(proto)DUNE computing update

Ken Herner (with material from Andrew Norman and Heidi Schellman)

WLCG GDB

17 October 2018

Reminder: Plan is to put a future huge LAr detector “DUNE” in the Homestake Gold Mine

Underground Laboratory

DUNE Far Detector site

- Sanford Underground Research Facility (SURF), South Dakota
- Four caverns on 4850 level (~ 1 mile underground)

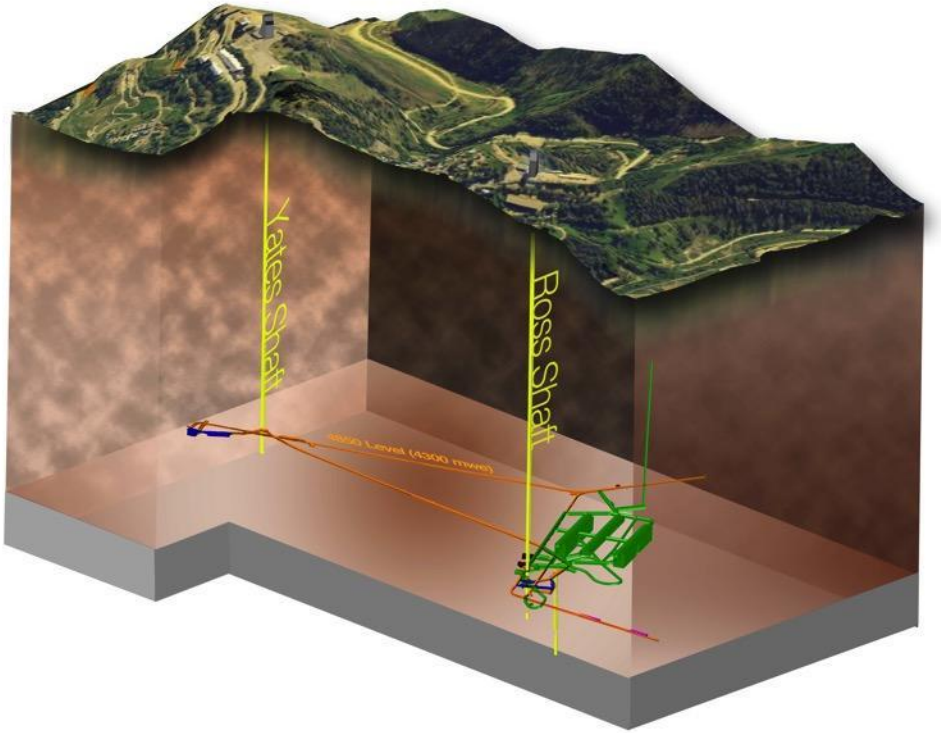

Full DUNE is 4 of these – filled with liquid Argon

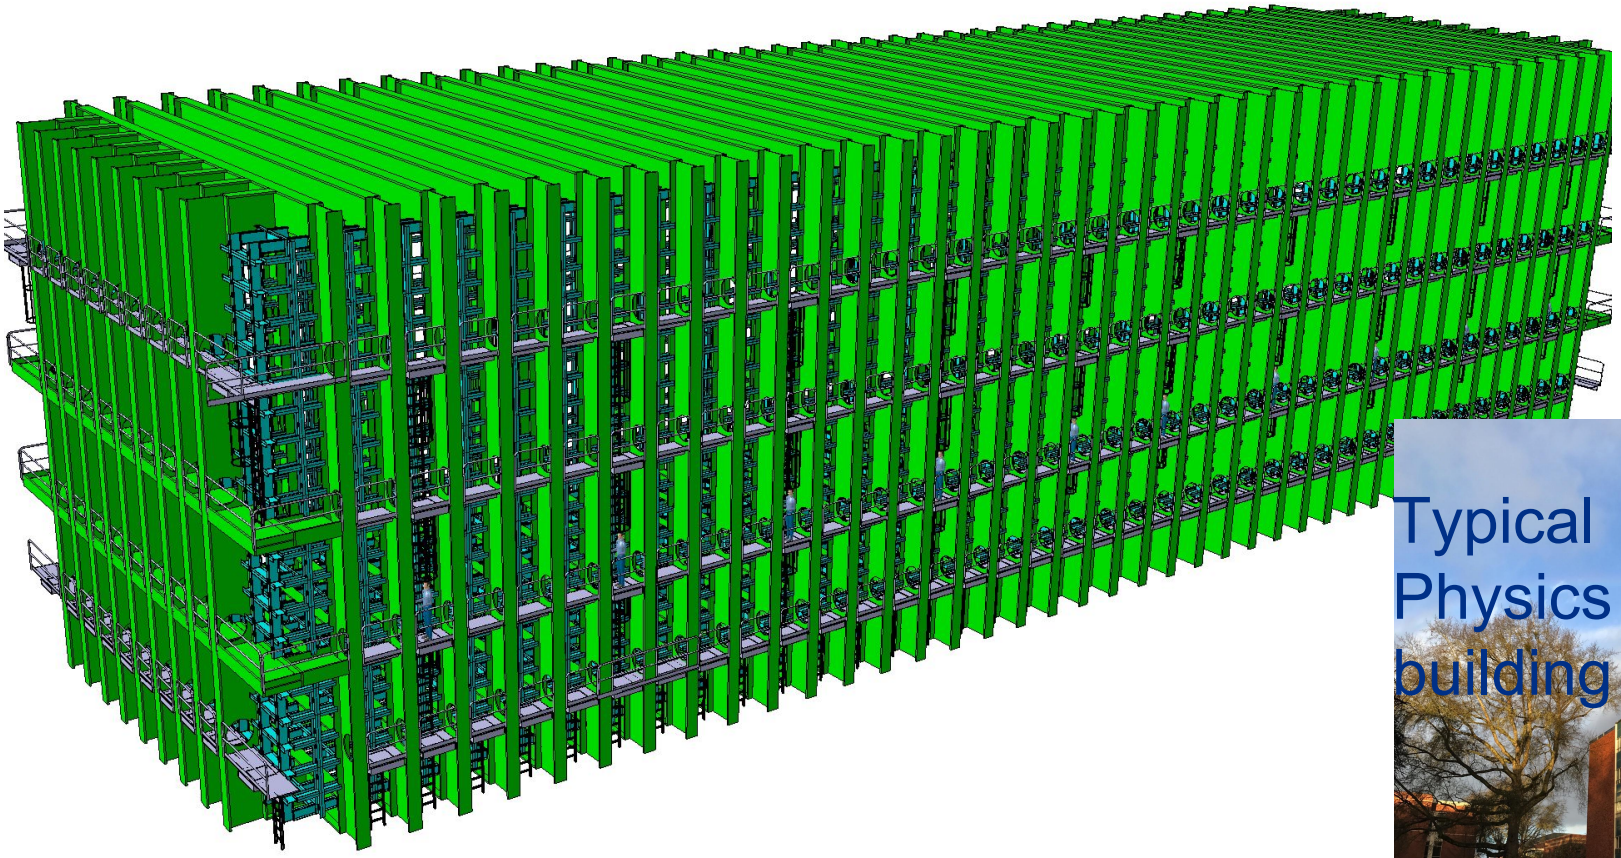

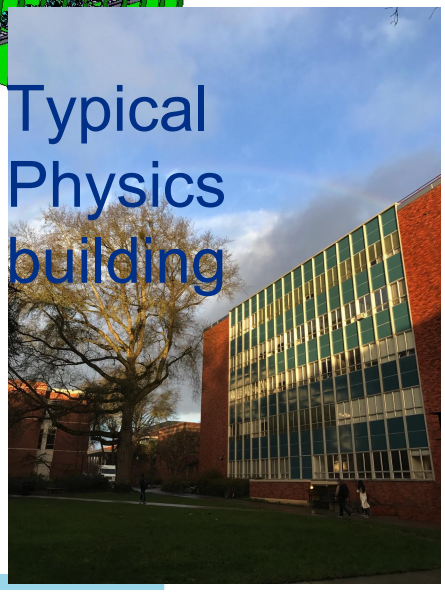

External Dimensions: 19.1m x 18.0m x 66.0m

protoDUNE: the here and now

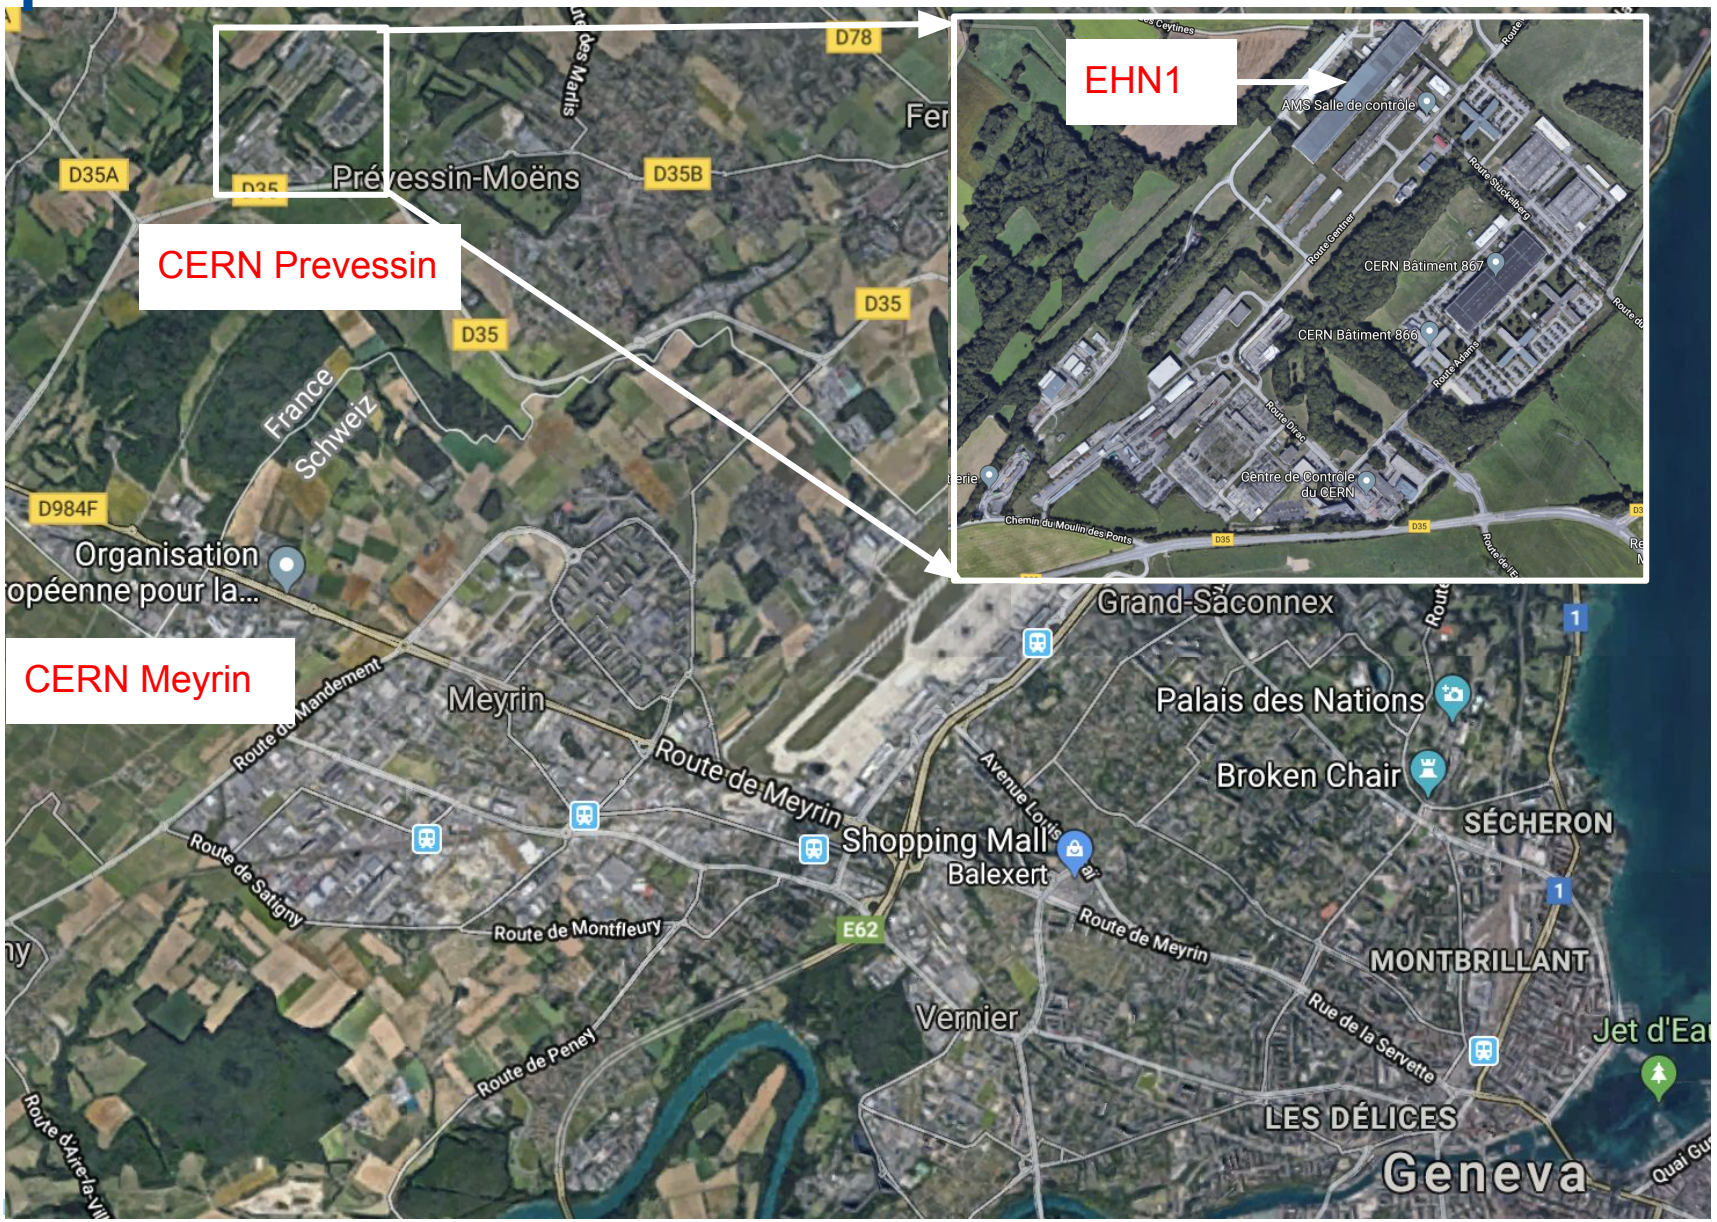

protoDUNE: Inside EHN1

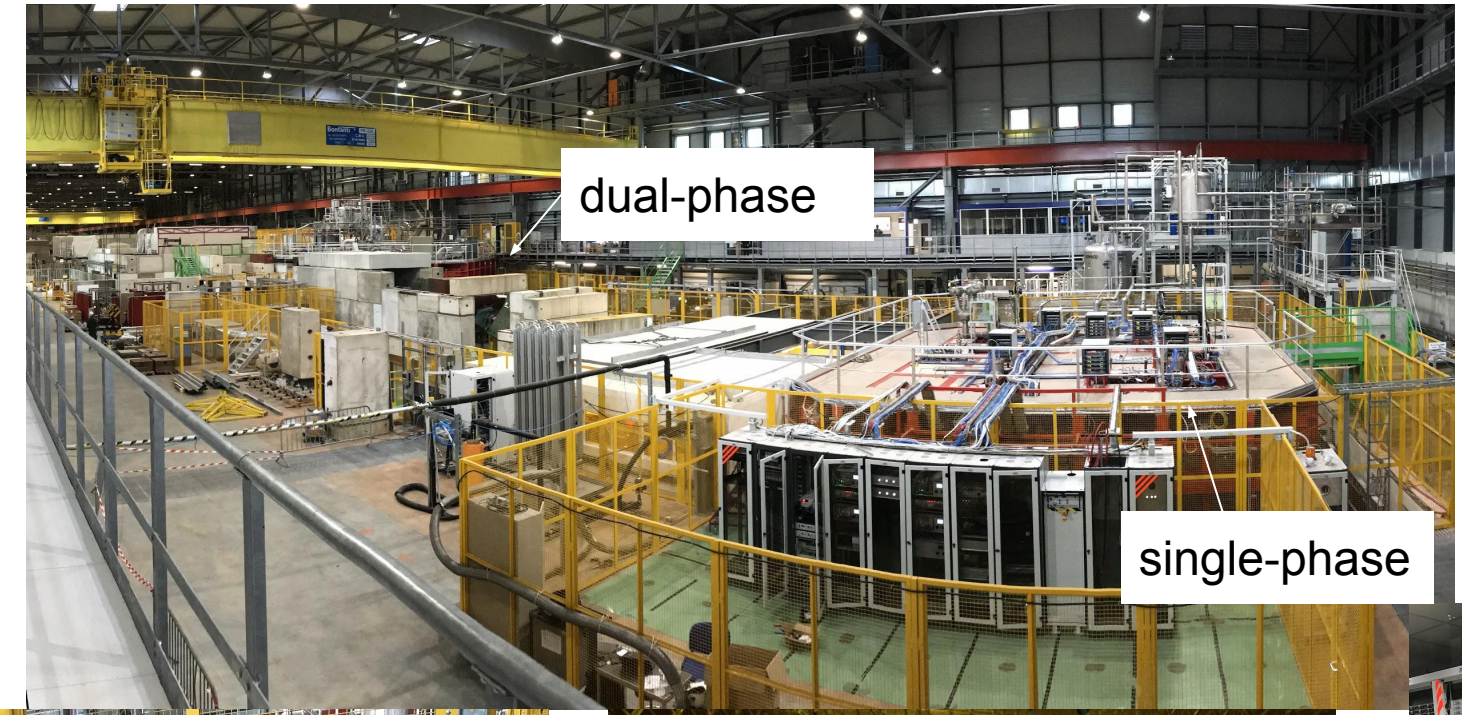

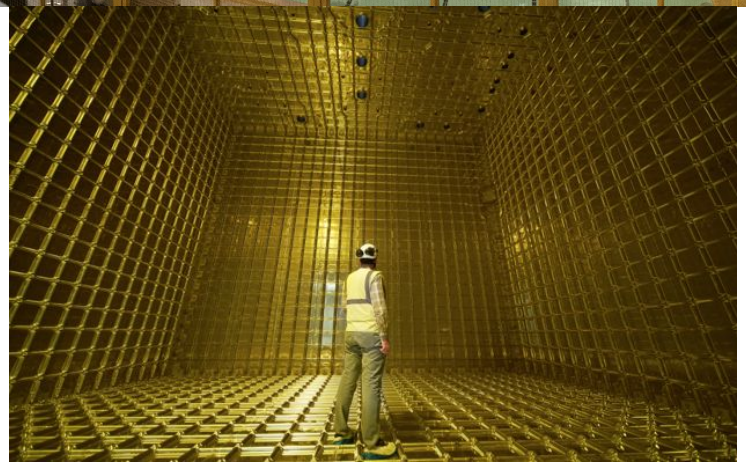

The computing problem

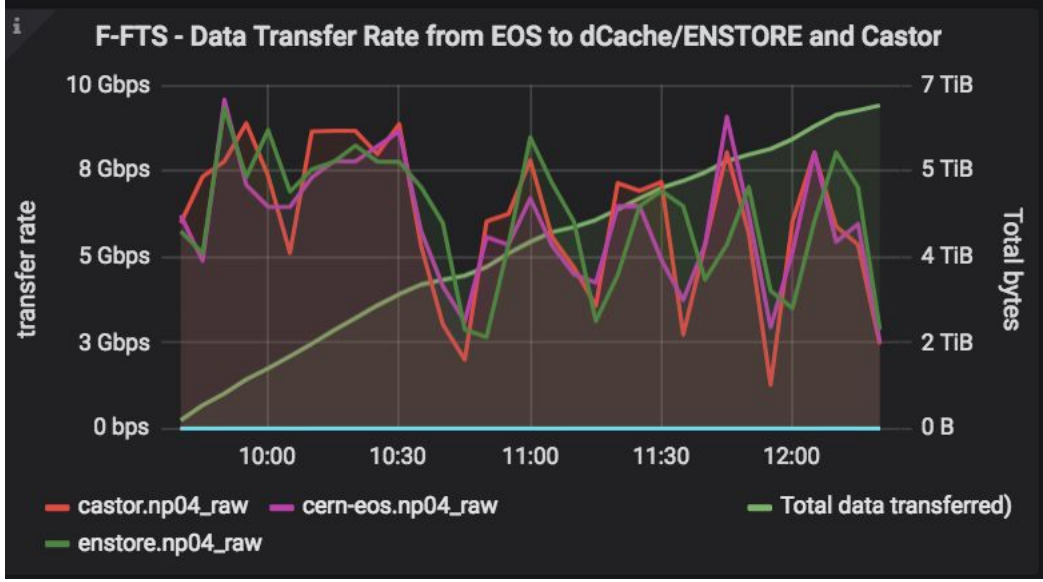

3 hrs of raw data transfers from CERN->FNAL 5 Oct

- ProtoDUNE can write up to 2-3 GB/sec of data when running at 25 Hz beam.
- Expect 2-5 PB of raw data from the 6 weeks of running this fall + more from additional cosmic running + ProtoDUNE-DP
- Big DUNE will write 10-30 PB/year from the far detectors and probably more than that from the near detector.
- Instantaneous rates are already >= CMS/ATLAS.
- Integrated rates in 2026 will be where ATLAS/CMS are now

What is the production group doing?

Two main efforts right now: protodDUNE data keepup processing (first priority) and MC production

Keepup: run raw data decoding and hit reconstruction (not tracking yet; will be added soon).

protoDUNE SP by the numbers

as of 17 Oct 2018

Raw data: 1298 TB
Raw “physics” data: 471 TB
Raw physics data fraction processed: 73.6% (complete through Oct. 10)
Total reconstruction output size: 186 TB

