


Production Software for sPHENIX

Links

- Meeting Link
 - <https://bnl.zoomgov.com/j/1604686747?pwd=cWJtem92VDE0ZG5FVTh5T28zMnE1QT09>
 - Meeting ID: 160 468 6747
 - Passcode: 155261
- [sPHENIX production workflow slide from Chris 20230328](#)
-  sPHENIX Distributed Computing Update

 NPPS 20230913

Production Observations

Meetings

Meetings are open to anyone interested. They are announced on [mattermost](#) and on [an open mailing list for distributed software](#).

- **Meeting Monday Oct 13, 2025 1 PM EDT**
 - Online Production:
- **Meeting Monday Sep 22, 2025 1 PM EDT**
 - Online Production:
 - Working on reproducing the chain from STREAMING event through SEED for earlier runs, so far 66456 through 68000
 - Problem: No, or almost no jobs are actually running.
 - Tentative explanation: Despite the online run monitor showing full transfer,

| Run Table | | | | | | | | |
|-----------|---------|-------------------------|-------------------------|------------------------|------------|---------|-----------|------------|
| Run Numt | runtype | Begin Time | End Time | | Run Length | In HPSS | In Lustre | First Seg. |
| 67963 | physics | 06/20/2025, 09:19:07 AM | 06/20/2025, 10:19:03 AM | <div><div></div></div> | 59.9 mins | 100% | 100% | 100% |
| 67962 | physics | 06/20/2025, 08:42:11 AM | 06/20/2025, 09:17:25 AM | <div><div></div></div> | 35.2 mins | 100% | 100% | 100% |
| 67961 | physics | 06/20/2025, 08:20:22 AM | 06/20/2025, 08:40:32 AM | <div><div></div></div> | 20.2 mins | 100% | 100% | 100% |

there are actually **no or not enough files available**

| RawdataCatalog=> select filename,status from datasets where runnumber=67962 and filename not like '%seb%' limit 10; | |
|---|--------|
| filename | status |
| physics_intt7-00067962-0000.evt | 0 |
| TPC_ebdc07_0_physics-00067962-0000.evt | 0 |
| TPC_ebdc07_1_physics-00067962-0000.evt | 0 |
| TPC_ebdc08_0_physics-00067962-0000.evt | 0 |
| TPC_ebdc08_1_physics-00067962-0000.evt | 0 |
| TPC_ebdc19_0_physics-00067962-0000.evt | 0 |
| TPC_ebdc19_1_physics-00067962-0000.evt | 0 |
| TPC_ebdc20_0_physics-00067962-0000.evt | 0 |
| TPC_ebdc20_1_physics-00067962-0000.evt | 0 |
| physics_intt1-00067962-0000.evt | 0 |

- Across a large range:

```
RawdataCatalog=> select runnumber,daqhost,status from datasets where runnumber between 67001 and 68000 and daqhost not like 'seb%'
and status!=0 and daqhost !='gl1daq'; runnumber | daqhost | status
-----+-----+-----
67939 | ebdc11:1 | 1
67447 | ebdc01 | 1
67448 | ebdc01 | 1
67449 | ebdc01 | 1
67450 | ebdc01 | 1
67451 | ebdc01 | 1
67452 | ebdc01 | 1
67456 | ebdc01 | 1
67457 | ebdc01 | 1
67458 | ebdc01 | 1
67466 | ebdc01 | 1
67467 | ebdc01 | 1
67470 | ebdc01 | 1
67447 | ebdc01:1 | 1
67449 | ebdc01:1 | 1
67450 | ebdc01:1 | 1
67456 | ebdc01:1 | 1
67458 | ebdc01:1 | 1
67466 | ebdc01:1 | 1
67719 | ebdc13 | 1
67764 | ebdc13 | 1
67597 | ebdc13:0 | 1
(22 rows)(!) ←
```


- Between 68001 and 69000, there are 270k available files, but I'd like to understand the above.

- Calo Fitting
 - It's working. Todo: Needs to be changed to only use the hcals, seb16 and seb17.
 - Monitoring and judicious resubmission allows me to refine their RAM needs. So far, indications are 2.5GB is often not enough, 4GB usually is.
- Rolling streaming cosmics:
 - Currently have 2000 held jobs that exceeded 18 GB RAM. I will save these (as pickle? JSON? For some reason the size is the same) in a way that they could be resubmitted and kill them. Same on a smaller scale for other jobs that keep exceeding "reasonable" RAM allocations
 - Whether we'll ever revisit them is a different matter, but if so, the information will be there.

- AOB

• Meeting Monday Sep 8, 2025 1 PM EDT

- Online Production:
 - Progress on monitoring; memory control. Example plot for held SEED jobs:



- majority is concentrated in the lower left, where memory exceeded the originally allotted RequestMemory
- most of the rest is held even though the memory request was updated
 - note that this doesn't show those jobs who successfully went back to running after adjustment
 - "Reason: Job was held for being idle after running"
 - Some fiddling with qedit and release will probably handle this.
 - Otherwise, I now have the machinery to clone the job and rerun it "from scratch"
- Next focus for framework integration
 - 1. Automatically produce fully transferred runs in full
 - For cosmics as well? Yes → Double-check
 - 2. CALO
 - 2a Combining
 - 2b Calofitting
 - 3. TRKR_TRACKS, JETS?
 - don't prioritize. Underlying calibrations need so much work it wouldn't be meaningful.
- Disk space discussion

• Meeting Monday Aug 18 2025 1 PM EDT

- Online Production:
 - Releasing held jobs with new memory requirements:

```
condor_qedit -const 'JobBatchName=="main.DST_TRKR_SEED_run3auau_new_newcdbtag_v001" && JobStatus==5' RequestMemory='8096MB'
condor_qedit -const 'JobBatchName=="main.DST_TRKR_SEED_run3auau_new_newcdbtag_v001" && JobStatus==5' PeriodicHold '(NumJobStarts >= 2 && JobStatus==1)'
condor_release -const 'JobBatchName=="main.DST_TRKR_SEED_run3auau_new_newcdbtag_v001" && JobStatus==5'
```

- Overwriting existing jobs, for example for special runs that require all segments, use `--force`, `-f` to delete existing files and wipe databases

```
create_submission.py --rule DST_STREAMING_EVENT_run3physics --config
/sphenix/u/sphnxpro/mainkolja/ProdFlow/short/run3auau/v001_combining_run3_new_nocdbtag.yaml --runs 72199 --no-onlyseg0 --force
```

- (still testing not in main yet)

- Discussion: Removing used DSTs (currently lustre is 87% full).
 - This should for the time being be triggered manually
 - AND require that the files are on HPSS? **NO**
 - Keep combined DSTs. Delete CLUSTER once SEED is produced. For the purposes of QA, histograms from SEED are all that's needed, DSTs can go as well.
- AOB
 - Just noticed:

```
du [...]
28M /.../DST_STREAMING_EVENT_ebdc39_run3auau_new_nocdbtag_v001-00072199-00008.root
du -sh /[...]/DST_STREAMING_EVENT_ebdc39_run3auau_new_*72199*|wc -l
1045
```

- Do we want this many small files?
- Check file size during spider! Copies are going poorly.
- Pick a recent fully transferred run and test full run production. E.g. 72552, 72592
 - How long do they run, how long do the segment0 ones take?
- Thomas has issues with a cosmics run - he'll scrub and resubmit it, then we'll troubleshoot on MM

• Meeting Monday Aug 11 2025 1 PM EDT - skipped

-

• Meeting Monday Aug 4 2025 1 PM EDT

- Online Production:
 - On-going combination and cluster, seed jobs are mostly smooth. One snag in STREAMING combination that shouldn't occur again.
 - Run TRKR jobs on older runs? Currently the lower bound is 71000
 - Yes, starting at 70000 and see how that goes.
 - How to signal runs that are transferred in full (e.g. for QA) and should be produced in full.
 - 2+ runs per run (or per fill, at least second run of the fill)
 - Use the query ("edit" in the kebab menu) from [Jin's Grafana page](#)

- Somewhat related request from Thomas regarding TRKR downstream: "Do you know if there's a straightforward way to change the requirement for the job to run to be only that the MVTX and INTT have all their files instead of all tracking detectors? It seems like a bunch of the runs that Greg wanted produced got skipped over because they were missing TPC files and he said that a lot of them will either be missing the TPC or be very sparse in the TPC, so it's ok to still produce them without it."
 - Current status: Required are MVTX and INTT + 30 TPC streams (ignoring TPOT, it's almost always there anyway AFAIK)
 - Decision: For **cosmics only**, relax even further and run for any run that has any streaming file available.

● **Meeting Monday Jul 28 2025 1 PM EDT**

- Online Production:
 - Fixing xferslot bug that stalled production (only needed for combination; other jobs using it block available job slots for no reason)
 - Discussion of calorimeter production, cosmics needs
- Disk space: 14PB (17%) available.
- Joe has created a readme overleaf document, we'll try and work on it as an md on github.
- AOB
 - Monitoring - how many events have been produced? Lost?

● **Meeting Monday Jul 21 2025 1 PM EDT**

- Online Production:
 - Discussion/review of methodology to prioritize new runs for combination (and other?) jobs.
 - Group condor submission files by runnumber (instead of blindly by chunks)
 - Segments: Ideally, if more than seg0 exists, then use it all.
 - Submission is (eventually) separated from condor file creation.
 - This allows:
 - Order by latest runs first
 - Create / submit only a capped number of jobs (instead off 700k+ idle jobs in the queue)
 - Caveat: If this equilibrates/saturates, there's a chance older runs never get produced. Unlikely to become a problem, and can always be kicked off by hand, but something to be mindful of.
 - Other changes:
 - Phase out DST lake. Seemed a good and tidy idea, but with simple and unavoidable snags it can lead to 10s of millions of files in one directory which brings lustre to a screeching halt.
 - Implement a switch for special jobs (cosmics, Vernier scan, ...) to re-enable original combining logic instead of only using segment0
 - Also needs a well thought out **--force** switch that not only overrules prod db/file db status but makes sure there's no job or pre-existing file hiding out there that will overwrite the forced production later. Easy for the current Vernier scan request. Not too bad once we transition out of the DST lake, but ideally, it needs a mechanism to also kill existing condor jobs. Probably feasible with prod db changes during the standalone submission process.
 - TRKR, SEED jobs: Want to use the logic currently in dev.
- Offline:
 -

● **Meeting Monday Jul 14 2025 1 PM EDT**

- Combining scheme discussion: the current scheme to wait for all files from a subsystem for a given run to be on lustre just isn't working.

Proposal: Push the processing for the first segment (or more?) to just get the q/a going to make sure we don't take garbage. This does need a change in approach (right now incomplete runs do not get put into the raw data catalog)

 - for 80% of the offline QA we really only need the 0th segment. The problem is that we need the 0th segment for all servers
- Online:
 - CLUSTER experienced a large number of held jobs over the weekend. condor logs are inconclusive, but I suspect 12 GB RAM may be not enough. Running smaller batches now for debugging
 - SEED workflow works, awaiting signoff for auto-production
 - CALO-like jobs follow the same job pattern as the two above, working now on turning them on.

● **Meeting Monday Jul 7 2025 1 PM EDT**

- Disk space

```
sphnxpro@sphnxprod03 mainkolja> df -h /sphenix/lustre01/
Filesystem      Size  Used Avail Use% Mounted on
10.42.34.179@tcp:/sphnx01  79P   65P   14P   83% /direct/sphenix+lustre01
```
- Online
 - Last week, revamp, refactor, and unification in db naming scheme between **dsttype** (DST_TRKR_CLUSTER), **dataset** (run3cosmics), and **tag**, i.e. the triplet new_nocdbtag_v001.
 - TRIGGERED (Calo combining) is running smoothly
 - Updated code from JaeBeom exists; it should integrate invisibly when merged.
 - STREAMING (Tracking combining) is putting a strain on the system. It seems files are produced faster than they can be registered and moved into place
 - Led to an out-of-memory error when the list of files to process reached ~20M. Alleviated with a crude chunking mechanism (**find ...| head -n 500000**)
 - pipe into a large file and only run find again once that's exhausted.
 - Best angle of attack: Consolidate db write in chunks. Only about 30% of the time sink, but the largest part that can be optimized.
 - TRKR_CLUSTER: 110k DSTs and histos produced.
 - Can be used to develop TRKR_SEED jobs.
 - Some issues:

- Naming convention was inverted - fixed now. Existing files can be used but will be supplanted by correctly named ones.
- Some 700k jobs were in the queue. This doesn't seem right, most likely explanation is a bugged check against the production db
 - dataset db, not prod db, but correct. Caused by the above refactorization.
 - First impression is that this does not affect event combining, but I will triple-check.
- Are 16 GB needed? Joe will check
- Since we're already using the RAM, could multithreading help throughput? Set nCPU to 4 and see what happens.

- Offline
 - Track based DSTs for 40 runs complete, and they look better than the QM production for the “golden run” so those will remain

• Meeting Monday Jun 30 2025 PM EDT

- Disk space:

| Filesystem | Size | Used | Avail | Use% | Mounted on |
|---------------------------|------|------|-------|------|--------------------------|
| 10.42.34.179@tcp:/sphnx01 | 79P | 66P | 13P | 84% | /direct/sphenix+lustre01 |

 - Significant cleanup over the last week. But it will remain an issue.
 - Cosmics etc. deleted and stoppedIf there are specific runs or run ranges you would like to have remain and are critical for ongoing analysis, please let me know as soon as possible and we will reproduce them

- Online:
 - Dealing with a huge data influx.
 - Requested cosmics: 6M files slowly moving into place
 - Rolling physics:
 - Triggered under control, currently 6k files.
 - Streaming:
 - Histos tbd
 - Not helped by Lustre problems.
 - Funky behavior of `lfs find`:

```
lfs find /sphenix/lustre01/sphnxpro/production/dstlake/run3auau/physics/DST_STREAMING_EVENT -maxdepth 1 -type f |wc -l
```

`104507`
 - Good! spider is keeping up - however, this find command took more than an hour to complete.
 - I just tried regular find, that is plenty fast - but not recommended for lustre?
 - Priorities:
 1. Make sure raw production does not pile up. Current transfer rate of ca. 50 files/second, barely keeping up, if at all. Current time sinks:

| | | | | | |
|--------|---------|-------|---------|-------|---|
| 50000 | 483.598 | 0.010 | 483.598 | 0.010 | {built-in method posix.rename} |
| 100070 | 196.009 | 0.002 | 196.212 | 0.002 | {built-in method posix.stat} |
| 100000 | 163.950 | 0.002 | 163.950 | 0.002 | {method 'execute' of 'pyodbc.Cursor' objects} |
| 100000 | 122.116 | 0.001 | 122.116 | 0.001 | {method 'commit' of 'pyodbc.Cursor' objects} |

- two avenues:
 - bundle db access into chunks.
 - investigate if posix.stat is caused by `shutil.move`; it's not invoked explicitly. Instead, try `os.rename()`. However, "shutil.move simply calls os.rename in most cases". Have to see if true.
 - Altogether ca 100% speedup possible; after that, explore multiple transfer (spider) hosts and speak with Jane.
- 2. TRKR_CLUSTER test run ran successfully. Start and monitor rolling production
- 3. From this template, derive all other downstream jobs (particularly CALOFITTING, SEEDs)
- 4. Automate loop better, monitor disk space, ...
- Copy input onto the local disk, check if successful, no more worry on that side.
- Check how many files we actually produce (over say 24h)
- What files need to be kept? prdfs?
 - CALOFITTING is end product, potentially much smaller Jet DSTs

• Meeting Monday Jun 16 2025 3PM EDT

- Online:
 - Running combination loops for physics, cosmics, beam, calib
 - ~100 jobs currently held

| — Schedd: sphnxprod03.sdcc.bnl.gov : <130.199.6.180:9618?... @ 06/16/25 14:52:03 | | | | | | |
|--|---|------------|------|------|------|------|
| OWNER | BATCH_NAME | SUBMITTED | DONE | RUN | IDLE | HOLD |
| sphnxpro | kolja.DST_STREAMING_EVENT_run3auau_new_nocdbtag_v000 | 6/13 15:20 | 706 | 385 | — | 46 |
| sphnxpro | kolja.DST_TRIGGERED_EVENT_run3auau_new_nocdbtag_v000 | 6/13 15:30 | 180 | 147 | — | 2 |
| sphnxpro | kolja.DST_TRIGGERED_EVENT_run3cosmics_new_nocdbtag_v000 | 6/16 13:15 | 579 | 1018 | — | — |
| sphnxpro | kolja.DST_STREAMING_EVENT_run3cosmics_new_nocdbtag_v000 | 6/16 13:20 | 3647 | 727 | — | 84 |

- Some final touches on prod db, then focus is on CLUSTERING
- (as template; once it works the rest is simple)
- Change of datasets table columns proposal
 - Lookups are too complicated and rely on like ‘%’ in any lookup script, how about
 - dataset (run1auau, run2pp, run3auau, run3cosmics,...), currently primary key - change it, or better name?
 - Dst type (DST_CALO, DST_TRACKS,...)
 - Tag (ana492_2025p003_v005)
 - others?
 - Added status (default = 1), currently 1: on disk, 0: only in hpss
 - Took 4 hours to initialize to 1

- Meeting Monday Jun 09 2025 1PM EDT
 - Data readiness
 - New triggered event combining took care of a lot of cosmic problems (all sebs started with the wrong event, discarding the complete run, handled now)
 - One more low hanging fruit today (rollover of 32 bit beam clock leads to mismatch), then reprocessing also of the run2 data
 - Seb19 (II1) not handled - II1 needs special data packet
→ IGNORE seb19 in production
 - Streaming and triggered combination currently controlled by hand (physicsmode='beam').
 - Found the first bug of the run, the spider works but didn't use the correct final directory.
 - Downstream, currently kicked off by hand by Chris as well
 - Shuhang asked for CALOFITTING ← fixed ana build and dbtag?
→ Not needed, but better to avoid future confusion
 - Lamination fitting macro/workflow
 - space charge corrections - run on output of CLUSTERing → second priority for downstream automatization (after calofitting)
 - Lustre cleanup status
 - Other things we can safely remove?
 - PRDFs in /sphenix/lustre01/sphnxpro/fromhpss (event combining production of ~40 tracking runs finished). du claims 1.5PB in fromhpss/physics
→ Keep for now, revisit next week
 - HPSS: **Bandwidth is limited**, more likely than not saturated by raw data alone
 - Run2pp streaming production
 - Cluster DSTs for ~40 runs produced under ana489_2024p020. Next up seeding, but requires calibrations. Work ongoing

- Meeting Monday Jun 02 2025 1PM EDT
 - Expectations for beam
 - Collisions expected for Friday
 - Readiness
 - Production loops running on phnxsub01 and phnxsub02
 - Tracking workflows from sub01, calorimetry from sub02
 - Streaming event builders
 - runtime in ('physics', 'beam', 'cosmics', 'calib', 'dryrun', 'line_laser'). (clustering removes line_laser from the list).
 - 4000MB memory
 - 100000 events per output segment → TODO move to 10k. Setup separate rule.
 - Triggered event builders
 - runtime in ('cosmics', 'physics', 'beam', 'dryrun')
 - calib??
 - 2048MB memory
 - 100000 events per output segment [20k “overflow” ... but I think we no longer use that parameter... TBD]
 - TODO: shutdown loops. Switch to Kolja's system.
 - Everyone should check prodflow to confirm that macros are up to date
 - New workflow(s) defined which will run after a job exits the running state. i.e. job finishes or fails.
 - Uses the production DB as input
 - Log files will be automatically compressed in order to reduce the footprint of the logfiles on disk.
 - At some point we switched off the xfer slot mechanism for both event builders. Re-enable?
 - Cleanup of past datasets (space on data02, lustre)

```
$ df -h /sphenix/data/data02/ /sphenix/lustre01/
Filesystem                Size      Used Avail Use% Mounted on
/gpfs/mnt/gpfs02          251T    232T   20T   93% /gpfs/mnt/gpfs02
10.42.34.179@tcp:/sphnx01  79P     74P   5.0P   94% /direct/sphenix+lustre01
```

Cleaning up 25 TB from /sphenix/data/data02/sphnxpro/testbed
Compressing several smaller (100's GB) directories under data02

Biggest datasets (in TB):

| | | |
|----------------------|--|------------------------|
| ana450_2024p009 | | 1502.3982888045590000 |
| ana462_2024p010_v001 | | 1790.0161645180220000 |
| ana484_nocdbtag_v001 | | 1907.8998904436430000 |
| ana468_2024p012_v001 | | 1924.9627285188490000 |
| ana462_nocdbtag_v001 | | 3010.6784400156690000 |
| ana446_2024p007 | | 4104.0289177335990000 |
| new_nocdbtag_v003 | | 4540.9539311057830000 |
| new_nocdbtag_v004 | | 7399.6319854158110000 |
| new_nocdbtag_v000 | | 7452.5182023704490000 |
| mdc2 | | 17191.9586880897960000 |

Question: what can be removed? Needs input from collaboration.

- AOB
-
- Meeting Monday May 19 2025 1PM EDT
 - Dry run data production
 - Required update to the full naming convention

- (dsttype)_(daqhost)_(serverid)_(runtype)_(build)_(cdbtag)-(run)-(segment)
 - Applied to the streaming data. Even the intt and mvtx will be saved with server ID = 0
 - Triggered workflows will also need to set a server id... (else we have two naming conventions)
- Each of the 48 ebdc prdfs is now processed by its own job.

Run summary

| dstgroup | runtype | nsubmitted | nstarted | perc | nran | perc | nfailed | perc | nheld | perc | nfinished | perc |
|--------------------------|----------|------------|----------|------|------|------|---------|------|-------|------|-----------|------|
| DST_STREAMING_EVENT_intt | run3auau | 599 | 599 | 100 | 599 | 100 | 0 | 0 | 0 | 0 | 599 | 100 |
| DST_STREAMING_EVENT_mvtx | run3auau | 407 | 407 | 100 | 407 | 100 | 36 | 9 | 0 | 0 | 371 | 91 |
| DST_STREAMING_EVENT_ebdc | run3auau | 2823 | 2823 | 100 | 2823 | 100 | 6 | 0 | 0 | 0 | 2695 | 95 |

94 ebdc jobs still running. 28 jobs have gone on hold. 6 failed. So 94+28+6 = 128 matches the 2823 - 2695 = 128 jobs “not finished”.

Table broken down by [host and server](#)

- Overall job efficiency looks correct...
 - I don’t see the same number of jobs submitted per ebdc and per ebdc server ID... Seems to be three groups. 71 and 54 jobs, 27 and 26 jobs, 37 and 31 jobs. Normal?

| | | | | | | | | | | | | |
|------------------------------|----------|----|----|-----|----|-----|---|---|---|---|----|-----|
| DST_STREAMING_EVENT_ebdc10_0 | run3auau | 71 | 71 | 100 | 71 | 100 | 0 | 0 | 0 | 0 | 65 | 92 |
| DST_STREAMING_EVENT_ebdc10_1 | run3auau | 54 | 54 | 100 | 54 | 100 | 0 | 0 | 0 | 0 | 54 | 100 |
| ... | | | | | | | | | | | | |
| DST_STREAMING_EVENT_ebdc21_0 | run3auau | 27 | 27 | 100 | 27 | 100 | 0 | 0 | 0 | 0 | 27 | 100 |
| DST_STREAMING_EVENT_ebdc21_1 | run3auau | 26 | 26 | 100 | 26 | 100 | 0 | 0 | 0 | 0 | 26 | 100 |
| .. | | | | | | | | | | | | |
| DST_STREAMING_EVENT_ebdc13_0 | run3auau | 37 | 37 | 100 | 37 | 100 | 0 | 0 | 0 | 0 | 36 | 97 |
| DST_STREAMING_EVENT_ebdc13_1 | run3auau | 31 | 31 | 100 | 31 | 100 | 0 | 0 | 0 | 0 | 31 | 100 |

→ reflects issue with transfer of files to sdcc

- Need to setup the triggered workflow to append the server ID=0 as well (eg seb17_0), otherwise it will be complicated to setup the input queries for the jet workflows.
- The htcondor python API is throwing a deprecation error when performing a schedd query.
 - Submission is unaffected
 - Workaround is to switch to the new htcondor2 module for queries
 - Will split off as a separate process the code which monitors condor for jobs going on hold.
- TODO: combine runtype line laser (not downstream)
- 12k triggered job event builders on prod01 from last years run w/ old naming convention
- Expectations for beam... May 29th
- Rerunning run2pp clustering seeding.
 - Tried, but clustering jobs 12GB in memory after 30min. Debugging. Issue with CDB loading in chn map. Then continuous growth.
 - Need to debug b/c this will be a bigger issue with AuAu
- AOB
- **Meeting Monday May 12 2025 1PM EDT**
 - Dry run

- 48 “dryrun” runs taken Friday-Saturday, first run # 64058 last run # 65105.
- Calorimeter processing is done, but 91 calofitting jobs evicted and put on hold due to exceeding 2G memory limit (out of a bit over 800 jobs total).
 - (incomplete, needs job selection) snippet for memory usage: `condor_history -format "%d." ClusterId -format "%d\t" ProcId -format "%d\n" MemoryUsage`
 - https://portal.osg-htc.org/documentation/htc_workloads/submitting_workloads/monitor_review_jobs/
 - We will need this for the future - this should become part of the regular production running
 - best as a separate daemon
- 38 streaming event building jobs still running. Few downstream tracking segments were processed. They are not really needed?
 - TPOT jobs are not being submitted. Reason TBD (hopefully in the next hour...)
- Data volumes adjusted by reducing calorimeter thresholds and turning off TPC ZS to be similar to what we expect for AuAu collisions. Daq ran at 9kHz, which is about twice what we had last year.
-
- TPC issues
 - Data is being written by the daq as tpc10_0_line_laser-00064099-0000.evt files.
 - The production script does need to be updated to minimally extract the tpc10_0 and map it into the tpc file list.
 - For each EBDC, there are two sequence of files, e.g.
 - **Serverid 0 list:** SELECT runnumber, hostname, serverid, sequence, filename FROM filelist WHERE (runnumber = 64097 AND hostname NOT LIKE '%:%' AND hostname LIKE '%ebdc%')

| runnumber | hostnam | serverid | sequence | filename |
|-----------|---------|----------|----------|--|
| 64097 | ebdc00 | 0 | 0 | /bbox/bbox0/W/tpc/dryrun/TPC_ebdc00_0_line_laser-00064097-0000.evt |
| 64097 | ebdc01 | 0 | 0 | /bbox/bbox1/W/tpc/dryrun/TPC_ebdc01_0_line_laser-00064097-0000.evt |
| 64097 | ebdc02 | 0 | 0 | /bbox/bbox2/W/tpc/dryrun/TPC_ebdc02_0_line_laser-00064097-0000.evt |
| 64097 | ebdc03 | 0 | 0 | /bbox/bbox3/W/tpc/dryrun/TPC_ebdc03_0_line_laser-00064097-0000.evt |

- **Serverid 1 list:** SELECT runnumber, hostname, serverid, sequence, filename FROM filelist WHERE (runnumber = 64097 AND hostname LIKE '%ebdc%' AND serverid = 1)ORDER

BY hostname

| runnumber | hostname | serverid | sequence | filename |
|-----------|----------|----------|----------|--|
| 64097 | ebdc00:1 | 1 | 0 | /bbox/bbox0/W/tpc/dryrun/TPC_ebdc00_1_line_laser-00064097-0000.evt |
| 64097 | ebdc01:1 | 1 | 0 | /bbox/bbox1/W/tpc/dryrun/TPC_ebdc01_1_line_laser-00064097-0000.evt |
| 64097 | ebdc02:1 | 1 | 0 | /bbox/bbox2/W/tpc/dryrun/TPC_ebdc02_1_line_laser-00064097-0000.evt |
| 64097 | ebdc03:1 | 1 | 0 | /bbox/bbox3/W/tpc/dryrun/TPC_ebdc03_1_line_laser-00064097-0000.evt |
| 64097 | ebdc04:1 | 1 | 0 | /bbox/bbox4/W/tpc/dryrun/TPC_ebdc04_1_line_laser-00064097-0000.evt |
| 64097 | ebdc05:1 | 1 | 0 | /bbox/bbox5/W/tpc/dryrun/TPC_ebdc05_1_line_laser-00064097-0000.evt |

- This breaks the naming convention expected by the production scripts.
- TPC files are not being added to the file list used by the streaming event builder.
- Kaboom
 - 100k events every 2 s... I thought there was protection in the jobs against running w/ only GL1.
- Identified a possible issue w/ the production status tables... using *int* rather than *bigint* to track event count for each run. Probably not an issue there. But potential issue with the unique ID of the

- Update on cosmics reproduction
- Spider to populate filecatalog (discussion)
 - Naive `ls` shouldn't be used on lustre - replacement `lfs find`? Will check with Jane
 - Concept: write to name.root.partial, only change the name when file is correctly transferred.
- There is something very wrong with this approach, new segment after 2sec is not sustainable:
 - This job is writing 2 files every second. It's already on segment 30k+. And I'm seeing similar output rates on all of our recent jobs. Both before and during the non zero suppressed data taking. We're running 1k streaming event builders plus downstream processing. So... 500Hz (building) + 500Hz (cluster) + 500Hz (seed) + 500Hz (track) = 2kHz of DSTs being written to disk. Writes to the production DB will be something like 500 Hz + 10* (1500 Hz)... ~ 15kHz.
-
- New calorimeter jet workflow
- AOB
- Offline pp production with DST_JET and DST_JETCALO with ana484 2024p018_v001.
 - 2 held jobs on sphnxprod01 600+ held jobs on sphnxprod02, similar amount of submissions on each node.
 - Should be a representative sample of runs on each submission host
 -
 -

● Meeting Monday May 05 2025 1PM EDT

- ...
- Switch to new naming convention (stream name in lowercase) for event builders etc...

Production=> select id,dstfile,status,started,ended,nevents from production_status where dstfile like 'DST_TRIGGERED_EVENT_TEST%';

| id | dstfile | status | started | ended | nevents |
|---------|--|----------|---------------------|---------------------|----------|
| 8543167 | DST_TRIGGERED_EVENT_TEST_seb16_run3auau_new_nocdbtag_v000-00062512-00000 | finished | 2025-05-03 00:25:09 | 2025-05-05 07:05:56 | 39470510 |
| 8543169 | DST_TRIGGERED_EVENT_TEST_seb18_run3auau_new_nocdbtag_v000-00062512-00000 | finished | 2025-05-03 00:25:31 | 2025-05-04 02:47:59 | 39470510 |
| 8543168 | DST_TRIGGERED_EVENT_TEST_seb17_run3auau_new_nocdbtag_v000-00062512-00000 | finished | 2025-05-03 00:25:15 | 2025-05-05 06:03:11 | 39470510 |
| 8543170 | DST_TRIGGERED_EVENT_TEST_seb20_run3auau_new_nocdbtag_v000-00062512-00000 | finished | 2025-05-03 00:25:31 | 2025-05-04 01:37:38 | 39470510 |

- Should be ready to put into production in the triggered workflow.
 - Tested through calofitting for one run.
 - Streaming should be ready to go as well (just not tested beyond seeing output of the builder).
 - Should discuss how to place this into production.
 - This will create files with identical contents in the streaming workflows... i.e.
DST_STREAMING_EVENT_INTT0_run3auau-run-seg will be identical to
DST_STREAMING_EVENT_intt0_run3auau-run-seg.
 - Do we want to elevate the version number in the downstream jobs? New CDB tag on the horizon? Presume we should break at the last submitted cosmics run.
- This required
 - updates to the stageout filename matching
 - passing runtime explicitly through the query
 - some straightforward modifications to the input query
- Breaks the regex used in building the PFN list from the logical file list. (No more need to build PFN list, so will retire that feature).
- Also may have broken the regex which is used in querying condor to flag and reset held jobs. TBD.
- Also now using the hostinfo table to build the required list of DSTs for downstream processing and the run table to select runs.
 - The most straightforward rewrite of the original queries turned out to be too slow. Had to rewrite the input queries using joins and a few stages of filtering to get the performance we need.
- Input queries for the event builders select on the run number and daq host. Any filename / dataset name will be picked up for processing from the daq.

