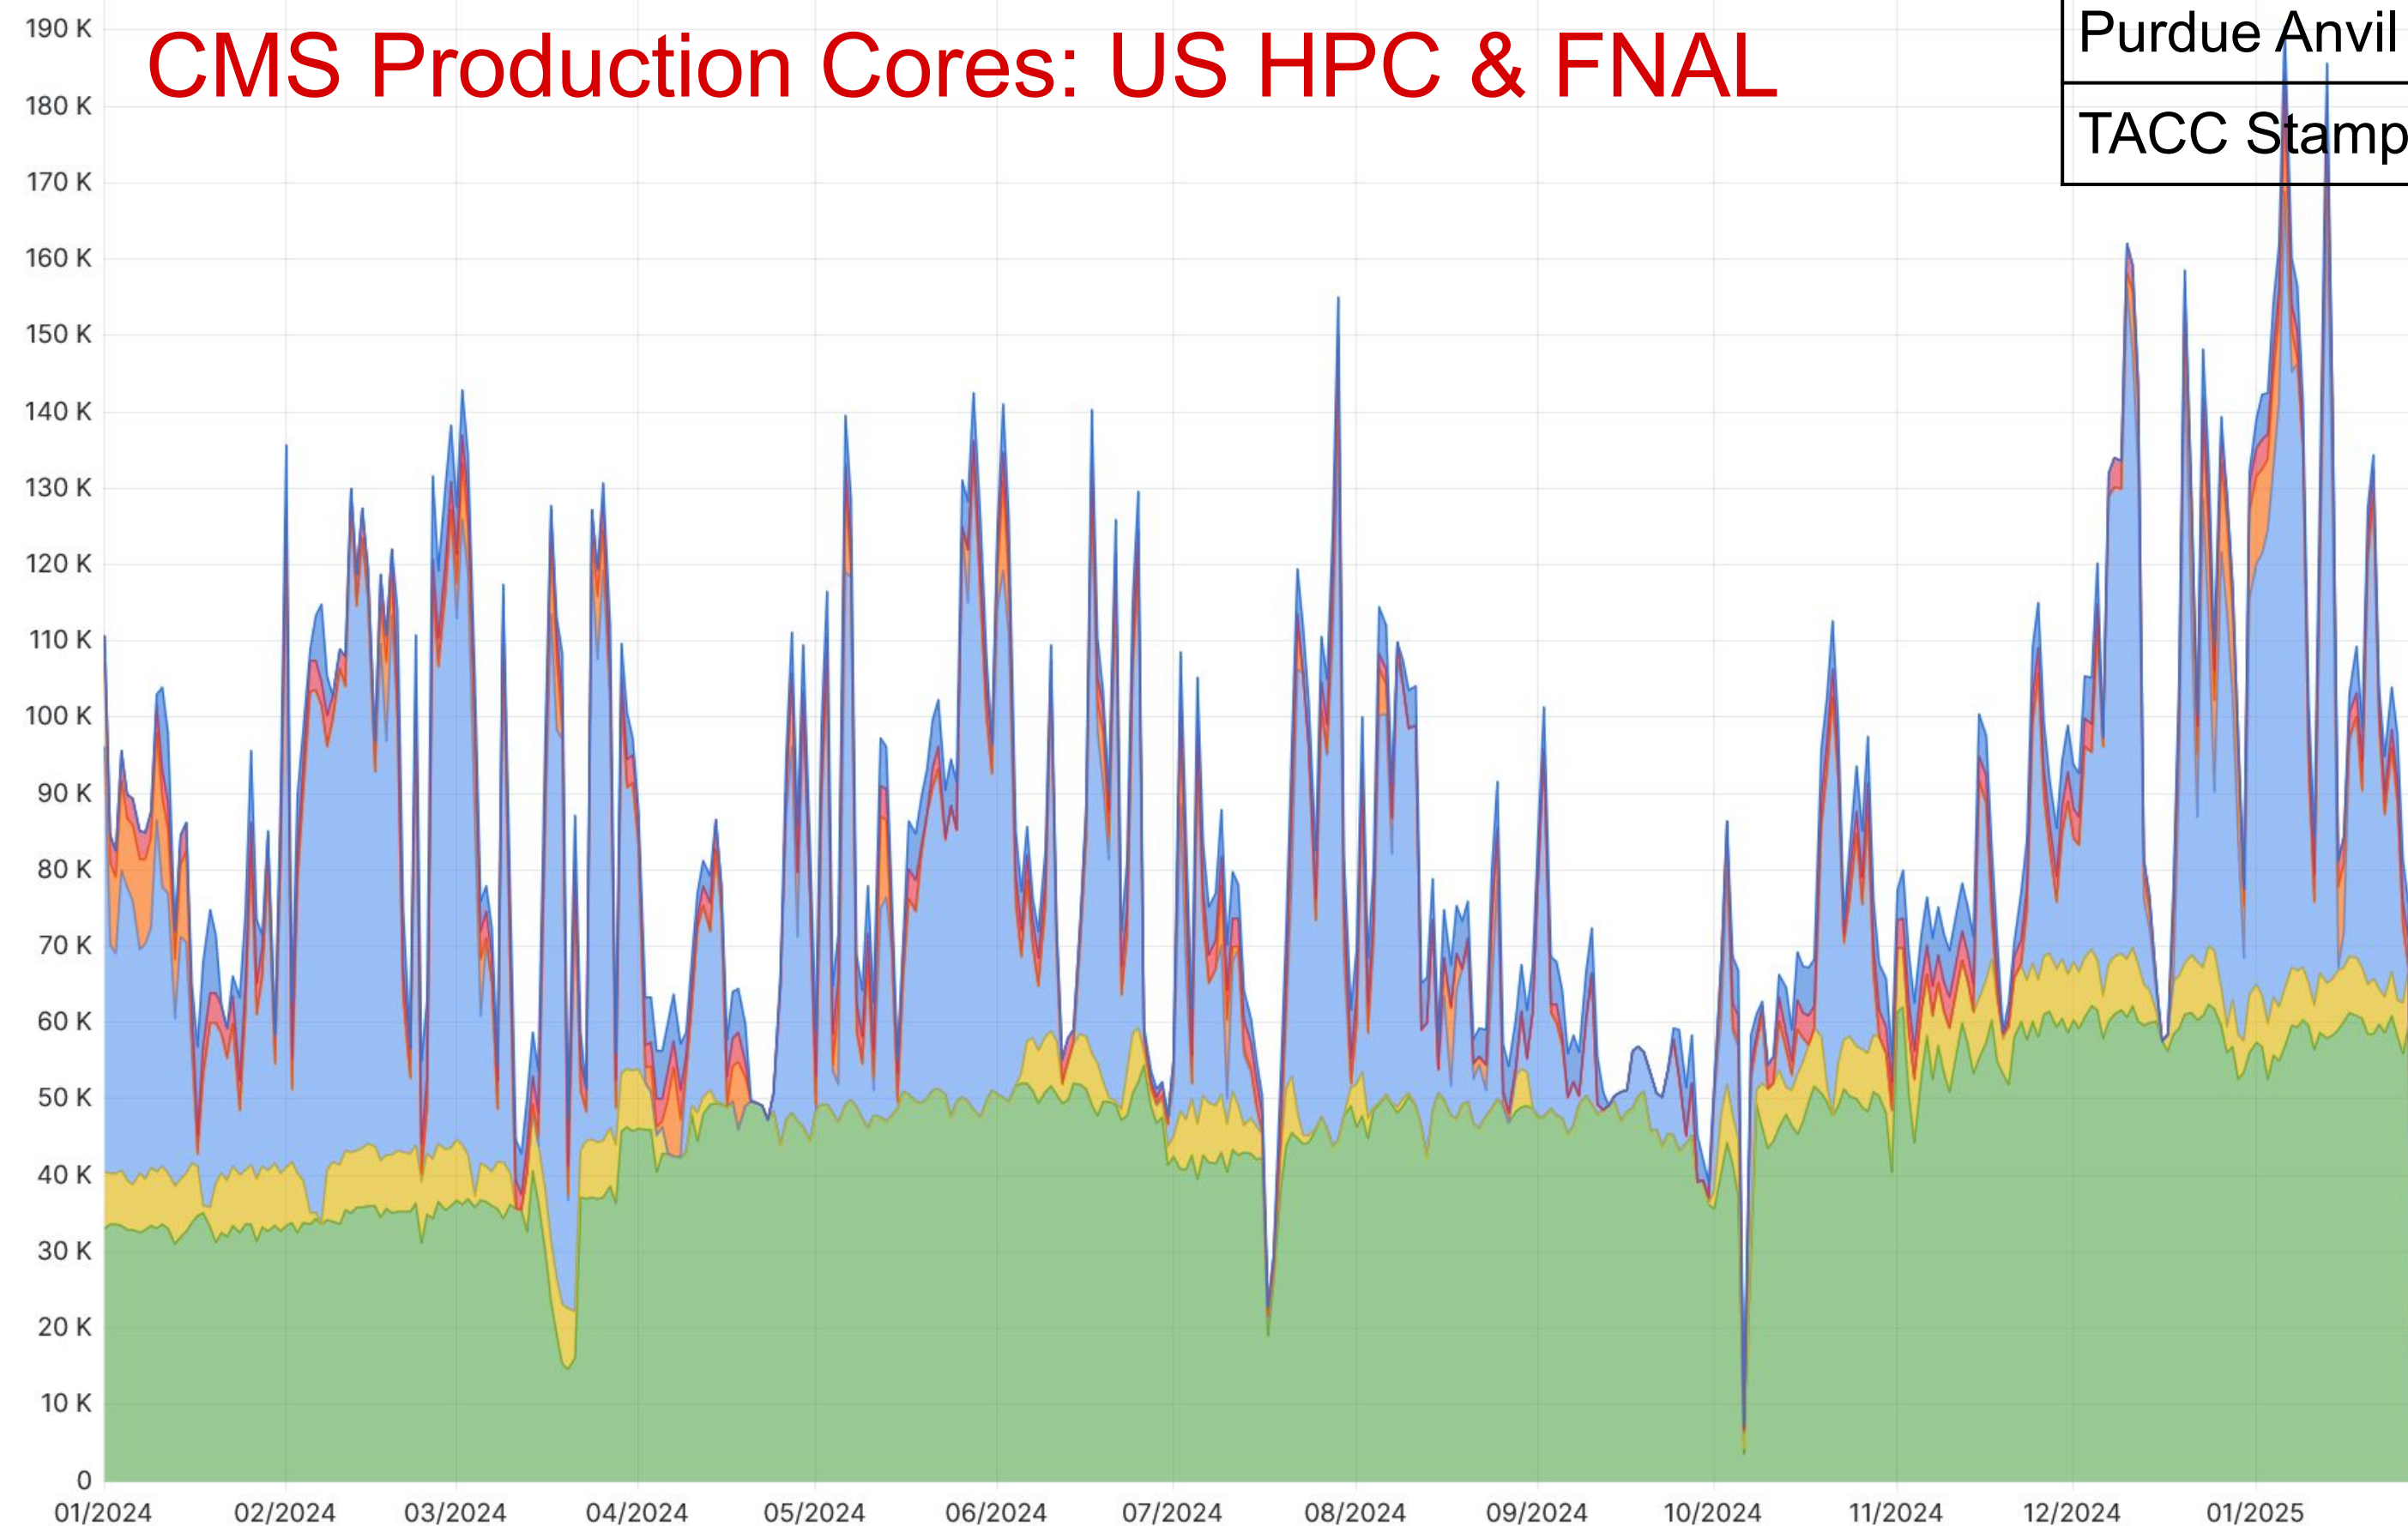
BACKUP

CMS HEPCloud Usage - Last 12 months

- Comparable to FNAL CMS T1 in scale (but not continuously available/used, utilization fluctuates)

HPC Resource	Allocation (core hours)	Allocation Period	%Used*
NERSC Perlmutter (DOE)	77M CPU	Jan 2025 - Jan 2026	13%
TACC Frontera (NSF)	36M	Jun 2024 - May 2025	75%
PSC Bridges-2 (ACCESS)	23M	Oct 2024 - Sep 2025	39%
SDSC Expanse (ACCESS)	23M	Oct 2024 - Sep 2025	44%
Purdue Anvil (ACCESS)	23M	Oct 2024 - Sep 2025	94%
TACC Stamped3 (ACCESS)	1M	Oct 2024 - Sep 2025	0%