Raw data

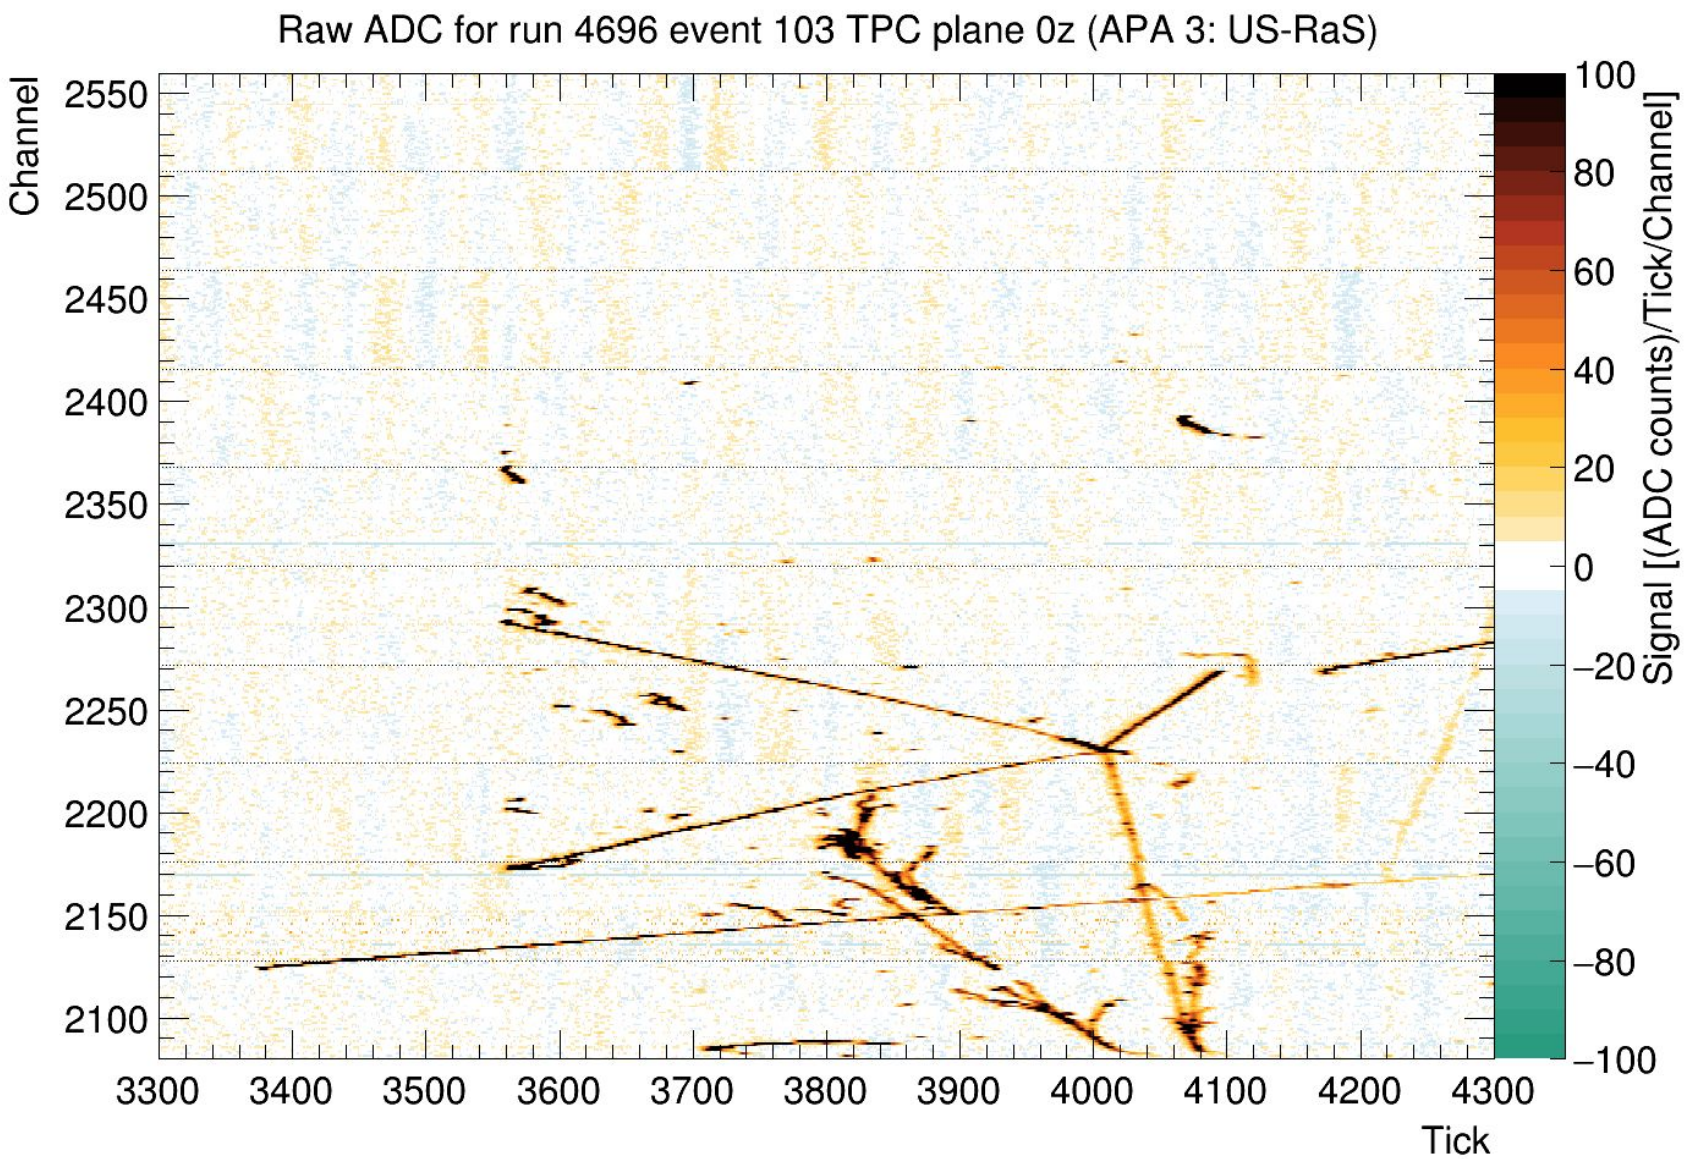
Raw ADC for run 4696 event 103 TPC plane 0z (APA 3: US-RaS)
Channel
2550
2500
2450
2400
2350
2300
2250
2200
2150
2100
3300
3400
3500
3600
3700
3800
3900
4000
4100
4200
4300
Tick
100
80
60
40
20
0
−20
−40
−60
−80
−100
Signal [(ADC counts)/Tick/Channel]

protoDUNE SP Keepup Processing

Processing all kinds of runs (test, noise, physics)

Able to get lots of jobs running, but reconstruction takes a while for physics runs

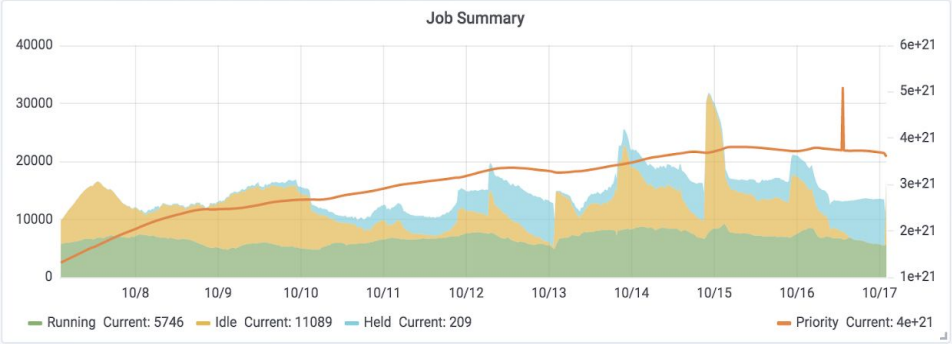

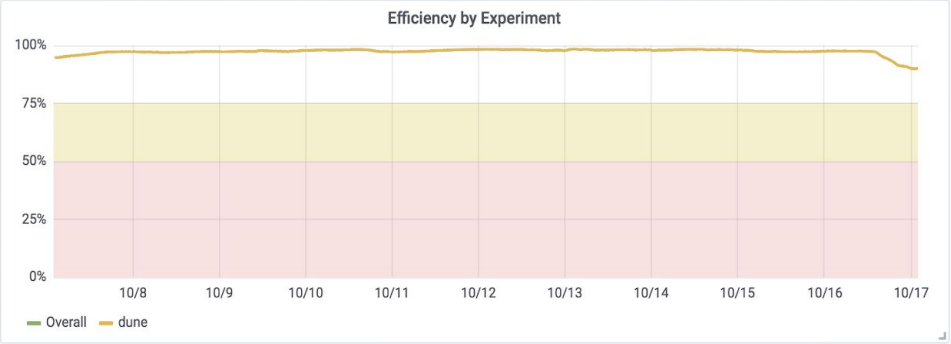

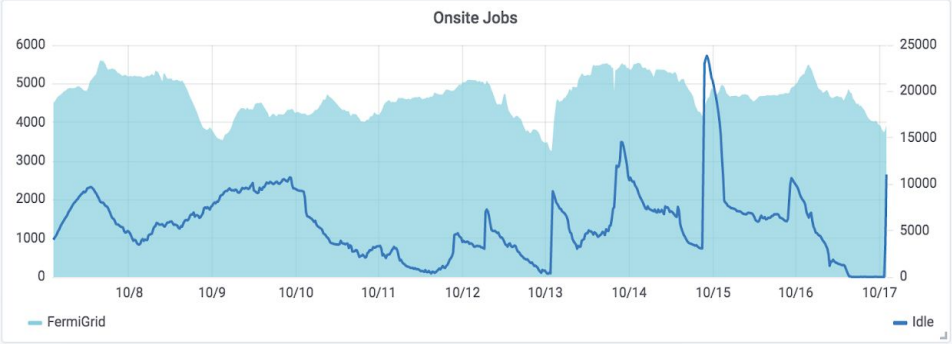

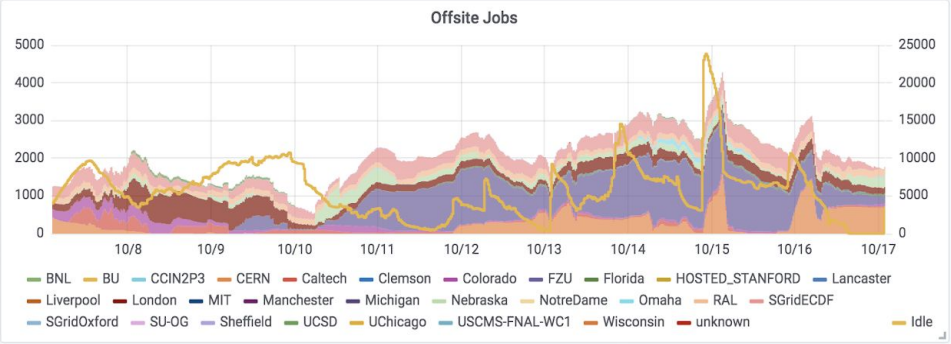

DUNE uses the **larsoft** suite (http://larsoft.org/) for reconstruction and analysis. Common art-based framework for LArTPC detectors (e.g. Microboone, ICARUS). art is a CMSSW fork.

Keepup Processing (2)