- Run Summary

| Prior week | | | | | | | | | | | | | |
|--------------------------|----------|------------|----------|------|--------|------|---------|------|-------|------|-----------|------|-------------|
| dstgroup | runtype | nsubmitted | nstarted | perc | nran | perc | nfailed | perc | nheld | perc | nfinished | perc | |
| DST_STREAMING_EVENT_INTT | run3auau | 10373 | 10277 | 99 | 10277 | 99 | 0 | 0 | 0 | 0 | 9944 | 96 | |
| DST_STREAMING_EVENT_MVTX | run3auau | 5163 | 5127 | 99 | 5127 | 99 | 49 | 1 | 80 | 2 | 4822 | 93 | |
| DST_STREAMING_EVENT_TPOT | run3auau | 1073 | 1031 | 96 | 1031 | 96 | 1 | 0 | 1 | 0 | 852 | 79 | |
| DST_STREAMING_EVENT_TPC | run3auau | 28523 | 28025 | 98 | 28025 | 98 | 4 | 0 | 405 | 1 | 25337 | 89 | |
| DST_TRKR_CLUSTER | run3auau | | 264136 | | 263671 | 100 | 263671 | 100 | 1853 | 1 | 135 | 0 | 258059 98 |

| | | | | | | | | | | | | | |
|--------------------------|-------------|----|--------|--------|-----|--------|-----|-------|-----|-----|-----|--------|-----|
| DST_TRKR_SEED | run3auau | | 259783 | 259541 | 100 | 259541 | 100 | 0 | 0 | 71 | 0 | 259444 | 100 |
| DST_TRKR_TRACKS | run3auau | | 259506 | 258551 | 100 | 258551 | 100 | 159 | 0 | 6 | 0 | 258347 | 100 |
| DST_TRIGGERED_EVENT | run3auau | | 222723 | 206328 | 93 | 206328 | 93 | 8720 | 4 | 138 | 0 | 174729 | 78 |
| DST_CALOFITTING | run3auau | | 172382 | 172185 | 100 | 172185 | 100 | 6075 | 4 | 67 | 0 | 165704 | 96 |
| DST_CALO | run3auau | | 171736 | 171734 | 100 | 171734 | 100 | 4216 | 2 | 19 | 0 | 167415 | 97 |
| DST_JETS | run3auau | | 75119 | 75107 | 100 | 75107 | 100 | 75104 | 100 | 3 | 0 | 0 | 0 |
| DST_TRIGGERED_EVENT_SEB | run3auau | 20 | 20 | 100 | 20 | 100 | 0 | 0 | 0 | 20 | 100 | | |
| DST_STREAMING_EVENT_TEST | run3auau | 14 | 14 | 100 | 14 | 100 | 0 | 0 | 0 | 14 | 100 | | |
| DST_TRIGGERED_EVENT_TEST | run3auau | 4 | 4 | 100 | 4 | 100 | 0 | 0 | 0 | 4 | 100 | | |
| DST_STREAMING_EVENT_TPC | run3calib | | 1301 | 1300 | 100 | 1300 | 100 | 598 | 46 | 14 | 1 | 688 | 53 |
| DST_TRKR_CLUSTER | run3calib | | 215 | 215 | 100 | 215 | 100 | 0 | 0 | 0 | 0 | 215 | 100 |
| DST_TRIGGERED_EVENT | run3cosmics | | 223693 | 223278 | 100 | 223277 | 100 | 8531 | 4 | 316 | 0 | 190073 | 85 |
| DST_CALOFITTING | run3cosmics | | 185700 | 185450 | 100 | 185450 | 100 | 7472 | 4 | 369 | 0 | 177345 | 96 |

To date

| dstgroup | runtype | nsubmitted | nstarted | perc | nran | perc | nfailed | perc | nheld | perc | nfinished | perc |
|--------------------------|-------------|------------|----------|------|-------|------|---------|------|-------|------|-----------|------|
| DST_STREAMING_EVENT_INTT | run3auau | 2953 | 2920 | 99 | 2920 | 99 | 0 | 0 | 0 | 0 | 2707 | 92 |
| DST_STREAMING_EVENT_MVTX | run3auau | 2192 | 2165 | 99 | 2165 | 99 | 25 | 1 | 80 | 4 | 1886 | 86 |
| DST_STREAMING_EVENT_TPOT | run3auau | 99 | 97 | 98 | 97 | 98 | 0 | 0 | 1 | 1 | 71 | 72 |
| DST_STREAMING_EVENT_TPC | run3auau | 3900 | 3863 | 99 | 3863 | 99 | 1 | 0 | 405 | 10 | 1179 | 30 |
| DST_TRKR_CLUSTER | run3auau | 23674 | 23647 | 100 | 23647 | 100 | 1840 | 8 | 0 | 0 | 21302 | 90 |
| DST_TRKR_SEED | run3auau | 22903 | 22896 | 100 | 22896 | 100 | 0 | 0 | 0 | 0 | 22896 | 100 |
| DST_TRKR_TRACKS | run3auau | 22896 | 22884 | 100 | 22884 | 100 | 0 | 0 | 0 | 0 | 22884 | 100 |
| DST_TRIGGERED_EVENT | run3auau | 22156 | 22155 | 100 | 22155 | 100 | 1091 | 5 | 0 | 0 | 18613 | 84 |
| DST_CALOFITTING | run3auau | 14596 | 14594 | 100 | 14594 | 100 | 247 | 2 | 0 | 0 | 14347 | 98 |
| DST_CALO | run3auau | 15289 | 15289 | 100 | 15289 | 100 | 0 | 0 | 0 | 0 | 15289 | 100 |
| DST_JETS | run3auau | 15058 | 15058 | 100 | 15058 | 100 | 15058 | 100 | 0 | 0 | 0 | 0 |
| DST_TRIGGERED_EVENT_SEB | run3auau | 20 | 20 | 100 | 20 | 100 | 0 | 0 | 0 | 0 | 20 | 100 |
| DST_STREAMING_EVENT_TEST | run3auau | 14 | 14 | 100 | 14 | 100 | 0 | 0 | 0 | 0 | 14 | 100 |
| DST_TRIGGERED_EVENT_TEST | run3auau | 4 | 4 | 100 | 4 | 100 | 0 | 0 | 0 | 0 | 4 | 100 |
| DST_STREAMING_EVENT_TPC | run3calib | 264 | 264 | 100 | 264 | 100 | 0 | 0 | 0 | 0 | 264 | 100 |
| DST_TRKR_CLUSTER | run3calib | 40 | 40 | 100 | 40 | 100 | 0 | 0 | 0 | 0 | 40 | 100 |
| DST_TRIGGERED_EVENT | run3cosmics | 21128 | 21128 | 100 | 21128 | 100 | 1092 | 5 | 0 | 0 | 17585 | 83 |
| DST_CALOFITTING | run3cosmics | 13576 | 13575 | 100 | 13575 | 100 | 235 | 2 | 0 | 0 | 13340 | 98 |

(18 rows)

- cups.py and/or stageout has been broken since last week. Only the last DST segment is actually copied to lustre.
 - https://chat.sdcc.bnl.gov/sphenix/pl/5qmjzd8tqtgx9kmegtcg5ropgo
 - Tracking productions
 - The raw hit production for run2pp is ongoing. Should finish mid week. Need to check if files are in the file catalog, Chris mentioned there were a nontrivial number that were not
 - At that point we will submit clustering jobs. ~75k jobs
 - Analyzers working on tpc calibrations before seeding jobs start, but that will take over >½ the farm for ~4 days
 - Need to reprocess sizeable chunk of cosmics data from the last month due to event combiner issues in silicon
 - Lustre I/O questions
 - The tracking event combining currently writes 100k events per file. We will never be able to process 100k events through jobA in any reasonable amount of time. I believe the change was made to reduce writing to lustre frequently - solutions?
 - 10k is probably the maximum
 - Job A 15s/event → 10k events @ 48h (best case). Probably x2 off of the best case.
 - Aim for 24h jobs. Job A stage the input files if practicable. 2500-5000 events. Run2AuAu should be informative (30s/event there).
 - 3k segments in pp. Nominal expectation is 54M events in 1h (1 run)
 - “Dry-run” on Thursday/Friday
 - Confirm w/ daq what the runtype is.
 - TODO... Ping Dmitry A.
 - AOB
 - Cleanup obsolete productions
- Meeting Monday Apr 28 2025 1PM EDT
 - Old Business
 - HOLDS are now marked in the production status table and *cleared*...
 - See summary below (or the 4/14 meeting until I post it...) for the number of jobs held compared to the total number submitted.
 - Qualitative observation is that (except for the mvtx) we have a reasonable rate of jobs going on hold.
 - We have also been able to recover some held jobs by resubmitting with higher memory allowance
 - See “message” field in the production status table for hold reason
 - Online Production
 - DST_TRIGGERD_EVENT_run3auau and cosmics switched to raw data catalog as input. DST_STREAMING query is ready... just need to integrate.
 - Summary... Percentage of MVTX and TPC jobs going on hold has increased. (May simply require resubmitting w/ higher memory allowance. TBD).

| Past Week | | | | | | | | | | | | | |
|--------------------------|-------------|------------|----------|------|-------|------|---------|------|-------|------|-----------|-------|----|
| dstgroup | runtype | nsubmitted | nstarted | perc | nran | perc | nfailed | perc | nheld | perc | nfinished | perc | |
| DST_STREAMING_EVENT_INTT | run3auau | 3186 | 3186 | 100 | 3186 | 100 | 0 | 0 | 9 | 0 | 2915 | 91 | |
| DST_STREAMING_EVENT_MVTX | run3auau | 910 | 910 | 100 | 910 | 100 | 50 | 5 | 165 | 18 | 677 | 74 | |
| DST_STREAMING_EVENT_TPOT | run3auau | 351 | 351 | 100 | 351 | 100 | 0 | 0 | 9 | 3 | 230 | 66 | |
| DST_STREAMING_EVENT_TPC | run3auau | 8815 | 8814 | 100 | 8814 | 100 | 0 | 0 | 1771 | 20 | 7017 | 80 | |
| DST_TRKR_CLUSTER | run3auau | 77347 | 77345 | 100 | 77345 | 100 | 3 | 0 | 108 | 0 | 77006 | 100 | |
| DST_TRKR_SEED | run3auau | 77636 | 77634 | 100 | 77634 | 100 | 0 | 0 | 29 | 0 | 77605 | 100 | |
| DST_TRKR_TRACKS | run3auau | 77557 | 77554 | 100 | 77554 | 100 | 1 | 0 | 2 | 0 | 77550 | 100 | |
| DST_TRIGGERED_EVENT | run3auau | 85572 | 71976 | 84 | 71976 | 84 | 5238 | 6 | 65 | 0 | 60097 | 70 | |
| DST_CALOFITTING | run3auau | 61914 | 61912 | 100 | 61912 | 100 | 609 | 1 | 52 | 0 | 60035 | 97 | |
| DST_CALO | run3auau | 60647 | 60645 | 100 | 60645 | 100 | 27 | 0 | 14 | 0 | 60441 | 100 | |
| DST_JETS | run3auau | 60534 | 60532 | 100 | 60532 | 100 | 60529 | 100 | 3 | 0 | 0 | 0 | |
| DST_STREAMING_EVENT_TPC | run3calib | 324 | 324 | 100 | 324 | 100 | 47 | 15 | 1 | 0 | 276 | 85 | |
| DST_TRKR_CLUSTER | run3calib | 37 | 37 | 100 | 37 | 100 | 0 | 0 | 0 | 0 | 37 | 100 | |
| DST_TRIGGERED_EVENT | run3cosmics | 83711 | 83708 | 100 | 83708 | 100 | | 5000 | 6 | 35 | 0 | 70688 | 84 |
| DST_CALOFITTING | run3cosmics | 70450 | 70329 | 100 | 70329 | 100 | 325 | 0 | 17 | 0 | 69670 | 99 | |

| Since Online Began | | | | | | | | | | | | | |
|--------------------------|-------------|------------|----------|------|--------|------|---------|------|-------|------|-----------|------|--|
| dstgroup | runtype | nsubmitted | nstarted | perc | nran | perc | nfailed | perc | nheld | perc | nfinished | perc | |
| DST_STREAMING_EVENT_INTT | run3auau | 9870 | 9728 | 99 | 9728 | 99 | 31 | 0 | 26 | 0 | 9211 | 93 | |
| DST_STREAMING_EVENT_MVTX | run3auau | 4914 | 4861 | 99 | 4861 | 99 | 121 | 2 | 284 | 6 | 4420 | 90 | |
| DST_STREAMING_EVENT_TPOT | run3auau | 1009 | 969 | 96 | 969 | 96 | 2 | 0 | 13 | 1 | 800 | 79 | |
| DST_STREAMING_EVENT_TPC | run3auau | 27050 | 26589 | 98 | 26589 | 98 | 11 | 0 | 2146 | 8 | 24405 | 90 | |
| DST_TRKR_CLUSTER | run3auau | 241503 | 241065 | 100 | 241065 | 100 | 13 | 0 | 135 | 0 | 237648 | 98 | |
| DST_TRKR_SEED | run3auau | 237737 | 237502 | 100 | 237502 | 100 | 0 | 0 | 71 | 0 | 237405 | 100 | |
| DST_TRKR_TRACKS | run3auau | 237392 | 236449 | 100 | 236449 | 100 | 159 | 0 | 6 | 0 | 236245 | 100 | |
| DST_TRIGGERED_EVENT | run3auau | 201595 | 185201 | 92 | 185201 | 92 | 7538 | 4 | 138 | 0 | 157140 | 78 | |
| DST_CALOFITTING | run3auau | 158875 | 158680 | 100 | 158680 | 100 | 5840 | 4 | 67 | 0 | 151499 | 95 | |
| DST_CALO | run3auau | 157294 | 157292 | 100 | 157292 | 100 | 4216 | 3 | 19 | 0 | 152894 | 97 | |
| DST_JETS | run3auau | 60870 | 60858 | 100 | 60858 | 100 | 60855 | 100 | 3 | 0 | 0 | 0 | |
| DST_STREAMING_EVENT_TPC | run3calib | 1037 | 1036 | 100 | 1036 | 100 | 598 | 58 | 14 | 1 | 424 | 41 | |
| DST_TRKR_CLUSTER | run3calib | 175 | 175 | 100 | 175 | 100 | 0 | 0 | 0 | 0 | 175 | 100 | |
| DST_TRIGGERED_EVENT | run3cosmics | 203426 | 203011 | 100 | 203010 | 100 | 7373 | 4 | 316 | 0 | 173204 | 85 | |
| DST_CALOFITTING | run3cosmics | 173061 | 172812 | 100 | 172812 | 100 | 7250 | 4 | 369 | 0 | 164876 | 95 | |

- New DST_TRIG event builder / workflow...
 - Simplified event builder / much faster / no recoveries yet. Once an seb goes out of sync it is dropped. Tower building works based on it. (And is identical). Two more mods incoming. Fitting step update coming. 20M events/h. Fitting step... reads all event files. Should not see infinite loops. Downstream fitting is robust
 - The workflow is essentially the same as what we use in streaming.
 - Submit jobs by subsystem (SEB01 to SEB20) across all input segments in the run (plus gl1)
 - Script will accumulate the runlist and run the macro.
 - Jobs will periodically write their own output.
 - Follow the streaming naming convention... e.g.
 - DST_STREAMING_EVENT_INTTn_run3aua etc...
 - DST_TRIGGERED_EVENT_SEBnn_run3auau
 - ... or do we want to convert SEBnn to a human readable name, e.g. EMCnn, HCALnn, etc?
 - No, so the daq hostlist can be used to derive filenames and check for missing files
 - Where will the SEBs be merged? At the calo fitting stage? Or do we keep the parallelization until later?
 - At the fitting stage
 - Macro/code is undergoing final test today
 - Anticipate test jobs in online production by end of Monday. In production Tuesday (?) (Not accounting for the Scotty principle).
 - However... query seems to be a bit on the slow side... may not be well optimized at the start. O(500ms/run).
 - Any interest in resubmitting the full cosmics dataset? Or just from a recent run?
 - 62456 is the first run after the HCal timing issue resolved
 - MBD production... at what stage is this done?
 - Incoming change. TPC (and TPOT?) EDBC will be writing two files. So number of files will increase. edbc_0 or _1 or ...
 - edbc00 ... edbc00_0 all same hostname
 - SLURP updates
 - We can now import data from multiple databases when running the input query. e.g. if we are using the raw data catalog for input, we can query the daq db for the list of runs, and the number of events per daqhost. The selection which is returned will result in a temporary table being created when we run the input query against the raw data catalog. We can use this temporary table to help run our selections.
 - What is going on with those 3+k Dsts which are not in the file catalog?
 - This may/may not be related to why in the initial attempt at a larger production we had >22k raw hit DSTs but when we tried submitting the clustering jobs only 5k jobs went through
 - Listed on mattermost. April 24th. 20 runs of raw hit DSTs.
 - TODO... update stageout to be silent (no logging) until both copy back and file catalog. Thinking is that cups got interrupted when trying to write to stdout right after copying the file to lustre, but before it updated the file catalog.
 - QA
 - Every few seconds the qa segfaults on sphnxdev01 in libqadrawmvtx, dmesg -T:
[Mon Apr 28 10:11:11 2025] root.exe[4181593]: segfault at 0 ip 00007fa4709bfe34 sp 00007ffde99238a0 error 4 in libqadrawmvtx.so.0.0.0[7fa4709ba000+12000] likely on CPU 86 (core 18, socket 1)
 - Joe: I can look into it
 - Status of 1TB calo log (/sphenix/user/sphnxpro/htmllogs/makehtmlcalo.log)?
 - AOB
 - Meeting Monday Apr 21 2025 1PM EDT
 - File Transfer
 - \$ find /sphenix/data/data02/sphnxpro/production/run3cosmics/cosmics/new_2025p000_v000/DST_CALOFITTING/run_* -name *err | xargs grep -l quota|wc -l → 59645
 - Clean up data02 further, delete everything (logs, hists) older than this year, only keep in the future things related to publications
 - Offline:
 - Who is running from sphnxprod01 (runy2fitting.sh, ca. 7k running, 10k total)? Why?
 - (sphnxprod01) PPG-03 asked for reprocessing of some run2auau runs. Event combining and waveform fitting are close to complete. DST_CALO next.
 - (sphnxprod01) run2pp calorimeter prdfs copied from HPSS previously not processed. Event combining done. Surprisingly few jobs on hold < 2% (the same for run2auau). But for run2auau a lot of waveform fitting jobs failed (not happening before?)
 - Hold: used to be 20ish%? Maybe follow up about the reason → quality of raw data improved, may well be the explanation

- Also, problematic jobs removed manually
 - Waveform fitting: seems to be caused by disconnect to the job (node crashed), O(3k)!
 - Ex. From Shuhang:
- /tmp/production/run2auau/physics/ana468_2025p001_v001/DST_CALOFITTING/run_0005480000054900/DST_CALOFITTIN
G_run2auau_ana468_2025p001_v001-00054863-00793.condor
- Event combining exhibited similar behavior last week, ana479 streaming production has many failures and jobs that seem to just have crashed/stopped running. Related to /data02/ issue?
 - Did many nodes die? `dmesg -T` on an affected node would reveal more
 - Joe "I believe I submitted those jobs under sphnxprod02 and the job number was 15799 in case you want to check those out"
- Online Cosmics Production
 - Switched to tag 2025p001 from run 61988.
 - Remaining jobs from prior runs using 2025p000. Second loop setup running 1 every 5 min.
 - Production status table is now in good shape to do job accounting.
 - Removed duplicate entries.
 - Resubmitted job now clears the job it replaces.
 - No jobs running over the weekend (not even event combining). Not sure what is going on as files are transferred to sdcc
 - 200 runs behind, need to figure out what's happening from logs in tmp
 - Slurp updates
 - Automatic masking of detectors which are known to be off in development.
 - Requires a scheme to make information from the daq db (hostinfo table) available to support the input query against the file catalog.
 - Similar scheme needed to support raw data catalog queries, where we need the run table from daqdb to accelerate the queries.
 - Expect to have these in place by the end of the week.
 - AOB
 - Why do jobs stay in Idle for hours while only 8k slots are claimed?
 - Ex.: sphnxuser08 jobid: 71079.0
- **Meeting Monday Apr 14 2025 1PM EDT**
 - File Transfer Issues
 - Online Cosmics Production
 - Reproduction of data prior to 3/31 appears to have finished. Dedicated loop has been stopped
 - Production loop has been stable with no unplanned pauses.
 - I've resubmitted many of the held jobs which failed due to memory issue. Many appear to have succeeded.
 - There have been quite a few jobs which tripped the periodic hold condition.
 - It looks like condor does try to restart some jobs that go on hold.
 - I have not seen any instances of jobs running beyond the 10 day limit.
 - Jobs that have the periodic hold set have mostly been recovered.
 - But we still see jobs that are going on hold that do not (easily) recover. Specifically the MVTX.
 -

| Online Production Jobs | | | | | | | | | | | | | |
|--------------------------|-------------|------------|----------|------|--------|------|---------|------|-------|------|-----------|--------|-----|
| dstgroup | runtype | nsubmitted | nstarted | perc | nran | perc | nfailed | perc | nheld | perc | nfinished | perc | |
| DST_CALO | run3auau | 102565 | 102561 | 100 | 102561 | 100 | 7055 | | 7 | 46 | 0 | 95459 | 93 |
| DST_CALOFITTING | run3auau | 170821 | 164137 | 96 | 164137 | 96 | 9396 | | 6 | 738 | 0 | 150256 | 88 |
| DST_JETCALO | run3auau | 3037 | 2798 | 92 | 2798 | 92 | 2793 | | 92 | 0 | 0 | 5 | 0 |
| DST_JETS | run3auau | 1593 | 1590 | 100 | 1590 | 100 | 1568 | | 98 | 6 | 0 | 16 | 1 |
| DST_STREAMING_EVENT_INTT | run3auau | 6491 | 6482 | 100 | 6482 | 100 | 32 | | 0 | 35 | 1 | 6401 | 99 |
| DST_STREAMING_EVENT_MVTX | run3auau | 7438 | 7426 | 100 | 7426 | 100 | 52 | | 1 | 2102 | 28 | 5229 | 70 |
| DST_STREAMING_EVENT_TPC | run3auau | 28148 | 28126 | 100 | 28126 | 100 | 13 | | 0 | 1121 | 4 | 26967 | 96 |
| DST_STREAMING_EVENT_TPOT | run3auau | 549 | 549 | 100 | 549 | 100 | 2 | | 0 | 0 | 0 | 547 | 100 |
| DST_TRIGGERED_EVENT | run3auau | 120873 | 117201 | 97 | 117201 | 97 | 2614 | | 2 | 831 | 1 | 96217 | 80 |
| DST_TRKR_CLUSTER | run3auau | 316779 | 315327 | 100 | 315273 | 100 | 14125 | | 4 | 404 | 0 | 274707 | 87 |
| DST_TRKR_SEED | run3auau | 204856 | 204574 | 100 | 204574 | 100 | 977 | | 0 | 2727 | 1 | 200836 | 98 |
| DST_TRKR_TRACKS | run3auau | 179555 | 178617 | 99 | 178617 | 99 | 657 | | 0 | 82 | 0 | 177839 | 99 |
| DST_STREAMING_EVENT_TPC | run3calib | 24 | 24 | 100 | 24 | 100 | 0 | | 0 | 0 | 0 | 24 | 100 |
| DST_TRKR_CLUSTER | run3calib | 20 | 20 | 100 | 20 | 100 | 0 | | 0 | 0 | 0 | 9 | 45 |
| DST_CALOFITTING | run3cosmics | 88506 | 86709 | 98 | 86709 | 98 | 15876 | 18 | 1744 | 2 | 69078 | 78 | |
| DST_TRIGGERED_EVENT | run3cosmics | 59254 | 58852 | 99 | 58851 | 99 | 1993 | | 3 | 1430 | 2 | 46919 | 79 |

- Slurp / Ramenya improvements
 - The production loop catches ctrl-c interrupts. Initiates a 20 min pause, which will resume if not additional action taken. User can terminate or resume the loop during the 20min window.
 - Jobs which went on hold can be resubmitted by hand with higher memory request using `kaedama.py -mem 8192MB -config ... -rule ... -runs ... -unblock held -mark-held-jobs -clear-held-jobs`
 - The clear-held-jobs instructs kaedama to remove the jobs from the condor queue
 - The mark-held-jobs should be issued so that any jobs that have gone on hold since the last iteration of the production loop get flagged in the production status before being cleared. TODO: clear should imply mark.
 - Added support for runs with nonstandard sets of detectors... e.g. to run clustering w/out TPC17, INTT3 and MVTX1... `kaedama -mask TPC17 INTT3 MVTX1 -rule DST_TRKR_CLUSTER_run3auau ...`
 - ... works by removing from a list of required subdetectors
 - ... will only pick up jobs which exactly match the requirements (in this case all subdetectors except those in the mask list)
 - Plan is to automate this so that when runs are taken in a special condition, the downstream processing can account for it. (e.g. missing hosts and/or entire detectors off and known to be off.)
 - hostinfo table on daqdb

New triggered event combing is getting ready

Running 2 jobs - cemc and all others too slow for cemc (100k events per hour), jobs run over all events
No error recovery yet - jobs terminate gracefully on event misalignments. (saving everything up to this event) to select candidates for development.
Will be split into packet ranges (single or multiple input files - looking into flexible naming convention)
Subsequent processing needs to change their packet lookup
Will run under separate production system


- AOB
 - Naming convention for DST_*_run3auau vs DST_*_run3cosmics...
 -
 - (DMA) Work on Jet QA/production macros delayed due to personal travel. Aiming for PRs this week...
- Meeting Monday Apr 07 2025 1PM EDT
 - Online Production
 - Reproduction of data prior to 3/31 commenced.
 - Online production loop paused Friday night while the disk and filecatalog were cleaned up.
 - Second production loop was set up to manage the reproduction.
 - Online production loop was not brought back online until Saturday.
 - Jobs seem to be culminating
 - Online production loop was killed (excessive memory consumption) on Saturday night.
 - Issue was the JOB (dsttype,build,tag,ver,run,segment) : STATUS map built in kaedama
 - Required to determine if the proposed JOB has already been submitted to kaedama or not
 - Original implementation retrieved and held in memory all of the columns for each existing JOB
 - Reducing the query to just the proposed JOB and status (e.g. dropping the input filelist) solved the issue
 - We were routinely keeping 50k cores fed over the weekend. We are down to around 30k (9AM)... 20k (nowish).
 - Partly b/c the reproduction is culminating (seeing O(1) job submitted every 2m cycle)
 - Partly b/c jobs were being resubmitted when they were going on hold...
 - When the hold state was polled from condor and added to the production status table, the corresponding 'hold' entry in kaedama's default blocklist was not updated. Fixed.
 - DST_TRKR_CLUSTER jobs are now set to 8GB.
 - Between that and fixing the systems hold issue, only seeing 16 jobs go on hold in the last day....
 - Offline production
 - Missing runs from run2pp transferred to lustre / including replacing corrupted files. Able to rerun pp.
 - Compare run table in daqdb to the number of events produced as recorded in the datasets / file catalog
 - New CPUS...
 - prod01 prod02 ... 3 more requested
 -
 - Old Business
 - INTT calibration workflow...
 - Jet QA macros...
 - AOB

- Meeting Monday Mar 31 2025 1PM EDT
 - INTT timing calibration workflow
 - Simple module that can be added on to end of streaming event combining.
 - Produces some output - can we orchestrate this output? Where do other calibration artifacts go?
 - (Do we want to setup a calib/ (cdb/) directory next to the log/ and hist/ on data02? dst/ on lustre?)
 - Plan for calib or cdb directory.
 - CALIB_ ... for the output naming convention
 - If this is just a single number ... end product is an entry in the INTT database rather than a cdb file.
 - Here is the module :
<https://github.com/SPHENIX-Collaboration/coresoftware/blob/master/calibrations/intt/inttcalib/InttCalib.cc>
 - Online Production
 - Run number to switch to tag 2025p000: 59641
 - Streaming event builders are running back to run 57600.
 - How to handle jobs before the switch over? For now I am adding a second production loop so that runs prior to 59641 are completed under the 2024p017 tag.
 - Need to add a max-run option to the run cursor argument...
 - Action item: suspend the loop for 2024p017.
 - /direct/sphenix+u/sphnxpro/slurp/slurpprod/run3auau_production_loop.yaml
 - /direct/sphenix+u/sphnxpro/slurp/slurpprod/run3auau_production_loop2.yaml
 -
 - Hit 50k concurrent jobs on 3/26 w/out significant scaling issues observed (save for race between condor and slurp discussed below...) But ran straight into one of the corollaries to Murphy's law...
 - Over the weekend we saw DB timeouts when building the DST_STREAMING jobs.
 - Run cursor strategy for the event builders should be a ratchet. Once we see a run we are never interested in building it again.
 - Differs from downstream production, where the run cursor needs to accept new jobs from running event builders and upstream codes.
 - Streaming data setup and running through track reconstruction. Inputs accepted from the physics, beam and cosmics pipelines.

```
-- Schedd: sphnxprod02.sdcc.bnl.gov : <130.199.6.175:9618?... @ 03/26/25 22:52:52
```


| OWNER | BATCH_NAME | SUBMITTED | DONE | RUN | IDLE | HOLD |
|----------|--|------------|-------|-------|------|------|
| sphnxpro | DST_TRIGGERED_EVENT_run3cosmics_new_nocdbtag_v000 | 3/24 22:27 | 1256 | 1366 | 316 | 110 |
| sphnxpro | DST_TRKR_SEED_run3auau_new_2024p012_v000 | 3/24 23:28 | 7037 | - | - | 19 |
| sphnxpro | DST_STREAMING_EVENT_INTT0_run3auau_new_nocdbtag_v000-singlestreams | 3/26 03:24 | 151 | 38 | - | - |
| sphnxpro | DST_STREAMING_EVENT_MVTX0_run3auau_new_nocdbtag_v000-singlestreams | 3/26 05:09 | 67 | 6 | - | 11 |
| sphnxpro | DST_STREAMING_EVENT_MVTX3_run3auau_new_nocdbtag_v000-singlestreams | 3/26 05:42 | 42 | 2 | - | 1 |
| sphnxpro | DST_TRKR_CLUSTER_run3auau_new_2024p017_v000 | 3/26 19:38 | 28880 | 12655 | 8586 | 480 |
| sphnxpro | DST_TRIGGERED_EVENT_run3auau_new_nocdbtag_v000 | 3/26 19:49 | 1902 | 1008 | 1767 | 53 |
| sphnxpro | DST_CALFITTING_run3auau_new_2024p017_v000 | 3/26 19:53 | 727 | 1630 | 155 | 7 |
| sphnxpro | DST_CALFITTING_run3cosmics_new_2024p017_v000 | 3/26 19:55 | 20139 | 1119 | 300 | 1689 |
| sphnxpro | DST_TRKR_SEED_run3auau_new_2024p017_v000 | 3/26 20:20 | 16941 | - | 4097 | 3059 |
| sphnxpro | DST_CAL0_run3auau_new_2024p017_v000 | 3/26 20:31 | 409 | 192 | 33 | 39 |
| sphnxpro | DST_STREAMING_EVENT_INTT6_run3auau_new_nocdbtag_v000-singlestreams | 3/26 20:54 | 29 | 4 | - | - |
| sphnxpro | DST_TRKR_TRACKS_run3auau_new_2024p017_v000 | 3/26 21:21 | 9974 | - | 8111 | 26 |
| sphnxpro | DST_3ETS_run3auau_new_2024p017_v000 | 3/26 21:34 | 241 | - | 6 | 6 |
| sphnxpro | DST_TRKR_CLUSTER_run3auau_new_2024p012_v000 | 3/26 22:42 | 7659 | 2337 | - | - |

Total for query: 48628 jobs; 0 completed, 0 removed, 23371 idle, 19757 running, 5500 held, 0 suspended
Total for sphnxpro: 48628 jobs; 0 completed, 0 removed, 23371 idle, 19757 running, 5500 held, 0 suspended
Total for all users: 48628 jobs; 0 completed, 0 removed, 23371 idle, 19757 running, 5500 held, 0 suspended



- Triggered data setup and running through jet production. Two workflows implemented. One accepting inputs from the physics, beam and cosmics pipelines. A second one accepting only cosmics, and being processed only through a custom fitting stage for hcal calibrations.
 - Production loop is running on a 10 min cycle.
 -
 - Offline Production
 - Streaming
 - Running track fitting production now, almost done. Last before Quark Matter...(?)
 - Slurp improvements. and etc...
 - The DB tag can be set at the top level configuration file for the production loop. When a new DB tag is set, jobs should begin at the current run cursor (to be tested).
 - The run cursor is now based on the parent jobs (the ones producing input for a given workflow). But it can be the case that there is no parent, e.g. the event builders. Also there are cases where parents finish during execution of the submission script. So the earliest running child will be used if its run is lower than the parent run.
 - The time required to create and submit jobs has been significantly decreased, especially when submitting jobs immediately downstream of the streaming event builder.
 - Primary gain was b/c we were able to remove significant number of DB queries in building the jobs.
 - Should reduce the race issue, where the jobs can start and finish on the worker node before kaedama has a chance to mark them as 'submitted'.
 - Triggered event builders now use the database to obtain their input filelists and event ranges
 - AOB
 - Updates to Jet QA almost complete; expect PRs on macros and production area later this week!
 -

-
- **Meeting Monday Mar 24 2025 1PM EDT**

- Run 25 first beam (?)
 - Any additional workflows to be added?
 - Attach all workflows to 'beam', physics... data, cosmics, etc...?
 - No lag... but...
 - Standard DST_TRIGGERED... for testing
 - Raw data catalog? 'Almost'
 - Aiming for first beam by Wednesday . 
 - [More details from 9 o'clock meeting](#)
- DSTs and Datasets to be retained / deleted
 - 2 recent calorimeter productions. older productions. no QM productions to be deleted.
 - plan to push everything w/ ana build and tag into hpss from '23.
 - data from '24 ... better accounting, looking to retain only what is needed to support existing papers.
 - 20 distinct productions identified.
 - datasets w/ 'new' disguised as ana build not yet cleared off... triggered was ana446 (?)
 - ana387_2023p003
 - ana388_2023p004
 - ana391_2023p006
 - ana393_2023p006
 - ana395_2023p007
 - ana399_2023p008
 - ana402_2023p009
 - ana403_2023p010
 - ana403_2023p011
 - ana407_2023p013
 - ana410_2023p014
 - ana412_2023p015
 - ana427_2024p005
 - ana439_2024p007
 - ana441_2024p007
 - ana444_2024p007
 - ana446_2024p007
 - ana449_2024p008
 - ana450_2024p009
 - ana451_2024p009
 - ana453_2024p008
 - ana457_2024p008
 - ana459_nocdbtag_v001
 - ana461_2024p009_v001
 - ana461_nocdbtag_v001
 - ana462_2024p010_v001
 - ana462_2024p012_v001
 - ana462_nocdbtag_v001
 - ana463_nocdbtag_v001
 - ana464_2024p011_v001
 - ana464_nocdbtag_v001
 - ana466_2024p012_v001
 - ana466_2024p015_v001
 - ana466_nocdbtag_v001
 - ana468_2024p012_v001
 - ana468_2024p012_v002
 - ana470_2024p012_v001
 - ana472_2024p013_v001
 - ana472_2024p014_v001
 - ana472_2024p014_v002
 - ana472_2024p014_v003
 - ana472_2024p015_v001
 - ana472_2024p015_v002
 - ana472_2024p016_v001
 - ana473_2024p016_v001
 - ana473_nocdbtag_v001
 - Joe will provide list of tracking 'test' productions that can be removed.