CMS Production Cores: US HPC & FNAL

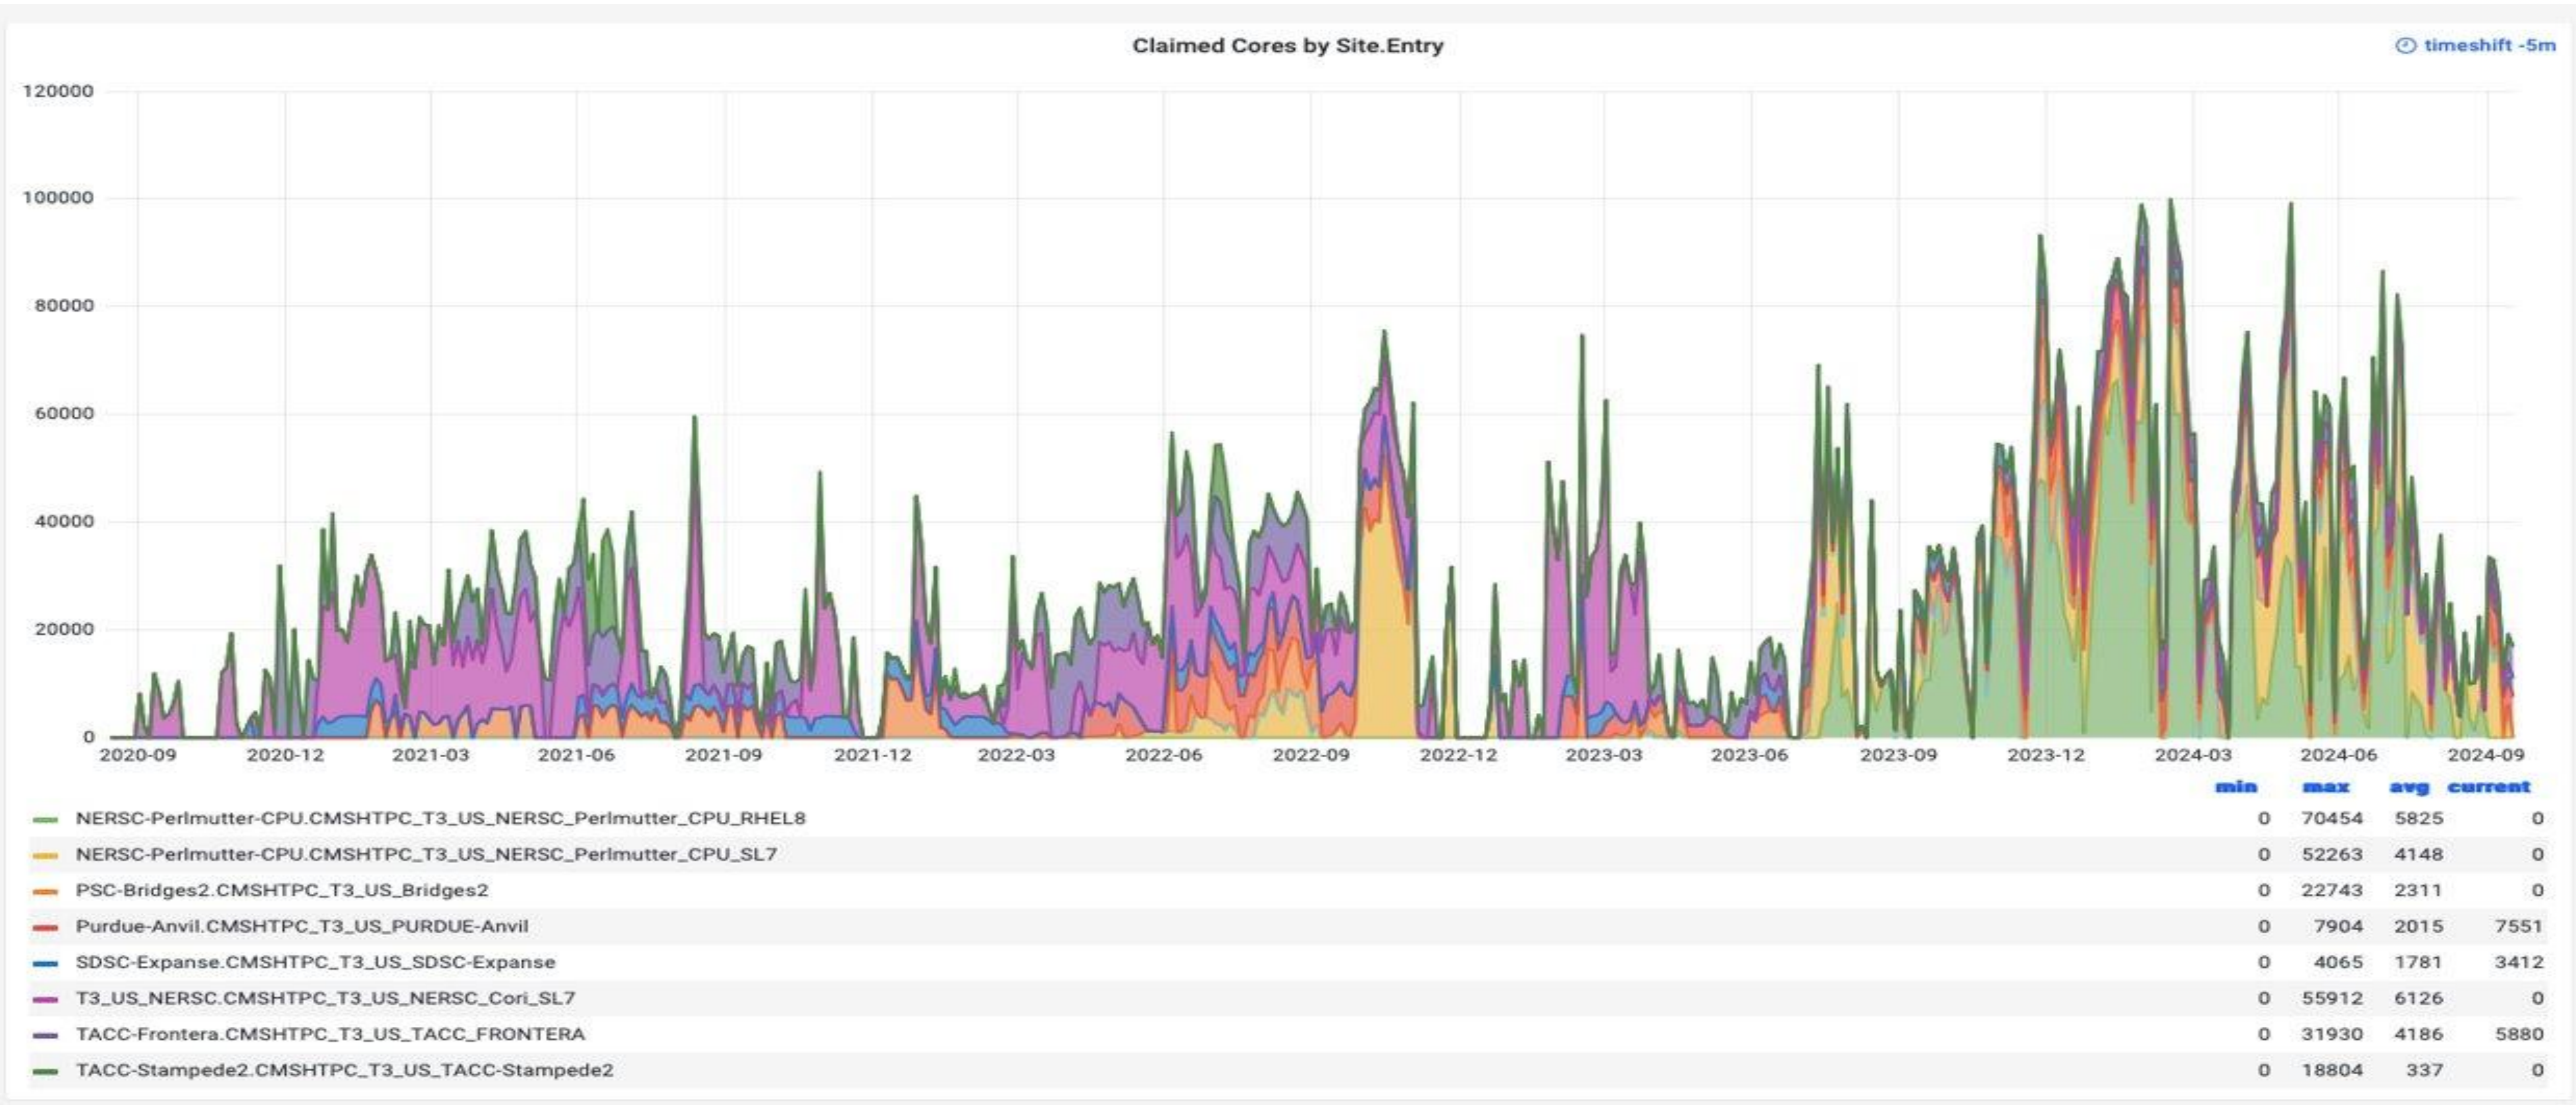


Currently Active Allocations

Name	Mean
T1_US_FNAL	45.9 K
T3_US_Anvil	4.31 K
T3_US_NERSC	26.7 K
T3_US_PSC	2.15 K
T3_US_SDSC	2.68 K
T3_US_TACC	4.08 K
TOTAL	85.9 K

CMS HEPCloud Usage 2020-Sep24

- Average of ~27k cores DC over last 4 years. Peaks of >100K cores!



DUNE/FIFE HEPCloud Usage 2020-Sep24

- Usage shown from NOvA, Mu2e, and DUNE experiments
 - Not shown is direct SLURM submission from other smaller projects