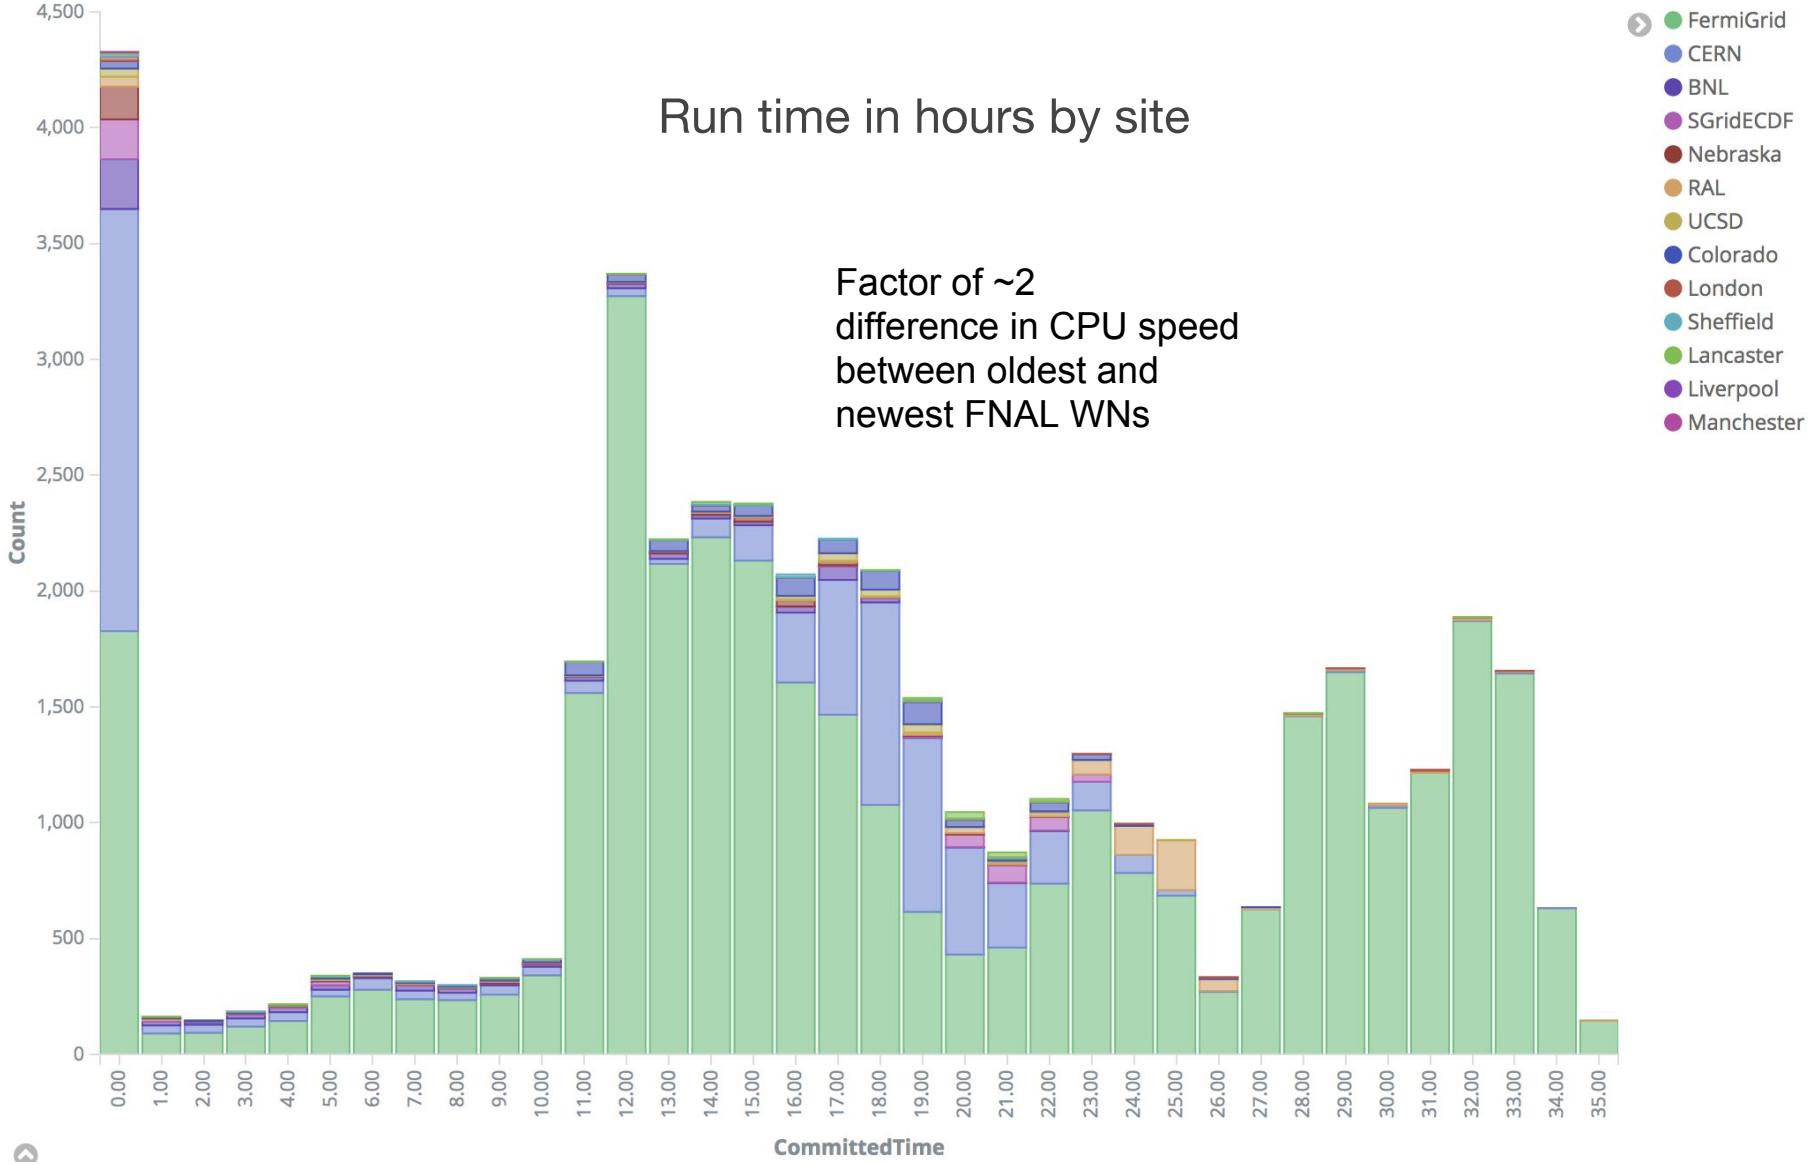

Keepup Processing (3)

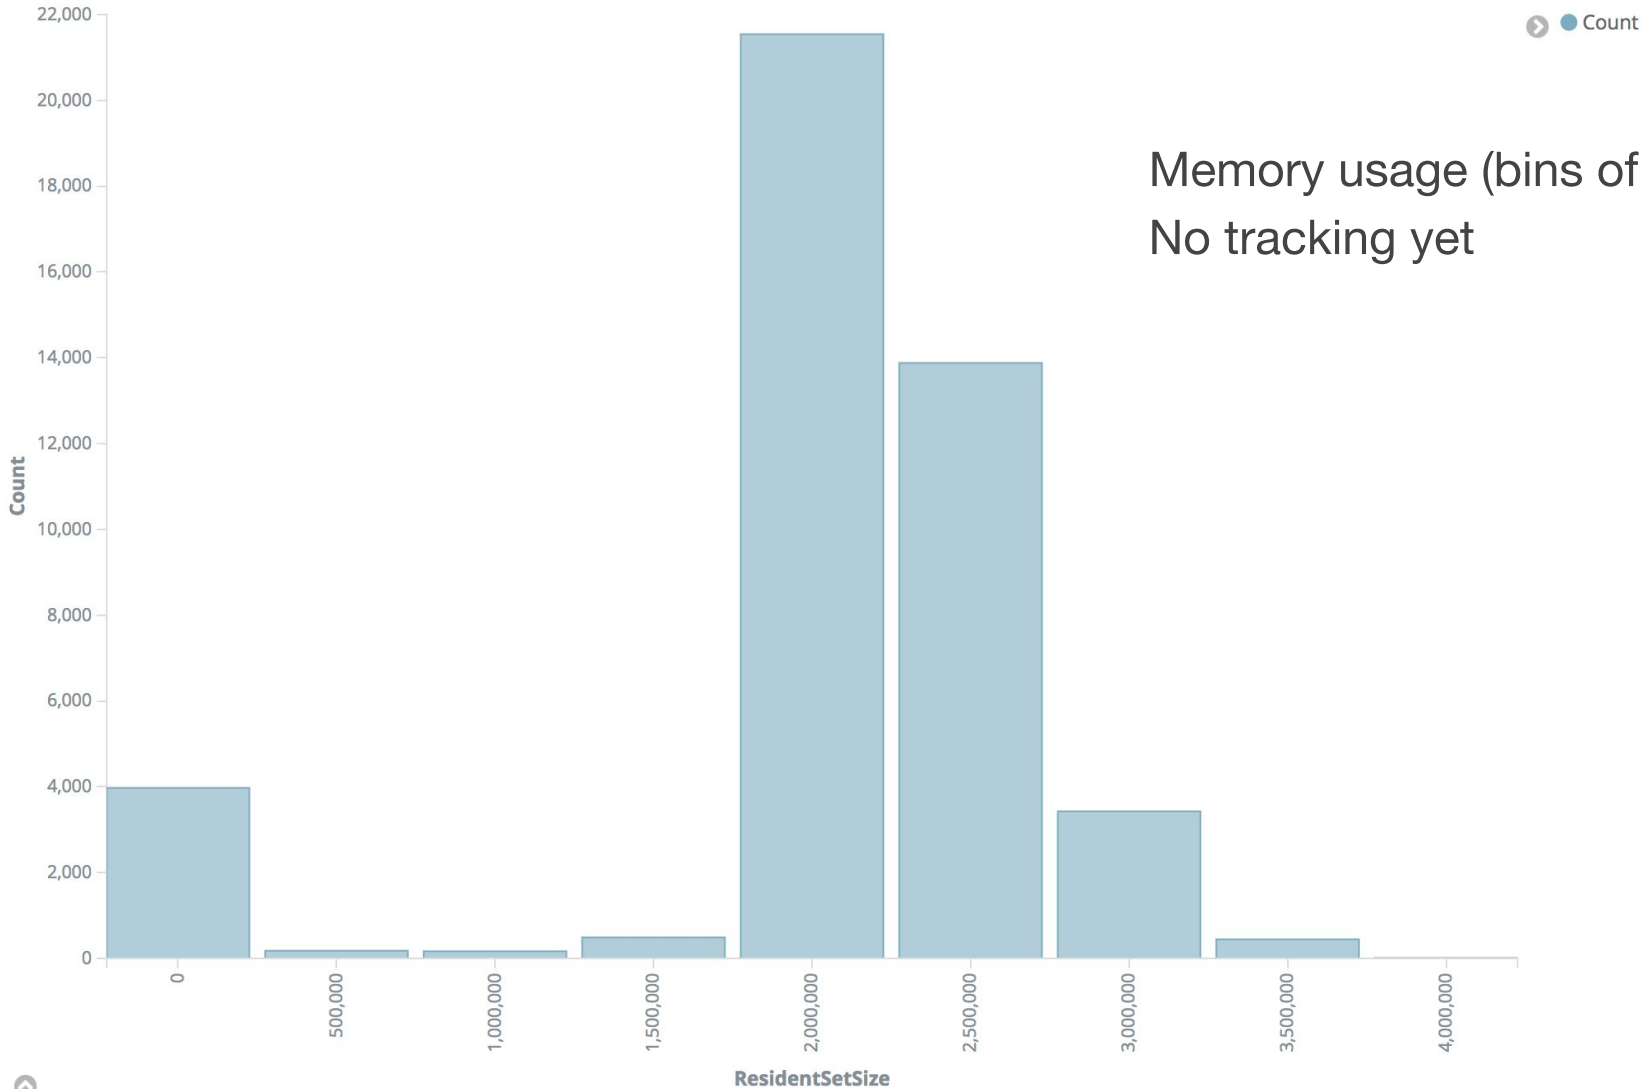

Memory usage (bins of 500MB)

No tracking yet

Production keepup processing- POMS setup

POMS is a FNAL-designed workflow management system; evaluating it and others for long-term needs

https://pomsgpvm01.fnal.gov/poms/campaign_stage_info?campaign_stage_id=1775

Apps | D0 Pages | STT Pages | News | DataPreservation | SCD | DarkSide | FIFE | OSG | DESGW | Imported From Fir... | DUNE | Other Bookmarks

Configure Work
Compose Login/Setup
Compose Job Type
Compose Campaign Stages
Config File Templates

Jobs
All (Last Day)
Failed_By Exit Code and Campaign Stage Name (Last Day)
Failed_By Exit Code and Campaign Stage Name (Last Week)
Failed_By Exit Code and Node Name Last Day)

Fail	6.39 K	3%
Held	7.82 K	4%
Success	182 K	93%

idle jobs - 37
running jobs - 1,313
max idle time - 3 day

cpu efficiency - 31.47%
memory efficiency - 21.05%

statistics for last 30 days

Reports/Status

- Production Shifter Page (Landscape)
- POMS Campaign Stage(landscape)
- Campaign Stage Stats (Landscape)
- Submissions (Landscape)
- Submission Time Bars (1wk)
- Campaign Stage Submission Files

Actions

- Edit Campaign Stage
- Launch Campaign Stage Jobs Now
- Launch Campaign Stage Test Jobs Now
- Kill running/idle Jobs for Campaign Stage
- Hold running/idle Jobs for Campaign Stage
- Release held Jobs for Campaign Stage
- Schedule Future Job Launches
- Mark Campaign Stage Inactive

Campaign Stage

Name: protoDUNE_SP_keepup_run4544
Experiment: dune
Dataset: np04_raw_keepup_run4544_v2
Software Version: v07_06_00
VO Role: Production
Param Overrides: [
['-Osubmit.N=', '250']
['-Oglobal.version=', '%(version)s']
['-Oglobal.quals=', 'e17:prof']
['-Oglobal.basename=', '%(dataset)s']
['-Oglobal.fclfile=', 'protoDUNE_SP_keepup_decoder_reco.fcl']
['-Oglobal.datastream=', 'noise']
['-Oglobal.tardir=', '/pnfs/dune/resilient/users/dunepro']
['-Oglobal.tarfile=', 'protodune-sp-keepup.tar.gz']
['-Osam_consumer.limit=', '1']
['-Osubmit.cpu=', '1']
['-Osubmit.expected-lifetime=', '36h']
['-Osubmit.memory=', '3500MB']
['-Osubmit.dataset=', '%(dataset)s']
['-Osubmit.email-to=', 'kherner@fnal.gov']
['-Oglobal.outdir=', '/pnfs/dune/scratch/dunepro/dropbox/np04/output']
['-Oglobal.logdir=', '/pnfs/dune/scratch/dunepro/beam_prep/logs']]
cs_split_type: drainingn(250)
cs_split_dimensions: None

Job Type

Name: protoDUNE_SP_keepup
Launch/Setup: fife_launch
Definition Parameters: [
['-c', '/dune/app/home/dunepro/protodune-sp/beam_prep/keepup_template.cfg']]
Input Files Per Job: 0
Output Files Per Job: 0
Output File Patterns: %root
Creator: kherner
Created: 2018-09-06 03:57:18.200709-05:00
Updater: kherner
Updated: 2018-10-05 23:17:51.423225-05:00

Initial keepup submissions are 3500MB memory and 36h run time

Production keepup processing- POMS setup

https://pomsgpvm01.fnal.gov/poms/campaign_stage_info?campaign_stage_id=1775

Apps | D0 Pages | STT Pages | News | DataPreservation | SCD | DarkSide | FIFE | OSG | DESGW | Imported From Fir... | DUNE | Other Bookmarks

Configure Work
Compose Login/Setup
Compose Job Type
Compose Campaign Stages
Config File Templates

Jobs
All (Last Day)
Failed_By Exit Code and Campaign Stage Name (Last Day)
Failed_By Exit Code and Campaign Stage Name (Last Week)
Failed_By Exit Code and Node Name (Last Day)

Fail	6.39 K
Held	7.82 K
Success	182 K

statistics for last 30 days

Reports/Status

- Production Shifter Page (Landscape)
- POMS Campaign Stage(landscape)
- Campaign Stage Stats (Landscape)
- Submissions (Landscape)
- Submission Time Bars (1wk)
- Campaign Stage Submission Files

Campaign Stage

Name: protoDUNE_SP_keepup_run4544
Experiment: dune
Dataset: np04_raw_keepup_run4544_v2
Software Version: v07_06_00
VO Role: Production
Param Overrides: [
['-Osubmit.N=', '250']
['-Oglobal.version=', '%(version)s']
['-Oglobal.quals=', 'e17:prof']
['-Oglobal.basename=', '%(dataset)s']
['-Oglobal.fclfile=', 'protoDUNE_SP_keepup_decoder_reco.fcl']
['-Oglobal.datastream=', 'noise']
['-Oglobal.tardir=', '/pnfs/dune/resilient/users/dunepro']
['-Oglobal.tarfile=', 'protodune-sp-keepup.tar.gz']
['-Osam_consumer.limit=', '1']
['-Osubmit.cpu=', '1']
['-Osubmit.expected-lifetime=', '36h']
['-Osubmit.memory=', '3500MB']
['-Osubmit.dataset=', '%(dataset)s']
['-Osubmit.email-to=', 'kherner@fnal.gov']
['-Oglobal.outdir=', '/pnfs/dune/scratch/dunepro/dropbox/np04/output']
['-Oglobal.logdir=', '/pnfs/dune/scratch/dunepro/beam_prep/logs']]
cs_split_type: drainingn(250)
cs_split_dimensions: None