- Online Production
 - Hcal cosmics... are setup and running through waveform fitting.
 - No need to attach ‘beam’
 - QA files still don’t appear to be produced. Pinged Martin about this but apparently it has not been resolved, will double check...
 - Tracking cosmics
 - Up to clusters seems to be fine. Track seed DSTs also being produced, so machinery seems to be in place
- Offline Productions
 - Streaming
 - Will re-run track fitting DSTs tonight with new Acts geometry fix
 - Started raw hit production last week over larger set of full runs on SDCC. Probably those will finish mid-end week
 - Triggered
 - Quality checks... which runs went into the calorimeter reconstruction? (No selection criteria beyond the 1k event min). Runs ≥ 47289 .
- SLURP updates
 - Making use of condor class ads to attach meta information to the jobs.
 - DSTTYPE, DATASET, RUNNUMBER, SEGMENT, CUPSID (the unique ID of the production status row which created the job)
 - Enables a fast query of the condor schedd by setting a constraint
 - During the production loop, we can query the condor schedd with a narrow constraint to identify jobs which are on hold, and then set the production status accordingly.
 - Sets the message field to contain the HoldReason
 - Sets the ended timestamp to the date/time when the job went on hold
 - Current implementation only checks for held jobs when new jobs are available for submission.
 - Email people about holds?
 - CLear holds?
 - We can also query what runs are currently being processed. Allows us to advance the run cursor, making our production loop DB queries more efficient.
- New CDB instance
- New condor pool, should be transparent to the end users. Prod02 setup to submit to this.
- AOB
- **Meeting Monday Mar 17 2025 1PM EDT**
 - Readiness for Beam & Online Cosmics production
 - Production loop has remained stable / no exceptions encountered.
 - <https://chat.sdcc.bnl.gov/sphenix/pl/ebbihu1r1pdg3dpahoidfmo7a>
 - Production failed after 58059
 - Streaming event building and cluster finding are active. Production cursor is sitting at run 58059 for each of them.
 - O(10s) of held event builder jobs / 16 cluster finder jobs
 - Clustering jobs more memory intensive than previously observed. (Alma 9?)
 - Readiness for further downstream workflows?
 - TODO: Setup and run through tracking. Setup and run calorimeter loop.... primarily test of offline QA. Expect jet finding updates to come along for the AuAu data.
 - TODO: Migrate off of sub05 → Onto sphnxprod01. (After handing off list of held jobs to Joe).
 - Triggered production
 - Run2pp reproduction is finished. DST_CALO not to be produced before QM. No new large scale productions anticipated.
 - Looking at the cups statistics files, I don’t see any instances where we failed to update the file catalog on the first attempt for the event builder jobs. (Checking downstream jobs as well...)
 - Believe that event builder jobs exceeding “few” hours are spinning. Once one segment fails, do the remaining segments also have problems?
 - Streaming production
 - Nothing major. Will repass over seeds and tracks this week
 - Slurp updates
 - Now with a *production (run) cursor*. A run number, pointing to the last run with jobs left to be submitted, is tracked for each production. Kaedama (and ramenya) can reference the cursor when specifying a run range.
 - `-runs cursor #` restricts query to `run>=cursor`
 - `-runs cursor N #` restricts query to `run>=cursor` and `run<=cursor+N`
 - Updated cups to verify that the file staged out has the same size as the file on the local disk. Will try up to 5 times (with 2 min delay) to copy on failure. Tracks statistics in the cups statistics file.
 - Have a plan B in place... if we see instabilities in database writes, write the dbquery to a file and execute at a later date.
 - Starting to provide ‘canned’ input queries. For most workflows, user provides the DST type(s) and dataset(s) to be processed as parameters, and the system will use a default DB query. For more complex processing, full query can still be specified..
 - AOB
- **Meeting Monday Mar 10 2025 1PM EDT**
 - Condor and Lustre and ... problems
 - Power outage started at approx 9:42 AM on Friday.
 - ~~■ Python / kernel bugs / knocking out individual nodes ???~~
 - Lustre issues
 - “The bug seems to affect the locking mechanism when multiple processes write to the same directory in lustre...”
 - Does this mean that lustre works lock-free if we write to independent directories?
 - ~~● Could introduce a segment group (e.g. seg_00400_00200) analogous to the run group (run_00052400_00052500) if useful...~~
 - Costin built a version of the lustre software with bug fix. 100 nodes available. Plan A if they survive. Plan B would be to revert.
 - Trying now with Shuhang restarting production.

- Need to know by this evening where we are.
 - Immediate needs... target machines with requirements. SDCC condor flag?
 - New submission node incoming.
 - ~~Write to data02???~~ Insufficient space.
 - Cosmics run support? Next TPC gain scan is when they need cosmics back in production. Suspend until Joe gives the word.
 - Once we have the 50/50 split we can target machines with cosmics. Next couple of days.
 - 50k concurrent jobs on phnxsub05. Shuhang is only seeing 15k or so. On sub05. (Competition with user jobs?).
 - Memory resources....
 - Fitting jobs 1GB. Combining 2GB.
- Restart of the production system... TBD. On pause for a couple of days.
 - Online restart
 - Will restart from the “current” cosmic ray run and then back fill, so that we don’t hit the full backlog at once
 - Will add in the triggered event building at this point... (and downstream workflows?)
 - Triggered event restart
 - Had a quick look at the production status on Friday, and it appears that many of the production jobs had reached a finished state before the power outage on Friday.
 - Any objections to starting downstream production. Nope.
 - Streaming restart
 - (Joe) tried to resubmit some seed production but half of the jobs die quickly. We think it is because there are calibration files produced with root 6.32 and the production runs in the sl7 container with root 6.26
 - Another strange thing - production with ana466 and new db tag (using CDB file made with root 6.26) all blow up in memory. 5 weeks ago we ran this production with the same ana build with no issues. *What has changed* (running in a container? Multithreading not handled well)?
 - Need to produce these data ASAP for quark matter, we are out of time. Can we run in a9 natively? We have been developing on the successful ana466 cluster DSTs for a month in alma9 without an issue
 - [Yes... need to remove the container in the jobwrapper.sh.](#)
 - ... need new ana build ... native alma9 ...
 - ... useful to ditch SL7. ... calo production seems to be ok ...
 - GSL version SL7 vs alma9. Lower energy events fitting differences in calos.
- SLURP
 - The production loop will have a “run cursor” this week. Jobs will be submitted for any run which is larger than the cursor. When jobs are submitted, the cursor advances to the last run submitted.
 - Planning to add a “heartbeat” timestamp. This will be the timestamp of the last write to the “message” column in the production status table. For streaming event builders, this will record the timestamp of the last DST staged out.
 - It would also be useful to add a column in the datasets table, which would be filled with the unique ID of the production job which created it.
- AOB
 - no-dbinput for calo builders...
 - 1.4 thousands of jobs on hold from crashed machines...

● Meeting Monday Mar 03 2025 1PM EDT

- Online Cosmics Production
 - ...
- Triggered Data Production ana462 2024p010 ver001
 - ...
 - Plan to wait on simulation jobs after simulations
 - Shuhang to take things after the event building
- Streaming Production
 - Seeding production this week aiming for D0 mass plot by QM
- Online Cosmics Production
- Slurp updates...
 - Stupid useful prompt update... displays the branch and state (green is clean) of slurp and ./ProdFlow. If either is red, think twice about submitting production jobs.

```
jwebb2@ssh02:/home/sphenix/u/jwebb2
sphenxpro [slurp:run3auau-cosmics-dev] [prodflow:online_streaming_run3auau_ana468_2024p012_v001.test] [slurptest/]> emacs -nw slurp/setup
sphenxpro [slurp:run3auau-cosmics-dev] [prodflow:online_streaming_run3auau_ana468_2024p012_v001.test] [slurptest/]> emacs -nw ProdFlow/run3auau/DST_STREAMING_run3auau_new_2024p012_v001.test
sphenxpro [slurp:run3auau-cosmics-dev] [prodflow:online_streaming_run3auau_ana468_2024p012_v001.test] [slurptest/]> cd slurp; git stash; cd -
Saved working directory and index state WIP on run3auau-cosmics-dev: f44c21d Trap signals and update status.
/direct/sphenix+u/sphenxpro/slurp/slurptest
sphenxpro [slurp:run3auau-cosmics-dev] [prodflow:online_streaming_run3auau_ana468_2024p012_v001.test] [slurptest/]> cd ProdFlow/; git stash; cd -
Saved working directory and index state WIP on online_streaming_run3auau_ana468_2024p012_v001.test: f0f4dea ...
/direct/sphenix+u/sphenxpro/slurp/slurptest
sphenxpro [slurp:run3auau-cosmics-dev] [prodflow:online_streaming_run3auau_ana468_2024p012_v001.test] [slurptest/]> █
```

- Production needs from PPGs:
 - **ppg02** (dN/deta with INTT): data are already produced, simulations running, no new production needed, problem: INTT dead map no the same in data and simulations.
 - **ppg03** (dE_T/deta): no reply
 - **ppg04** (jet structure in AuAu): No need for centralized production. Trying to get running embedding of pythia jet10 & jet30 events into locally produced calorimeter DSTs from 2024 Au+Au run 54912.
 - **ppg05** (tracklet flow in AuAu): no reply

ppg07 (neutral mesons): no reply

ppg08 (di-jets): no data/simulation needs

ppg09 (jet structure): no data/simulation needs

ppg10 (jet structure): no data/simulation needs

ppg11 (neutral mesons): need to finish calorimeter production, simulations are the same as the jet group (MB, JetXX, PhotonXX triggers)

ppg12 (isolated photons): data: finishing calorimeter waveform fitting (ana468); simulation: samples of jet5 and, with lower priority, jet15; only need the DST_TRUTH and G4Hits produced.

- Balance of production vs user slots?

- AOB
 - 20 PB free on lustre now...

- **Meeting Monday Feb 24 2025 1PM EDT**

- Triggered Data Production ana462 2024p010 ver001
 - No evidence that we lost anything to the DB upgrade
 - 125515 in finished state
 - 210 jobs with no file catalog entry. All ran from sub05 w/in ~2 hours on 2/12 - 2/13
 - The jobs die when the file is being copied out to disk. There is no exception thrown.
 - There is no attempt to update the file catalog, b/c we don't reach that part of the code.
 - All such jobs failed on alma9 nodes.
 - Rest look normal
 - 10668 in failed state
 - No file catalog entries. No output files. Some log files exist.
 - Failure reasons TBD
 - 151302 in submitted state (job never advanced to 'running' 'finished' or 'failed')
 - jobs did not go on hold at this level
 - no file catalog entries as expected
 - 3051 log files.
 - 264 jobs wrote output. They each have job logs.
 - Notable: run 53050 failed on 500 of ~560 segments.
 - Waveform fitting was performed with ana462.
 - Accounting TBD
 - Friday's mattermost chat noted ana463 or later preferred. Ran with ana468.
 - Observed ½ of all jobs going into the failed state...
 - Do we have a reason for this? Whose jobs is it to find out?
 - Sasha looked at this. Doesn't look like memory overflow. Some jobs went past 10 GB and succeeded. Log files truncated in many cases.
 - Wait for response from ...
 - Maybe a condor config issue. Once SDCC reconfigs ... may be able to restart.
- Streaming Production
 - Running seed DST production now. Will aim for a track DST production at the end of the week, trying to finish up some calibrations in push for QM
 - No real issues noticed so far, jobs churning along
 - Additional Runs should show up 'at some point soon'
- Online Cosmics Production
 - Online production loop up and running.
 - System submitted run 56938. MVTX run only. Issues with MVTX software indicated in log file.
 - No GL1 (but file was present), gl1 file for run 56938 does not contain any events
 - still running off of the daq db
 - Odd issue noticed when I restarted the processing loop. The production_status table is out of sync with the table structure expected by slurp. An extra 'revision' column was added. It should have been and was added to production_setup.
 - Issue is fixed and submission loop is running now. 30 min delay between loop iterations.
- New common slurp installation
 - mkdir blah
 - cd blah
 - source /sphenix/u/sphnxpro/slurp/slurp/setup
 - checkout prodflow
 - run. No more need to checkout slurp.
 - touch .slurp/testbedto enable testbed mode... (may become default).
- AOB
 - Move of SL7 hosts to Alma9 should start today (disabling condor submission from the interactive SL7 hosts, draining the condor jobs)
 - Submission to SL7 still possible from phnxsub01-04, please don't use those for production

- **Meeting Monday Feb 10 2025 1PM EDT**

- Triggered Data Production ana462 2024p010 ver001 status

| dstname | nruns | min_run | max_run | njobs | status | duration |
|---|-------|---------|---------|--------|-----------|-----------------|
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 567 | 49400 | 53879 | 32531 | running | |
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 918 | 47201 | 53879 | 159123 | submitted | 00:00:00.777778 |
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 852 | 49400 | 53880 | 115965 | finished | 02:46:56.379313 |
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 395 | 49402 | 53875 | 9268 | failed | 06:14:55.791756 |
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 130 | 47601 | 48398 | 21429 | started | |
| (5 rows) | | | | | | |
| dstname | nruns | min_run | max_run | njobs | status | duration |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 441 | 49400 | 53880 | 14429 | running | |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 50 | 49432 | 53854 | 148 | submitted | 00:00:06 |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 843 | 49400 | 53880 | 98966 | finished | 05:26:35 |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 2 | 50898 | 52018 | 2 | failed | 00:02:30 |

- Backlog of 159k jobs remain, with run 47201 the first submitted run.
 - Down to 3600 jobs... submitted primarily from phnxsub02, 03 and 04.
- 21k jobs in ‘started’ state were not passed their input lists. To be restarted.
- There is a class of jobs which go on hold on phnxsub05 which *should be recoverable*...
 - Are terminated early for reasons which are not clear
 - fun4all exits with 0 exit code
 - condor logs have been missing since we restructured the output directories
 - output DST truncated
 - log files consistent with job being terminated during the stageout script, possibly b/c GetEntries.C ‘blows up’ when it hits a truncated root file.
 - Ran single such job interactively
 - Normal termination of the job
 - Output DST same size as other segments / not truncated
 - Will cleanup the production status table / filecatalog and resubmit on phnxsub05
- TODO:
 - df .
 - Check if user jobs are outcompeting on 2,3,4
 - check for blackhole nodes
- Streaming Data Production
 - Ana464 production
 - Raw hit DSTs are probably close to being finished
 - CDB update over the weekend, will start ana464 2024p012 clustering jobs today
 - Need to produce PPG02 cluster DSTs off of special branch in ProdFlow. Assuming we can not run this concurrently with the above production (i.e. need to run special ppg02 production and let those jobs drain out, then switch branches to master and run regular jobs)
 - Watch shifts start next week - how do we setup online productions so that they are ready to go when cosemics start later in the week?
 - event builder + downstream workflow will be stream parallel
 - aim for production loop up and running for next week’s cosmic data
 - will need to discuss ultimately what the division of farm resources between production and users...
 - bottleneck from last run should be fixed... able to transfer full data... may be disk constrained.
 - plan to run tracking through job A... then raw data will be cleared from disk.
 - need to verify TPC quality in each run / do not need full run to do so
 - event builder in a monitoring mode...
 - Productions that can safely be deleted
 - Run2pp
 - Anything under new*
 - Ana464_2024p011_v001 TRKR_CLUSTER DSTs only (i.e. nothing else from ana464 can safely be deleted)
 - Ana461_2024p009 and nocdbtag v001 can be deleted
 - Run2auau
 - Anything under new*
 - ana463_nocdbtag_v001
- Data carousel ... missing datasets to be restored to disk
- HPSS ... stage with directory structure ... filecatalog will have locations ...
 - Existing PRDFs staying in place
 - ...
 - Input managers now raw data catalog aware...
 -
- Next version of slurp...
 - Sufficiently stable that we should consider moving to using a central installation.
 - Testbed setup allows us to isolate testing and development from it
 - Have a version which moves the functional bits into a python package, enables a central install (e.g. ~sphnxpro/slurpinst)
- AOB
 - sPHENIX indico page... / calendar
 - Announce to sphenix software
 - Next meeting Feb. 24.
- **Meeting Monday Feb 03 2025 1PM EDT**
 - Triggered Data Production ana462 2024p010 ver001 / ver002
 - Reproduction started from event building
 - Jobs are limited to 3 days run time. So far I don’t see any jobs running past the 3 day limit.
 - Current status (which does not reflect held jobs) below
 - ... snapshot is Monday 11 AM
 - There are low rates of issues with database connections indicated
 - ‘Running’ should be read as the sum of running and held state.
 - I need to update the code which flags for held jobs to account for the fact that we are running across 4 different submission hosts.
 - Plan to flag these as ‘evicted’ in the proddb.

- Oddly... if I query the collector for the list of schedd's... I do not see phnxsub05? (But I get phnxsub01, and it is the only one showing held jobs...
- We are holding down ~10k slots for the event building.
 - Backlog of 228k jobs waiting to run. Currently have xferslots=1 on the event building jobs spread over phnxsub02,03,04 and 05.
- We can begin downstream jobs ...
 - Run with ana462 2024p010 v002 OR a new build or tag?

| dstname | nruns | min_run | max_run | njobs | status |
|---|-------|---------|---------|--------|-----------|
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 201 | 50547 | 53875 | 7008 | failed |
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 536 | 50545 | 53880 | 69686 | finished |
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 1375 | 47403 | 53879 | 228361 | submitted |
| DST_TRIGGERED_EVENT_run2pp_ana462_nocdbtag_v001 | 352 | 50536 | 53879 | 21313 | running |
| totals | 1997 | 47403 | 53880 | 326368 | - |
| (5 rows) | | | | | |
| dstname | nruns | min_run | max_run | njobs | status |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 39 | 50548 | 53854 | 369 | submitted |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 207 | 50548 | 53880 | 1591 | running |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 526 | 50550 | 53880 | 66352 | finished |
| DST_CALOFITTING_run2pp_ana462_2024p010_v001 | 1 | 50898 | 50898 | 1 | failed |
| totals | 532 | 50548 | 53880 | 68313 | - |

- Streaming Data Production
 - Tested and merged in segmentized QA production for event combining
 - Updated AuAu event combining macro/script to run parallel streams
 - PPG02 special production
 - PPG02 aiming for preliminary for QM. They need a specialized production run with certain settings in the clusterizer
 - What is the protocol for this? I am starting the raw hit production for the AuAu run today. I suggested they make a macro starting with the production job0 macro with their settings in place
 - Do we put this on a branch and run it from the branch?
 - This will likely not be the last time some PPG needs a specialized production, so we should discuss a long term strategy
- AOB
- Meeting Monday Jan 27 2025 1PM EDT
 - Production ana462 2024p010 ver001
 - Triggered pp reproduction started on Friday.
 - DST_CALO production jobs are complete. One job needs to be recovered (test setting cups).
 - The production script was setup so that the filenames are based on the input filenames. Will rename and update the catalog to correct this... (will write a production workflow to do this).
 - Can be done in sql on the cmd line
 - Useful if someone could take a look at the log files and data and confirm that the correct CDB tag was picked up.
 - /sphenix/data/data02/sphnxpro/production/run2pp/physics/ana462_2024p010_v001/
 - /sphenix/lustre01/sphnxpro/production/run2pp/physics/ana462_2024p010_v001/
 - ana 446 was a ‘new’ build so... it was still a bit of a moving target.
 - ‘scrap’ and full rerun v002.
 - ... new calofitting incoming, so event building can/should proceed.
 - Do we need to rerun event building from 446. Yes. Not 100% sure the 446 build stayed consistent.
 - Still have probs with inf loops, wrong lsat event, etc....
 - Set a 3d limit if practicable.
 - DST_JETCALO nearly complete.
 - AuAu production running... actually... ran. 3k jobs from the requested run lists completed fairly quickly.
 - ... @5k events per job... seems a little light...
 - 10 runs AuAu.
 - Streaming still iterating on understanding data
 - Slurp updates
 - In production mode, require the local repo’s HEAD to point to a valid commit on the same branch in remote.
 - Applied in the ProductionSystem/Run2Trigger area for this production
 - . Not (yet) in the streaming area.
 - Note... diff is the correct tool here, not the hash. Ideally track only significant changes. (no whitespace, eg...)
 -
 - TODO:
 - Triggered event building
 - AOB
 - DST_TRIGGERED_EVENT_run2pp_ana451_2024p009 production ??? still running/on hold. clean it up? YES.
 - DST_STREAMING_EVENT_*_run2pp_ana459_nocdbtag_v001 also clean up?
 - Joe - Yes, we can remove it
 - Data preservation ...
- No Meeting Monday Jan 20 2025
 - Lab Holiday
- Meeting Monday Jan 13 2025 1PM EDT
 - Status of / plan to finish ana451 production
 - ???
 - Next fixed build production
 - Plan to run under alma9 and deal with any issues we find

- - Need to start a single streaming event combining fixed build tracking production ASAP.
 - There are weird memory issues on alma9 that I'm trying to debug, but they are very nontrivial
 - ...
 - MBD vertex object copied twice... should be fixed...
 - Event combining does not need a new ana build...
 - ana.459 nocdbtag ver001
 -
 - Single stream production... setup a testbed and production area asap w/ new fn convention.
 -
 - Single stream support... “dev” branch.
 - Show stopper last week
 - DST_TRKR_SEED jobs are submitted multiple times. (2x for each input segment)
 - Reason
 - Laminations_DST_TRKR_CLUSTER_run2pp_new_2024p009-00053880-00008.root files are cataloged
 - stageout.sh uses regex to extract dsttype and dataset from the filename... and it is flexible, so that we can match things like HIST_ ... CDB_ ... and Laminations_ ...
- ```
regex_dsttype_run="([A-Z]+_[A-Z_]+[a-z0-9]+)_([a-z0-9]+)_(202[3456789]p[0-9][0-9][0-9]|nocdbtag)-([0-9]+)-([0-9]+)"
```
- The *dsttype* is matched in these cases as DST\_TRKR\_CLUSTER\_run2pp b/c it ignores the leading mixed-capitalization string...
  - Why I ran into this? Tried to better optimize the input DB query... replacing a wildcard search on the filename with a query on the dataset and dsttype...
- ```
select
    'filecatalog/datasets' as source,
    runnumber,
    segment,
    filename as files,
    filename || ':0:' || events as fileranges
from
    datasets
where
    dsttype='DST_TRKR_CLUSTER_run2pp' and dataset='new_2024p009'
order by runnumber
```
- Workaround is to add ‘events>0’ to the where clause...
 - Need to establish (revisit...) a convention for non-DST files
 - Test jobs running... will merge the dev branch after meeting.
 - Not really QA not really DST ... this is a calibration object...
 - Decision... loosen the naming convention... any word groups before the build / tag.
 - New naming convention.
 - Adds a ‘version’ number to the filename.
 - Version indicates a meaningful change in production conditions which is not tracked as part of the library builds or the cdb tags.
 - i.e. a change in the production macros
 - DST_TRKR_SEED_run2pp_ana432_2025p001_*v001*
 - Version number is set in the yaml file.
 - If not set the version number will be 000.
 - Version 000 is reserved for new builds only.
 - New build can only produce files with version 000.
 - slurp should be able to process inputs with no version number
 - Will require an update to all stageout.sh scripts... (will take this opportunity to consolidate to a common... or at least fewer... stageout.sh scripts).
 - Aiming for this to be done by COB (but likely tomorrow)
 -
 - AOB
 - **Meeting Monday Jan 06 2025 1PM EDT**
 - Readiness for next fixed build production(s) Topics of discussion...
 - Infrastructure
 - Space on lustre / conditions DB / ...
 - Freed up 1PB, 3PB currently free... (need to revisit directory structure)
 - FC and dataset table should be in sync with whats currently on disk.
 - Run 1 data needs to be dealt with... discussion w/ collaboration.
 - Development before break / testbed area can be cleaned up for tracking.
 - Goal 10% of tracking data for QM. Need to clean of ____ PB.
 - 50 best runs.
 - No reason to keep raw data on disk that isn't going to be used ... (confirm its on HPSS.)
 - Need data carousel to restore files from HPSS.
 - Calorimeter side... main part of run 1 is calorimeter.
 - Planning to delete everything with ‘new’.
 -
 - No open PR's... Updates to calorimeter workflows?
 - Jet skim macro incoming
 - No open PR's...Updates to tracking workflows?
 - Jet QA upstream PRs... but not on critical path.

- Event combining expected to utilize single stream combining... Code should be in the prodflow area.
- Next production will be under alma 9
 - ana build 458 is built under alma9.
 - run native (test under container).
 -
- Updates to production system...
 - slurp supports running mixed input / output tags. e.g.
 - DST_STREAMING_EVENT_new_nocdbtag → DST_TRKR_HIT_new_2024p009
 - In principle you should be able to mix any set of inputs (i.e. cluster and seed with different build/tag can be used as input to track finding). But this has not been tested. THIS HAS TO BE TESTED.
 - ramenya supports a new loop steering file...
 - each rule that is defined can be executed with different arguments
 - allows us to run multiple workflow definition files in one loop...
 - i.e. tracking workflows, calorimter workflows, etc... can be mixed in a single production loop
 - fine grained tuning of max number of jobs to submit in each pass (i.e. submit up to 100 streaming event builders, 10 triggered event builders, 42 cluster finders, etc...)
 - cups will now log failures in database access. It will log the update it was trying to perform, the number of attempts it made, the timestamps (before/after) the update, and the last exception returned by pyodbc if any. The log is written out beside the job log to DST_blah.dbstat.
 - Changes for single-stream job processing.
 - Have merged the changes to slurp needed for single-stream processing. Code is tested against current workflows. Ready to merge / test against latest single-stream workflows.
- Files / datasets inconsistencies
 - Files made it in, datasets didn't
- Preparation for TPC cosmics data taking
 - TPC plans to start cosmics next week - need a production pipeline in place for this as they want to do gain balancing
 - Will also need laser hit/cluster production in place when watch shifts start, but that is on a O(weeks) timescale
 - Full event combining. Data zero suppressed? (Either or).
 - ... setup single streaming ... R&D ... debug it ...
- Discussion and AOB
 - Continue output file (and filepath) naming discussion
 - Deleting single files not efficient, need to be bunched in subdirs for easy rm -rf without chance to remove other files/dirs
 - Means runwise dirs are lowest subdir, tag next level, dst type, production - did i miss anything?
 - Unique filename conventions to delete via matching in DB tables
 - Cosmics right now under run2auau
 - Need another field in database table (pp,auau,cosmics,...)?
 - Versioning of a given DST (inputs changed, macros changed, ... etc...)
 - <build>_<tag>_<version>
 - Need to add some version tag to tag input files
 - Scenario: combining is rerun, but subsequent calo pass uses previous ana_cdb tag

```
_default_filesystem = {          // run2pp      physics      anaXXX_2024p007      run_X_Y      DST_STREAMING_EVENT
    'outdir' :                  "/sphenix/lustre01/sphnxpro/production/${runtype}/${runname}/${build}_${tag}/run_${rungroup}/{leafdir}"
    , 'logdir' : "file:///sphenix/data/data02/sphnxpro/production/${runtype}/${runname}/${build}_${tag}/run_${rungroup}/{leafdir}/log"
    , 'histdir' :                "/sphenix/data/data02/sphnxpro/production/${runtype}/${runname}/${build}_${tag}/run_${rungroup}/{leafdir}/hist"
    , 'condor' :                  "/tmp/production/${runtype}/${runname}/${build}_${tag}/run_${rungroup}/{leafdir}"
}

TODO:  runname/runtype/build_tag/leafdir/run_X_Y
```

● Meeting Monday Dec 16 2024 1PM EDT

- Stream Parallel Test Production
 - Large scale (100 run) jobs started last week. Aiming to wrap by end of next week.
 - Stress test of decoder.
 - MVTX fixes this weekend... small batch to be run to demonstrate error handling improves things.
 - Close to readiness for a full production
 - Need to carefully check output filesize.
 - Nearly ready to move this out of the testbed... estimating first of year for physics quality production.
 - 'Full' production constrained by most of data in hpss...
 - New job 0... (combined hit unpacking and clustering)
 - New input query... working to first order, but ...
 - Issue in unpacker, ... possibly resolved by recent commit
 - Mixing builds in a workflow chain... still requires 'direct_path: /path/to/*/directories/of/*all*/inputs'
- Fixed Build Calorimeter Production
 - run2pp event building. Job submitted and quit... reason TBD. Add **—no-dbininput** to kaedama.
 - run2auau 20 runs requested. and submitted.
 - Request for MBD triggered events only. i.e.. production of trigger specific DSTs... to be discussed in other forum..
- nocdbtag
 - Output of event builders (and any other workflwo) that doesn't requires CDB should have the tag 'nocdbtag'
 - Flag to disable cdb lookup
 - Calorimeter event builder... doesn't event access cdb...
 - Latest ana build...
 - Action item... calorimeter production fixed build ana.453 ... alma 9 .. expect ~today tomorrow.

- Macro with git commit IDs...
 - Alma 9 support integrated [PR63](#)
 - DONE: update workflows to leverage... mostly requires removing the ‘executable’ from the condor job description.
 - <https://github.com/sPHENIX-Collaboration/ProdFlow/pull/43>
 - Support for stream-parallel production is ready to go
 - <https://github.com/klendathu2k/slurp/pull/64/>
 - Will merge after workflows are updated for Alma 9 (ideally before Monday’s meeting)
 - Default directory structure is set as per last weeks discussion, except that the last directory follows DST_TRKR_CLUSTER_run2pp.
 - To be fixed.
 - ...
 - Slurp has been updated with alma 9 and stream-parallel processing support. Prodfow workflows should all be alma9 capable now.
 - Simple tests...
 - DST_STREAMING_EVENT run2pp ana444 2024p007 submits under SL7 and A9 ... and runs ...
 - DST_TRKR_HIT run2pp ana444 2024p007 submits under SL7 and A9 ... and runs ...
 - DST_TRIGGERED_EVENT run2pp ana444 2024p007 submits under SL7 and A9 ... issue with building inlist. TBD
 - DST_TRIGGERED_EVENT needs to be enabled for dbinput (still passes input files via arguments list.)
 - runs w/ –no-dbinput flag...
 -
 -
 - AOB
 - Add SaveGitTags.C to the production macros to save git commit ids in DSTs
 - Status of missing DB inserts
 - Status of nocdbtag
 - Request from MVTX group to reprocess cosmics (55000+ runs) with ana.453 (event combining only)
 - Code builds under alma 9. Soon will be able to submit native to alma 9.
 - Two ways... 1) alma9 native, 2) alma9 container.
 - ... observation that code runs faster in containers...
- **Ad hoc meeting 12/11**
 - Thanksgiving test reasonably successful
 - MVTX being fixed
 - TPC is the biggest timing issue
 - Would like to get a start sooner rather than later ...
 - Raw data catalog based query failed...
 - Raw data catalog was implemented after their stable development branch was established
 - significant divergence between the isolated development branch and the main branch...
 - quickest path forward to a test production is to use the daqdb, but remove the condition that all files must be transferred. Limit runs processed to those which are known to be fully transferred.
 - Bringing the streaming changes onto the main branch on the task list following alma9
 - Clustering inputs are to be aggregated by segment... i.e. all files with common run and segment number fed into the f4a macro.
 - Need to double check that the TPOT number of events per output segment is the same across the production
 - SQL query should be straightforward... needs just to group by runnumber, segment.
- **Meeting Mon Dec 09 2024 1 PM EDT**
 - Running jobs with an ana build / cdb tag that differs from the input jobs
 - There is currently a limitation in slurp. When building the list of input files (the physical file list) from the file catalog, it is presumed that the input files and output file names have a 1:1 mapping... i.e. the build and tag do not change.
 - We lookup all files that have the dataset ‘{build}_{name}’
 - This limitation can be removed... but carefully
 - There is a workaround for this. The physical file list can be built using a direct filesystem lookup. This is not ideal, but can be used for the moment.
 - The calorimeter workfest has started (this is why I want this to be done now)
 - Tracking reproduction status
 - Inputs are the DST_TRKR_HIT w/ ana 441 cdb tag 2024p007
 - DST_TRKR_CLUSTER: 17695 submitted 745 still running
 - DST_TRKR_SEED: 16950 submitted 4450 still running 7714 held
 - DST_TRKR_TRACKS: 4785 submitted 26 running
 - Directories:
 - /tmp/production/
 - /data/data02/sphnxpro/production
 - /lustre01/sphnxpro/production
 - physics/run2pp/DST_TRKR_CLUSTER_run2pp/ana449_2024p008/”run range”
 - TODO: swap build w/ DST type
 - TODO: strip run2pp from the ...
 - TODO: physics/run2pp/ana449_2024p008/”run range”/DST_TRKR_CLUSTER
 - QA check
 - Runs w/out tracks and seeds
 - Runs with x100 smaller output files (no tracks / seeds)
 - Tracks to emcal clusters no match / completely random
 - Suspicion that there is a clustering issue / downstream jobs fail
 - Calorimeter production issues

- Default dbinput dbrange...
 - calorimeter scripts were not setup for that
 - ... direct_path needs to be setup for the pfn list building
 - Calibration output filename is built with a replace rather than building from the arguments
 - run2pp / run2auau
 - edbc-by-edbc
 - Jobs seems to be crashing w/out logfile to debug / ...
 - Possibly a typo early in the script
 - Possibly other crash and condor reroutes the log file... condor_hist –long
 - New features ... git pull
 - Directory paths are now set by default. Workflows can (and should) remove the filesystem block
 - Testbed areas replace production with production-testbed
 - If provided, a filesystem block in the yaml file will override in part or in full the system default.
 - Next up is alma9. Plan is to implement a general wrapper that gets passed the name of the script to run, and then passes down the argument list.
 - ...
 - AOB
- **Meeting Mon Dec 02 2024 1 PM EDT**
 - PR#39 Fix issue with build configuration
 - <https://github.com/sPHENIX-Collaboration/ProdFlow/pull/39>
 - Production readiness...
 - Issue with segment number... internal data structures / synchronization issue / hardcode seg=0
 - Event builder macro change (?)
 - Code change (?)
 - Directory structure...
 - ☒ ~~TODO: YAML files should implement the new directory structure.~~
 - ☒ ~~TODO: cosmics/ directory should → streaming/~~
 - ☐
 - Plan is to run streaming production w/ the current workflow, switch to EDBC-by-EDBC when available.
 - Production will be the done via the alma9 submission
 - ☒ ~~TODO: Containerize the production scripts for alma9~~
 - small increase in throughput expected
 - jobA jobC reproduction... rerun now. Event builder will be rerun later.
 - ☒ ~~TODO: Containerize the production scripts for alma9~~
 - ☒ ~~TODO: YAML files should implement the new directory structure~~
 - Joe ran large scale edbc-by-edbc ... things don't look so bad. Pretty close to running a full production, aiming towards end of the year for a full tracking dataset production.
 - Order of 2 weeks. Currently 80-90% success rate.
 - Lustre disk space. 8PB free. (14% availability).
 - Database should be cleaned when deleting files.
 -
- **Meeting Mon Nov 25 2024 1 PM EDT**
 - odbc.ini status
 - Status of current productions ...
 - run2pp reproduction ana448 2024p008 (jobA jobC...)
 - Updating ProductionSystemRun2pp (pulling latest slurp and ProdFlow)
 - Will run with latest changes (discussed below) to test at scale
 - Submitted seed finding pass.
 - Two observations
 - Post condor submit, the update which writes each condor job Id to the production status table is expensive. (17k jobs → 17k updates...). Need to move the update out of the loop if practicable.
 - We (currently) have to use direct filesystem lookup to build the physical file names for the inputs to these jobs.
 - the file catalog query which builds the PFN lists presumes that the input and output datasets match.
 - ana449_2024p008 seeds are built from the ana441_2024p007 clusters
 - ana449_2024p008 tracks are built from the ana441_2024p007 clusters and ana449_2024p008 seeds
 - calorimeters
 - ... disk space discussion ...
 - 6PB of calorimeter DSTs
 - Deleted the “new” DSTs (calo). New from august/sep streaming DSTs. DB should be deleted as well.
 - Chris will delete.
 - Raw data 2023?
 - Need for filepath convention
 - x3 reduction from combined to waveform fits
 - run2pp reproduction will wait until we confirm space is available. Test job of 1-2 runs may be submitted prior.
 - SLURP updates [... git pull ...]
 - Alma9 readiness... phnxsub05 ... everyone should be able to log on
 - Modified setup script to setup aliases to system python (python 3.9)
 - Installed missing packages (for user sphnxpro)
 - Jobs successfully generated and passed to condor... (condor was inop during tests).

- Runlist Selection Options...
 - Users can specify
 - –runs SINGLE_RUN
 - –runs FIRST_RUN LAST_RUN
 - –runs run0 run1 ... runN
 - –runlist <whitespace delimited text file>
 - Runlist can also be forced as part of the input query
 - ... where runnumber in (run0, ..., runN)
 - Runlist can also reference another table (w/in the same database being accessed in the query)...
 - ... where runnumber in (select runnumber from goodrun_table where ...)
 - ... where runnumber not in (select runnumber from badrun_table where ...)
 - Believe this covers the use cases outline last week
- Made –dbinput option default
- RawdataCatalog can be specified for the input query in slurp. (db: raw).
 - Tested, works with streaming event builder workflow.
 - runlist table is a superset of the daqdb filelist table, so... created an alias in slurp that makes this transparent. All workflows remapped to RawdataCatalog after the next git pull.
- Output Consolidation (work in progress)
 - Defining a master configuration file.
 - Toplevel directories defined in master config.
 - Default directory structure will be generated by slurp during submission.
 - YAML file may override the defaults.
- Better timeout handling during job submission.
 - Input query will now be attempted 10 times, randomized delay between attempts.
 - Connection established at the point of the query.
- Single streaming tests
 - The script to run the macro was wrong, was not setting the ana build correctly and was not passing the file lists properly to the macro
 - I am trying to fix this now. I still do not understand some behavior of the kaedama.py script - sometimes running it on a single stream/ebdc submits jobs for every single stream/ebdc, and sometimes it does not. Not sure exactly what I'm doing wrong
- AOB
- **Meeting Mon Nov 18 2024 1 PM EDT**
 - Focus on immediate needs, cleanup can be done once all this is implemented:
 - Multiple sources for run numbers
 - Daq DB (physics)
 - Add query to triage table (remove no good runs)
 - Calibration status (tbi)
 - Text files
 - Special table with selected runs
 - How do we submit run ranges (probably cannot just submit all runs, command line?)
 - How do we run this in a cron which keeps filling the queue?
 - Information about runs comes from the daq DB, if relevant for the production - it needs to fetch them from there
 - Streaming flag
 - Zero suppression
 - Strobe length
 - ...
 - Source for files is the RawdataCatalog
 - Files in Datasets will get a status flag
 - >0 good for processing -
 - 1 = in lustre, higher numbers reserved for whatever
 - 0 in hpss
 - <0 bad
 - Detailed Feedback and fixing of problems
 - Get info on problematic files (not “5 runs didn’t process”)
 - Use production to patch up problems
 - Patch up and update md5 sum during production?
 - Many streaming files have 0xFFFFFFFF
 - Some bogus first, last event numbers
 - Bogus sizes
 - first/last bco needs filling
 - Move to production DB, ditch production_status DB
 - Output location consolidation
 - Top dir in lustre: /sphenix/lustre01/sphnxpro/production
 - Then production (run2pp, run2auau, run2cosmic,<calibration>)
 - Then tag (build_cdb)
 - Then DST type
 - Then runrange (100's)
 - Top dir in gpfs: /sphenix/data/data02/sphnxpro/production
 - Same subdir structure as in lustre

- Move to Alma9
 - 36k cores moved already... not a great success
 - Workflow production scripts have a significant amount of duplicated code... stagein of support files, building file lists, etc... Grouping these into functions within a single 'slurpinit.sh' script that will simplify maintaining the production scripts.
 - And simplify refactoring of (e.g.) odbc.ini
 - Thomas and Joe tried running single stream production with limited success - some observations
 - Success of script seemed to depend on run number (?)
 - 52907 did seem to resubmit, files were being produced as of yesterday
 - Joe started an ana445 yml file to test with the new macro. Blindly just changed the ana441->ana445 - TPC jobs appear to be running, but no silicon/TPOT jobs (?)
 - Possibly run doesn't have the TPOT / Si
 - New plans to discuss:
 - Job 0 is going to absorb the hitset unpacking. Event combining will run 38 jobs (one for each prdf) and then job0 will read in all the output files and produce cluster DSTs. Added new [macro](#) on Friday
 - Jin made the TPC decoder packet dependent. So we could further parallelize by splitting the 24 EBDC jobs into 48 packet jobs (2 packets per EBDC). Would require development on production workflow side to supply packet ID to macro
 - Other news - Tony will send Chris a new alignment file today, so we will want to update to 2024p008
 - ...
 - AOB
- **Meeting on Stream Parallel Workflow Wed Nov 13 2024 10AM**
 - Work within the testbed environment
 - log in to phnxsub04 ...
 - **Setup submission directory**
 - /direct/sphenix+u/sphnxpro/ProductionSystemTestbed/slurp-singlestream
 - Workflow directory ... (actually a symbolic link)...
 - /direct/sphenix+u/sphnxpro/ProductionSystemTestbed/slurp-singlestream/ProdFlow
 - Both directories setup on a single stream processing branch.
 - Single stream event builder is in
 - ProdFlow/run2pp/DST_STREAMING_EVENT_run2pp_ana441_2024p007.yml
 - Rule is DST_STREAMING_EVENT_run2pp_streams

```
./kaedama.py -rule DST_STREAMING_EVENT_run2pp_streams
--config ProdFlow/run2pp/DST_STREAMING_EVENT_run2pp_ana441_2024p007.yml
--runs first [last]
--streamname TPC{id}
--streamfile TPC_ebdc{id}_physics
--dbinput
```

 - See also the *submit_stream_parallel.sh* script
 - DST_TRKR_HIT_run2pp_streams rule is setup to match the output of the stream parallel builder...
 -
 - Output goes in /sphenix/lustre01/sphnxpro/testbed...
 - ...
 - Resubmission options for kaedama...
 - If a job will overwrite a file in the filecatalog the system will not submit the job:
 - -r will override this behavior
 - If the production status table contains an entry for the job you are trying to submit, it will be blocked from submitting
 - -u *status* will unblock jobs with the given status... status can be submitted, submitting, started, running, finished or failed.
 -
 - Multiple sources for run numbers
 - Daq DB (physics)
 - Add query to triage table (remove no good runs)
 - Calibration status (tbi)
 - Text files
 - Special table with selected runs
 - How do we submit run ranges (probably cannot just submit all runs, command line?)
 - How do we run this in a cron which keeps filling the queue?
 - Information about runs comes from the daq DB, if relevant for the production - it needs to fetch them from there
 - Streaming flag
 - Zero suppression
 - Strobe length
 - ...
 - Source for files is the RawdataCatalog
 - Files in Datasets will get a status flag
 - >0 good for processing -
 - 1 = in lustre, higher numbers reserved for whatever
 - 0 in hpss
 - <0 bad

- Detailed Feedback and fixing of problems
 - Get info on problematic files (not “5 runs didn’t process”)
 - Use production to patch up problems
 - Patch up and update md5 sum during production?
 - Many streaming files have 0xFFFFFFFF
 - Some bogus first, last event numbers
 - Bogus sizes
 - first/last bco needs filling
- Move to production DB, ditch production_status DB
- Output location consolidation
 - Top dir in lustre: /sphenix/lustre01/sphnxpro/production
Then production (run2pp, run2auau, run2cosmic,<calibration>)
Then tag (build_cdb)
Then DST type
 - Top dir in gpfs: /sphenix/data/data02/sphnxpro/production
 - Same subdir structure as in lustre
- **Meeting Mon Nov 04 2024 1 PM EDT**
 - Year 2 AuAu jet macro: <https://github.com/sPHENIX-Collaboration/ProdFlow/pull/30>
 - Macro is merged. Will setup and run under testbed environment today / tomorrow
 - Setting up an alma 9 Testbed Environment
 - phnxsub05 / tmux session will be sPHENIX-Alma9Testbed
 - SLURP requirements
 - Will share the standard testbed database / directory structure
 - ... wrapper script tpo run in container TBD...
 - DST_TRIGGERED_EVENT_run2pp accounting
 - Identified source of the difference between the 32.6B event vs 19.2B event starting event samples...
 - Emma starts with a runlist where all calorimeter modules are present, with no restriction on the existence of the GL1 file or requirement that all files have been transferred to SDCC
 - The (online) production rule builds a runlist where *any* calorimeter module is present, the GL1 file is present, and all files have been transferred to SDCC.
 - ...
 - Planning for next set of fixed build productions
 - Plan to use the new raw data file catalog as input.
 - Production rules will be similar to the online production rule... but the check on the ‘transferred to sdcc’ field is no longer needed.
 - Expect to segment the triggered data into fixed-event segments. (100k events/output files?)
 - Full or partial reproduction? (Leaning toward full b/c we didn’t process everything and we should do a final processing on alma9).
 - Expect to run multiple passes with last on alma9.
 - Expect to run stream-parallel streaming event builders? ... (any progress on timing?)
 - One improvement this past week. (Runaway container in GL1).
 - TODO: Get these fully operational. Run hit unpacker on each output then combine at job0 ...
 - TODO: Setup test production through job0 with job0 taking single segment / job. INTT will need to have more events per segment...
 - Job0 can be independently clustered. Final combination could be done at the seeding stage. TBD.
 - ...
 - Addressing / monitoring database stability issues
 - Previously noted a 0.5% failure rate in registering output with the file catalog.
 - Upgrade planned for cups:
 - file catalog and production status updates will have a uniform number of retries at random intervals
 - statistics will be collected on the number of retries, date/time of successful and failed attempts, record the exception encountered
 - RUN, SEGMENT, DSTTYPE, DATASET, the update type, START time, END time, N tries, Exception
- Run2 cosmics
- ...
- AOB
 - Field off AuAu MBD calibrations in the database? Runs 54279, 54280 and 54281.
 - Wiki page / production documentation
 -
- **Meeting Mon Oct 28 2024 1 PM EDT**
 - Run2AuAu online
 - DST_TRIGGERED_EVENT_run2auau_new_2024p007
 - 1968 jobs still running this morning (all restarted after the 7day expiration). I’ve removed them from the queue (and saved the condor_q to a file)
 - DST_STREAMING_EVENT_run2auau_new_2024p007
 - 62 jobs past expiration (removed)
 - 136 running 78 on hold
 - Still see downstream tracking jobs starting / running
 - Most of the data seems to be through the pipeline
 - Chatted with Tom Smith... expectation is that when we upgrade to newer version of condor w/ alma 9...time limit gets attached to the job, and we don’t need to play the periodic hold game.

- TODO: Confirm that we can place things on hold
- RAW data catalog is ready... will handle files which are incorrectly flagged in daq DB as not on sdcc. In future we will use this as the master... so plan to drop the daqdb for fixed build and online productions.
- Run2pp
 - Held DST_TRKR_SEED jobs from last week raising the memory 2GB to 4GB
 - 2331 jobs finished (~ 80%)
 - 19 appear to still be running (checked 1 log file, so 1 is certainly running)
 - 491 jobs went back on hold
 - Few condor logs checked... memory issues.
 - Strong suspicion that these are noisy TPC events which cause the problem and exhaust the memory
 - ... progress expected ~week
 - DST_TRIGGERED_EVENT_run2pp_ana437_2024p007
 -

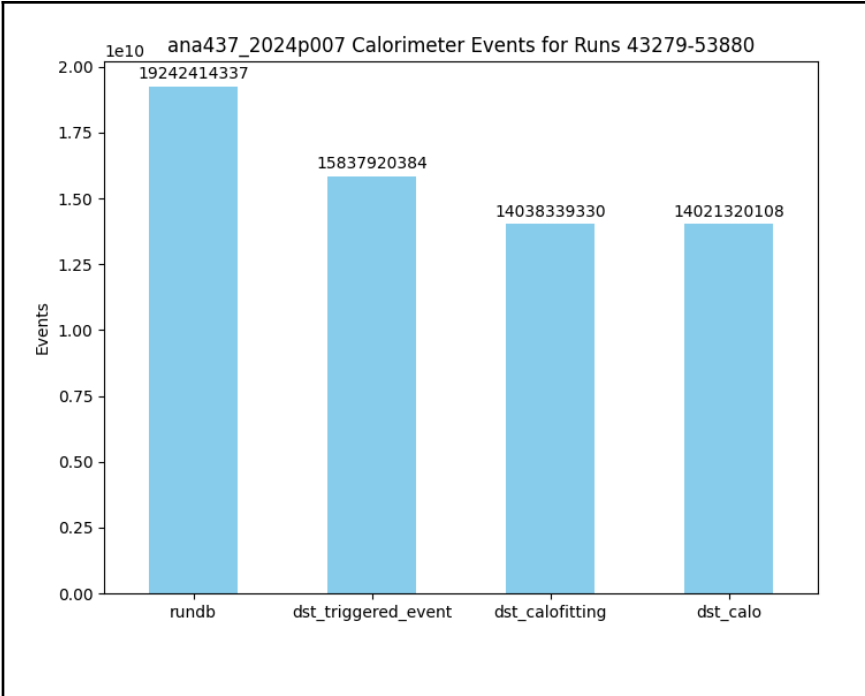


Figure 1

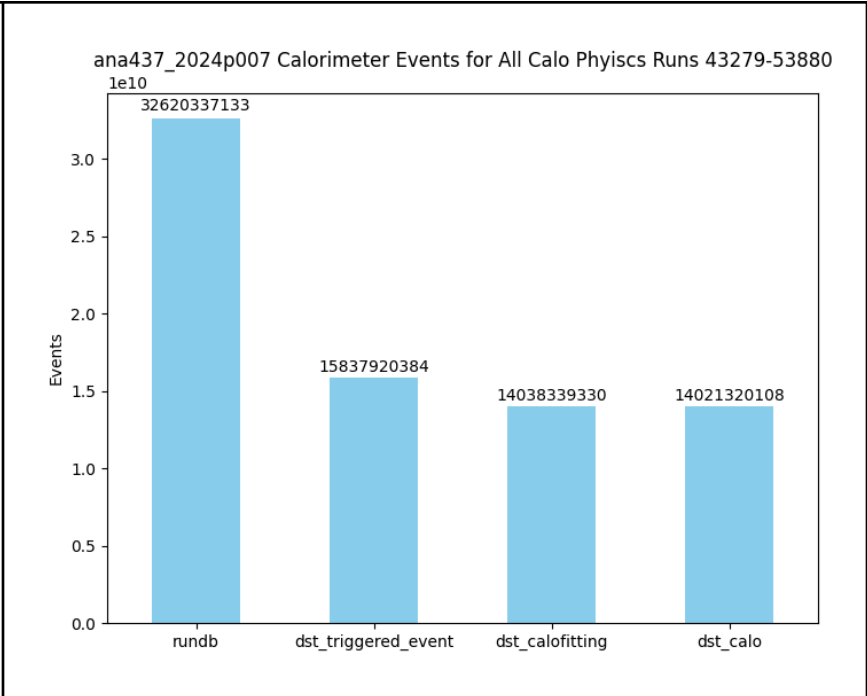


Figure 2

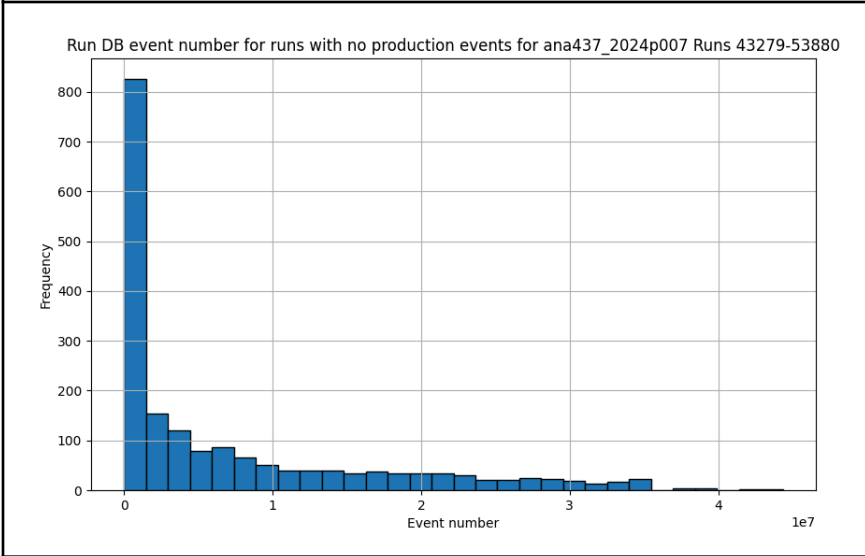


Figure 3





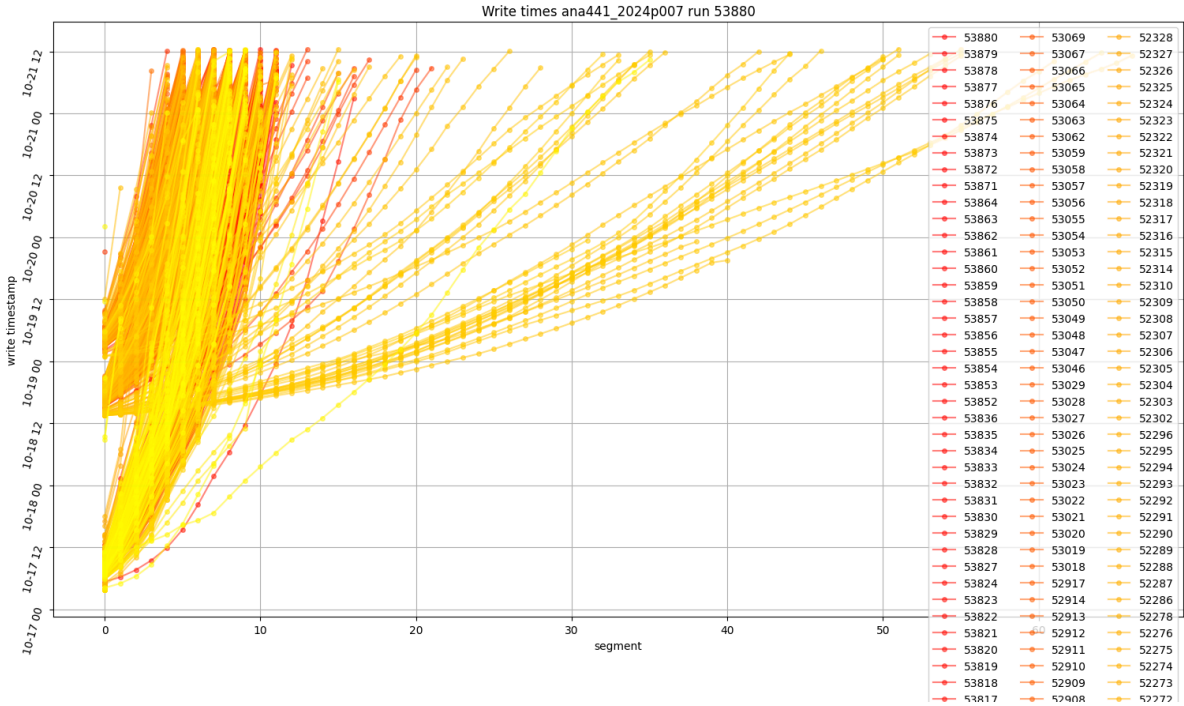
Figure 4

-
- ana441_2024p007 and ana439_2024p007 cleanup
 - DST_STREAMING_EVENT processed before first streaming runs / macro not properly configured for the earlier runs
 - Downstream processing was properly limited to the run range
 - Noted some small discrepancies in the file catalog. 0.005 of files were not properly registered in both the files and datasets table.
 - I have an update to the `cups.py` script which (a) makes the filecatalog update atomic... both will succeed or fail; (b) applies a uniform `retry` scheme if we can't update to the DB. Retries 10 times, with random delay between (Ntry+1)*10 seconds between attempts.
 - Should deploy with the next large scale fixed build production
- Run2 Cosmics (Shuhang)
 - Setup for development under the testbed area
 - DST_TRIGGERED_EVENT_run2cosmics
 - TODO: apply similar naming scheme in the streaming data ...
-
- AOB
 - Year 2 AuAu jet macro: <https://github.com/sPHENIX-Collaboration/ProdFlow/pull/30>
 - Runs under ana443 (or new)
 - TODO: run jet finder over run2auau online production
 -

- Meeting Mon Oct 21 2024 1 PM EDT
 - Online AuAu wrap up / plans for fixed build production



- Setting up area for fixed build production. Initially under the testbed environment.
- Jet QA PR over the weekend... plots are showing up.
 - Macro change TBD. pp / auau flag and trigger list.
- Resubmit after macro change ... triggered data ... jet production
- TODO: event accounting (i.e. emma's script).
- Run2pp production
 - Held jobs... 50% seeding (big busy single events). 50% hit unpacking. few% others. 5k+ jobs total.
 - TODO: repeat with 4G
 - 142161 total TRKR_HIT jobs submitted ... so 3-ish % fail rate
 - Switch to a new memory for the production rule and resubmit these.
 - Thomas Marshall production deputy for streaming productions
- Run2pp ana437_2024p007 calorimeter production finished for missed final runs (~13k jobs).
 - 25% of event building jobs slowed down to a crawl and eventually were evicted.
 - But waveform fitting and reconstruction steps went through essentially without losses.
 - Biggest calorimeter dataset so far / about 1/2 of the data
 - TODO: event accounting
- DST_STREAMING_EVENT_ana441_2024p007 status...
 - 3845 jobs submitted
 - 1530 have reached 400k+ events processed
 - 2097 event builders still running running [~800 3d rest 5d]
 - ... additional runs fixed in file catalog?




- Stream-parallel streaming data production
 - Some timing information.
 - Run 53880 whole stream production vs stream parallel processing...
 - TPOT first to finish with 38M events (failure is marked b/c log file is large since each event number is printed out).
 - TPC streams are the slowest, and therefore drive the total time of the whole stream production.
 - Parallel streaming builds the TPC events 10x faster than the full stream processing.
 - But uses >24x as many nodes to do so
 - So... parallel building is a factor 2.4x *slower*
 - Snapshot below is taken at 2.5 days (9AM-ish this morning).
 - Extrapolating the parallel stream TPC jobs... we need to run the jobs for 16 days to build this run to completion IF the time per segment is linear (see below).

FileCatalog=> select id,dstfile,status,nevents,now()-started,message from production_status order by id desc limit 40;


| id | dstfile | status | nevents | ? | column? |
|--------|---|---------|----------|------------------------|---------|
| 171277 | DST_STREAMING_EVENT_run2pp_ana441_2024p007-00053880-00000 | running | 60000 | 2 days 12:30:47.962271 | |
| 171276 | DST_STREAMING_EVENT_TPOT_run2pp_ana441_2024p007-00053880-00000 | failed | 38361009 | 2 days 12:32:50.962271 | |
| 171275 | DST_STREAMING_EVENT_TPC23_run2pp_ana441_2024p007-00053880-00000 | running | 580000 | 2 days 12:32:50.962271 | |
| 171274 | DST_STREAMING_EVENT_TPC22_run2pp_ana441_2024p007-00053880-00000 | running | 580000 | 2 days 12:32:50.962271 | |
| 171273 | DST_STREAMING_EVENT_TPC21_run2pp_ana441_2024p007-00053880-00000 | running | 600000 | 2 days 12:32:51.962271 | |
| 171272 | DST_STREAMING_EVENT_TPC20_run2pp_ana441_2024p007-00053880-00000 | running | 570000 | 2 days 12:32:51.962271 | |
| 171271 | DST_STREAMING_EVENT_TPC19_run2pp_ana441_2024p007-00053880-00000 | running | 570000 | 2 days 12:32:51.962271 | |
| 171270 | DST_STREAMING_EVENT_TPC18_run2pp_ana441_2024p007-00053880-00000 | running | 570000 | 2 days 12:32:51.962271 | |
| 171269 | DST_STREAMING_EVENT_TPC17_run2pp_ana441_2024p007-00053880-00000 | running | 510000 | 2 days 12:32:50.962271 | |
| 171268 | DST_STREAMING_EVENT_TPC16_run2pp_ana441_2024p007-00053880-00000 | running | 520000 | 2 days 12:32:50.962271 | |
| 171267 | DST_STREAMING_EVENT_TPC15_run2pp_ana441_2024p007-00053880-00000 | running | 570000 | 2 days 12:32:51.962271 | |
| 171266 | DST_STREAMING_EVENT_TPC14_run2pp_ana441_2024p007-00053880-00000 | running | 590000 | 2 days 12:32:51.962271 | |
| 171265 | DST_STREAMING_EVENT_TPC13_run2pp_ana441_2024p007-00053880-00000 | running | 590000 | 2 days 12:32:51.962271 | |
| 171264 | DST_STREAMING_EVENT_TPC12_run2pp_ana441_2024p007-00053880-00000 | running | 570000 | 2 days 12:32:51.962271 | |
| 171263 | DST_STREAMING_EVENT_TPC11_run2pp_ana441_2024p007-00053880-00000 | running | 570000 | 2 days 12:32:51.962271 | |
| 171262 | DST_STREAMING_EVENT_TPC10_run2pp_ana441_2024p007-00053880-00000 | running | 570000 | 2 days 12:32:51.962271 | |
| 171261 | DST_STREAMING_EVENT_TPC09_run2pp_ana441_2024p007-00053880-00000 | running | 580000 | 2 days 12:32:51.962271 | |
| 171260 | DST_STREAMING_EVENT_TPC08_run2pp_ana441_2024p007-00053880-00000 | running | 600000 | 2 days 12:32:51.962271 | |
| 171259 | DST_STREAMING_EVENT_TPC07_run2pp_ana441_2024p007-00053880-00000 | running | 580000 | 2 days 12:32:51.962271 | |
| 171258 | DST_STREAMING_EVENT_TPC06_run2pp_ana441_2024p007-00053880-00000 | running | 560000 | 2 days 12:32:51.962271 | |
| 171257 | DST_STREAMING_EVENT_TPC05_run2pp_ana441_2024p007-00053880-00000 | running | 530000 | 2 days 12:32:52.962271 | |
| 171256 | DST_STREAMING_EVENT_TPC04_run2pp_ana441_2024p007-00053880-00000 | running | 600000 | 2 days 12:32:52.962271 | |
| 171255 | DST_STREAMING_EVENT_TPC03_run2pp_ana441_2024p007-00053880-00000 | running | 560000 | 2 days 12:32:52.962271 | |
| 171254 | DST_STREAMING_EVENT_TPC02_run2pp_ana441_2024p007-00053880-00000 | running | 560000 | 2 days 12:32:52.962271 | |
| 171253 | DST_STREAMING_EVENT_TPC01_run2pp_ana441_2024p007-00053880-00000 | running | 610000 | 2 days 12:32:52.962271 | |
| 171252 | DST_STREAMING_EVENT_TPC00_run2pp_ana441_2024p007-00053880-00000 | running | 560000 | 2 days 12:32:52.962271 | |
| 171251 | DST_STREAMING_EVENT_MVTX5_run2pp_ana441_2024p007-00053880-00000 | running | 33270000 | 2 days 12:32:52.962271 | |
| 171250 | DST_STREAMING_EVENT_MVTX4_run2pp_ana441_2024p007-00053880-00000 | running | 28910000 | 2 days 12:32:52.962271 | |
| 171249 | DST_STREAMING_EVENT_MVTX3_run2pp_ana441_2024p007-00053880-00000 | running | 30410000 | 2 days 12:32:52.962271 | |
| 171248 | DST_STREAMING_EVENT_MVTX2_run2pp_ana441_2024p007-00053880-00000 | running | 28240000 | 2 days 12:32:52.962271 | |
| 171247 | DST_STREAMING_EVENT_MVTX1_run2pp_ana441_2024p007-00053880-00000 | running | 34180000 | 2 days 12:32:52.962271 | |
| 171246 | DST_STREAMING_EVENT_MVTX0_run2pp_ana441_2024p007-00053880-00000 | running | 26530000 | 2 days 12:32:52.962271 | |
| 171245 | DST_STREAMING_EVENT_INTT7_run2pp_ana441_2024p007-00053880-00000 | running | 1970000 | 2 days 12:32:52.962271 | |
| 171244 | DST_STREAMING_EVENT_INTT6_run2pp_ana441_2024p007-00053880-00000 | running | 2070000 | 2 days 12:32:52.962271 | |
| 171243 | DST_STREAMING_EVENT_INTT5_run2pp_ana441_2024p007-00053880-00000 | running | 2040000 | 2 days 12:32:52.962271 | |
| 171242 | DST_STREAMING_EVENT_INTT4_run2pp_ana441_2024p007-00053880-00000 | running | 1960000 | 2 days 12:32:52.962271 | |
| 171241 | DST_STREAMING_EVENT_INTT3_run2pp_ana441_2024p007-00053880-00000 | running | 2100000 | 2 days 12:32:52.962271 | |
| 171240 | DST_STREAMING_EVENT_INTT2_run2pp_ana441_2024p007-00053880-00000 | running | 1940000 | 2 days 12:32:52.962271 | |
| 171239 | DST_STREAMING_EVENT_INTT1_run2pp_ana441_2024p007-00053880-00000 | running | 1820000 | 2 days 12:32:52.962271 | |
| 171238 | DST_STREAMING_EVENT_INTT0_run2pp_ana441_2024p007-00053880-00000 | running | 1710000 | 2 days 12:32:53.962271 | |

(40 rows)



- Nonlinearity?
 - Not likely to be memory leaks. Memory fragmentation?
 - Try running without output....
 - Code change to print event # and timestamp....
- Jet meeting ...
 - track jet QA expecting end of week. DST_TRKR_TRACKS → jet finding QA histograms
- 2024 HCAL calibrations...pp ... latest ana437_2024p007
 - cosmics first up... cosmics are mia
 - TODO: restart the cosmics builder. Test first and point Blair to them.
- AOB
- Meeting Mon Oct 14 2024 1 PM EDT
 - DST_TRIGGERED_EVENT_ana437_2024p007
 - Missing several downstream products ...
 - Running kaedama... over the range 43263 to 53880
 - 328 waveform fitting jobs submitted
 - >850 calib jobs submitted + 270 jobs submitted + ...
 - ... looks like we did not run the jet production (or most of the jet production).... but many (most) of these jobs are exceeding 2GB of memory when they run. Need to check for/fix memory leaks before proceeding with the jet production
 - Missing piece closure test.
 - Pp Streaming with fixed build... ana439_2024p007. new?
 - Why do we only have 8 runs with >=40 segments (400k events)?
 - run2auau status / discussion
 - Production restarted for streaming event builders + downstream tracking from run 53881 onwards w/ correct auau macros
 - Production restarted for triggered event builders + downstream jobs... run 54100 onwards... w/ correct auau macros
 - Why is it run from a dev branch, what happens to PR's which use the master as basis?
 - ... I have to cherry pick the relevant commits and merge into the online branch. Would prefer commits to the online that target the online. and we can merge to master at later point.
 - The `online_triggered_event_run2auau_new_2024p007_v1.0` ... is meant to apply and track any and all changes that we apply during the online production, and allows us to keep this separate from the master branch so that they do not impact any fixed build productions we may have going on concurrently.
 - Action: merge all onto the master branch. future fixed productions should be done from a branch.