Immediate Dependencies

protoDUNE_SP_keepup_run4544
Submissions 911 Located 0

Recent Launch Outputs

- 20181008_033144_kherner
- 20181008_033009_kherner
- 20181008_031134_mengel
- 20181008_013920_mengel
- 20181008_013212_mengel
- 20181008_001132_mengel
- 20181007_234506_mengel
- 20181007_233055_mengel
- 20181007_230022_mengel
- 20181007_220020_mengel
- 20181007_210024_mengel
- 20181007_162956_mengel
- 20181007_160525_kherner
- 20181007_160007_mengel
- 20181007_153949_mengel
- 20181007_151830_mengel
- 20181007_142516_mengel
- 20181007_130026_mengel
- 20181007_111845_mengel
- 20181007_054518_mengel
- 20181006_184102_kherner
- 20181006_181101_kherner
- 20181006_174101_kherner
- 20181006_171101_kherner
- 20181006_164101_kherner

Automated Recovery job launches

Recovery jobs submitted with 48h run time and 5000 MB mem

Initial submissions are 250 jobs at a time every 30 minutes (tunable)

Current status: https://pomsgpvm01.fnal.gov/poms/campaign_stage_info?campaign_stage_id=1775

Keepup outputs

Keepup runs the raw data decoder and reconstruction through hit finding (including photon detector bits)

Saves outputs to separate data tiers (decoded-raw and hit-reconstructed) in SAM (not using Rucio yet; highly likely in the future)

Analyzers can start from these files for tracking studies; no need to run expensive hit-finding many times

Do not save the raw data branches in outputs (saves a factor of ~2 in space)

Going Global with User Jobs

- International collaborators often bring additional computing resources to bear
- Users want to be able to seamlessly run at all sites with unified submission command.
- Following OSG setup process makes it easy to have sites around the globe communicate with a common interface, with a variety of job management systems underneath.
- Integration times as short as 1-2 weeks; all accessible via standard submission tools. **Record set-up time is just 2 hours!**
- DUNE requires a few CMVFS repositories: dune.openesciencegrid.org, fermilab.opensciencegrid.org, larsoft.opensciencegrid.org
- Local batch system (HTCondor, CREAM, ARC, etc.) does not really matter. **Use glideinWMS pilots (same as CMS) to run the user jobs**

DUNE jobs at non-US sites, past 3 months

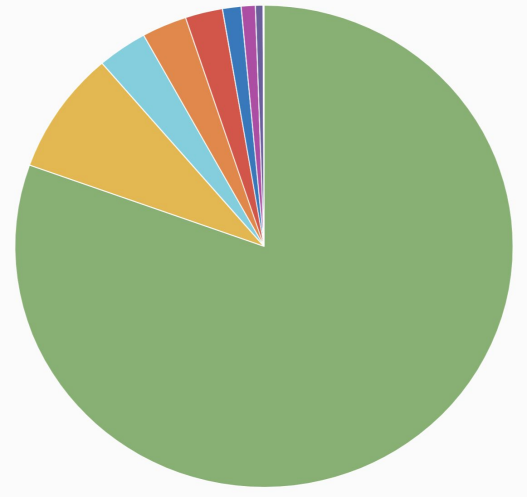

	total
CERN-PROD	385 K
RAL-LCG2	39.6 K
UKI-SCOTGRID-ECDF	15.5 K
UKI-SOUTHGRID-OX-HEP	14.01 K
UKI-LT2-IC-HEP	11.61 K
FZU	5.78 K
UKI-NORTHGRID-LANCS-HEP	4.39 K
UKI-NORTHGRID-SHEF-HEP	2.366 K
UKI-NORTHGRID-LIV-HEP	269
UKI-NORTHGRID-MAN-HEP	22
IN2P3-CC	1.0

Keepup dataflow

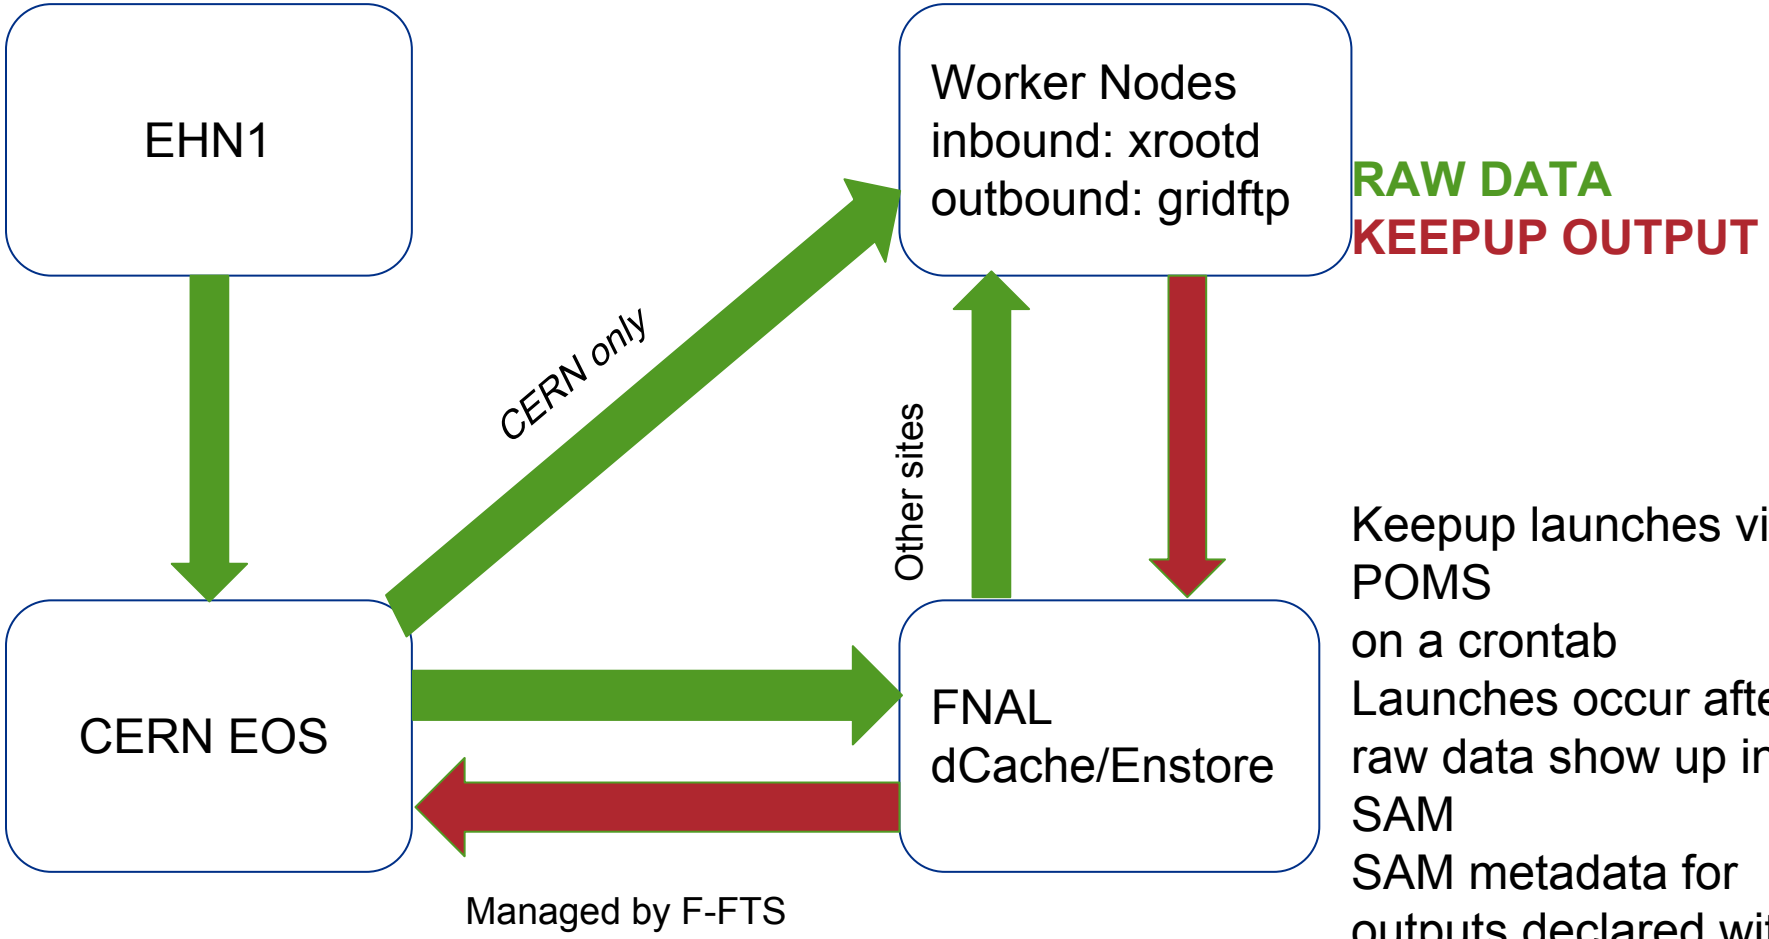

Keepup launches via POMS
on a crontab
Launches occur after raw data show up in SAM
SAM metadata for outputs declared within the job

protoDUNE Cross Labs Monitoring

- protoDUNE detectors now running
  - Significant computing infrastructure at both CERN and FNAL (buffer disk, long-term storage, reconstruction and analysis resources)
- Developed a unified monitoring system to show data from both laboratories side-by-side:
  - Covers data transfer rates, storage health and usage, job submission and status