- Offline Production
 - Goal - auto pilot
 - Give list of runs (the whole range) and just monitor progress, no calling submissions by hand
 - Cron jobs (stable against submission host crashes/reboots)?
 - Current...
 - ramenya submit RULEDEF.yaml --rules DST_STREAMING DST_TRIGGERED ...
 - Monday 8AM end of AuAu
 - Switch to preparation for fixed build production of the complete dataset.
 - ... MBD calibration ... waiting on Mickey to sign off so that I can kick off the jobs
 - → 53881 (AuAu begin) to 54100 mbd calibrations to be done
 - AOB
- Meeting Mon Oct 7 2024 1 PM EDT
 - run2auau status / discussion
 - ...
 - First usable collisions this morning. 7 AM triggered detectors taking data under ‘beam’ runtime. Several runs worth.
 - Transfers not running so data not going to SDCC.
 - ... full runs to be transferred ...
 - Plan to push to lustre everything, data rate allowing
 - ... calorimeter workflows once data available ...
 - run 54117 is available... /sphenix/lustre01/sphnxpro/physics/emcal/beam/*54117* and replace “emcal” with other subsystem names
 - 54094 appears to be the first one
 - New build... 2024p007
 - run2pp production(s)
 - DST_TRIGGERED ...
 - Calorimeter reconstruction of the full run2pp dataset with ana437_2024p007 tag finished in a few days.
 - Fast submission rate, 25k jobs submitted in less than a minute.
 - Jobs idle on condor for many hours.
 - Deputy – Shuhang
 - Production QA
 - ana437 recorded 30B events / production reached 18B events / 14B made it through to event combining / 13 waveform / 12B final calibrated DSTs...
 - 30B → 18B issue is that data is not all on disk
 - 20-30% software failures on event combining
 -
 - DST_STREAMING ...
 - Physics quality only from 51429 to 53880... all other runs will need to be custom jobs by hand.
 -
 - Friday’s production attempt failed w/ many event builders timing out during a database query to get their input parameters.
 - Restarted on Sunday evening...
 - Added protection in the script... if the input file lists contain no file, we mark the run as failed and return an exit code 111 to distinguish the failure mode
 - We feed in 50 jobs during each iteration of the loop, with 10 minutes between iterations. This should limit the load on the database
 - Downstream jobs (which take only one input file) will have those files passed on the command line, not via the production database table. Up to 5k jobs start at once. 10 minute loop.
 -
 - Monday morning status
 - slurp improvements ...




- significantly improved the submission times
 - a little better optimized for building physical file lists from the filecatlog
 - much better optimized for building physical file lists from filesystem lookup
 - time to submit 3k runs / streaming event builders about 2min....
 - consequently, 3k streaming jobs start up in quick succession and try to obtain input files from the production status database... database is not happy, requests time out and builders get empty file lists.
- AOB
-
- **Meeting Mon Sep 30 2024 1 PM EDT**
 - Online production status
 - Failure modes
 - Talked to Tom S. about restarting jobs. After looking at some examples he could not offer a solution immediately, said he will talk to condor developers.
 - DST_TRIGGERED_run2pp_ana437_2024p007 problem (NOT put on hold) jobs example: run 47198 works well until segment 155, then only segments 163, 168, 191, 205 (all short, 2-3k events, but good for waveform fitting). Missing segments have segfaults (e.g. seg 165), short segments have “MBD event stack is empty, we are done” e.g. seg 168.
 - Jobs put on hold (few percent, seemingly random) produce normal size root files, log files have lots of “G11in: Found dropped Event at event 7821872 with Packet Number: 3176524637” meggages, they appear in successfully finished jobs too.
 - Likely indicates incomplete runlist. **TODO: integrate from start file to end of file list.**
 - DST_CALOFITTING_run2pp_ana437_2024p007 all jobs have “ls: cannot access HIST_*.root” error message
 - Few percent get evicted and put on hold without root or log file output. Condor logs have a shadow exception “Assertion ERROR on (result)”, memory usage ~750Mb, image size 1465716. Example: run 47063 segment 78.
 - Short event builder segments finish waveform fitting successfully.
 - **TODO: Spam the data02 disk with log files**
 - DST_TRIGGERED_run2pp_ana437_2024p007 Production Status
 - 149841 event builder jobs finished
 - 2209 runs / 14.9B events processed.... (should have 20B events)
 - This accounts for all of the data which is in the file catalog *and* on disk.
 - 35693 event builder jobs are marked as failed (Fun4All returned nonzero exit code)
 - 36039 event builder jobs are marked as running...
 - 36017 jobs were removed because they ran in excess of 12h
 - ~4k were put on hold yesterday 7-8AM, discussion below...
 - 22 jobs finished and staged output, but we did not get the ‘finished’ update in the DB
 - 106606 waveform fitting jobs finished
 - 1983 runs / 10.6B events processed
 - No failures
 - 25k jobs are still running on condor
 - 18k jobs were placed on hold yesterday around 7-8AM
 - HPSS still staging to lustre...
 - TODO: run as a production loop.... once / 6h
 - Event builder jobs...
 - daqdb has 226828 entries
 - status table has 149443 entries... 77385 jobs difference... (but may be integrating over too many daqdb entries... TBD)
 - file catalog has 116778 entries
 - datasests has 116785 entries
 - Event accounting in <https://github.com/sPHENIX-Collaboration/ProdFlow/tree/master/RHIC2023/CALOR/macros>
 - Large number of jobs went on hold 9/29 between 7-8AM
 - ~4k event builders which had been running for 7+ hours.
 - 18375 waveform fitting jobs were on hold when I checked in the afternoon. Approx 200 of these were already on hold.
 - We have had an additional 21 jobs go on hold since yesterday
 - Condor logs do not indicate memory issues. We only see that the jobs failed the periodic hold condition.
 - Looking at the class ads, the jobs did fail the PH condition
 - They had been executing for anywhere from 1 - 3d
 - I don’t see any anomalies when looking at the number of jobs that executed on each node
 - Does not appear to have any correlation with the length of time the job was running. Lower numbered clusters / longer execution time are represented, as are higher numbered clusters / shorter execution time.
 - No correlation with run numbers
 - Copied .../sphenix/data/data02/sphnxpro/SchedLog.20240930T082118
 - **TODO: Ping Tom Smith with this... ask to retain more?**
 - Long executing queries...
 - If, for example, I were to submit the full DST_TRIGGERED_run2pp_ana437_2024p007 data set, I would get the following error.
 - pyodbc.Error: ('40001', '[40001] ERROR: canceling statement due to conflict with recovery\nDETAIL: User query might have needed to see row versions that must be removed;\nError while executing the query (1) (SQLExecDirectW)')
 - TODO: ping Dmitry
 - because (1) the query to build the jobs is somewhat complicated and (2) there is a significant amount of data to process. but...
 - We don’t really expect people to be deleting entries in the daqdb or the file catalog... at least, not with any great frequency.
 - Once there is a timeout... it will consistently throw the same exception on the same connection. But a subsequent connection seems to succeed.
 - Tracking is ready for a fixed build reproduction of all existing data on sdcc. Not going to be the whole data set but that is fine, we need something stable. Run a small test production on a single run to produce a few segments for us to examine first?
 - Build: ana.439
 - CDB tag: 2024p007
 - Start production from run #:
 - Full ‘standard’ production. (And test the single readout card)

- Preparation for AuAu running later this week/early next week
 - Collisions expected Wednesday night. Trackers turned on Friday earliest.
 - Last pp run #.... [53880](#)
 - TODO: Workflows setup for zero and nonzero suppressed data... (e.g. how many events).
 - How to setup production system to mirror last august... quick turnaround time ...
 - Main problem was non-zero suppressed data...
 - Need to separate zero and non-zero suppressed data
- JetQA / HTML merged.
- AOB / Discussion

● Meeting Mon Sep 23 2024 1 PM EDT

- Production status
 - Held jobs / open issues / etc...
 - Dozen failed over the past week. Mix of 7 day holds, and stalled jobs. Few jobs hit the 7 day mark and restart. (Have seen individual jobs restart several times, if not caught). Need to involve Tom Smith / SDCC. How to prevent it from starting in the first place.
 - *Parallel Processing triggered event builders*
 - Running under testbed
 - Only processes runs which have been completely transferred
 - Calo runs 80-90k in single input segment
 - Segments each run into 100k event chunks, with last segment allowed to be 120k.
 - Each job's event range is determined at submission time and passed to the production script through condor
 - No additional load on DB from running jobs
 - Tested with O(5k) jobs from run range 50000 to 50500:
- The graph shows I/O throughput over a 5-hour period. The y-axis represents throughput in GB/s, ranging from 0 to 400. The x-axis shows time from 15:00 to 20:00. The 'read' throughput (green line) starts low, rises sharply to a peak of about 340 GB/s at 16:30, then gradually declines with some fluctuations, ending around 100 GB/s. The 'write' throughput (orange line) remains very low throughout the entire period, staying below 20 GB/s.
- O(1k) jobs continued to run beyond 4h. Chris identified code issues, library to be patched.
 - Event realignment, ... buffering, clock counters, vector comparison, ...vector not being cleared at the proper place.
 - Setting up to put this into production...
 - Build is ana.437
 - CDB tag at 2024p007
 - Produce up to calo fitting.
 - Calibrations to be used at the analysis stage
 - Start during the day. Loop in SDCC (both DB and condor expertise). Testing limited to O(100) jobs. O(1 run).
 - 300 GB/s is > max
 - Aim for a 9AM start tomorrow.
 - Any limitations would be managed by condor resources. TBI.
 - Cleanup...
 - Same naming convention.
 - ODBC driver update
 - Dmitry identified issues with timeouts due to old ODBC driver. Magic entry in odbc init file... latest odbc driver should work alot better w/ the psql databases.
 - ... TODO ... restart the production loops.
 - AOB
 - Discussion

● Meeting Mon Sep 16 2024 1 PM EDT




- Held and stalled (running, no output) jobs
 - EMCAL jobs on hold primarily 7 day limit / Few jobs unexplained
 - Fixing up daq db. Few runs marked as not transferred, ... auto loop should pick them up.
- Jet QA up and running from run 52332.
 - Update? Expecting a PR by end of week for the HTML side of QA. Output looking reasonable.
- Migration of production job definitions / scripts / macros to sPHENIX repo
 - We had established such a repo back in the PanDA era:

- <https://github.com/sPHENIX-Collaboration/ProdFlow>
- Propose to create a “run2pp” directory in this repo
 - run2pp/ – top level contains the job description “yaml” files
 - run2pp/TrackingProduction – tracking workflow scripts and macros
 - run2pp/CalorimeterProduction – &etc...

- Parallel Processing...
 - Very little code changes. Skipping through files relatively fast now.
 - Figure out which segment the starting file is in. That’s the first file in the file list. Add he rest.
 - Needs a macro change. First, last event inclusive. [A,B] processes all events event>=A and event<=B.
 - event alignment is not yet implemented.
 - Working just for the calorimeters at the moment.
 - Streaming needs a bit more work... events → bunch crossing counters.
 - may be advantageous to run multiple streams... b/c allows us to run 24 2GB jobs vs one big one.
 - Next production pass... will have multiple thousands of jobs running concurrently.
 - Calorimeter aiming for QM... abstracts 31st Oct

- Single stream on testbed for run 52345 ...



- “Full stream” productions exceeded the 7 day condor limit
- Single stream jobs running for three days now... and still writing output
- Next step is to feed into downstream tracking workflows.


- ODBC
 - Driver has alot of overhead.

- AOB
- Discussion

• Meeting Mon Sep 09 2024 1 PM EDT

- Recent production system updates (by date desc)
 - Nifty graphical command line utility...
 - Soft cap on the number of jobs the production system will place on the condor schedd
 - Reduce load on production DB replicas by passing the row ID as a job parameter [not yet merged]
 - Use production DB to pass long filelists to the jobs / make default
 - Calorimeter builders switch to ana.435 from run 52332
 - If a partial run is copied (i.e. single segment) it will run. If a full run is copied, it will run. n.b. negligible(?) possibility that a full run could be picked up as a partial run depending on timing of file transfers.
 - If this becomes an issue, we may need to add a ‘queued_for_sdcc_transfer’ flag to the DAQ db. (Or make transferred_to_sdcc an enumeration of ‘false’, ‘queued’, ‘true’).

- Snapshot at 9AM this morning




- We have been experiencing significant number of idle jobs for last week or so
 - Users need a magic line... limiting them to 23k jobs... so far working well. Snapshot may have caught right after the production loop submitted jobs.
 - STAR jobs interfering with jobs starting
 - No more problems over the weekend
- Streaming event builder jobs are no longer going on hold. Elimination of condor copy back seems to have solved these issues.
- More holds on triggered event builders
 - Currently 19 calorimeter builder jobs on hold b/c they exceeded the 7d limit...
 - Runs 51908 51914 51915 52106 52321 ...
 - TODO... detect when a job stalls out...
 - Looking at the log files these ^^^ jobs did not write any output after 2-3 days.
 - Run 51905 is sitting at 15+ days... snuck by our periodic hold condition. Condor job 45268.0 (which I just removed)

- Another class of 9 jobs went on hold after ~10 h - 2d. Appear to be memory issues.
- /tmp/calologs/ana435_2024p007/run_00052600_00052700/DST_TRIGGERED_EVENT_run2pp_ana435_2024p007-00052686-00000.condor

- Jet Production workflow
 - Workflow implemented and test runs made under the testbed
 - No DST output, so stageout was failing, but this was expected
 - More concerning is that there is no histogram output from the macros
 - **[Derek]** macro was missing a couple headers/shared objects for the ZDC (that’s where the “Missing Zdcinfo…” errors were emanating from)
 - [PR open!](#)
 - Go back to the high lumi zero xing angle period.
- Working on single stream event builder workflow...
 - Basic idea is to parameterize the streaming workflow...

```
for i in 00 .. 23; do
./kaedama.py --runs 52345 --config
production-rules/DST_STREAMING_EVENT_run2pp_ana435_2024p007.yaml --rule PHYS_DST_ST
REAMING_EVENT_run2pp --streamname TPC${i} --streamfile TPC_ebdc${i}_physics
done;
```

and etc&




- Originally expediting event building. Not the long term solution. Could be used in general regardless of skipping around. File management is the downside. Could consider event combining and hit unpacking run in the same job.
- Parallel reading is coming along. Not mutually exclusive...
- Discussion on how to move forward...
- Raw data catchup is still ongoing...
 - Data from June being transferred.
 - Heads up... expect low run numbers to be propagating through the system.
- Old business / discussion
- AOB

• **Meeting Mon Aug 26 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261

- Status of fixed build calorimeter production
- Current state of the infrastructure
 - Lustre and HPSS simultaneous transfers too much bandwidth
 - To HPSS 2G/s faster... decision to stick to direct HPSS transfer only
 - Now need to stage to disk... data carousel upcoming...
 - HPSS cache disk i/o not designed for simultaneous in out
 - Already 2x higher than the 10GB/s sustained rate (remedied by transfers only in between fills)
 - May need to be selective about runs to process / fractions of runs to process
 - Volume of TPC volume is 5x to 7x initial plans (offset by CAD...)
 - Lustre, gpfs, condor, databases,
 - Readiness to restart? *Good to go.*
- Testbed environment established

- Running under sphnxpro on phnxsub04, under tmux session sPHENIX-Testbed
 - Filecatalog and production status tables on sphnxhpssdbmaster.
 - /direct/sphenix+u/sphnxpro/ProductionSystemTestbed
 - main branch:
 - Verified streaming and triggered workflows operate in the testbed environment
 - By default, dst, histograms and log files go under a ‘testbed’ directory structure
 - rsync-stagein branch:
 - ... modifies workflows to stage in their own support files (eg macros, production scripts, cups, ...)
 - Ready to test
 - Status of (solution for) held jobs
 - Status of simultaneously open DB connections... does not scale.
 - TODO: move the TESTBED output up... in hierarchy.
 - TODO: job timing queries / plots
 - ... Joe will take a look at log files ...
 - No longer need ‘dominoes’ (30 min) level of delivery.
- Discussion of procedures
 - Deliverables
 - Calorimeters... run over everything that we have. Seems to be running well.
 - Tracking...
 - Old runs @ 100 event chunks
 - Need to get into a mode with useful output. Reproduce old runs ? All runs (with new or fixed?) ... or continue with just incoming runs?
 - *Focus on the new data with turnaround time sufficient to validate quality of the data*
 -
 - Resubmission of previous DST productions
 - ...
 - Integration of new workflows
 - Propose that all changes go onto a separate branch of slurp-examples (or other, better named repo...)
 - Merge into the test environment and test (CI/CD pipelines)...
 - Then integrate into the production environment
 - Do we need a separate ‘fixed build’ production environment from the ‘online’ and ‘testbed’ environments?
 - Jobs from last ~5 days not running
 - From calos, no histogram files produced since 51825 (probably from lustre/cvmfs move issues). For trackers, no histogram files produced since 51617 (jobs indicate submitted but not running?)
 - TODO: Clear existing outputs / logs and resubmit.
 - Need to be able to do an intelligent ‘clear’ command.
 - Including TPC drift velocity calibration into production workflow
 - Not a trivial workflow.
 - 1) Whole tracking with initial guess, 2) extract drift velocity, 3) rerun tracking
 - Requires jet qa, tracking and calorimeter production
 - Input is streaming tracking production and calo production. Default tracking... script to do the fitting. Good fit results in a CDB tree. Failed fit goes to another fit attempt....
 - IN: DST_TRKR_TRACK from current workflow
 - IN: DST_CALO_CALIB from current workflow
 - OUTPUT: root file → CDB tree
 - Reprocess jobs A and C. Thinking to overwrite original outputs? Encourage a new filename, and delete the intermediate file.
 - Production macro repo
 - There is in general confusion about where to go to learn about the production macros, which are not (necessarily) the same as in sPHENIX-Collaboration/macros.git
 - TODO: Establish a production area w/in the sPHENIX git repository.
 - Open PR’s
 - <https://github.com/klendathu2k/slurp-examples/pulls>
 - Jet QA
 - Production macro → input file list, output histograms, uses calo DSTs as inputs. Will be using tracks + calorimeters along the way.
 - Discussion
 - Re-processing of older runs with ana430_2024p007. Plot made using “ls”.



- Done, except for 45800-45900 run range (“fit data is empty” error), some missing waveform fits for later runs
- Some DSTs are in the wrong directories. (Filecatalog?)
- Calibrations in database?

- Lack of disk space.
 - Resubmission (--resubmit) should wipe old files regardless of directories. We need to remove the feature so that the system always blocks on the existence of a file.
 - AOB
- **Meeting Mon Aug 19 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Current situation
 - Recovering from an issue that occurred starting yesterday afternoon where jobs were being submitted multiple times
 - This does not appear to be every run since yesterday (statement needs confirmation)
 - This caused problems with the database servers. [Restart was required]
 - Temp redirect everything through the write DB / now back online.
 - This does not appear to be reproducible (at low loads AND using only the write DB)... streaming event builder is running since 1135AM (thanks Sasha), with no evidence of multiple submits.
 - 2nd issue are jobs that get resubmitted rather than placed on hold
 - A first test will be do we see issues after Chris switches back to the read DB?
 - Over the past week+
 - 100 events/streaming builder output file seems to be stressing the system and revealing sub optimal code. Likewise, the database size (7M entries) is getting prohibitive to do exhaustive queries... FC DOS...
 - Database servers are being overloaded
 - We are presenting more input files to the downstream processing chains in 1 hour than we can submit in 1 hour
 - Potential for a race condition when database takes a long time to insert the production status (or worse fails, but this has not been directly observed).
 - Some issues / fixes that have been applied
 - Make sure that queries have a limit to the number of rows returned (where possible).
 - Eliminate as many queries that don't use indexed columns as possible
 - DST name is not indexed → Use RO replicas to lookup the uniq ID of all rows with a given run, segment. Then select in python code the ID with the corresponding DST name. Write is based on the unique ID. Future improvement will pass the unique ID down to the job through its argument list.
 - Kaedama can exit early if it has no candidate jobs to build
 - Can prune the candidate output list, removing any candidate that already has a file on disk
 - Use bulk insert into the production_status DB
 - Time to submit 1k jobs went from 2min to 35s.
 - Future
 - Achival DB table for production status
 - Move higher number of events/job. Aiming for <1-2h for the downstream tracking jobs.
 - Condor holds issue
 - Two jobs timed out... one was resubmitted, the other not. Shadowlog may be informative.
 - Argument list exceed due to linux limitation
 - False positives... todo: test if a nonzero exit results in a hold.
 - Discussion
 - TEST production system (which is subject to the loads involved).
 - Path forward
 - Restart from 51588
 - Tracking for priority list zero suppressed studies
 - Launch if not done
 - AOB
 - Log files... error message local variable r not defined
 - Next week Xudong
- **Meeting Mon Aug 12 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Current configuration of the production system
 - DST_STREAMING, and detector specific DSTs
 - slurp-examples/sPHENIX/RAMENYA_LOOP_PHYSICS_2024.yaml
 - 100 events per file, ana build = new, cdb tag at 2024p002
 - fast production mode disabled at 10AM this morning
 - Downstream tracking jobs
 - ana build = new, cdb tag at 2024p007
 - DST_TRIGGERED
 - production-rules/DST_TRIGGERED_EVENT_run2pp_ana430_2024p007.yaml
 - 10k events per file, ana build 430, cdb tag at 2024p007
 - Downstream reco jobs
 - Waveform fitting, calibration corrections
 - ana build = 430, cdb tag at 2024p007
 - MBD_CALIBRATION
 - ana build = new, cdb tag at 2024p005
 - TODO: double check tpc calibration data type...
 - TODO: triggered histograms
 - O(500) jobs steady state, with O(10k) downstream production periodically. Some instability in DB, needs characterization, instrumentation, testing...
 - Discussion of recent issues
 - Held jobs (not false positives...)
 - Few jobs per day (not the 10k's from last week).
 - Data 02 filled up, causing shadow exceptions when condor was failing to copy back log files and DSTs.... many jobs went on hold as a result.
 - Transfer of the stdout and stderr files that condor caches on the worker node is under the control of the "output" and "error" submission lines. I

- I incorrectly thought that condor copy back had been disabled by adding the command `transfer_output_files = ""`... see below
- ...
- Introduced a new requirements expression on Wednesday, requiring that the NumJobStarts is undefined or zero.
 - Initial submission test worked fine. Jobs submitted and completed.
 - Subsequent job submissions, both production loop and manual, froze during the call to condor submit.
 - Reverted to the old (default) requirements expression.
 - Will revisit later today / tomorrow.

- Recent Improvements / fixes
 - Eliminated the copy back of the condor.stdout and condor.sterr files.
 - Solved issue where condor copies back files... `transfer_output_files = ""`
 - The condor submission language has some subtleties...
 - Passing an empty string like `variable = ""` results in the variable not being passed at all
 - Passing as `variable = ""` results in the variable being set to ""
 - Passing as `variable = ""` results in the desired behavior...
 - Using more selective query when testing the production status DB for existing jobs that satisfy a given submission request
 - Results in significant speedup of each iteration of the submission loop
 - Better logging... kaedama now creates a dated log file for each rule that it submits, under a directory structure...

```
[phnxsub02] $ ls /tmp/kaedama/kaedama/*/2024-08-12.log -1
/tmp/kaedama/kaedama/BEAM_DST_STREAMING_EVENT_run2pp/2024-08-12.log
/tmp/kaedama/kaedama/BEAM_DST_STREAMING_EVENT_run2ppfast/2024-08-12.log
...
/tmp/kaedama/kaedama/PHYS_DST_TRIGGERED_EVENT_run2pp/2024-08-12.log
/tmp/kaedama/kaedama/PHYS_MBD_CALIBRATION/2024-08-12.log
```

- Jet production/QA
 - Link to macros:
 - [“Jet production” macro](#)
 - [Common Jet QA methods](#)
 - Action item: base macros on filelists, and build configuration, rule, submit, ...

- Discussion on how to ingest the MBD calibrations
 - meta-database needs to be implemented, api for querying which runs have which calibrations, ...

- Status of open action items
 - ☐ Vernier scan run 48029, crashed after 7.8M out of ~22M events
 - ☐ Multiple problems... main one is GL1 was not stable. Gaps where 20k GL1 events are missing. Essentially the run was not good. New VS run to be done.
 - ☐ Ana430 reprocessing older runs (43-45k run range) which failed due to MBD, maybe re-try with MBD calibrations in place?
 - ☐ Restart failed runs
 - ☐ *TRKR_SEED has no build_dbtag directory*
 - ☒ ~~Stop streaming fast production: output of full/fast production identical in size – the same?~~

- Discussion
- AOB

• Meeting Mon Aug 05 2024 1 PM EDT [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261

- Parallel event building [Joe]
 - Pushed some code that allows us to run the streaming event combining and trkr hit set unpacker jobs in parallel, on a PRDF by PRDF basis. Dramatically improves speed because we can process O(40) files in parallel
 - Problem - DSTs can not be hadd-ed together. So they need to be combined a different way (e.g. to not break downstream workflows for users)
 - Discussed developing a new workflow in the production that takes advantage of this, and then switching from old to new so that we don't break any existing functionality
 - End goal move people away from the DST_STREAMING_EVENT files... and work from the hits. Or from single merged file... or ...
 - Naming convention:
 - DST_STREAMING_EVENT_run2pp → DST_STREAMING_MVTX5_run2pp etc...
 - Should be able to stagein these inputs

- Production 2024p007 [Jason]

```
-- Schedd: phnxsub02.rcf.bnl.gov : <130.199.6.55:9618?... @ 08/05/24 10:36:16
OWNER  BATCH_NAME                                SUBMITTED   DONE    RUN    IDLE   HOLD   TOTAL JOB_IDS
...
sphnxpro DST_TRIGGERED_EVENT_run2pp_ana430_2024p007  7/31  21:08    1596    654     _    24    2274  29979.3 ... 30289.373
sphnxpro DST_CALO_run2pp_ana430_2024p007           8/2    00:44    3128     _     2    3130  30400.1230 ... 30473.322
sphnxpro DST_CALOFITTING_run2pp_ana430_2024p007    8/2    02:41    2693     _     2    2695  30431.779 ... 30434.91
sphnxpro DST_TRKR_HIT_run2pp_new_2024p007          8/2    13:55   11956     _    19   11975  30584.21 ... 31053.30
sphnxpro DST_TRKR_CLUSTER_run2pp_new_2024p007      8/2    15:59    188     _     9     197  30612.43 ... 31066.34
sphnxpro DST_TRKR_TRACKS_run2pp_new_2024p007       8/3    13:41    119     _    19     138  31051.2 ... 31069.15
sphnxpro DST_TRKR_SEED_run2pp_new_2024p007         8/3    14:14     82     2     4     88  31062.10 ... 31913.16
...
```


- Reprocessing all runs in the ‘physics’ era for triggered readout with fixed build ana.430
 - also running new... single pass calo reconstruction
- Online processing of the streaming readout switched to 2024p007 from run 49703 with new
- Fixed build production of the triggered detectors is continuing
 - First test of the 2 step fitting, calibrated reconstruction workflow
 - The online production has not been switched yet
 - At the moment we are running 654 event builders
 - Calorimeter fitting has 2 jobs on hold...
 - Production was paused on Friday to switch to stage in of job inputs... not unpaused
 - So... starting to recover. Submitting fitting jobs by hand, will relaunch the full workflow after the meeting.
- Triggered production thus far did not overstress the farm. Observed O(10k) concurrent jobs from this production within few hours after submission.
- New triggered production
 - Switch it off. Calorimeters are sufficiently stable.

- Summary of job failures over the weekend, for details see email to the production list [Sasha]
Two recent groups of failed runs:
 - 1) “False negatives”, both TRKR and CALO jobs, not many, all from Saturday.
 - 2) Old calo reco runs in the 43k-44k run range reprocessed with ana430_2024p007 tags. All segfault at event 1001.

- Slurp Updates [Jason]
 - Downstream jobs are now staging in their inputs
 - Resolution of the copy back issue (?)
 - condor copy back can be disabled by passing an empty string for the output file list
 - issue was in slurp where we check that condor parameters are specified in the yaml file before we pass them to condor
 - we were testing that each parameter evaluated to “True” rather than testing that the parameters was not “None”
 - thus far I don’t see any root files, or temporary job files, being copied back. TBD.

- Discussion / AOB [All]
 - Double slashes in file names... in filecatalog? In log files? Minor fix.

• Meeting Mon July 22 2024 1 PM EDT [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261

- Issue with large log files in event builder productions
 - daq error issued on every hit
 - messages should now be gone
- Additional issue with dat02 stability
 - SDCC global limit on inodes in gpfs / fixed (increased # of inodes...
 - Need to resubmit impacted jobs...
 - Need to identify run range and resubmit
- TODO: drop the 1GB filter on log size
- TODO: odbc.ini file ... should not be local
- ... Feature updates
 - Tracking size of the log files in the production_status table column = logsize
 - Not working yet... module pyodbc cannot be loaded this one time that I run cups...
 - Adding a flag when size exceeds 10MB
 - Fixed issue with the ‘held job’ monitor. Now records jobs going on hold that never reach the ‘start’ state, e.g. when stdout.log cannot be written to.
- Issue with new periodic hold ,
- We have jobs triggering the periodic hold condition after few hours...

| id | dstfile | status | flags | start | end | message |
|---------|---|--------|-------|---------------------|---------------------|--|
| 3424414 | DST_MVTX_EVENT_run2pp_new_2024p002-00048864-00000 | failed | 5 | 2024-07-21 02:25:42 | 2024-07-21 07:23:48 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3424413 | DST_MVTX_EVENT_run2pp_new_2024p002-00048862-00000 | failed | 5 | 2024-07-21 02:25:22 | 2024-07-21 03:13:06 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3424412 | DST_MVTX_EVENT_run2pp_new_2024p002-00048861-00000 | failed | 5 | 2024-07-21 02:25:09 | 2024-07-21 02:43:04 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3424411 | DST_MVTX_EVENT_run2pp_new_2024p002-00048859-00000 | failed | 5 | 2024-07-21 02:25:09 | 2024-07-21 02:58:04 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3411070 | DST_MVTX_EVENT_run2pp_new_2024p002-00048826-00000 | failed | 5 | 2024-07-20 15:01:25 | 2024-07-20 23:19:43 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3411069 | DST_MVTX_EVENT_run2pp_new_2024p002-00048825-00000 | failed | 5 | 2024-07-20 15:01:23 | 2024-07-20 19:23:13 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3405780 | DST_MVTX_EVENT_run2pp_new_2024p002-00048824-00000 | failed | 5 | 2024-07-20 11:54:10 | 2024-07-20 12:07:54 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404663 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048743-00447 | failed | 5 | 2024-07-20 10:32:09 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404662 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048743-00344 | failed | 5 | 2024-07-20 10:32:09 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404528 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02835 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404527 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02792 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404526 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02761 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404525 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02862 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404519 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02859 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404518 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02808 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404509 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02800 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404508 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02811 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404504 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02836 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404502 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02777 | failed | 5 | 2024-07-20 10:32:03 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |
| 3404479 | DST_TRKR_CLUSTER_run2pp_new_2024p005-00048660-02797 | failed | 5 | 2024-07-20 10:32:02 | 2024-07-20 10:35:49 | "The job attribute PeriodicHold expression (NumJobStarts >= 1 && JobStatus == 1) evalu |

- - 7 days ... need to stop wasting the time
 - Eviction → false restart breaks the 7 day limit
 - Need a cleanup script to hold these jobs
 -
- Possible Lustre Issue? DaqDB issue? Transfer script?
 - During submission loop I noticed part of an error message from kaedama... which indicates a mismatch between the list of files on lustre and the list provided by the daqdb query.


```
Trying rule ... PHYS_DST_STREAMING_EVENT_run2pp 2024-07-22 10:45:48
Trying rule ... PHYS_DST_TPC_EVENT_run2pp 2024-07-22 10:50:28
Trying rule ... PHYS_DST_TPOT_EVENT_run2pp 2024-07-22 10:51:59
('GL1_physics_gl1daq-00048990-0000.evt', '/sphenix/lustre01/sphnxpro/physics/GL1/physics//GL1_physics_gl1daq-00048990-0000.evt')
('TPOT_ebdc39_physics-00048990-0000.evt', None)
('GL1_physics_gl1daq-00048991-0000.evt', '/sphenix/lustre01/sphnxpro/physics/GL1/physics//GL1_physics_gl1daq-00048991-0000.evt')
('GL1_physics_gl1daq-00048991-0001.evt', '/sphenix/lustre01/sphnxpro/physics/GL1/physics//GL1_physics_gl1daq-00048991-0001.evt')
('GL1_physics_gl1daq-00048991-0002.evt', '/sphenix/lustre01/sphnxpro/physics/GL1/physics//GL1_physics_gl1daq-00048991-0002.evt')
('TPOT_ebdc39_physics-00048991-0000.evt', None)
Trying rule ... PHYS_DST_INTT_EVENT_run2pp 2024-07-22 10:53:47
Trying rule ... PHYS_DST_MVTX_EVENT_run2pp 2024-07-22 10:55:32
Trying rule ... DST_TRKR_HIT_SET_physics 2024-07-22 10:57:55
Trying rule ... DST_TRKR_CLUSTER_SET_run2pp 2024-07-22 11:23:11
Trying rule ... DST_TRKR_SEED_SET_run2pp 2024-07-22 11:40:06
Trying rule ... DST_TRKR_TRACKS_SET_run2pp 2024-07-22 11:56:13
```

- When I submitted these jobs by hand... 1.5h later... the job submission proceeded w/out an issue.
-

- Mickey says the MBD calibration output is good and we should start the calibration workflow running. (Attached from run # **48900+** ... will run based on earlier as *needed/requested*)
 - Expect this to be running today
 - ROOT files containing MBD calibration trees. MBD info... duplicate calorimeter stuff... etc... May be able to trim it down.
- Preparation for TPC tests later this week - need everything to be on autopilot for a variety of workflows
 - Streaming builder (full and subsystem) are on autopilot
 - Cosmics should be on autopilot as well (aim for now-ish)
 - Test runs today in various daq configurations.
 - On friday... want fast turnaround time
 - Requires an additional production loop
 - 5-10 min runs / time to process full run? TBD.
 - Will test full workflow beforehand...
- Discussion
- AOB