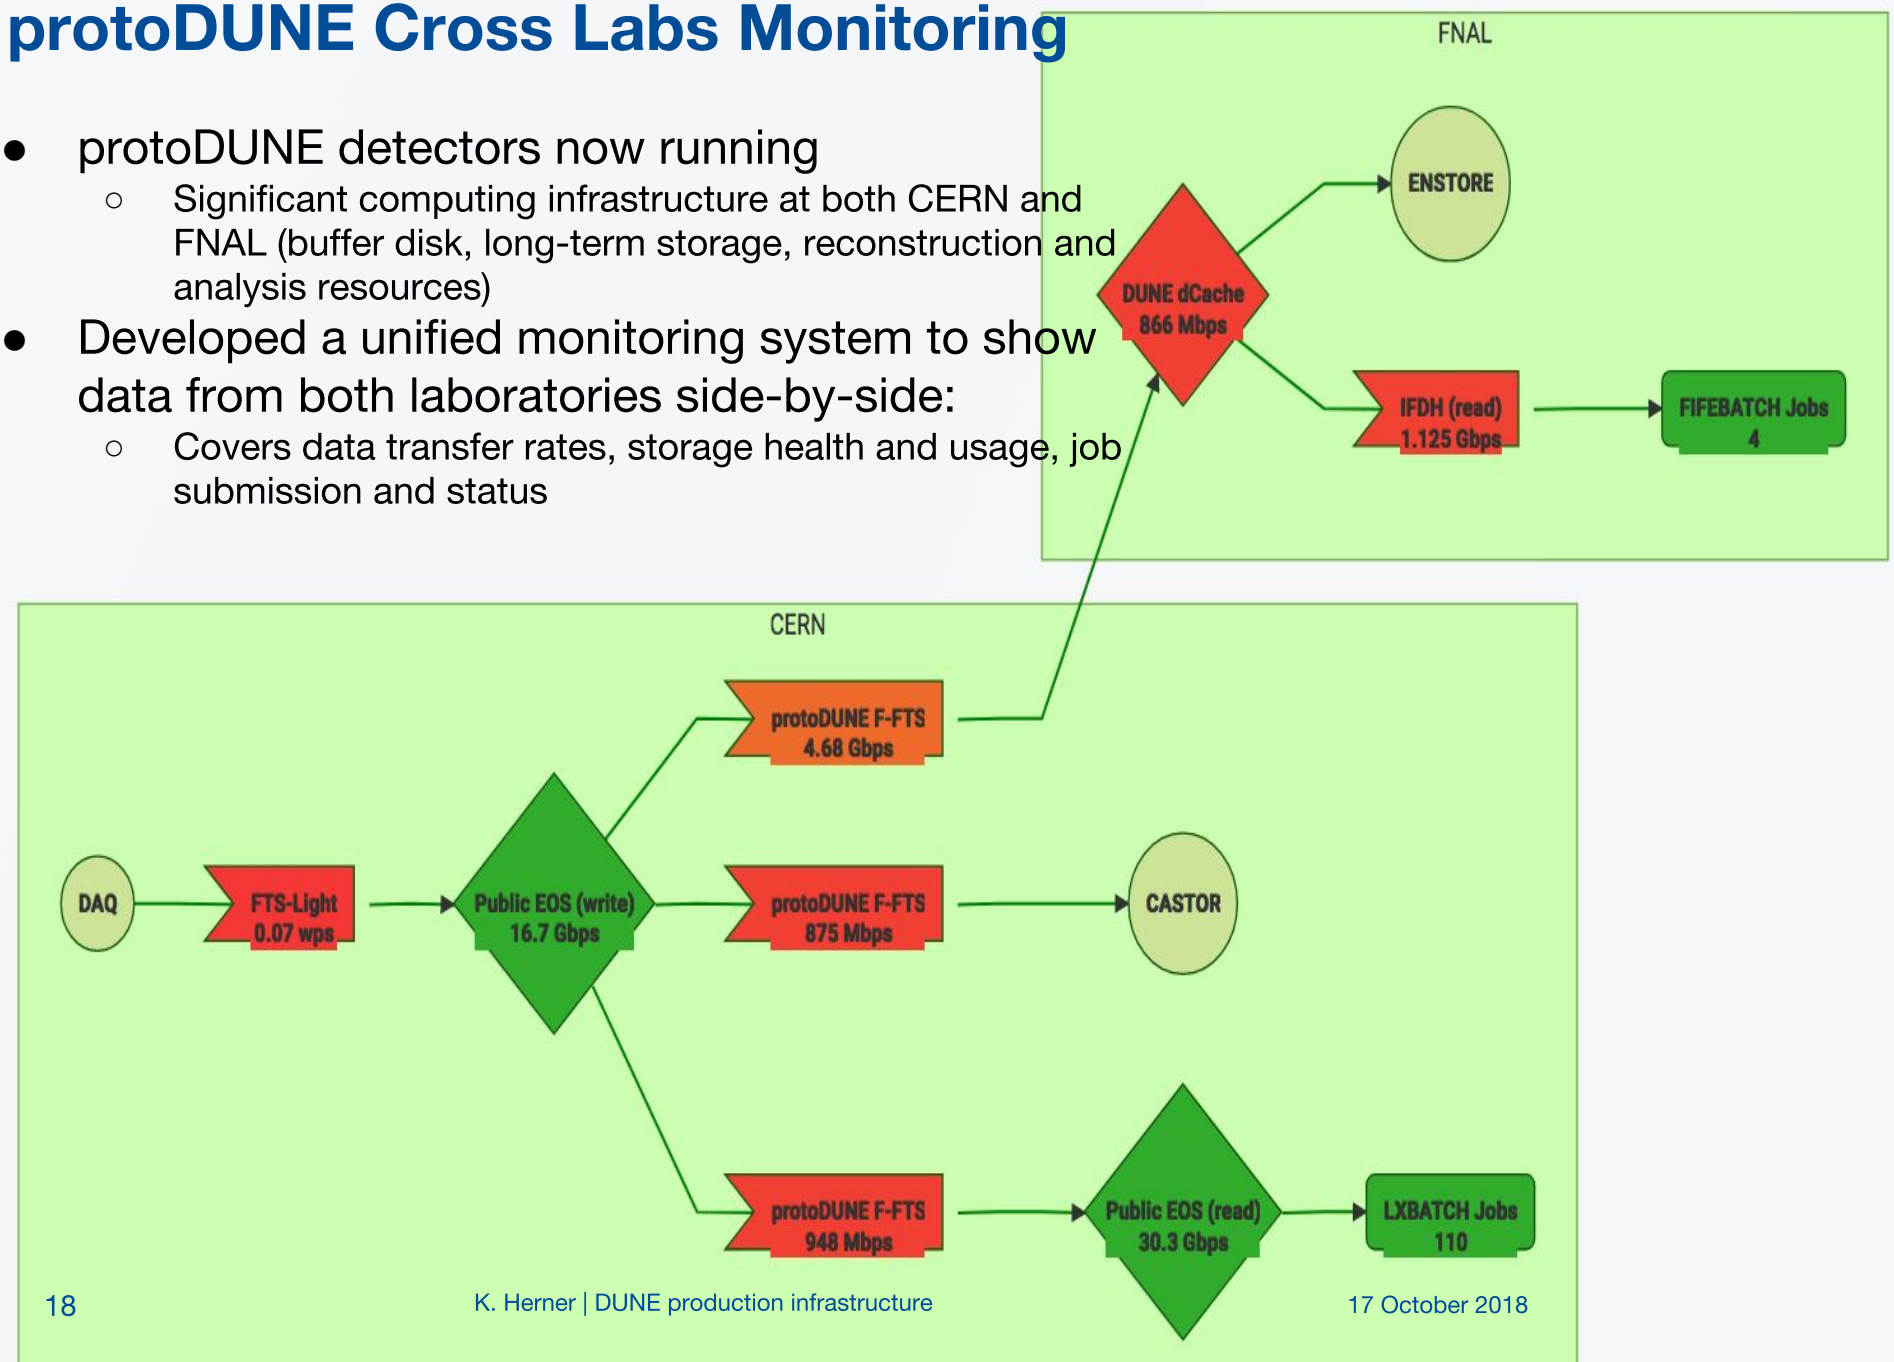

Central MC Production

Also running MC production (most recent campaign is MCC11)

Use POMS to set up infrastructure

Run samples once centrally; distribute outputs to entire collaboration

MCC11 jobs only, July 23 to October 8 (right: no FNAL, no CERN)

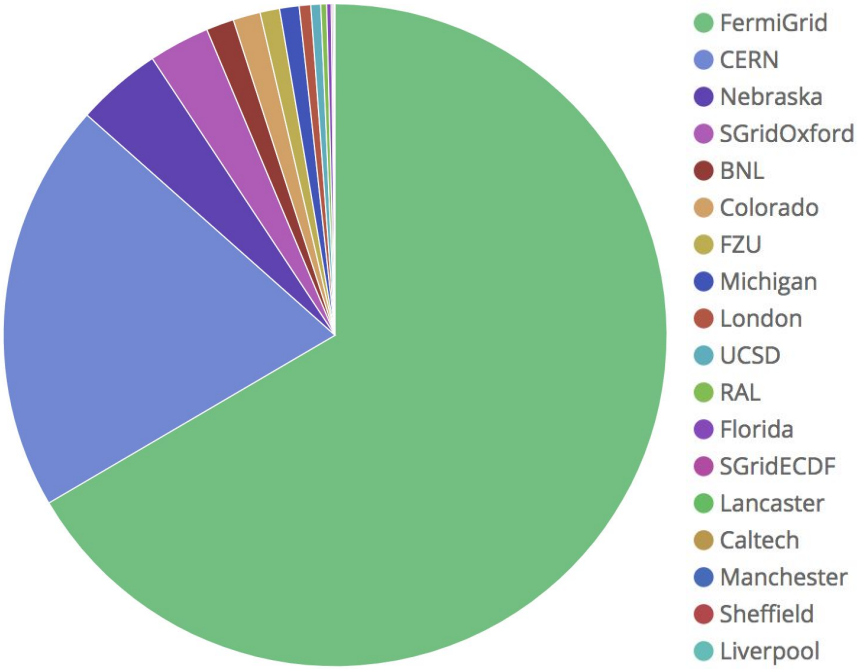

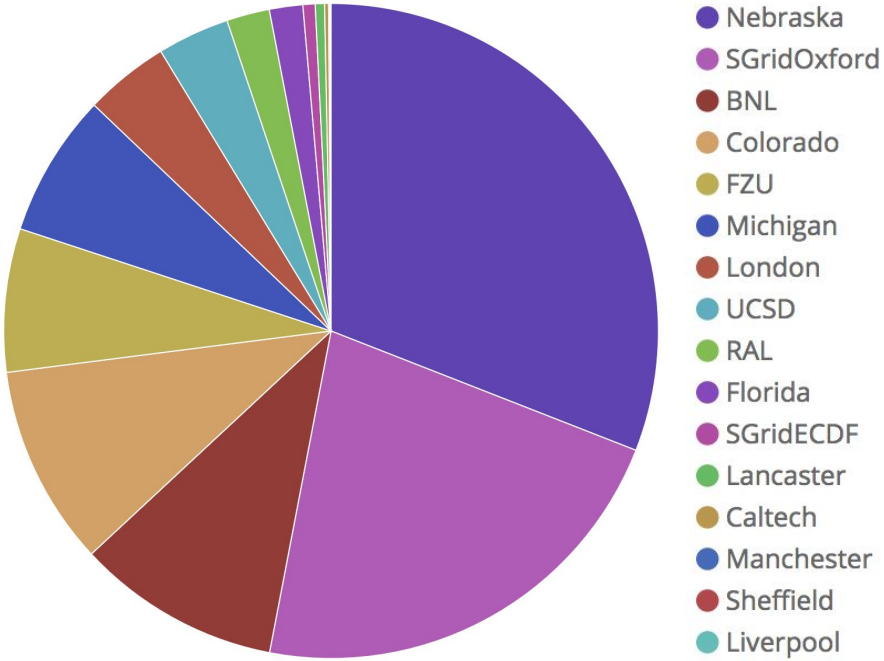

Taking the next steps

Low-hanging fruit in the short term:

1) Increase job distribution: move all sites to multicore setup (tracking can use multiple cores), higher memory limits and longer pilot maximum run times. (DP MC may come in at some point too).
2) Integrate site SEs starting with UK sites:
   a) stream inputs locally (or from CERN)
   b) b) copy outputs locally, set up some FTS back to FNAL

Longer term:

Develop computing consortium

Evaluate technologies and make long-term choices (data management, workflow management, software, etc.)

Adapt to changing infrastructure landscape

Future production dataflow?

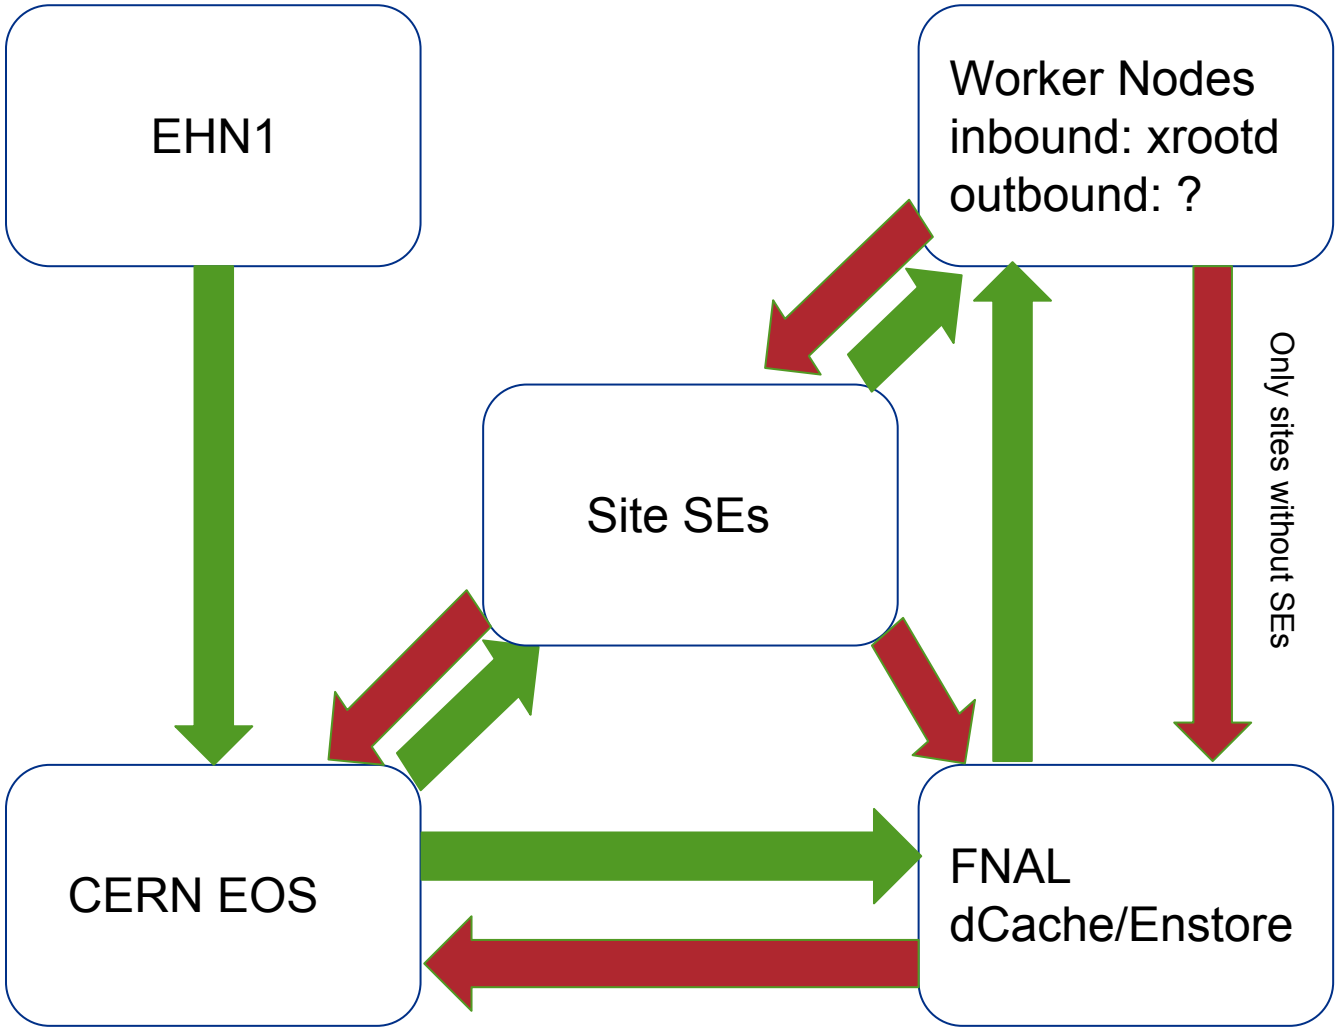

RAW DATA
KEEPUP OUTPUT

Not all paths may be here for all workflows: determined by replication/retention policies