●

- Meeting Mon July 15 2024 1 PM EDT [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261

- Production / Development summary
 - Reprocessed runs 44300-46999

```
$ psql -d ProductionStatusTest -h sphnxproddbreplica.sdcc.bnl.gov -U argouser --command "select
dstname,count(run) count,sum( case when status='failed' then 1 else 0 end ) as failed,su
m(case when flags=5 then 1 else 0 end) as evicted,sum( case when status='running' then 1 else 0 end ) as
running from production_status where run>=44300 and run<=46999 and id>2500000 group by dstname;"
```

| dstname | count | failed | evicted | running |
|---|-------|--------|---------|---------|
| DST_CALO_run2pp_new_2024p004 | 58152 | 9098 | 7619 | 0 |
| DST_MBD_CALIBRATION_run2pp_new_2024p002 | 5 | 0 | 0 | 0 |
| DST_MBD_CALIBRATION_run2pp_new_2024p005 | 5 | 0 | 0 | 0 |
| DST_MVTX_EVENT_run2pp_new_2024p002 | 375 | 256 | 256 | 2 |
| DST_STREAMING_EVENT_run2ppfast_new_2024p002 | 881 | 124 | 121 | 0 |
| DST_STREAMING_EVENT_run2pp_new_2024p002 | 881 | 512 | 503 | 2 |
| DST_TPC_EVENT_run2pp_new_2024p002 | 230 | 183 | 182 | 2 |
| DST_TPOT_EVENT_run2pp_new_2024p002 | 660 | 129 | 129 | 0 |
| DST_TRIGGERED_EVENT_run2pp_new_2024p003 | 666 | 511 | 445 | 2 |
| DST_TRKR_CLUSTER_run2ppfast_new_2024p004 | 27435 | 31 | 0 | 0 |
| DST_TRKR_HIT_run2ppfast_new_2024p004 | 27448 | 1 | 0 | 0 |
| DST_TRKR_SEED_run2ppfast_new_2024p004 | 27380 | 3243 | 16 | 0 |
| DST_TRKR_TRACKS_run2ppfast_new_2024p004 | 26819 | 2725 | 0 | 0 |

- Reproducing without introducing a new build or cdb tag is disruptive to analyzers with our current procedures.
 - reproduction... clone the cdb tag
- Switched to new_2024p005
 - INTT hot map / calorimeter hot towers / micky’s beam beam calibrations
 - Starting with run 48099
 - Directory structure for output and log files now includes the build and cdb tags
 - *in principle can include anything in the SPhenixJob class / arguments to the production script...*
 - Histogram directories not changed so as not to disrupt QA... can make this change now if desired
- Micky provided an updated MBD calibration workflow... tested over the weekend (see the table above...). May be ready to place into production.
- Data set status table

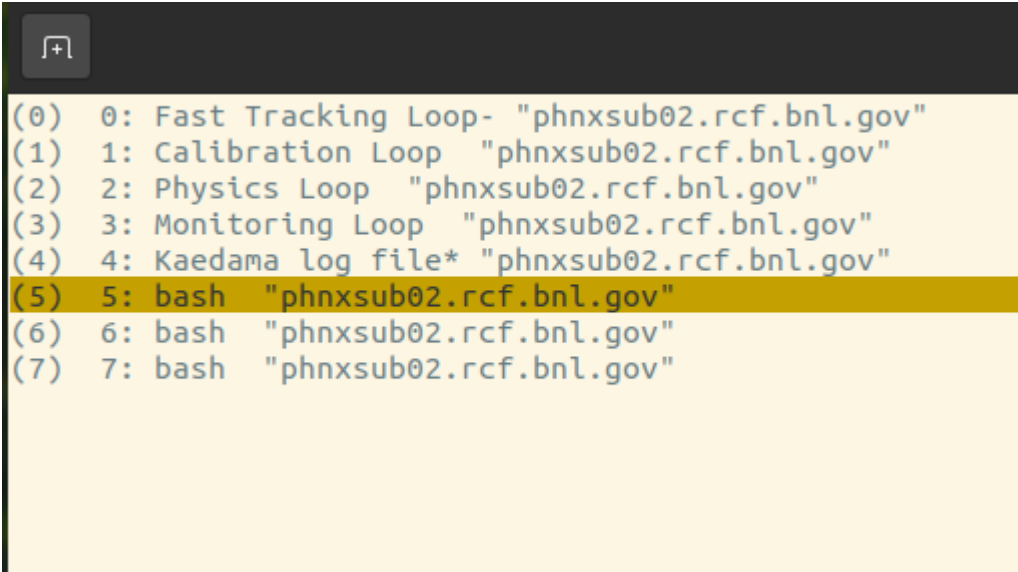
| dstname | run | nsegments | created | updated | finalized | status | blame |
|---|-------|-----------|---------------------|---------------------|---------------------|-----------|-------|
| DST_STREAMING_EVENT_run2ppfast_new_2024p002 | 48359 | 50 | 2024-07-15 13:15:34 | 2024-07-15 13:26:16 | 2024-07-15 13:26:21 | finalized | cups |
| DST_STREAMING_EVENT_run2ppfast_new_2024p002 | 48358 | 8 | 2024-07-15 12:56:17 | 2024-07-15 13:12:19 | | finalized | cups |
| DST_MVTX_EVENT_run2pp_new_2024p002 | 48356 | 24 | 2024-07-15 12:17:07 | 2024-07-15 13:37:20 | | finalized | cups |
| DST_STREAMING_EVENT_run2ppfast_new_2024p002 | 48357 | 50 | 2024-07-15 12:17:07 | 2024-07-15 12:26:11 | 2024-07-15 12:26:13 | finalized | cups |
| DST_INTT_EVENT_run2pp_new_2024p002 | 48357 | 273 | 2024-07-15 12:14:27 | 2024-07-15 13:37:24 | | finalized | cups |
| DST_INTT_EVENT_run2pp_new_2024p002 | 48356 | 245 | 2024-07-15 12:14:27 | 2024-07-15 13:37:34 | | finalized | cups |
| DST_TPOT_EVENT_run2pp_new_2024p002 | 48356 | 1057 | 2024-07-15 12:13:27 | 2024-07-15 13:37:47 | | finalized | cups |
| DST_TPOT_EVENT_run2pp_new_2024p002 | 48357 | 1179 | 2024-07-15 12:13:27 | 2024-07-15 13:37:47 | | finalized | cups |

- Tracks when a new dataset is created, updated (file added), finalized, and total # of files added.
 - Records the ‘blame’ of who updates it.
 - Working for tracking workflow... still debugging for calorimeter.
- Introduce new periodic hold expression...
 - Old: PeriodicHold = (NumJobStarts >= 1 && JobStatus == 1)

- New: PeriodicHold = ((NumJobStarts >= 1 && JobStatus == 1) || (NumJobStarts >= 2 && JobStatus == 2))
 - ... to stop jobs from getting restarted
 - Introduced after run 48369 / cluster 24751
 - Readiness for first fixed-build production?
 - Analysis ongoing on 2mo old data
 - Recent ana build should be acceptable
 - Aim for streaming readout production... repass of event building
 - **ana.427 2024p005**
 - First “physics” run
 - Streaming only (calors to be done with new ana build)
 - Recent job failures due to failures to transfer files
 - Jobs go on hold b/c unable to transfer files
 - Condor is still trying to copy back some files (eg input.list)
 - Discussion
 - AOB
- **Meeting Mon July 8 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Joe fixed an issue with a corrupted run number for INTT. ~300 failed INTT_EVENT jobs resubmitted on Friday around noon, going back to 06/21, run range 44300-46914. 22170.* on condor. As of Monday morning, half of the jobs are still running, but those that finished had finished successfully.
 - What to do with duplicate entries?
 - TODO... clean all jobs in this range and resubmit (streaming, tracking...)
 - Macro updated for calorimeter reconstruction on Saturday around noon, including ZDC now.
 - Macros updated for tracking reconstruction to expand tracking QA.
 - Tracking chain is broken (TRKR_HIT production fails) because of MVTX QA using ABORTEVENT in case of nodes missing, so no output is produced since approximately run 46200. This is a general behavior of all (?) QA modules. PR made to fix MVTX QA.
 - Jobs can run for 7 days now.
 - More jobs are restarted instead of being put on hold after eviction?

TODO: Update pathname convention: /path/to/directory/BUILD_TAG/RUN_RANGE
 - **Meeting Mon July 1 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Joe fixed MVTX memory leak on Thursday, ~300 jobs re-submitted for both MVTX_EVENT and STREAMING_EVENT.
 - MVTX jobs get slower with time from 3-4 min/seg to over 1hr /seg. ~80k events processed in 3 days and then evicted.
 - Very few MVTX fail again but hit 3 day limit, but many STREAMING fail again after several hours (mostly not status 5)
 - check recent mvtx jobs...
 - Jobs are still hitting the 3 day running limit, and some still keep running instead of being put on hold (several per day).
 - increase limit to 7 days
 - Reconstruction production moved to db tag 2024p004 on Friday.
 - For the calorimeters it is a slight change in the waveform template that seems to have little effect on the reco energy but reduces the chi2 for well behaving towers.
 - For the trackers there is a fix of the MVTX alignment.
 - Very approximately, this new db tag is applied to runs after ~46800.
 - Forgot to change db tag for streaming reconstruction input files, tracking production stopped after HITS step. Fixed today.
 - Condor “transfer slots” problem on Saturday, jobs idling for many hours, Jason fixed it by turning it off by setting *xferslots* to 0
 - 5 digits in segment name
 - new event combiner since Saturday packetwise clock alignment. Old problems gone. Still some work to be done. 47100+
 - re-process old failed event combining jobs with new code.
 - Deeper_production_information wiki page created: https://wiki.sphenix.bnl.gov/index.php?title=Deeper_production_information
 - **Meeting Mon Jun 17 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Production System Status
 - Streaming Detector Specific Event Builders
 - **MVTX : 592 / 595 jobs held** (go on hold after some successful processing)
 - INTT: 16 / 599 jobs held
 - TPOT: 3/ 658 jobs held
 - TPC: 6 / **230** jobs held...
 - There is mismatch between the # of TPC only workflows and the full streaming workflows. Not understood yet.
 - Disconnect between MVTX group and production
 - Decoder w/ only MVTX runs fast and not memory hungry, but... adding GL1 shows memory issues...
 - Their tests work... but not useful for production.
 - condor polling every 30 min for held jobs
 - job status updated to failed w/ flags=5 for jobs in the held state.
 - occasionally htcondor and/or db connection fails. code will skip the current iteration.
 - Jobs resubmitted past the three day limit... ⇒ need a solution ⇐
 - the condition for job going on periodic hold is not respected at a high frequency

- xferslots mechanism implemented for event builders
 - '2' setting, so expect one builder per node
 - ought to reduce the IO load on the nodes
- noticed a new issue in the production loop
 - receiving a timeout for the 'sh' module to execute the production script
 - raises an exception that ends the production loop
 - added exception handling. issue is reported on screen, and processing continues.
- Ran over the weekend (longer?) with `neventsper=1000` in the physics streaming workflow
 - This was introduced b/c event builder jobs were taking too long, jobs were failing before first output, & etc...
 - We are seeing 20k segments on this setting, too many (10k limit based on the naming convention)
 - Reverted to 10k / output file
 - Move to 2k
 - Focus on updating the naming convention
- System is running on phnxsub02 under a tmux session `sPHENIX-Integration`
 - Four workspaces are setup running the physics, beam/calib, fast physics and condor monitoring loops.
 - The kaedama logfile is being tailed on the 5th workspace



-
- What (if any) additional job metrics do we need to be keeping track of? e.g.
 - Event processing rate? IO (stageout) rate?
 - Hostname
- Noticed a class of jobs in the production system flagged as 'finished' but which staged no output files
 - An example log file...
/sphenix/data/data02/sphnxpro/streaminglogs/run_00046500_00046600/DST_STREAMING_EVENT_run2pp_new_2024p001-00046568-0000.out
 - Fun4All crashed... possibly on first event?
 - But root's exit status was zero... ???
 - "Do not panic" ... file did not have any events ...

psql -d ProductionStatusTest -h sphnxproddbreplica.sdcc.bnl.gov -U argouser --command "select dstfile,status,started at time zone 'utc' at time zone 'edt' as start,ended at time zone 'utc' at time zone 'edt' as end,cluster||'.'||process,flags as clusterprocess,message from production_status where dstfile like 'DST_TRKR_HIT%_run2pp_new_2024p002%' and flags=5 order by id desc;"

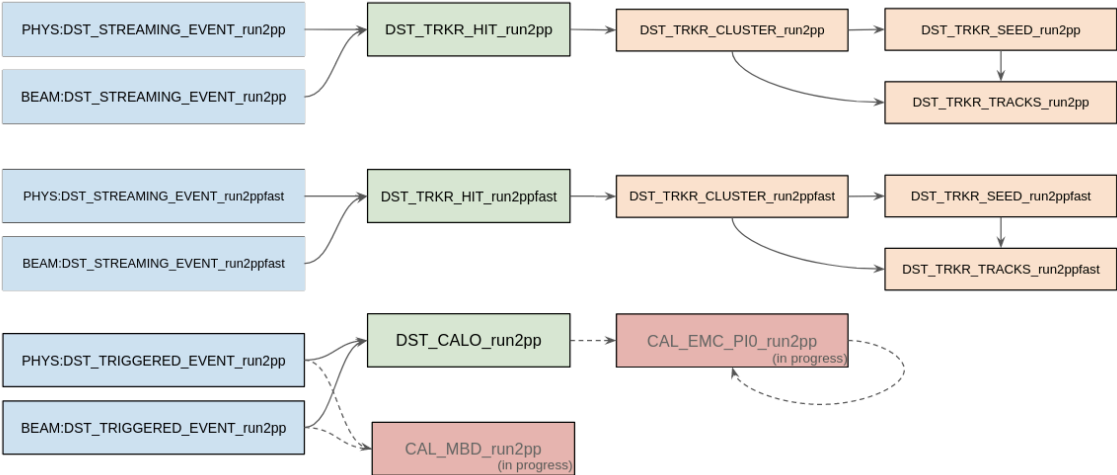
- Organizational Matters
 - Service taskers (shift replacement)
 - production deputies ... need to define roles, documentation, etc... date
 - Task list
 - Informs documentation and infrastructure needs
 - Aspirational
- Next week...
 - Sasha should organize
- Discussion
- AOB
 - First version of resync is working... starting up after this meeting
- **Meeting Mon Jun 17 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Calorimeter productions DB issue
 - SLURP provides both the analysis build and global cdb tags as arguments to the production run script, which are supposed to be forwarded to the production macro
 - When the Y2 calorimeter production macro was integrated, it was not modified to accept the global cbd tag as an argument. Thus, we ran with the ProdA_2024 tag instead of the 2024p001 tag as intended.
 - Macro was updated on Friday and jobs resubmitted going back ~2 weeks, starting with run 44277. Run against global cdb tag of 2024p002. (Event builder remains at p001).
 - Approx 75k jobs were submitted.
 - When I looked at the jobs, I saw up to 6500 slots in use concurrently (presumably farm occupied with simulations and/or user jobs)
 - Backlog of jobs was finished Sunday AM. No jobs failed or held.
 - We can submit older jobs if needed
 - DST_CALO_run2pp_new_2024p002 is stable
 - 80426 have been submitted (80k+ finished)
 - 250 jobs currently running

- File transfer issues
 - Inform the shift crew
- O(10) jobs / day go on hold b/c of memory excursions... likely beam clock related
 - Tracking ... everything other than fast tracking... memory limits reached... possibly map not being cleared.
 - Fix pushed for TPOT... but others may still be needed.
- Reorganization? / Transition to Formal Production Scheme
 - Establish a clear line of communication between “customers” and production team
 - Communication should go through the production manager
 - Reduce number of channels (random MM channels → Commissioning_Data_Production)
 - Role of the production mailing list?
 - Low traffic... announcements of new productions
 - How to handle calibrations/builds?
 - Documentation for future deputy production managers?
 - Documentation updates for users in the wiki?
 - Where to find logfiles
 - 2024 productions
 - Locate DSTs
 - Do we need deputies at the current time (and what would they do)?
 - Joe and Blair are effectively already in this slot / would be useful to get people appointed to offload that role
 - Need eyes on log files / build up the expertise to keep an eye on event building, understand what is going on when jobs go on hold...
 -
- Operational plans moving forward
 - Transition to production db (from production status test)
 - Integration of workflows
 - Scope of the automatic production loop
 - One macro to rule them all
 - Unify the macro API
 - DB tag a function argument w/ sensible defaults
- Calorimeter Calibration Discussion
 - Remember that slurp operates in a fire-and-forget mode...
 - (Set of inputs) → Production Rule → (Set of *condor jobs*)
 - Once the production system submits the jobs, there is no additional following.
 - Up to the jobs themselves to communicate back their status and fill the file catalog
 - We cannot determine, based on examination of the file catalog, whether a given run has been completely processed. (Event builder may still be running and creating segments).
 - We could add an arbitrary text flag to the datasets table, which would enable the event builder job to indicate when the run is completed. (Event builder would update all entries as “done” ... or ... “finished nsegment/ntotal” or ...
 - Calorimeter production seems to be fairly stable.
- Increase fast production N total and / or N per job?
- HCAL cosmics... opportunistic 1 day / week.
-
- **Meeting Mon Jun 10 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Need to deal with held jobs (most exceeding 3 days, first held jobs was submitted 5/18)
 - and jobs which fail the periodic hold and get resubmitted forever wasting resources:


```
condor_q | grep -v ' H ' | head -30 | awk '{print $1" "$5}'
```

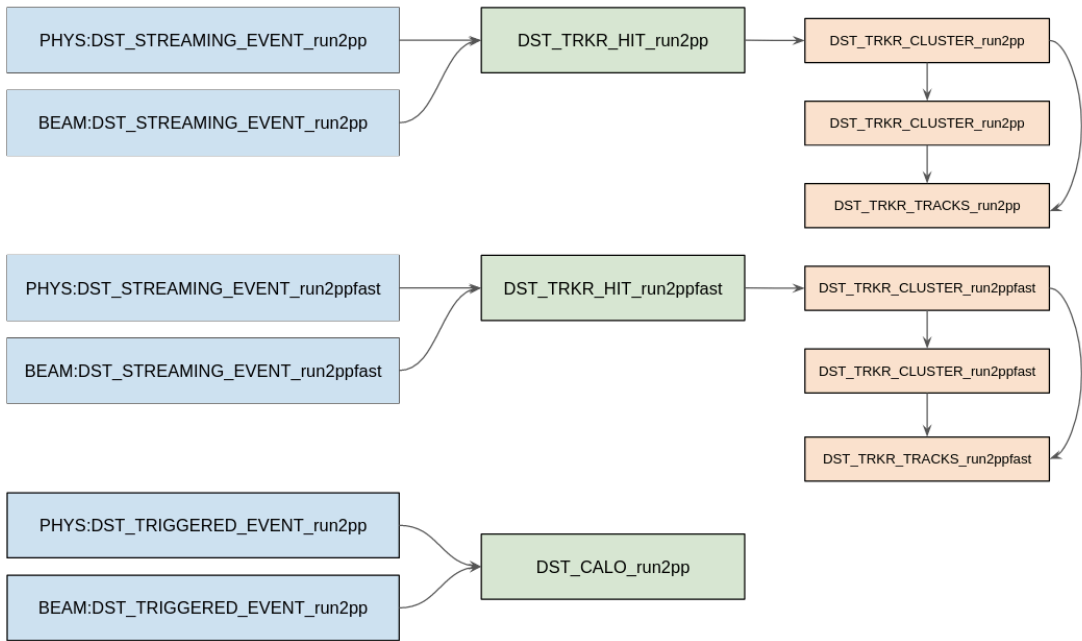
| ID | ST |
|--------|-------------|
| 2427.0 | 20+11:53:26 |
| 5541.0 | 13+14:12:51 |
| 5584.1 | 13+11:06:20 |
| 5999.0 | 12+07:50:11 |
| 6813.1 | 10+01:08:40 |
| 7360.0 | 8+11:17:12 |
| 7857.0 | 7+08:17:18 |
| 7926.0 | 7+03:44:55 |
| 7926.1 | 7+03:44:50 |
| 8003.0 | 6+23:57:01 |
| 8003.2 | 6+23:57:07 |
| 8423.0 | 5+19:34:29 |
| 9398.0 | 3+05:17:22 |
| 9398.1 | 3+05:17:22 |
| 9427.0 | 3+04:09:23 |
| 9430.0 | 3+03:53:41 |
- What is going on with the calorimeter productions? (14k jobs at this point... 130PM on Saturday...)
 - (note ... paused submission of calo y2 at 1247 on 6/10)
 - Possible calibration issue? Possible issue with (non) zero suppressed data?
 - Chris inspected the log file... Every event MBD printout
- Need to deal with release management... Currently running everything against NEW. Exposes us to potential for unvetted code updates. Weekly build? QA of “new” release? ... To be discussed in software meeting.
- TPC Calibration Runs

- 45131, 45134, 45137 Question from Friday Night
 - Last two ran through w/out issue. First did not start b/c daqdb indicates that not all files have transferred
- Suspect jobs 44701 44702 44703 44715 44718
 - Report that only 2 of 10M events were produced
 -
- Reproduction Plan
 - Scope... everything since physics declared to be reprocessed. (Some data before physics may be useful to look at, but not top priority).
 - Time to migrate to the ProductionStatus, rather than the ProductionStatusTest, DB
 - Will submit against a new DB tag, so that we have an independent namespace
 - Will update to the development version supporting the run-range naming convention ABC_XXX_yyy_<run1>-<run2>-<seg>-<iteration>
 - Once the system is running against the official DB, few policies will be put in place
 - Ensure traceability of changes
 - No changes within the running directory *other than* the files which specify which workflows are running at any one time
 - Use pull requests to update production codes (e.g. macros, runscripts, yaml configuration files)
 - Changes to the production system itself should be vetted separately, then applied via PR
 - Library release policy... Use an ana build. (Fast should stay new. (??))
 - Timescale? End of week.
 - Two stage shift to new system?
 - Dependencies / show stoppers?
 - Waveform processing
 - Old files? Wipe them out. All of them.
 - Staged approach... may be running duplicate event builders for O(1wk).
 -
- AOB
- Meeting Mon Jun 03 2024 1 PM EDT [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261
 - [New hit format...](#) from run 44477.
 - Paused production loops this morning. Will resume after appropriate build... 1508 parent and child jobs to (potentially) cleanup. Unpause and continue running...
 - Clock check errors since the beginning. Event combining triggered and scalers... As of now GL1 off by 1. Will impact anyone trying to combine triggered and streaming data. So at some point will need to reprocess. Current fix should take care of the additional GL1 event (but TBD (and hardcoded offset from Friday evening to fix... but the GL1 may have unfixed...))).
 - Plan will be to reprocess event combining and downstream processing (anyway)...
 - Need mechanism in place to parse log files.
 - Calorimeters ... no zero suppressed towers (status flag is not correctly set?)
 - Another event combining change coming along...
 - New reproduction on time scale of next week... against new DB tag 2024p???
 - useful point to switch to the ProductionStatus instance from ProductionStatusTest
 - Data Production Status
 - Downstream data production is acceptably robust. “Few” held jobs over O(100k)
 - Up to few hundred held DST_TRIGGERED and DST_STREAMING event builders. Out of O(100k)... Reason(s) TBD.
 - Production System Status and Work Plan
 - Tracking and calorimeter workflows are implemented.
 - Calibration workflows are next priority
 - MBD workflow is ready for review (Mickey pinged, will look into code this evening... noted some errors in log file, needs access to intermediate results, etc...)
 - Calorimeter is nearing readiness...
 - [Stage 1 histograms...](#)
 - Merging / fitting O(2d)

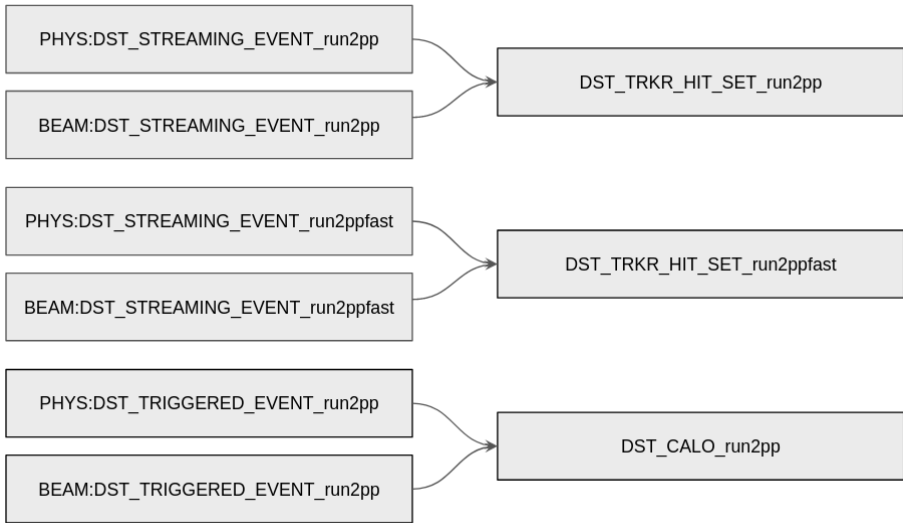


- Fast tracking... means two things (1) everything but TPC, and (2) small number of events with everything.

- Streaming Resubmission for runs [41620,41626] and [41967,41992]
 - 1 failed job. The rest appear to have ended normally (all jobs went on hold b/c the “BEAM” yaml file still instructed condor to copy back the log files). Downstream hit, seed, cluster and tracking jobs appear to have finished.
- Discussion
- AOB
- **No Meeting Mon May 26 2024 (Memorial Day Weekend)**
 - Currently running the following workflows... streaming up through full track reconstruction. Calorimeters through tower energy extraction.



- **Meeting Mon May 20 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261
 - Ran into an issue with running zero-suppressed TPC data through the production system
 - 100 events / output file set for the streaming event pipeline to support fast turnaround of TPC commissioning analysis
 - Okay for non-zero suppressed data, but zero suppressed data a bit like drinking from the fire hose...
 -
 - We can handle submitting O(1k) jobs concurrently in batch mode without an issue. O(100k) was (at a minimum) taking longer to submit than the delay in the submission loop.
 - Interactive submission will be needed, and we (may) need to clear the backlog of jobs prior to Friday.
 - Presumably we can run the backlog with more reasonable # of events / job
 - Weekend running
 - Enabled streaming and triggered workflows.



- Fast Production Pipeline: DST_STREAMING_EVENT_run2ppfast
 - Loop set on a 5 min delay between iterations
 - 1k event building for each physics, beam run
 - 100 events/per output file
 - Setup Friday 5PM after last minute requirements discussion with Jin and Chris
 - Checked in on Saturday and first wave of run2ppfast jobs had started but not yet finished.
 - Looked like some of the changes I had implemented in, e.g., scripts and macros had been reverted?
 - Fixed the issue and things worked. Sunday AM killed and resubmitted the initial 200 fast jobs which were still running.
 - Looking at mean time to complete the hit set... jobs run fast (<1min)... with current farm loading it takes a few minutes to start the jobs some times.


- | dstname | run | average_time_from_subnit |
|--------------------------------------|-------|--------------------------|
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43130 | 00:04:16.5 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43129 | 00:03:58.4 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43128 | 00:00:41.8 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43127 | 00:08:45.2 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43126 | 00:03:25.8 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43125 | 00:04:01.1 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43124 | 00:04:44.2 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43123 | 00:00:34.4 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 43122 | 00:03:24.1 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42953 | 00:00:43.5 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42952 | 00:00:52.3 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42951 | 00:04:40.6 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42944 | 00:04:19.4 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42943 | 00:03:54.9 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42942 | 00:00:44 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42941 | 00:00:43 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42940 | 00:06:14.6 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42939 | 00:06:02.4 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42938 | 00:06:04 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42937 | 00:00:45.5 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42936 | 00:00:45.3 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42921 | 00:00:43.3 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42910 | 00:01:07.7 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42909 | 00:01:07 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42906 | 00:00:35.7 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42904 | 00:00:31.7 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42903 | 00:00:31.4 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42891 | 00:00:57.1 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42874 | 00:00:59.8 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42872 | 00:00:54.4 |
| DST_TRKR_HIT_run2ppfast_new_2024p001 | 42832 | 00:00:51.7 |

- Normal Production Pipeline
 - beam and physics modes both feed into the same DST production pipelines
 - names of the files now follow a consistent naming convention regardless of data taking mode (beam vs physics).

- Friday processing calorimeters w/out gl1 and/or mbds...
 - all put oh hold (reason TBD) few diff runs.
 - finished / shadow exception copying outputs
 - some jobs did stage outputs
- Misc / Old / New Business
 - Run 10k events for trigger combining
 - What’s going on with the GL1 packet?
-
- QA jobs... histogram files. Extension is botched.
-
- AOB
- Discussion
- Meeting Mon May 13 2024 1 PM EDT [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261
 - Production Updates
 - Y2 Calorimeter Automated Production.
 - Initial startup 1130AM 05/09.
 - Observed issue with job duplication. 1135 duplicates inserted at 12:20 PM and 2:00 PM on 05/09.
 - Normal operation of the production loop until the duplicate jobs were submitted.
 - Dmitry doesn’t see any issue with DB replication during that time.
 - Number of inserts/updates < O(100k) are not a concern
 - Issue not observed again over 24h...
 - Possibly an issue with network saturation?
 - DB consistency check added to kaedama.
 - If the read instance of the DB shows the last unique entry ID less than the write instance, we log a warning and retry once/minute for 5 minutes. Then give up.
 - Ran O(10k) jobs for 24h...
 - Paused event loop when second issue discovered
 - File stageout failing b/c we matched the output file with a wildcard which was *a/so* picking up a histogram file produced by the macro. Resulted in the histogram file being overwritten instead of the DST copied to the destination.
 - Runs 41420 and 41568 [beam] Y2 macro.
 - reported running slow ... but running and creating output ...
 - Unusable for physics. (?) Exclude from future productions.
 - Problem may be at event combining stage... TBD.
 - Some runs before 40800 should be excluded as well
 - More investigation at event combining stage TBD.
 - DST issues
 - Missing GL1
 - Data will be reprocessed
 - Plan will be to get naming conventions correct across physics, beam, calib and cosmics.
 - ... streaming to follow ...
 - Global vertex map... mbd not being filled.
 - Condor Delays
 - Chris reported that condor stopped running new jobs over the weekend.
 - Typical delay between submitted and started state is O(1min), with current loading of the farm.
 - Looking at the submission delays over the weekend, I see
 - We have “normal” behavior until a point where the submission takes a significantly longer amount of time to begin running.
 - Below I tabulate the times, production types, delays and cluster/process # where the problems seemed to begin.
 - `select dstfile,status,submitted,started,started-submitted as DELAY,cluster,process,nsegments from production_status order by id desc;`

| dstfile | status | submitted | started | delay | cluster | process | nsegments |
|---|----------|---------------------|---------------------|----------|---------|---------|-----------|
| ... | | | | | | | |
| DST_BEAM_run2pp_new_2023p013-00042119-0000 | running | 2024-05-12 10:52:09 | 2024-05-13 07:02:49 | 20:10:40 | | 1257 | 0 1721 |
| ... | | | | | | | |
| DST_CALO_run2pp_new_2024p001-00042079-0003 | running | 2024-05-12 06:23:28 | 2024-05-12 10:42:49 | 04:19:21 | 1215 | 1 | 1 |
| ... | | | | | | | |
| DST_CALO_run2pp_new_2024p001-00041939-0018 | running | 2024-05-12 05:41:11 | 2024-05-12 06:16:09 | 00:34:58 | 1207 | 0 | 1 |
| ... | | | | | | | |
| DST_STREAMING_EVENT_run2pp_new_2024p001-00042083-0000 | finished | 2024-05-12 02:30:30 | 2024-05-12 05:12:55 | 02:42:25 | 1182 | 0 | 1 |
| ... | | | | | | | |
| DST_STREAMING_EVENT_run2pp_new_2024p001-00042074-0000 | running | 2024-05-12 00:12:55 | 2024-05-12 00:27:46 | 00:14:51 | 1165 | 0 | 19669 |
| ... | | | | | | | |
| DST_CALO_run2pp_new_2024p001-00041927-0015 | running | 2024-05-11 21:54:55 | 2024-05-12 00:02:51 | 02:07:56 | 1147 | 0 | 1 |
| ... | | | | | | | |
| DST_CALO_run2pp_new_2024p001-00041927-0012 | running | 2024-05-11 18:55:21 | 2024-05-11 19:31:00 | 00:35:39 | 1136 | 0 | 1 |
| ... | | | | | | | |
| DST_CALO_run2pp_new_2024p001-00041568-0120 | running | 2024-05-11 17:20:04 | 2024-05-11 18:16:01 | 00:55:57 | 1129 | 0 | 1 |

- I don't see any correlation with job type...
 - I note that we see this same issue this morning...
 - sdcc looking at short and long term solutions to this...
 - Currently no user jobs... so production job stalls seem odd.
 - sPHENIX in shared pool may not be appropriate
 - Streaming HIT macro... should be robust against missing inputs
 - Histogram output from Y2 macro...
 - Simple fix... but would prefer to not send to lustre. GPFS preferred.
 - Blair... macro just waveform fitting
 - Discussion
 - AOB
- Meeting Mon May 06 2024 1 PM EDT Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261
 - Production Updates
 - Post Mortem
 - Large log files filling up /var/spool...
 - /dev/mapper/sysvg-root 62G 24G 35G 41% /