All SE to SE copies managed by some form of FTS

Long-term: The Consortium Plan

- Forming a DUNE consortium of interested parties to work on DUNE strategy, assemble resources and make technical decisions
- Where it works, collaborate with the broader HEP and Astrophysics communities on common tools and infrastructure
  - Grid jobs
  - Data management (Rucio!)
  - Authentication!
  - New coding methodologies

The Goal(s)

- Record, reconstruct and analyze the ProtoDUNE data ASAP
- Prepare a CDR for full DUNE (x100 in size) which starts taking data in 2026
- Work with other experiments SKA/LHC/HyperK/LSST/SBND to have robust, global shared strategies

Current status of Consortium

- Many institutions are already contributing resources!!!!
- Setting up a directory of institutional contacts and resources
- Setting up a formal structure for technical decision-making
- Setting up a formal structure for resource contributions
  - Hardware (cores/GB)
  - Operations (people to keep it all running)

Three-fold way

Facilities	Technical	Operations
National infrastructure	DUNE standards	Lab support
Funding agencies	Common tools	Common costs
	New architectures	Collaborators

Many lessons to be learned from WLCG

Long term: software

Frameworks, reconstruction/tracking techniques, data management

Deep learning will play a critical role (already does) in particle ID and analysis

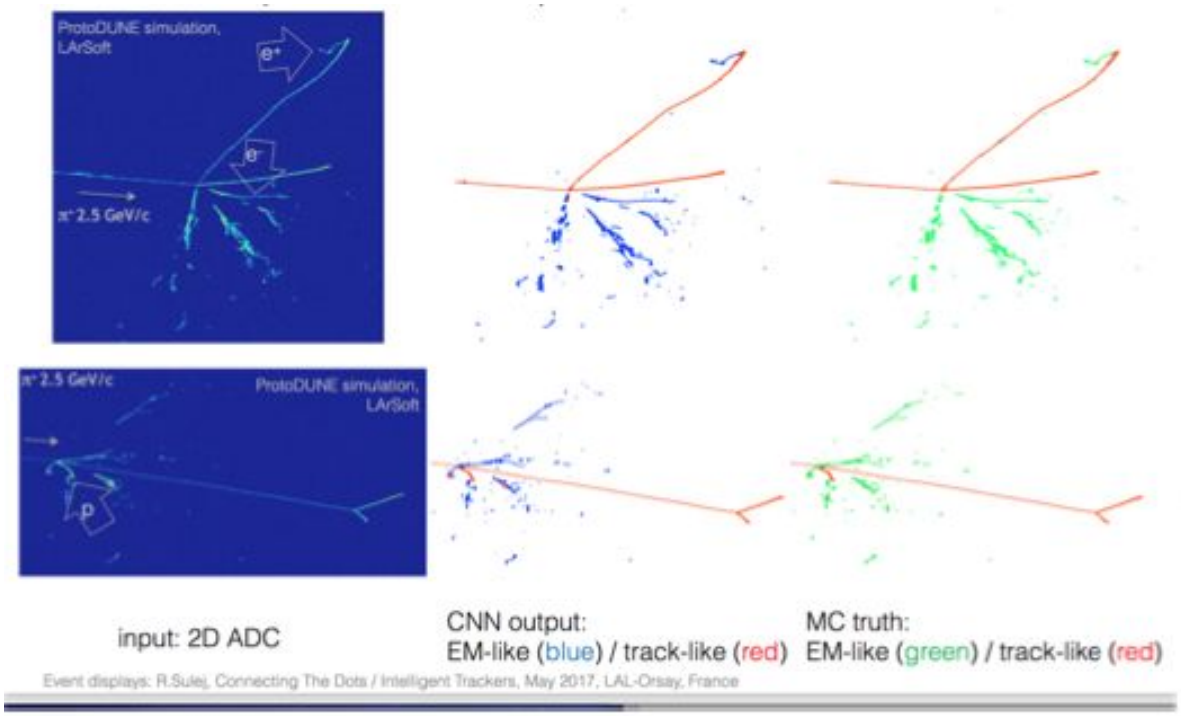

Works to separate out individual tracks and showers.

Allows better energy estimators

Need combo GPUs/CPUs at large scales

Long term: Infrastructure

DUNE will need resources of all types (HTC CPU, GPU, HPC)

Clear mandate from US DOE to make use of HPC and LCF machines

Art v3 is a step in the right direction (multiple events in separate threads)

Now is the time to re-think how we do things to fully leverage new resources

DUNE is still something of a green field in this regard!

CORI and Edison Supercomputer at NERSC were able to reproduce the 4-week long 2017 NOvA Analysis in a few hours

Over a million cores on CORI and Edison were provisioned and run through the **HEPCloud** facility in two analysis runs in May 2018.

Summary

Production group has been very busy the past few months with both protoDUNE data and MC processing.

Short-term modifications to improve throughput include modifications to the gWMS factory entries for some sites and integration of site storage elements

Currently running with beam (next week and one more week after that) and then another big MC push expected for the TDR in spring 2019. Plenty of work to do in the next 6 months!

WLCG offers many lessons on how to do international computing at scale

DEEP UNDERGROUND NEUTRINO EXPERIMENT

BACKUP

Production MC Workflow Submission Example

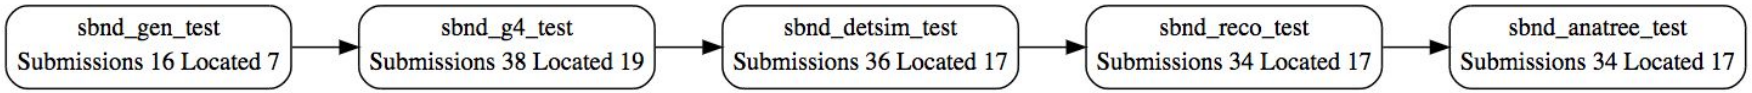

Test workflow for SBND

Campaign Stages

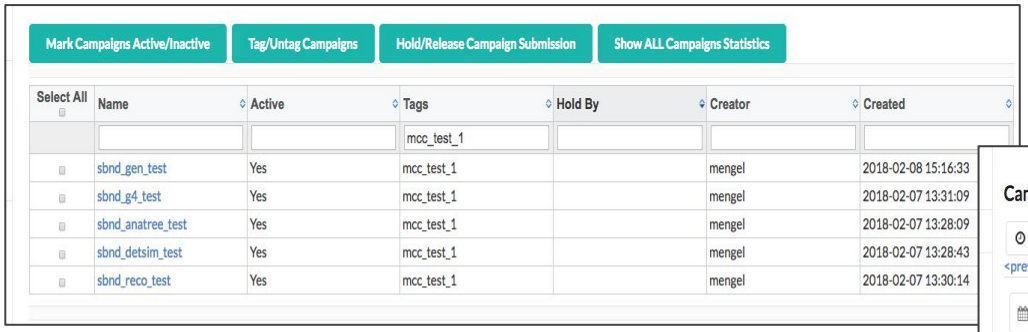

Mark Campaigns Active/Inactive | Tag/Untag Campaigns | Hold/Release Campaign Submission | Show ALL Campaigns Statistics

Select All	Name	Active	Tags	Hold By	Creator	Created
			mcc_test_1			
	sbnd_gen_test	Yes	mcc_test_1		mengel	2018-02-08 15:16:33
	sbnd_g4_test	Yes	mcc_test_1		mengel	2018-02-07 13:31:09
	sbnd_anatree_test	Yes	mcc_test_1		mengel	2018-02-07 13:28:09
	sbnd_detsim_test	Yes	mcc_test_1		mengel	2018-02-07 13:28:43
	sbnd_reco_test	Yes	mcc_test_1		mengel	2018-02-07 13:30:14

Campaign Results

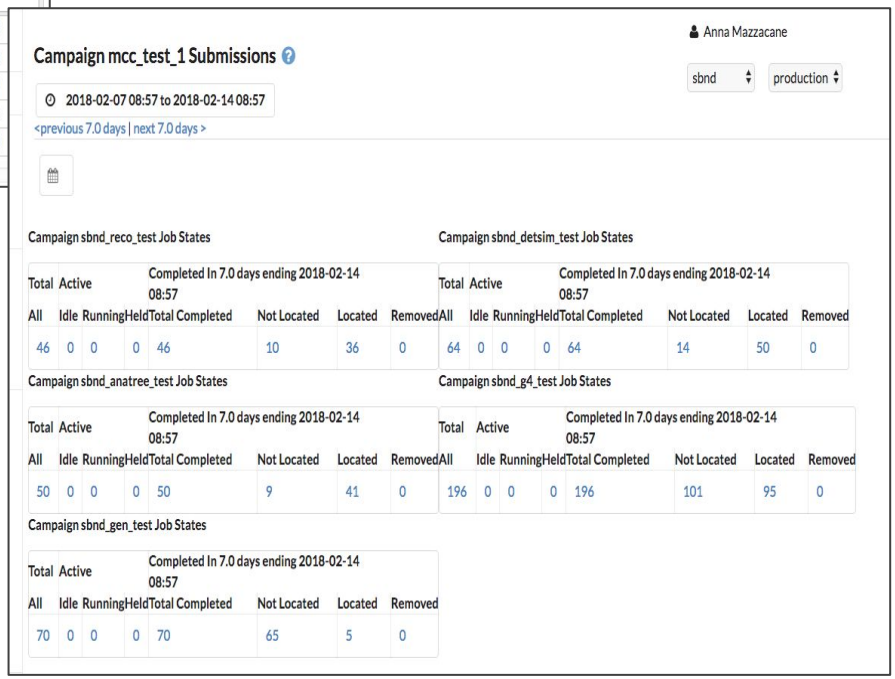

Anna Mazzacane

Campaign mcc_test_1 Submissions

sbnd | production

2018-02-07 08:57 to 2018-02-14 08:57

<previous 7.0 days | next 7.0 days >

Campaign sbnd_reco_test Job States

Total	Active			Completed In 7.0 days ending 2018-02-14 08:57			
All	Idle	Running	Held	Total Completed	Not Located	Located	Removed
46	0	0	0	46	10	36	0

Campaign sbnd_detsim_test Job States

Total	Active			Completed In 7.0 days ending 2018-02-14 08:57			
All	Idle	Running	Held	Total Completed	Not Located	Located	Removed
64	0	0	0	64	14	50	0

Campaign sbnd_anatree_test Job States

Total	Active			Completed In 7.0 days ending 2018-02-14 08:57			
All	Idle	Running	Held	Total Completed	Not Located	Located	Removed
50	0	0	0	50	9	41	0

Campaign sbnd_g4_test Job States

Total	Active			Completed In 7.0 days ending 2018-02-14 08:57			
All	Idle	Running	Held	Total Completed	Not Located	Located	Removed
196	0	0	0	196	101	95	0

Campaign sbnd_gen_test Job States

Total	Active			Completed In 7.0 days ending 2018-02-14 08:57			
All	Idle	Running	Held	Total Completed	Not Located	Located	Removed
70	0	0	0	70	65	5	0

- Each stage of the workflow takes a SAM dataset input with the output files of the previous stage. Often run gen+g4 together, or even gen+g4+detsim
- Stages can have different job resource requirements
- Each stage starts automatically when the previous one is finished
- Usually requesting 2 cores for reco stage

Auxiliary File Delivery via OSG's StashCache

Example of the problem: several neutrino experiments use O(100 MB)"flux files" for simulating neutrinos in rock, etc.

- Several files are used per simulation job
- Total input is few GBs
- Not well-suited to CVMFS for several reasons (each job uses random subset; can thrash worker node cache)
- Transferring all the time sub-optimal (plus no caching at all that way)

StashCache service developed and widely used by OSG to solve just such a problem:

- **Files replicated to regional caches, streamed via transparent xrootd connection from closest cache (geoIP) if not cached locally (small cache on the machine itself)**
- Can be layered on top of a dedicated CVMFS repository
- User sees a consistent POSIX file path across all sites, no need to worry about "how" the file gets there

Four FNAL experiments are using the service so far:

FIFE assists experiments with setting up and populating area.

StashCache (2)

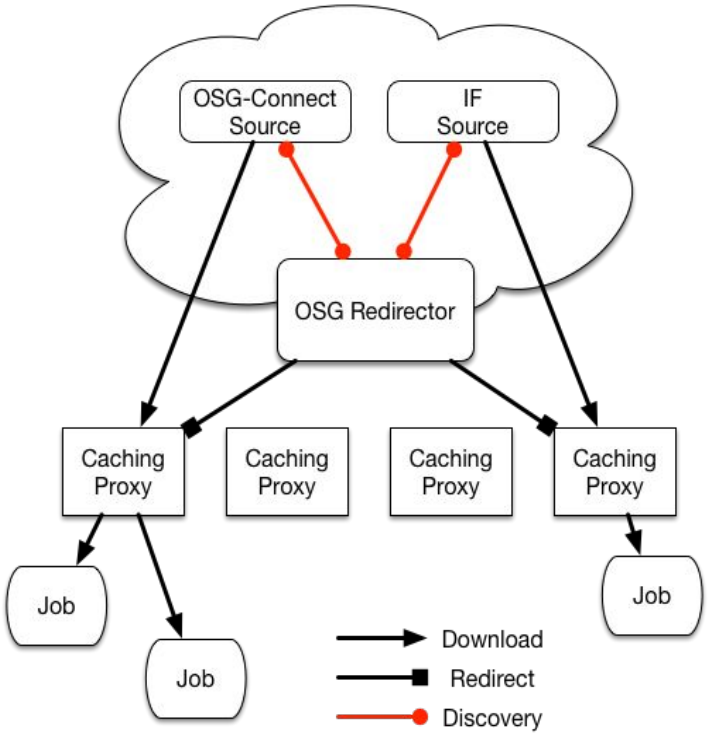

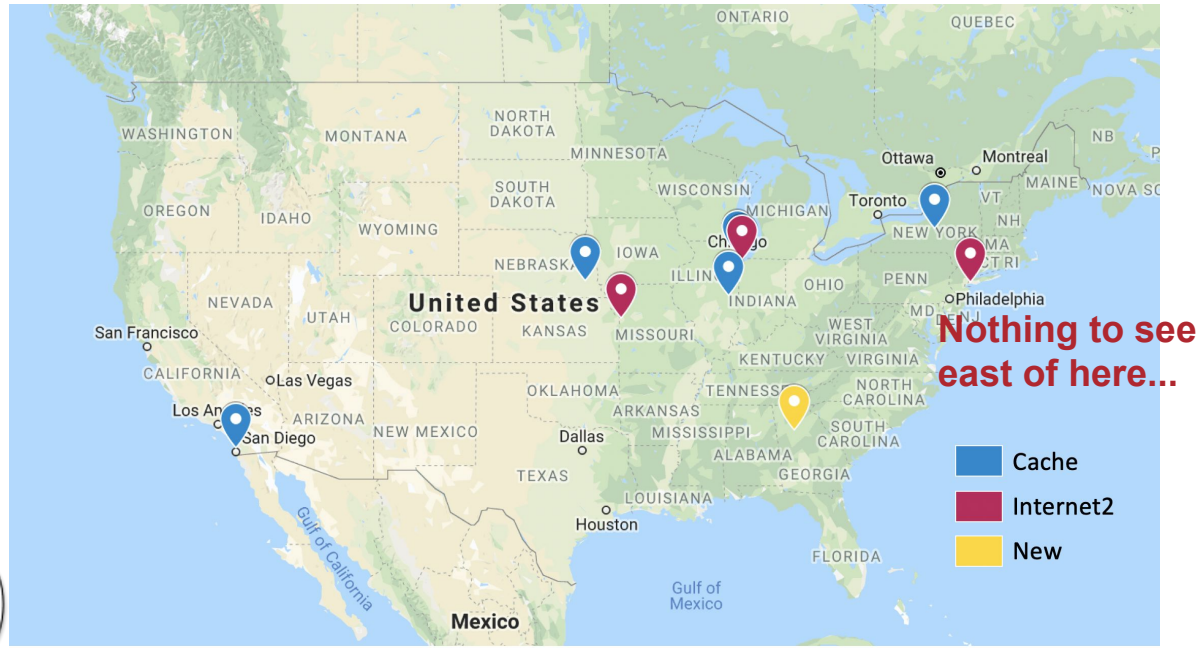

More info:

https://opensciencegrid.org/docs/data/stashcache/overview/
https://derekweitzel.com/2018/09/26/stashcache-by-the-numbers/

Off-line Production Shifter Dashboards

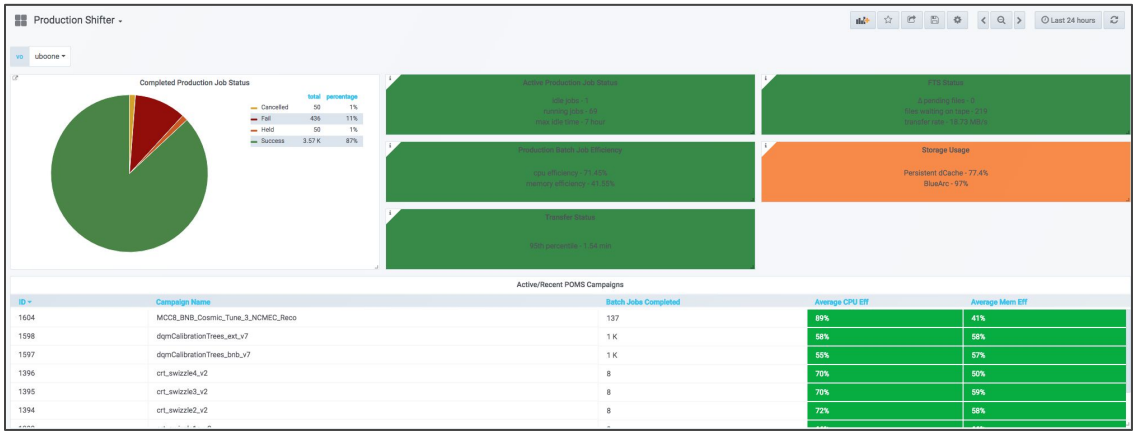

- Customizable dashboards for each experiment
- Display information experiments deem relevant to shifters in a single place
- Could drill down to more details

Efficient User Code Distribution

- Until January 2018 analysers were able to access their Bluearc code dev area
  - Was NFS mounted on all FermiGrid nodes.
- After Bluearc has been unmounted analyzers have started to use dCache scratch area to upload tarballs and stage them from dCache to worker nodes. It turned out to be a bad idea:
  - too much hammering on single pools when many jobs of the same type start
  - can completely block access to files owned by other users on a given pool
- Interim solution: stage tarballs in “resilient” dCache area (files replicated 20x on different pools):
  - Advantages: keeps the system running, doesn’t block users, more efficient.
  - Disadvantages: tarballs use 20x space; need careful cleanup policies
- Longer-term solution: looking into automated distribution via CVMFS.
  - Maybe should wait for CVMFS release that allows parallel publishing (should be released by the end of 2018)

Data management: FIFE SAM and FTS

SAM originally developed for CDF and D0; many FNAL experiments now using it

- A File metadata/provenance catalog
- A File replica catalog (data need not be at Fermilab)
- Allows metadata query-based “dataset” creation
- An optimized file delivery system (command-line, C++, Python APIs available)
- PostrgreSQL backend; client communication via http: eliminates need to worry about opening ports for communication with server in nearly all cases
  - Heretofore mostly followed the "send the data to the jobs" model. Enhancements underway to also make allowances for the "jobs to the data" school of thought

Simplifying I/O with IFDH

- File I/O is a complex problem (Best place to read? What protocol? Best place to send output?)
  - FIFE cannot force opportunistic sites to install specific protocols
- **I**ntensity **F**rontier **D**ata **H**andling client developed as common wrapper around standard data movement tools; shield user from site-specific requirements and choosing transfer protocols
- Nearly a drop-in replacement for cp, rm, etc., but also extensive features to interface with SAM (can use ifdh to pull in files associated with SAM project, etc.)
- Supports a wide variety of protocols (including xrootd); automatically chooses best protocol depending on host machine, source location, and destination (can override if desired)
  - **Backend behavior can be changed or new protocols added in completely transparent ways**
  - Special logic for automatically parsing Fermilab dCache and CERN EOS paths; shield user from having to know proper URIs

Full workflow management

- Now combining job submission, data management, databases, and monitoring tools into complete workflow management system
  - **Production Operations Management Service (POMS)**
- Can specify user-designed “campaigns” via GUI describing job dependencies, automatic resubmission of failed jobs, complete monitoring and progress tracking in DB
  - Visible in standard job monitoring tools

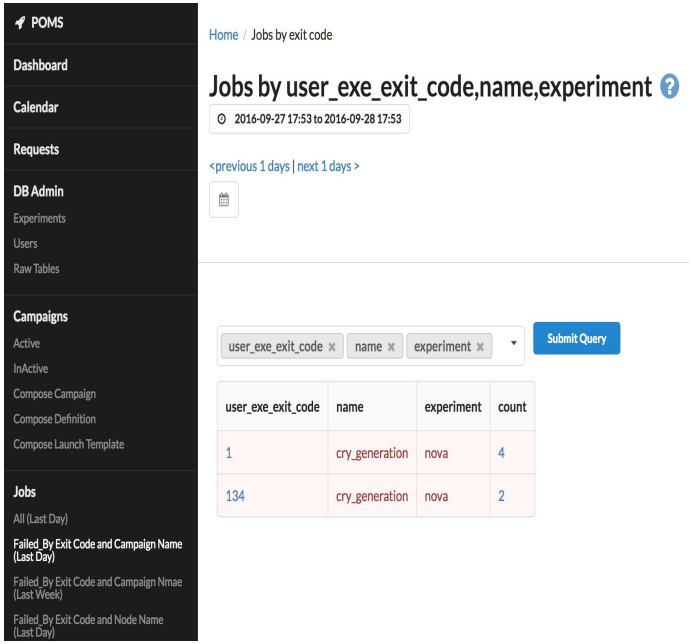

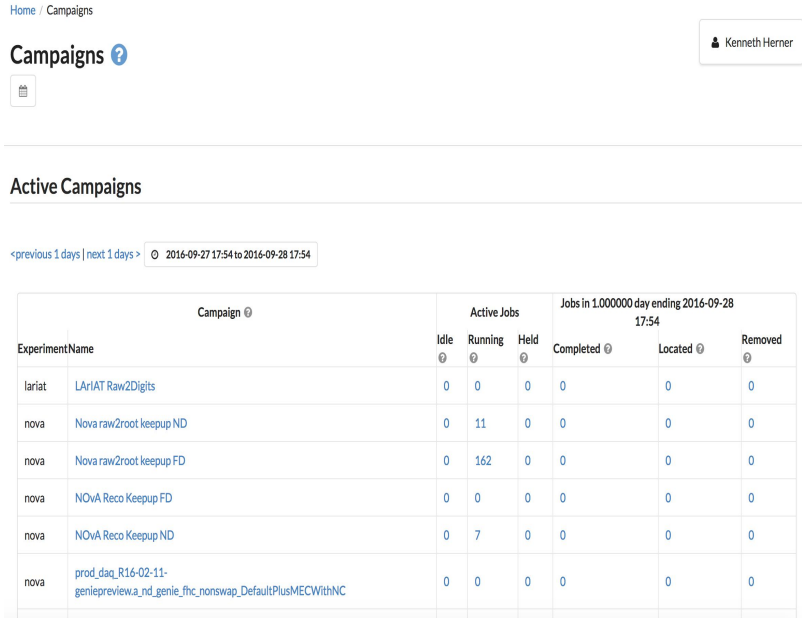

- Usable for production-level running and user analysis
- REST API for data I/O
- Command line tools for needed operations
- Supports POMS launching jobs, or experimenters launching jobs and using POMS only for tracking