- - Hardcoding truncation at 25MB for now
 - Other options...
 - Write back a small log file
 - Copy back full log directly with (gzip then) cp
 - Streaming Jobs held... (new copy back may not be working right)...
 - | | | | | | | | | | | | |
|-------|----------|-----|-------|------------|---|------|------|----------------|------|------------------------------|------------------------------|
| 592.0 | sphnxpro | 5/6 | 11:06 | 0+00:45:18 | H | 3800 | 4151 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 592.1 | sphnxpro | 5/6 | 11:06 | 0+00:20:19 | H | 3800 | 3907 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 592.2 | sphnxpro | 5/6 | 11:06 | 0+00:15:12 | H | 3800 | 4151 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 592.3 | sphnxpro | 5/6 | 11:06 | 0+00:20:16 | H | 3800 | 4151 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 592.4 | sphnxpro | 5/6 | 11:06 | 0+01:03:30 | R | 3800 | 3907 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 592.5 | sphnxpro | 5/6 | 11:06 | 0+00:45:20 | H | 3800 | 4151 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 592.6 | sphnxpro | 5/6 | 11:06 | 0+00:30:13 | H | 3800 | 4151 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.0 | sphnxpro | 5/6 | 11:06 | 0+01:03:30 | R | 3800 | 3663 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.1 | sphnxpro | 5/6 | 11:06 | 0+00:25:17 | H | 3800 | 3907 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.2 | sphnxpro | 5/6 | 11:06 | 0+01:03:30 | R | 3800 | 3907 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.3 | sphnxpro | 5/6 | 11:06 | 0+01:03:30 | R | 3800 | 3907 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.4 | sphnxpro | 5/6 | 11:06 | 0+01:03:30 | R | 3800 | 3418 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.5 | sphnxpro | 5/6 | 11:06 | 0+01:03:29 | R | 3800 | 3663 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.6 | sphnxpro | 5/6 | 11:06 | 0+01:03:30 | R | 3800 | 3418 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |
| 593.7 | sphnxpro | 5/6 | 11:06 | 0+01:03:29 | R | 3800 | 3663 | run_cosmics.sh | 1000 | DST_BEAM_run2pp_new_2023p013 | DST_BEAM_run2pp_new_2023p013 |

 - MVTX issue... May 1st data is corrupted... out of memory is likely
 - QA/QC
 - Fixed as of yesterday
 - Sympton? Jobs on hold (memory issue?) 0 run number (we do have a few of those in the logs)...
 - HARmonize our naming convention(s) TRIGGERED_RAW_run2pp, ... etc...
DST_BEAM→DST_STREAMING_RAW_run2pp, ...
 -
 - New workflows
 - MBD_CALIBRATION workflow test implementation submitted and it ran... will check log file to make sure code ran without issues... then setup.
 - PHYSICS event builder workflows defined for streaming and triggered detectors. Running in production loop.
 - DST_CALO_Y2 rule defined and added, taking input from triggered physics. In the loop. Producing DST_CALO_run2pp-<run>-<segment>.root
 - Tracker hits... upvoted
 - Etc...
 - AOB
 - - MVTX from affected runs. (Useful for INTT TPC alignment)... Run Range?
- Meeting Mon Apr 29 2024 1 PM EDT Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261
 - Start of Collisions
 - Started Saturday
 - but gl1 missing due to firmware

- still want data to be reconstructed...
- macro and code can handle this... slight code modification incoming
- discussion calibration calorimeters about not running w/ gl1
- event combining change
 - epd will be moved to an appropriate spot on the node tree
 - waiting for it to go through jenkins
- reprocessing of beam data that didn't go through following above change(s)

- Production Update
 - Setup a “beam” loop (s/cosmics/beam)
 - Issue #1... firmware update failed / jobs not created b/c gl1 files missing.
 - Issue #2... runlists are built with logic that depend on the filename... e.g. stripping the input filename around “cosmics” in order to determine the filelist to which it belongs. This logic should be changed.
 - Jobs are running

Summary of jobs which have reached staus='started'

| dsttype | num_jobs | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | last_job_started | last_job_finished | sum_events |
|---------------------------|----------|-------------------|---------------|-------------|--------------|------------|------------------|---------------------|---------------------|------------|
| DST_TRIGGERED_RAW_cosmics | 4 | 00:01:38 | 0 | 0 | 4 | 0 | 00:03:25.25 | 2024-04-29 15:07:56 | 2024-04-29 15:13:08 | 20003 |
| DST_TRIGGERED_RAW_beam | 8 | 00:28:30.5 | 0 | 5 | 3 | 0 | 04:32:26.333333 | 2024-04-29 10:07:02 | 2024-04-29 13:13:28 | 8657920 |
| DST_INTT_RAW_cosmics | 10 | 00:01:39 | 0 | 8 | 2 | 0 | 00:00:17.916667 | 2024-04-29 15:07:56 | 2024-04-29 15:08:54 | 0 |
| DST_INTT_RAW_beam | 11 | 00:15:44.454545 | 0 | 7 | 4 | 0 | 00:03:56.25 | 2024-04-29 16:07:26 | 2024-04-29 16:14:04 | 5000001 |
| DST_COSMICS_run2pp | 11 | 00:01:41.363636 | 0 | 8 | 3 | 0 | 00:00:17.363636 | 2024-04-29 15:07:56 | 2024-04-29 15:08:31 | 0 |
| DST_BEAM_run2pp | 10 | 00:15:46 | 0 | 6 | 4 | 0 | 00:03:46 | 2024-04-29 16:07:26 | 2024-04-29 16:14:39 | 1895001 |

Next submit: 3%|

- Running two production loops, 10 min between iterations
 - cosmics: | DST_COSMICS | DST_TPCCALIB | DST_COSMICS_INTT_RAW | DST_TRIGGERED |
 - beam: | DST_BEAM | DST_BEAM_INTT_RAW | DST_BEAM_TRIGGERED |
 - Incoming: LED_DATA for calors
 -

- Discussion of Failed Jobs
 - ... new standing item on the agenda ...

- Calorimeter Calibration Workflows
 - Use production_status as the authoritative table for inputs to jobs
 - allows us to form a query for input to the merge step... such as
 - select files from all jobs which succeeded if success rate > 90%
 - Run selection... doable.
 - Tack on tower building onto the event combining.

- AOB
 - *Tracker macro incoming (hit making)*

- **Meeting Mon Apr 22 2024 1 PM EDT [Meeting Link...](#)** Meeting ID: 160 468 6747 Passcode: 155261

- Production Update
 - Run 1 AuAu Calorimeter Production DST_CALO_run1auau_ana412_2023p015
 - Resubmitted 4/19 1230PM
 - No condor kaboom
 - Minor DB hiccup (ramenya production loop encountered a pydobc exception which implied it could not connect to DB... restarted loop w/out an issue).
 - Slurp file now contains a “fixit” rule, which submits only jobs where the output is missing.
 - 4/22 3051/55425 files returned. TBD.
 - Production Loop running
 - DST_COSMICS
 - DST_INTT_RAW
 - DST_TPCCALIB
 - DST_TRIGGERED_RAW
 - DST_TRIGGERED_cosmics added to production loop
 - Initially jobs ending w/ nonzero exit status...
 - Chris noted that this appears to be in the cleanup phase of Fun4all. Output is written and staged out... last output file may have been lost. Last event failed. Fixed.
 - Jobs began finishing w/out error at from 2024-04-20 15:10:16 UTC
 - Condor on phnxsub02 went down yesterday ~12 PM
 - Runs 40214 to 40240 “submitted” but never started
 - (40208 DST_TRIGGERED_RAW was held).
 - MVTX filling up disk on submission host as it copies back the ginormous error / log files
 - Looking back at the size of the logfiles which made it back...
 - 38401 is the first run where we see this issue
 - 2TB largest error file copied back... TODO cleanup
 - Error protection TBD... too big? 2MB. Head, tail and raise a flag in production_status.
- Slurp Update
 - modifying code to handle run-range naming convention.
 - DST naming convention
 - (XXX)_(YYYY)_(zzzz)-(nnnnnnnn)-(mmmm).ext
 - n = runnumber (08d)
 - m = segment number (04d)
 - Run range naming convention
 - (XXX)_(YYYY)_(zzzz)-(iiiiiii)-(jjjjjjj)-(nnnn)-(mm).ext
 - i = first run number (08d)
 - j = last run number (08d)
 - n = segment number (04d)

- o m = iteration (02d, optional)
 - Slurp is updated. cups needs to be updated so that we have status updates back to the production status table, stageout.
 - setting up calorimeter calibration workflow... TBD
 - reaction plane ...
- o Discussion
- o AOB

- Meeting Mon Apr 15 2024 1 PM EDT Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261
 - o Production Update
 - Production loop running on phnxsub04 from 4/12 at 1610
 - Jobs which have started since Friday (12AM UTC)

| | | | | | | | | | |
|--|----------|-------------------|---------------|-------------|--------------|------------|------------------|---------------------|---------------------|
| Summary of jobs which have reached staus='started' | | | | | | | | | |
| dsttype | num_jobs | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | last_job_started | last_job_finished |
| DST_TPCCALIB_run2pp | 1 | 00:00:30 | 0 | 0 | 1 | 0 | 00:05:10 | 2024-04-12 03:02:46 | 2024-04-12 03:07:56 |
| DST_INTT_RAW_cosmics | 144 | 00:03:35.194444 | 0 | 24 | 82 | 38 | 00:25:05.616667 | 2024-04-15 13:50:31 | 2024-04-15 14:10:39 |
| DST_COSMICS_run2pp | 266 | 00:00:41.804511 | 0 | 27 | 166 | 73 | 14:55:47.054394 | 2024-04-15 14:02:40 | 2024-04-15 12:04:08 |
| DST_CALO_run1auau | 107830 | 00:58:20.393202 | 0 | 6539 | 101291 | 0 | | 2024-04-13 04:23:36 | |

- Failed jobs seem to be mostly (maybe completely) data corruption
- For the first time, I am seeing run mis match flagged for GL1 files and MVTX inputs.
- [list of bad runs since Friday](#)
- Production loop is now processing cosmics, fast cosmics (no tpc) and tpc calibration runs
 - The fast cosmic workflow (aka INTT_RAW) launches when any of (INTT, TPOT, MVTX) and a GL1 file have been transferred to SDCC. Once the TPC files are transferred, the full cosmic workflow is run.
 - Implementing the INTT_RAW required an update to the stageout code to handle the (general) naming convention.
 - o This temporarily broke processing.
 - o Impacted runs cleared and resubmitted
- Calorimeter run1auau submitted 4/11 1700
 - 54k jobs with build ana.410, dbtag 2023p14 ...
 - impacted by the INTT_RAW issue (no output staged)
 - impacted by a condor hiccup ???
 - o large number of jobs ran into a shadow exception and failed.
 - o large number of jobs failed b/c of stale code in the creation of the output (log) directories. (Should have impacted 1% of jobs, but larger number impacted.
 - TODO: Review the logic here and make sure that data02 is not being hammered with attempts to mkdir...
 - resubmitted from phnxsub03 on 4/12 1530 with fix to the directory creation (but not necessarily the optimal fix)
 - again saw issues with shadow exceptions... 50% of jobs were failing
 - resubmitted from phnxsub03 on 4/13 ~10 AM or so.
 - o failed jobs only. staged in run-by-run to limit the number of concurrent job submissions.
 - o observed much higher rate of successful jobs, O(95%) of expected outputs
 - o As of Saturday afternoon approx 244 jobs were still running.
 - o As of today 11:30 AM, 191 jobs still running...
 - o There is a long tail to complete these jobs...
 - o Chris noted an issue on spool1315

HOWTO get list of missing outputs...

```
with
inputs as (
  select distinct filename as input, runnumber, segment
  from
    datasets
  where
    filename like 'DST_EVENT_auau23_ana399%'
),
outputs as (
  select filename as output, runnumber, segment
  from
    datasets
  where
    filename like 'DST_CALO_run1auau_ana410_2023p014-%'
)

select * from inputs left outer join outputs on inputs.runnumber=outputs.runnumber and inputs.segment=outputs.segment where outputs.runnumber is null;
```

2868 missing outputs, in 33 runs

| | | | | | | | | |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| 20859 | 20864 | 20868 | 20869 | 20877 | 20888 | 21199 | 21200 | 21626 |
| 21774 | 21776 | 21783 | 21923 | 21928 | 21933 | 21936 | 21939 | 21950 |
| 22055 | 22911 | 23050 | 23537 | 23702 | 23721 | 23724 | 23725 | 23726 |
| 23727 | 23740 | 23743 | 23745 | 23761 | 24316 | | | |

- Q's
 - INTT raw using new. filename like 'DST_INTT_RAW_cosmics_new%' or directories... or ...
 - Many single segments missing, but a couple (few) complete runs
 - Easy way to figure out ... open the file and do a T->Print()...
 - o double check INTT_RAW requires INTT...
 - o
- Incoming... trigger readout added to cosmics... new macro.
 - Synchronization object...

- o Slurp Update
- o AOB / Discussion

- Production Update
 - Production loop running cosmics production for 1+ week... modulo shutting it down / restarting for a few minor updates
 - Operation of the system appears to be stable ... over 16h period no noticeable memory consumption. (Design of the loop is that there is no saved state between iterations of the loop).

```
top - 10:43:23 up 107 days, 12:58, 8 users, load average: 0.03, 0.04, 0.05
Tasks: 346 total, 1 running, 343 sleeping, 2 stopped, 0 zombie
%Cpu(s): 0.2 us, 0.1 sy, 0.0 ni, 99.8 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 10250659+total, 40358372 free, 21501760 used, 40646456 buff/cache
KiB Swap: 32767996 total, 32721916 free, 46080 used, 74188368 avail Mem

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  %MEM    TIME+  COMMAND
 3111 root        0 -20 7653312   3.6g 930348 S   0.7   3.7   1872:53 /usr/lpp/mmfs/bin/mmfsd
563253 sphnxpro 20  0 5115116 29300 10500 S   0.7   0.0   16:31:02 python ./ramenya2.py submit slurp-examples/sPHENIX/DST_STREAMING_2024.yaml --nevents 0 --rules-file RULES --timestart 04/01/2024 --runs 35000 39999 --loop --delay 300
942618 sphnxpro 20  0 158104 2532 1540 R   0.7   0.0   0:00:18 top
  1 root        20  0 52172 3716 2208 S   0.0   0.0   32:12:65 /usr/lib/systemd/systemd --switched-root --system --deserialize 22
  2 root        20  0 0 0 0 S   0.0   0.0   0:02:62 [kthreadd]
  4 root        0 -20 0 0 0 S   0.0   0.0   0:00:00 [kworker/0:0H]
  6 root        20  0 0 0 0 S   0.0   0.0   15:13:30 [ksoftirqd/0]
  7 root        rt  0 0 0 0 0 S   0.0   0.0   2:25:39 [migration/0]
  8 root        20  0 0 0 0 S   0.0   0.0   0:00:00 [rcu_bh]
```

- Resubmitted several previously held runs
 - 35865 35866 35869 35875 35878 35881 35882 36234 36245 36246 36247 36252 36254 36255 36257 36258 36259 36260 36268 36298 36300 36301 36302 36306 36308 36309 36311 36313 36315 36316 36317 36318 36319 36320 36321 36322 36323 36325 36327 36328 36329
 - 4 of these are no longer valid jobs (probably trigger only)
 - Condor cluster 21161548
 - 36 jobs running, 1 finished.
- Several runs over the weekend that went into a failed state. Issue with INTT ‘losing the run number’, e.g.
\$ grep 'Single Streaming Input Manager INTT_0 reads' failed.logs/*
Single Streaming Input Manager INTT_0 reads run 1712113657 from file ...

| | | | | | | | | |
|-------|--|---------------------|-------|--|---------------------|-------|--|---------------------|
| 37223 | | 2024-04-07 23:08:39 | 37037 | | 2024-04-07 19:51:41 | 36993 | | 2024-04-07 18:39:20 |
| 37195 | | 2024-04-07 22:24:46 | 37035 | | 2024-04-07 19:46:14 | 36989 | | 2024-04-07 18:33:53 |
| 37192 | | 2024-04-07 22:13:49 | 37027 | | 2024-04-07 19:35:18 | 36972 | | 2024-04-07 18:06:37 |
| 37189 | | 2024-04-07 22:13:49 | 37024 | | 2024-04-07 19:29:51 | 36970 | | 2024-04-07 18:01:10 |
| 37186 | | 2024-04-07 22:08:20 | 37015 | | 2024-04-07 19:13:29 | 36968 | | 2024-04-07 18:01:10 |
| 37176 | | 2024-04-07 21:46:27 | 37008 | | 2024-04-07 19:02:34 | 36963 | | 2024-04-07 17:55:43 |
| 37160 | | 2024-04-07 21:24:33 | 37006 | | 2024-04-07 19:02:34 | 36958 | | 2024-04-07 17:44:49 |
| 37064 | | 2024-04-07 20:35:22 | 37004 | | 2024-04-07 19:02:34 | 36956 | | 2024-04-07 17:39:21 |
| 37057 | | 2024-04-07 20:24:26 | 37002 | | 2024-04-07 18:57:06 | 36944 | | 2024-04-07 17:23:00 |
| 37053 | | 2024-04-07 20:18:59 | 36999 | | 2024-04-07 18:50:26 | | | |
| 37052 | | 2024-04-07 20:13:31 | 36998 | | 2024-04-07 18:44:47 | | | |

- TPC and other subsystems?

- Slurp Update
- Calorimeter Calibration Workflow

- Two step, iterated workflow
 - Stage 0: Raw data combining / template fitting → DSTs w/ towers (uncalibrated)
 - In progress / nearly ready
 - Stage 1: Run calorimeter reconstruction / pi0 reconstruction on calorimeter event DSTs.
 - Inputs are chunks of 10-20 files from each run, generated from DB query to the file catalog *plus* a root file containing the calibration constants for the current iteration. All files matching the query must be in the file catalog.
 - CDB_CALO_PI0CALIB_ITER-<run0>-<runF>-<iter#>.root
 - Output is a set of histograms (tower-by-tower mass spectra?)
 - 1a template fit / signal extraction → move to stage 0
 - 1b pi0 reconstruction
 - Stage 2: Merge histograms and extract new set of calibration constants
 - Inputs will be the histogram files generated in stage 1 plus the calibration constants used (?)
 - Output new calibration constants file which initiates a new iteration
 - Termination condition:
 - EITHER difference between the input and output constants file is negligible
 - Output calibration constants are pushed the the CDB (may or may not automate?)
 - Stage output as CDB_CALO_PI0CALIB_DONE-<run0>-<runF>-<iter#>.root
 - does not match the input query, terminating the loop
 - OR max iterations reached
 - Output calibration constants are pushed the the CDB
 - Stage output as CDB_CALO_PI0CALIB_STOP-<run0>-<runF>-<iter#>.root
 - does not match the input query, terminating the loop
 - OR there is no need for the job to produce any output (or am I missing something here?)
 - The submission will end when the calibration does not push any output into the file catalog
 - Finished calibrations will be in the calibration status table
 - How to find failed calibrations?
 - As above - this needs the current value of the iteration to be known in the job
 - How to do that (coded into the filename)?

- Sliding window analysis

- Questions
 - How to update DB? Currently by hand.
 - What to do with failed jobs at stage 1? Possible that jobs fail.
 - Do you need to specify a good run list? (Run 23 yes. 24... may be more automated).
 - May be easier to specify a bad run list.
 - Planning to run concurrent with data taking.
 - How many files per job? 50-100 jobs/run
 - 1.296 Billion events at 24h running per day. Turnaround at 5 iterations. 5 h max job length.
 - Max number of jobs
 - What is the job length?
 - Currently est. at 15 sec/ 100K events
 - Est. 50h to run 1.296 B events once, so that is a minimum of 10 jobs.
 - stage 1: 1 day → 10 cpu days

- stage 2... fitting / hadding well below stage 1.
 - O(100) nodes concurrently devoted emc calibrations... may change based on below...
- size of inputs? data volume consideration for local copy. size constraint on number of jobs.
- ... emcal + vertex size
- Discussion
- AOB
-

Meeting Mon Apr 01 2024 1 PM EDT Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261

- Production Update
 - Production Loop running for DST_COSMICS from 4PM 3/29
 - Error in stageout script (hubris strikes again)... but good opportunity to test the “hot swap” capability of the system.

Active rules:
+-----+
| active rules |
+-----+
| DST_COSMICS |
+-----+

Summary of jobs which have not reached staus='started'
+-----+
| dsttype | prod_id | num_jobs | avg_time_to_submtt | min_time_to_submit | max_time_to_submit |
+-----+
| DST_COSMICS_run2pp | 58 | 1 | 00:00:00 | 00:00:00 | 00:00:00 |
+-----+

Summary of jobs which have reached staus='started'
+-----+
| dsttype | prod_id | num_jobs | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | min_job_duration | max_job_duration | sum_events |
+-----+
| DST_COSMICS_run2pp | 58 | 101 | 00:00:10.455446 | 0 | 55 | 46 | 0 | | | | 0 |
+-----+

Summary of jobs by condor cluster
+-----+
| dsttype | cluster | min_run | max_run | num_jobs | earliest_start | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | sum_events |
+-----+
| DST_COSMICS_run2pp | 21161497 | 36263 | 36266 | 17 | 2024-03-29 17:12:47 | 00:00:11.176471 | 0 | 17 | 0 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161498 | 36257 | 36257 | 1 | 2024-03-29 18:15:56 | 00:00:17 | 0 | 1 | 0 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161500 | 36257 | 36257 | 1 | 2024-03-29 18:56:43 | 00:00:18 | 0 | 1 | 0 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161501 | 36257 | 36257 | 1 | 2024-03-29 19:15:22 | 00:00:15 | 0 | 1 | 0 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161502 | 36263 | 36263 | 36 | 2024-03-29 20:00:00 | 00:00:13.3 | 0 | 36 | 0 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161503 | 36258 | 36258 | 1 | 2024-03-29 20:00:04 | 00:00:14 | 0 | 1 | 0 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161504 | 36466 | 36471 | 13 | 2024-03-30 01:55:19 | 00:00:30.530462 | 0 | 0 | 13 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161505 | 36466 | 36528 | 39 | 2024-03-30 02:00:16 | 00:00:12.2 | 0 | 1 | 29 | 0 | | 0 |
| DST_COSMICS_run2pp | 21161506 | 36268 | 36268 | 1 | 2024-03-30 03:44:52 | 00:00:11 | 0 | 1 | 0 | 0 | | 0 |
+-----+

Summary of failed jobs by run
+-----+
| dstname | run | segment | cluster | process | prod_id |
+-----+
+-----+

Next submit: 53%

- Fixed stageout script Sat 850AM. Cleared production status entries at 855AM. Did *not* terminate the submission loop (could have paused over tmux with ctrl-s?)

Active rules:
+-----+
| active rules |
+-----+
| DST_COSMICS |
+-----+

Summary of jobs which have not reached staus='started'
+-----+
| dsttype | prod_id | num_jobs | avg_time_to_submit | min_time_to_submit | max_time_to_submit |
+-----+
| DST_COSMICS_run2pp | 58 | 25 | 00:00:00 | 00:00:00 | 00:00:00 |
+-----+

Summary of jobs which have reached staus='started'
+-----+
| dsttype | prod_id | num_jobs | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | min_job_duration | max_job_duration | sum_events |
+-----+
+-----+

Summary of jobs by condor cluster
+-----+
| dsttype | cluster | min_run | max_run | num_jobs | earliest_start | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | sum_events |
+-----+
+-----+

Summary of failed jobs by run
+-----+
| dstname | run | segment | cluster | process | prod_id |
+-----+
+-----+

Next submit: 6%

- Proper response observed... (new production ID b/c updated code was submitted, jobs submitted, not yet picked up for production.

- SLURP Updates
 - Naming conventions
 - Output, log and condor directories now formatted as:
 - /PATH/TO/run_{floor(100*run/100):08d}_{100*ceil(run/100):08d}
 - not the most flexible syntax... \$(rungroup) condor macro is hard coded, applied at the file match. So decision earlier rather than later as to the number of runs to group in a single directory...
 - DST_TPCCOSMICS_* → DST_COSMICS_runNcollision_*
 - TODO: allow arbitrary DST_BLAH_BLAH_BLAH_ALLCAPS_runNcollision_*
 - DST_[QUALIFIERS_]+[type]_{tag1}_{tag2}-run-seg.ext
 - GL1, INTT, MVTX files are now passed to the event builder...
 - The mechanism *depends on* the naming convention for the files.
 - TODO:
 - Update at end of job is not being made (ended timestamp missing, nevents missing from production_status).
 -

- Recent fixes in new...
 - List of a few jobs which (quickly) went on hold over the weekend

| | | | |
|-------------|------------|---------------------------------|---|
| 21161502.5 | 0+00:05:12 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036298-0000 |
| 21161502.8 | 0+00:10:13 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036301-0000 |
| 21161502.9 | 0+00:05:13 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036302-0000 |
| 21161506.0 | 0+00:10:15 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036268-0000 |
| 21161507.0 | 0+00:20:11 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036257-0000 |
| 21161507.5 | 0+00:10:12 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036268-0000 |
| 21161507.10 | 0+00:10:13 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036298-0000 |
| 21161507.13 | 0+00:10:12 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036301-0000 |
| 21161507.14 | 0+00:05:12 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036302-0000 |
| 21161507.32 | 0+00:45:18 | DST_COSMICS_run2pp_new_2023p013 | DST_COSMICS_run2pp_new_2023p013-00036320-0000 |

- Resubmit 1140AM Monday... but, need to update the macro. TBD.
- Discussions
- AOB

Meeting Mon Mar 18 2024 1 PM EDT Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261

- Production Status
 - Calorimeter production ran / finished Friday night/saturday morning
 - TPC productions... ran into issue w/ extended input list of lists

- Operational Considerations
 - Need to keep a stable production area
 - Continual production (automated) should help here... notice if things go wrong faster
 - fixed stable branch
 - Updates to the production are need to go through an integration test
 - New codes to integration area
 - Run a standard set of production runs in new with recent DB tag
 - All jobs at small # of events
 - Small # of jobs with all events
 - If it passes, the integration area becomes the new production area
 - Production macros, scripts, definitions on their own branch
 - Ideally submit via a tmux session
 - \$ ssh sphnxpro@phnxsub02 -t tmux a -t sPHENIX-Production
 - (plan to create a bash script to initialize a session).
- Slurp Update
 - ProductionSystem directory created to support production running
 - No changes made directly here ... this will be overwritten with changes from the Integration directory
 - ProductionDefs holds the “slurp-examples”. A production branch is now created
 - Will plan to merge changes from git
 - ProductionSystemIntegration directory replicates the above. Changes will be deployed here before being copied over the production directories.
 - Following changes will be deployed and tested in this area, before
 - ramenya2.py can now read a list of rules from a text file...
 - Allows production manager to add new rules / comment out existing ones without shutting down the main loop
 - Fixed issue with calorimeter example... output manager is now created, and stageout script now executes cups to stage out results / copyscript.pl is now disabled.
 - TPC combiner workflow tested w/ new macro... but does not (yet) fill in the MVTX, INTT, GL1 file lists. TBD.
 - Naming conventions, directories, event totals...
 - TODO: new_2023p011 → dataset
 - TODO: output directory run_%08i_%08i etc...

| filename | runnumber | segment | size | dataset | dsttype | events |
|--|-----------|---------|-----------|------------|-----------|--------|
| DST_CALOR_auai23_new_2023p011-00022027-0013.root | 22027 | 13 | 105556295 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0126.root | 22027 | 126 | 105753057 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0167.root | 22027 | 167 | 105712649 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0160.root | 22027 | 160 | 105250183 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0068.root | 22027 | 68 | 104917060 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0004.root | 22027 | 4 | 104574867 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0071.root | 22027 | 71 | 103920559 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0028.root | 22027 | 28 | 104824419 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0024.root | 22027 | 24 | 104540251 | auau23_new | DST_CALOR | 123 |
| DST_CALOR_auai23_new_2023p011-00022027-0054.root | 22027 | 54 | 104507816 | auau23_new | DST_CALOR | 123 |

| lfn | full_file_path | size |
|--|---|-----------|
| DST_CALOR_auai23_new_2023p011-00022027-0004.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0004.root | 104574867 |
| DST_CALOR_auai23_new_2023p011-00022027-0119.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0119.root | 106804542 |
| DST_CALOR_auai23_new_2023p011-00022027-0005.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0005.root | 105106423 |
| DST_CALOR_auai23_new_2023p011-00022027-0157.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0157.root | 105114591 |
| DST_CALOR_auai23_new_2023p011-00022027-0153.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0153.root | 105021669 |
| DST_CALOR_auai23_new_2023p011-00022027-0138.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0138.root | 104028994 |
| DST_CALOR_auai23_new_2023p011-00022027-0140.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0140.root | 105254973 |
| DST_CALOR_auai23_new_2023p011-00022027-0149.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0149.root | 104278838 |
| DST_CALOR_auai23_new_2023p011-00022027-0156.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0156.root | 105067075 |
| DST_CALOR_auai23_new_2023p011-00022027-0139.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0139.root | 105763679 |
| DST_CALOR_auai23_new_2023p011-00022027-0084.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0084.root | 105833725 |
| DST_CALOR_auai23_new_2023p011-00022027-0113.root | /sphenix/lustre01/sphnxpro/slurptest/run_22000_22100/DST_CALOR_auai23_new_2023p011-00022027-0113.root | 104818976 |

- Cooldown expected 4/15 ... beam on or about 4/20
- AOB
 - Calorimeter cosemics incoming
 - New event builder macro incoming
 -

Meeting Mon Mar 18 2024 1 PM EDT Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261

- Production Update
 - First production run based on SLURP...
 - DST_TPCCOSMICS and DST_TPCCALIB (both running same event builder through run_cosmics.sh)

jwebb2@ssh02:/home/sphenix/u/jwebb2

Summary of jobs which have not reached staus='started'

| dsttype | prod_id | num_jobs | avg_time_to_submit | min_time_to_submit | max_time_to_submit |
|---------|---------|----------|--------------------|--------------------|--------------------|
| | | | | | |

Summary of jobs which have reached staus='started'

| dsttype | prod_id | num_jobs | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | min_job_duration | max_job_duration | sum_events |
|----------------|---------|----------|-------------------|---------------|-------------|--------------|------------|------------------|------------------|------------------|------------|
| DST_TPCCOSMICS | 31 | 39 | 00:00:10.487179 | 0 | 32 | 6 | 1 | | | | 0 |
| DST_TPCCOSMICS | 33 | 14 | 00:00:28.857143 | 0 | 13 | 1 | 0 | | | | 0 |
| DST_TPCCALIB | 32 | 1 | 00:00:16 | 0 | 1 | 0 | 0 | | | | 0 |
| DST_TPCCALIB | 34 | 23 | 00:00:20.956522 | 0 | 20 | 1 | 2 | | | | 0 |

Summary of jobs by condor cluster

| dsttype | cluster | min_run | max_run | num_jobs | earliest_start | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | sum_events |
|----------------|----------|---------|---------|----------|---------------------|-------------------|---------------|-------------|--------------|------------|------------------|------------|
| DST_TPCCOSMICS | 21161341 | 32549 | 33015 | 34 | 2024-03-13 19:40:50 | 00:00:08.470588 | 0 | 20 | 6 | 0 | | 0 |
| DST_TPCCOSMICS | 21161344 | 33316 | 33316 | 1 | 2024-03-14 06:13:50 | 00:00:41 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161348 | 33318 | 33318 | 1 | 2024-03-14 06:19:34 | 00:00:30 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161346 | 33320 | 33320 | 1 | 2024-03-14 06:20:46 | 00:00:04 | 0 | 0 | 0 | 1 | | 0 |
| DST_TPCCOSMICS | 21161347 | 33321 | 33321 | 1 | 2024-03-14 06:22:58 | 00:00:15 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161348 | 33322 | 33322 | 1 | 2024-03-14 06:30:16 | 00:00:31 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161350 | 32553 | 32564 | 3 | 2024-03-14 17:27:20 | 00:00:22 | 0 | 3 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161353 | 33006 | 33006 | 1 | 2024-03-14 20:53:23 | 00:00:22 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161364 | 33325 | 33326 | 2 | 2024-03-15 22:34:45 | 00:00:22 | 0 | 2 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161365 | 33325 | 33323 | 1 | 2024-03-15 22:43:02 | 00:00:08 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161366 | 33324 | 33338 | 3 | 2024-03-15 22:50:46 | 00:00:37.333333 | 0 | 3 | 0 | 0 | | 0 |
| DST_TPCCOSMICS | 21161367 | 33323 | 33332 | 4 | 2024-03-15 22:53:36 | 00:00:19 | 0 | 4 | 0 | 0 | | 0 |
| DST_TPCCALIB | 21161354 | 33314 | 33314 | 1 | 2024-03-14 18:00:14 | 00:00:10 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCALIB | 21161351 | 31701 | 33315 | 21 | 2024-03-14 17:27:20 | 00:00:19.761905 | 0 | 18 | 1 | 2 | | 0 |
| DST_TPCCALIB | 21161355 | 33408 | 33408 | 1 | 2024-03-15 00:13:36 | 00:00:40 | 0 | 1 | 0 | 0 | | 0 |
| DST_TPCCALIB | 21161360 | 32497 | 32497 | 1 | 2024-03-15 10:17:00 | 00:00:27 | 0 | 1 | 0 | 0 | | 0 |

Summary of failed jobs by run

| dstname | run | segment | cluster | process | prod_id |
|-----------------------------|-------|---------|----------|---------|---------|
| DST_TPCCALIB_new_2023p009 | 31715 | 0 | 21161351 | 1 | 34 |
| DST_TPCCALIB_new_2023p009 | 31722 | 0 | 21161351 | 2 | 34 |
| DST_TPCCOSMICS_new_2023p009 | 33320 | 0 | 21161346 | 0 | 31 |

Next submit: 38%

- Identified memory issues (overflowed the 2GB condor limit due to large number of spurious TPC hits)
- Monday 10:00 AM – removing TPC cosmics and TPC calibration jobs that were placed on hold...
 - calibration runs went to hold w/in 5min... increase to 16GB and resubmit...
 - these are hitting 32GB and going on hold...
 - most cosmics running > 25 min at 11AM
 - ...

jwebb2@ssh03:/home/sphenix/u/jwebb2

summary of jobs which have not reached staus='started'

| dsttype | prod_id | num_jobs | avg_time_to_submit | min_time_to_submit | max_time_to_submit |
|---------|---------|----------|--------------------|--------------------|--------------------|
|---------|---------|----------|--------------------|--------------------|--------------------|

summary of jobs which have reached staus='started'

| dsttype | prod_id | num_jobs | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | min_job_duration | max_job_duration |
|----------------|---------|----------|-------------------|---------------|-------------|--------------|------------|------------------|------------------|------------------|
| DST_TPCCOSMICS | 33 | 42 | 00:00:10.833333 | 0 | 42 | 0 | 0 | | | |
| DST_TPCCALIB | 34 | 20 | 00:00:15 | 0 | 20 | 0 | 0 | | | |

summary of jobs by condor cluster

| dsttype | cluster | min_run | max_run | num_jobs | earliest_start | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration |
|----------------|----------|---------|---------|----------|---------------------|-------------------|---------------|-------------|--------------|------------|------------------|
| DST_TPCCOSMICS | 21161370 | 33332 | 33332 | 1 | 2024-03-18 14:10:32 | 00:00:26 | 0 | 1 | 0 | 0 | |
| DST_TPCCOSMICS | 21161371 | 33323 | 33331 | 8 | 2024-03-18 14:15:24 | 00:00:03 | 0 | 8 | 0 | 0 | |
| DST_TPCCOSMICS | 21161372 | 33325 | 33325 | 1 | 2024-03-18 14:20:55 | 00:00:18 | 0 | 1 | 0 | 0 | |
| DST_TPCCOSMICS | 21161373 | 32564 | 33322 | 22 | 2024-03-18 14:26:04 | 00:00:12.045455 | 0 | 22 | 0 | 0 | |
| DST_TPCCOSMICS | 21161374 | 32549 | 32562 | 9 | 2024-03-18 14:31:13 | 00:00:05 | 0 | 9 | 0 | 0 | |
| DST_TPCCOSMICS | 21161377 | 32553 | 32553 | 1 | 2024-03-18 15:30:15 | 00:01:17 | 0 | 1 | 0 | 0 | |
| DST_TPCCALIB | 21161378 | 32497 | 33408 | 20 | 2024-03-18 15:45:02 | 00:00:15 | 0 | 20 | 0 | 0 | |

summary of failed jobs by run

| dstname | run | segment | cluster | process | prod_id |
|---------|-----|---------|---------|---------|---------|
|---------|-----|---------|---------|---------|---------|

next submit: 18%

- ... production status does not indicate when jobs go on hold b/c they are evicted... ~3 cosmics and all calibration jobs currently on hold.
-
-
- SLURP update
 - streaming workflows
 - On eviction, truncated job log is now available and copied back / condor file will now show eviction rather than failed copy
 - ramenya2.py
 - command line tool for submission, monitoring and management of jobs
 - *ramenya2.py submit --runs 0 9999999 --rules DST_TPCCALIB DST_TPCCOSMICS*
 - *ramenya2.py query --reports pending started clusters failed*
 - *ramenya2.py remove DST_TPCCALIB_NEW_2023p009 33320*
 - runs in a loop 5min (default) between iterations
 - when daq DB indicates that a run has had all files transferred to SDCC, work is launched
 - currently configured to process TPC cosmic and calibration runs
 - several canned queries of the production_status DB are implemented as part of the “query” subcommand
 - text based, color coded, ... failed jobs are highlighted in red
 - but not all failures are caught at this point... e.g. if job is evicted, this is not transmitted back to the production status table
 - remove sub-command cleans up the production status DB and the file catalog. Will add code to cleanup the files on disk as well.
 - When a run is cleaned up... it becomes available again for submission
 - Need to query condor and remove the held job
 - donburi.py
 - flask application ... primitive demonstration of a web application
 - FLASK_APP=donburi.py flask run
 - currently running on phnxsub02
 - lynx http://127.0.0.1:5000/clusters
 - lynx http://127.0.0.1:5000/failed
 - reports implemented in ramenya are exported
 - TODO: Run Range Convention... apply to output directory structure

we will most likely re-use the run range naming convention we used in PHENIX which looks like e.g. run_0000365000_0000366000 for files from run 365000 to run 365999

in perl this was like
my \$lower = \$runnumber - \$runnumber%1000;
my \$upper = \$lower + 1000;
my \$runpath = sprintf("run_%010d_%010d",\$lower,\$upper);
depending on how things go and how many files we get in those ranges we might have to go to %100

... changing this is easier said than done b/c of the way the directory paths are defined for condor (they use condor macros to specify the path)... but we have to create them before the jobs are submitted to condor (evaluate python expressions)... perhaps injecting a “precmd / precmdarguments” into the job file.

- Discussion
- AOB

Meeting Mon Mar 03 2024 1 PM EDT Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261

- Production Update (TPC cosmics)
- SLURP update
 - SLURP now passes PFNs as the \$(inputs) filed rather than LFNs. \$(ranges) will still be formatted as LFN:minEvent:maxEvent.
 - Could provide options to pass \$(inputs_lfn) or \$(inputs_pfn)...
 - Solved issue with condor files to /tmp

- A run will be submitted IFF all files in run satisfy the constraint: “having every(transferred_to_sdcc)”
- Utilizes the output manager’s closing script functionality
 - stageout script which should have only minor differences between production.... calls cups...
 - “cups” uses a simple “cp” to copy the files to lustre
 - md5sum before and after is required to be the same, otherwise no entry is made in the file catalog...
 - *May need to extend the production status table... this failure mode should be indicated.*
- Verified that calorimeter workflow still runs
- *ramenya* runs kaedam.py, so submission is correct. Polling output still points to the (deprecated) production_status tables in the filecatalog DB. To be fixed.
- Discussion
- AOB

Meeting Mon Mar 03 2024 1 PM EST [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261

- Production Update
 - Several missing runs @ event combining stage... zdc handling? tested 21520 (dropped zdc). Issue still not clear.
 - Streaming device production missing part is DB
- SLURP Update
 - Implemented cosmics workflow w/ copy back functionality
 - Input query (i.e. input file lists) can now be validated against a physical directory rather than the file catalog
 - Still need to add constraint that all files have been transferred
 - Test run 30117 shows files which have not been transferred
 - Does fun4all wait for the copy back process(es) to finish?
 - Have two conflicting tests at the moment...
 - Have not yet sorted out the location of the condor logs... when I set the condor files to be created in the /tmp area, it also tries (and fails) to write the output logs to those locations.
- Performance testing
 - Start to finish test O(10k) jobs for DB test, eg calorimeter production
 - New Calo production incoming timescale of next week
 - Dmitry A. would like advance notice 24h.
- Discussion
- AOB

Meeting Mon Feb 26 2024 1 PM EST [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261

- Production update
 - 2 productions
 - 1st had issues w/ database entries 2nd fine.
 - Setting dataset name... yaml configuration file. But also two different locations in copy script
 - Jobs placed on hold... copy back of log files. condor logs ought to be set to /tmp
 - 13 runs where DSTs no longer produced
 - event combining step did not produce output... tried to rerun 21520, no output again.
 - potential issue w/ zdc data
 -
- Slurp update
 - (nearly) Ready for the next release of *slurp*
 - Uses the new production DB
 - Separate read / write instances
 - There is a risk of repeating job submission if
 - kaedama is called with more frequently than the time to sync the replicas.
 - mitigated interactively b/c it does take O(1min) to go through the submission process
 - mitigated in batch b/c we will typically run w/
 - if the synchronization goes down
 - Use the write only until we see issues.
 - Added an “examples” repository
 - <https://github.com/klendathu2k/slurp-examples>
 - DST_EVENT and DST_CALOR rules in separate files, submit jobs as we have in the past
 - DST_EVENT_RANGES shows how to run with fixed nevents in each job
 - *Directing condor logs to /tmp...*
 - *rsync periodically next to job logs*
- Discussion
 - cups to take over file copy
- AOB
 - Cosmics data / time to hook up processing chain
 - Inputs, macro/script, ...
 - Stress test TBD

Meeting Mon Feb 12 2024 1 PM EST [Meeting Link...](#) Meeting ID: 160 468 6747 Passcode: 155261

- n.b. short meeting today as Sasha is at a conference.
- Production update
- Slurp update
 - Updated to allow event builder jobs to be submitted in run ranges.
 - Difficulty is that event builder jobs take a runlist as input.
 - Each row returned by the DB query results in one job submitted to condor...
 - So the DB query has to aggregate the files produced by the daq to a flat list for the given run
 - Which is then passed a comma-separated list of values to the payload script...
 - ... along with the first and last event in each file

- Input list is extended with the range of each file
- Passed down into the run script
- Run script adds each input file to the filelist based on whether the production event range has an event which falls in the range of the input file.
- ... need to make sure that the input files are sorted in ascending sequence order
- *Bottom line... slurp is setup now to break the event builder jobs up into event ranges... (modulo a proper sorting of the run lists).*

- Discussion
- AOB

Meeting Mon Feb 05 2024 1 PM EST Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261

- Production Update
- Slurp Update
 - Complete production (event builder + calorimeter) ran w/ automatic creation of the filelists
 - DAQ db and filecatalog must agree on the list of files available for a run
 - daqdb query restricts to the list of files needed in calorimeter production.
 - If the two agree, work is launched.

| | | | | | | | | | | | |
|--|---------|----------|-------------------|---------------|-------------|--------------|------------|------------------|------------------|------------------|------------|
| Summary of jobs which have reached staus='started' | | | | | | | | | | | |
| dsttype | prod_id | num_jobs | avg_time_to_start | num_submitted | num_running | num_finished | num_failed | avg_job_duration | min_job_duration | max_job_duration | sum_events |
| DST_EVENT_auiu23 | 8 | 174 | 00:01:50.3793 | 0 | 0 | 174 | 0 | 01:18:14.063953 | 00:11:15 | 05:32:39 | 42141254 |
| DST_CALOR_auiu23 | 9 | 55425 | 00:22:58.317943 | 0 | 0 | 55425 | 0 | 01:34:03.471194 | 00:00:12 | 02:40:05 | 41530694 |

- “dev” branch established in github to isolate development code from production. production system will utilize the main branch. Development on parallel branch until ready to deploy in production.
- Added a preprocessing step to kaedama.
 - Before running matches and building the jobs, a presubmit query and script may be applied.
 - Each row that is returned from the DB query is converted to a list of arguments to a script
 - User defined script is called for each row returned by the DB query

- General Discussion
- AOB

Meeting Mon Jan 29 2024 1 PM EST Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261

- Slurp Update
 - Added python’s logger (warnings, errors are now colorized and annotated with timestamps).
 - Event builder pipeline is now able to run automatically when files are added to the filecatalog...
 - Added a *runlist query* to the rule definition. When a runlist query is specified:
 - The daq db (replica) is queried for a list of all files (LFNs) which match the query.
 - List of files (LFNs) is checked against the filecatalog.
 - Llist of files (PFNs) is provided via condor as \$(inputs), which may be picked up by the production script.
 - Workflow will be dispatched IFF all files that the daq recorded in the DB are in the file catalog as well.
 - TODO verify checksums of the PFNs against the daqdb
 - If expected files are missing, a warning is issued and the given run is skipped
 - If there are files in the filecatalog (but not known to the daq DB) we issue a warning and continue
 - Incoming event builder will need to be launched over event ranges. One way to handle this would be to have a single process compute event ranges and input file lists, and then populate a table of the form

```
CREATE TABLE if not exists EVENT_BUILDER_RANGES (  
  id          serial          unique, -- unique identifier  
  run         int             not null, -- runnumber  
  segment     int             not null, -- segment number  
  first       int             not null, -- first event  
  last        int             not null, -- last event  
  seb01       text            not null, -- space separated list of LFNs  
  seb02       text            not null, -- space separated list of LFNs  
  seb03       text            not null, -- space separated list of LFNs  
  seb04       text            not null, -- space separated list of LFNs  
  seb05       text            not null, -- space separated list of LFNs  
  seb06       text            not null, -- space separated list of LFNs  
  seb07       text            not null, -- space separated list of LFNs  
  seb08       text            not null, -- space separated list of LFNs  
  seb14       text            not null, -- space separated list of LFNs  
  seb15       text            not null, -- space separated list of LFNs  
  seb16       text            not null, -- space separated list of LFNs  
  seb17       text            not null, -- space separated list of LFNs  
  seb18       text            not null, -- space separated list of LFNs  
# , ... or perhaps a single “inputs” list ...  
  foreign key (prod_id) references PRODUCTION_SETUP (id) ,  
  primary key ( id, run, segment, first, last ) )
```

-
- ..
- ..

- Discussion
- Test production using slurp
 - Sasha L. run test production last week
 - Three runs, DST_EVENT first, DST_CALOR next
 - After some confusion about correct macros and learning curve, test production completed successfully
 - For DST_EVENT input files are coming from pre-made text files.
 - No-slurp related: what to do with extra validation histogram files?
 - Minor stuff: failed runs had to be removed by hand from db, tags as input parameters, separate yaml files for DST_EVENT and DST_CALOR.
- TODO: Cleanup script...
- Run DB, Production DB, et... all ready to rock.
- AOB

Meeting Mon Jan 22 2024 1 PM EST Meeting Link... Meeting ID: 160 468 6747 Passcode: 155261

- Production System Update
 - Slurp “1.0”
 - Documentation
 - Consolidated the code. Rules are now defined in a single YAML file (greatly simplifying the business logic in kaedama).
 - snakemake considered...
 - but ultimately it relies on the existence of files on a filesystem to determine dependencies... for our purposes, a job running on condor is presumed to satisfy a dependency.
 - Run-by-run job status during the submission loop... jobs sorted by run and/or condor cluster.
 - Examples of event builder and calorimeter jobs passing back accumulated statistics
 - Ready for the experiment to setup and run a test production against current software
 - Quick setup / test of production QA
 - Fun4all will export a json file at the end of each run
 - Single production_quality table of the form

\d production_quality

| Table "public.production_quality" | | | | |
|-----------------------------------|-----------------------|-----------|----------|--|
| Column | Type | Collation | Nullable | Default |
| id | integer | | not null | nextval('production_quality_id_seq'::regclass) |
| dstname | character varying(32) | | not null | |
| run | integer | | not null | |
| segment | integer | | not null | |
| stat_id | integer | | not null | |
| qual | jsonb | | not null | |

Indexes:

"production_quality_pkey" PRIMARY KEY, btree (id, dstname, run, segment)

"production_quality_id_key" UNIQUE CONSTRAINT, btree (id)

Foreign-key constraints:

"production_quality_stat_id_fkey" FOREIGN KEY (stat_id) REFERENCES production_status(id)

select * from production_quality;

| id | dstname | run | segment | stat_id | qual |
|----|----------------------------------|-------|---------|---------|---|
| 3 | DST_CALOR_auau23_ana387_2023p003 | 22026 | | 0 | 136318 {"tag": "2023p003", "build": "ana387", "workflow": "calorimeter "} |
| 4 | DST_CALOR_auau23_ana387_2023p003 | 22026 | | 0 | 136318 {"tag": "2023p003", "build": "ana387", "workflow": "calorimeter "} |
| 5 | DST_CALOR_auau23_ana387_2023p003 | 22026 | | 0 | 136318 {"tag": "2023p003", "build": "ana387", "workflow": "calorimeter "} |

There are two syntaxes available for querying json data...

Postgres syntax for selecting json columns...

psql> select qual->'workflow' as workflow from production_quality where run=22026 and dstname like DST_CALOR%;

jsonpath expressions

psql> select jsonb_path_query("qual",'\$.workflow') as workflow from production_quality;

...

cups support for storing qa data / json format.
cups.py ... quality –qfile userMetaData.json

- Next priority is getting the automated submission of jobs in place.
 - Run DB to tell us what detectors were live in a given run, thus what set of files need to be present to consider a run’s dataset to be completely transferred to lustre.
- Discussion
 - event builder now takes back responsibility for updating the file catalog.
 - learning curve vs operational needs
- AOB


Meeting Mon Jan 08 2024 1 PM EST Room 1-224 New Meeting Link... cWJtem92VDE0ZG5FVTh5T28zMnE1QT09

- Torre - zoom link does not let me in. “If you are the meeting host, sign in to start the meeting 1604686747”
- Sasha - same for me
- Chris and I (jason) are on... hold on a sec... Try now (Tony was able to join).
- Sasha - I still get the same thing - it asks me to log in to BNL account
- What’s the passcode? Pass code is above (worked for me) but Chris has the numeric code below...
-
- Meeting ID: 160 468 6747
- Passcode: 155261
-
-
-
- Happy New Year
- Expected schedule and discussion of readiness plan
 - Beam in March
 - Initially support calorimeter and qa jobs

- March 4th CAD plan for cooldown (beam)? but sPHENIX wants ~April 1.
- Turn on... event building / reconstruction
- DB replication (run DB) ... need to hook into this for existing runs, etc...
- MBD/ZDC global reconstruction workflow can be setup
- Streamline / simplify workflow definition

- Summary of 12/20 SDCC-sPHENIX Meeting
 - Full day meeting between Chris and Martin, SDCC, Jason, to review progress, areas of responsibilities, etc...
 - Brief 1 page summary distributed
 - Updates to issues identified to be posted on the meeting agenda
 - <https://indico.bnl.gov/event/21477>

- Production System Update
 - Demonstration of complete workflow
 - Production script (ramenya) runs in a loop, executing the submission script (kaedama) every 5 minutes.
 - Event builder jobs are submitted immediately b/c their rule is satisfied (inputs available)
 - They run and register their outputs with the filecatalog as they are produced.
 - Calorimeter jobs are submitted when their rule is satisfied (event builder outputs are available)
 - Now w/ plots!
 - Scale test @ 174 runs... Plot shows the timestamp at which each job reports a given status (submitting, submitted, started, running, finished, failed). Transparent bars show the cumulative sums. 2 calor jobs failed out of 5690 submitted. [submitting, submitted, started, running, finished, failed]



- And with job summary output.
 - Repeat scale test with 2GB event builder inputs (thus 10x as many DST_CALOR jobs)...

| Summary of jobs which have reached staus='started' | | | | | | | | | | |
|--|----------|-------------------|------------|----------|-----------|---------|--------------|--------------|--------------|------------|
| dsttype | num_jobs | avg_time_to_start | nsubmitted | nrunning | nfinished | nfailed | avg_duration | min_duration | max_duration | sum_events |
| DST_EVENT_auiu23 | 174 | 00:02:45.178161 | 0 | 0 | 174 | 0 | 01:48:41 | 00:12:42 | 07:57:18 | 42376839 |
| DST_CALOR_auiu23 | 55777 | 02:42:43.534288 | 0 | 0 | 55775 | 2 | 01:21:39 | 00:00:19 | 02:46:09 | 0 |
| (2 rows) | | | | | | | | | | |

- (quick demo)
- Next steps
 - Setup a functional system for testing purposes by MM/DD focusing on QA-type jobs and calorimeter production
 - Date depends on run DB replication...
 - Build the event builder input runlists from the file catalog, but n.b.
 - Need to integrate with the run database to discover when the event builder inputs are ready...
 - The query needs to know whether all *necessary* files have been transferred from 1008 (involves detector state and/or knowledge of what has been shipped)
 - Leverage transferred flag in DB
 - Track total event counts per production (ought to be done by meeting...)
 - Pushed to file catalog and production status tables
 - Add a QA table which is filled on a per-job basis... from an output by fun4all, that contains QA information in e.g. json format.
 - Support for data taking / production by MM/DD
 - Implement the remaining workflows (event builder for tracking... tracking chain)
 - User friendliness
 - Plan to switch to snakemake to define production rules? Or refactor what we currently have. (Decision to Sasha).
 - Eye towards simplifying the definition of production rules vs the homebrew system in place

- Discussion

- AOB

Meeting Mon Dec 18 2023 1 PM EST Room 1-224 [New Meeting Link...](#)

- TODO: Zoom is requiring passcode (even though it is in url).
- Production System Update
 - Calorimeter workflow is working. 1 input : 1 output
 - Event builder workflow is working. Many input: Many output
 - Queries catalog for all raw event files for a given run, and passes the run# into the job
 - Job *currently* uses a prebuilt set of filelists for each run.
 - For production, will need to create these lists in flight... whether on the submission side or on the job side TBD.
 - Cataloging of output files still TBD. Once done, will test event builder → calorimeter workflow
 - System handles both 1:1 and MANY:MANY data production. (By extension, 1:MANY and MANY:1 as well).
 - Both small scale tests O(1 run)
 - Single application":
 - kaedama.py --rule {DST_CALOR,DST_EVENT} [--nevents=0 --resubmit --unblock state]
 - Command line option to specify the workflow (rule)
 - Optional the number of events (defaults to 0=all)
 - Optional to resubmit (overwrites files in the catalog)
 - Optional to unblock (ignores existing production status)...
 - n.b. currently this option invalidates all previous production status entries for the given DST.
 - Workflow management
 - Plan is to run kaedama periodically via cron job.
 - When new inputs in the file catalog satisfy the specified rule, production jobs will be launched.
 - Outputs will be registered in the file catalog by the job, and picked up as inputs on the next pass
 - ~~■ Expect to test builder → calorimeter production chain by Monday.~~
- SDCC Planned Downtime 12/19 (starts 6AM, but jobs start draining from condor 5PM 12/18 evening... earlier on long duration queues...) Expect 9PM 12/19 back to operations.
- SDCC+sPHENIX [planning meeting](#) 12/20
- sPHENIX update (?)
- Discussion
- AOB
- Next meeting (first of the year) Jan 8th,

Meeting Mon Dec 11 2023 1 PM EST Room 1-224 [New Meeting Link...](#)

- Production System Update
 - Discuss DB table implementation, job state tracking and implementation
 - Currently testing at the scale of a single run...
 - Readiness for test at scale
 - Plan to get this into production now. Will install system under sphnxpro account. Will work w/ Sasha on implementing the rule to run event building.
 - Will need to extend at some point to handle the multiple output files of some jobs.
 - [Link](#)
- sPHENIX Update
 - Time to cleanup lustre (rucio)... Vincent returns to BNL tomorrow... will coordinate.
- ...
- Discussion
- AOB

Meeting Mon Nov 27 2023 1 PM EST Room 1-224 [Meeting Link](#)

- Simple test jobs working... i.e. logfiles are produced locally and staged out at end of job. Testing calorimeter production shortly.... may have an inconsistency between setting an initialdir and transferring (log) files at the end of the job. TBD.
- Propose to create a DB table to track job status, which will persist past the end of the job as a record of what was run.
 - At least one table per production
- Chris...
 - calibrations DB meltdown. s/t changed... timeout not handled and job dies. Timeout not interpreted correctly on client side (dies instead of retries).
 - Event builder ... progress on triggering copy script when file is closed. Also capability to limit number of events go into output files (rather than filesize). Useful for bookkeeping. Plan for events 1 to N, N+1 to 2N, ...
 - Files will be written locally, then moved once closed. Available for launching downstream jobs before event builder job has finished.
 - # of output files can now be predicted based on input file largest event number.
 - streaming readout detectors... output mgrs create multiple output files. ... can start cosmic jobs running under new event builder scheme.
- AOB

Meeting Mon Nov 20 2023 1 PM EST Room 1-224 [Meeting Link](#)

- Initial thoughts / work on real time production system
 - Try to develop a general system, but time is short... impose sPHENIX policies where needed
 - Simple, lightweight, rules-based, minimize load on DB servers
 - Provide sensible defaults for most use cases
 - Flexibility to ~~get into trouble~~ do what you want
 - Scratch the surface and you know you're using condor

- [slide](#)
-
- ...
- Open discussion
- AOB

Meeting Mon Nov 13 2023 1 PM EST Room 1-224 [Meeting Link](#)

Thoughts:

- Need internal queuing - we cannot feed more than 200k jobs into condor even if we have 4 submission hosts
 - Though maybe scrap that - if we put the submission into a cron, we can just fill the condor queue and stop when we reach its limit. And then submit more on the next pass. That makes it easier to interrupt this in case things go south (no need to clean up internally queued jobs)
- Likely event building and reconstruction jobs are so different that it'll be hard to wedge them into the same frame
 - Reconstruction jobs (global, calo, tracking job0, jobA, jobC) have a similar structure (maybe they are candidates for a common approach)
- What kind of DB tables do we need?
 - Need to copy over existing run DB to get the list of raw data files for each run
 - filenames need to be parsed to assign them to correct list
 - Location comes from file catalog (todo: enter raw data files to file catalog)
- First goal - run the calorimeter production
 - Implement what is needed (access to file catalog, creation of file lists, job submission, entering resulting files in file catalog, exporting and saving condor logs, tracking production status, identifying crashed jobs)

Discussion

- Monitoring: condor_q ... send logs to gpfs... failed jobs on hold... (condor... on evict ship logs...) What capabilities can condor bring to the table to make our life easier? Eventually log stash.
- Condor support from sdcc.
- Pathname convention...
- pyodbc

Meeting Mon Nov 06 2023 1 PM EST Room 1-224 [Meeting Link](#)

n.b. Clocks move back in NY on Sunday Nov. 5th.

- RHIC schedule
 - No earlier than 3/3/24
- Old Business
 - Sending output to different RSEs during panda run.
 - Use case: Large (data) files to lustre / small (log) files to GPFS for optimization purposes
 - sPHENIX would need to support space tokens... does not currently exist in sP rucio. Interact w/ Vincent to find if this is possible.
 - Is the rucio data (local site)? mover exposed to the user? (So that we write to lustre once).
- PanDA
- Rucio
- SHREK
 - Minor update. Soft/hard limit on the total size of the submission package.
- Production
 - Latest production over the weekend, to create combined event files as 2GB rather than 20GB, failed.
 - Short duration jobs run successfully interactively
 - PanDA task terminates after about 1-1.5 h on the farm... the job itself seems to disappear sooner...
 - Logged in to worker node to watch a job fail... nothing obvious....
 - But directory disappears while job is still running. (I thought that condor did not instantly clean up the working directory?) ~few minutes
 - “Simple” workflow (no processing) runs fine (20min sleep).
 - When condor kills a job... pilot should have enough time to do a cleanup (including updating the real time log).
 - HOw to access the condor logs? On the CE. (?)
 - Job ID on the CE.
 - We need the condor user log from the CE...
 - Direct use of condor avoiding the CE entirely...
 - REquires 2-3 interactive submission hosts...
- AOB
 - Likely need a separate queue for event combining
 - Special treatment for event builder jobs, e.g. one EVB job per node.
 - Production mailing list requested (sphenix-production-l)

Meeting Mon Oct 30 2023 1 PM EDT Room 1-224 [Meeting Link](#)

- PanDA
 - jobReport.json file
 - Contains the userJobMetadata.json file ... e.g. <https://sphenix-panda.apps.rcf.bnl.gov/job?pandaaid=2283298>
 - Inserted into the panda DB... so (a) propagates to monitoring pages, (b) is queryable by the panda team
 - No access to PanDA internal DB from the sPHENIX side. Options are to (a) scrape the monitoring pages, (b) pull from rucio, (c) store into sPHENIX db, (d) ...
 - Decision. sphenix will add metadata to its own DB.

- Need log datasets (let's call the "logsets") to go to separate RSE (gpfs) because they are lustre-inefficient. (Can panda handle writing datasets to different RSEs).
- Misc fixes : monitoring on metadata and Md5 checksum
 - Pilot config needs to be merged at each upgrade to retrieve sPHENIX changes. Need a unit test for Md5 sum vs pilot upgrades...
 - Date to freeze PanDA?
 - Authentication Tokens need to be updated by date?
 - Condor upgrade date (after run please). (At least authentication).
- Rucio
 - Version for common sphenix setup?
 - Heizenbug... Multiple warnings when using rucio, e.g.
 - ~~InsecureRequestWarning: Unverified HTTPS request is being made to host 'sprucio01.sdee.bnl.gov'. Adding certificate verification is strongly advised.~~
- SHREK
- Production
 - Complete raw event and DST_CALOR production with DB tag 2023p001
 - Unit tests verify that each run/dataset has a consistent set of files

| | | |
|--|------|--|
| test_events_100_the_event_collection_should_be_filled ✓ | 9% | |
| test_events_101_every_run_should_have_file_replicas ✓ | 18% | |
| test_events_102_the_name_of_each_file_must_conform_to_the_naming_convention ✓ | 27% | |
| test_events_103_the_lowest_sequence_number_is_zero ✓ | 36% | |
| test_events_104_the_highest_sequence_equals_nfiles_minus_one ✓ | 45% | |
| test_calor_200_the_calor_collection_should_be_filled ✓ | 55% | |
| test_calor_201_every_run_should_have_file_replicas ✓ | 64% | |
| test_events_202_the_name_of_each_file_must_conform_to_the_naming_convention ✓ | 73% | |
| test_calor_203_the_lowest_sequence_number_is_zero ✓ | 82% | |
| test_calor_204_the_highest_sequence_equals_nfiles_minus_one ✓ | 91% | |
| test_joint_300_the_event_and_calor_collections_should_have_the_same_number_of_runs ✓ | 100% | |

- DST_CALOR reproduction with DB tag 2023p002
 - 6 runs succeeded under shrek/panda. Sasha reports O(dozen) success when running by hand.
 - Calibration issue on one run...
 - Sasa will try to track down discrepancy
 - Few notes
 - PanDA has fallen back to registering files with Adler32 checksum rather than md5
 -
 - Access to logfiles is tedious. The payload.stdout and payload.stderr should be pulled out, named according to the run and segment they contain and put into gpfs.
 - The logfile tar balls are currently copied to lustre which is not suitable for this kind of file
 - Do we need all the python stuff in those tar balls?
- AOB

Meeting Mon Oct 23 2023 1 PM EDT Room 1-224 [Meeting Link](#)

- PanDA
 - New feature to allow users to specify workflow name in pchain, deployed on the PanDA prod instance
 - Follow-up ...
 - jobReport.json file, if produced by job payload, is included in the job log file tarball that pilot uploads to Rucio. To display the contents of it on PanDA monitor, PanDA monitor needs to be able to download "rucio download" from Rucio.
 - Torre - Is this confirmed fact? I thought the json file was put into a database record, not just dumped as a file somewhere requiring Rucio to access the content
 - Torre - Tadashi has confirmed that the job report json is put in the panda DB and so is easily queryable. The 'other half' of the job of creating a custom job report json for sPHENIX's needs is to create a PanDA monitor module to present that info in the monitor as sPHENIX sees fit, at the job level and/or the aggregate summary level.
 - Token support for job submission* : tested successfully in the past, next step is to maintain valid tokens and make them accessible from harvester pods in OKD. Low priority so far, but to build PanDA images on newer OS (e.g. AL9), we need to move to newer HTCondor release, which supports token only submission. Will resume the thread with SDCC folks ...
 - Discussion of scheduling upgrade vs RHIC run.
 - Discussion of event accounting in the jobs. Propagation of information out to panda monitoring? rucio? sphenix database?

TODO – add to json file ... jobReport.json
 - Rucio
 - SHREK
 - Production
 - 10/09 production files produced in PanDA are "locked"... unremovable by sphnxpro account.
 - Created RHIC2023-sPHENIX-JUNK-DRAWER and moved datasets there
 - Question: Does PanDA protect the datasets it creates with... a replication rule? Or other mechanism? If so... is this the behavior we want for sPHENIX?
 - 10/14 Production...
 - identified 2 failure modes.
 - 14 runs which produce only single prdf output file. 21663 - 21719 (consecutive runs in the run list). These are reproducible.
 - 13 runs where there are gaps in the list of output files...
 - Fun4All produces N files with sequence numbers 0 to N-1. **We observe M<N files cataloged in rucio.**
 - Good news is this is recoverable. When I rerun the workflow on single run, all files are created/cataloged. (Bad news... difficult to debug b/c not reproducible).
 - The scratch space on lustre where the files are created is empty after the job... which is odd because...
 - The script does not remove the files... PanDA does this when it moves them to RSE.
 - The script only does a "mkdir"... no rm... into the scratch space.
 - (possibly I cleaned up by hand...)
 - ... and at the end of the job, ln -s so that PanDA can discover the files...



- Fun4All is trusted code... there are no gaps expected.
- 10/23 Production
 - Reproducing the calorimeter data ... and watching jobs / scratch space...
 - Class 1 above is reproduced. Jobs are ending (failing) with only a single PRDF output. (But the scratch space is not cleaned up...)
 - Sasha– probably an issue in the data and / or fun4all ... TBD.
 - Class 2 above is not reproduced. I do not (yet) see any runs with a gap in the PRDF output list. Error *detection* has been implemented in the workflows... Chris added verbosity flag so that Fun4All will output list of files it creates. I can capture / compare to list of files that will be presented to PanDA for archival / cataloging.
 - Testbed to speed up development
 - Use case? (1) Marshal raw event inputs into DS and run event building + ... workflow? Or (2) run workflow on existing built event datasets?
 - AOB
 - Event builder questions. Is it possible to increase the logging level so we can record when an event is processed and how much data is read in the condor logs?

Meeting Mon Oct 16 2023 1 PM EDT Room 1-224 [Meeting Link](#)

- PanDA updates
 - Feature request implemented in PanDA that allows user to give a meaningful workflow name in pchain. Being tested in the test PanDA instance. Some monitoring issues still to sort out but one can try it using the latest panda-client against the test PanDA instance now.
 - One potential problem is being discussed w.r.t. the newly created PQs for GPFS RSE: pilots can write to GPFS RSE now, but some configuration may prevent pilots from getting inputs from ANY RSEs, under investigation.
- Rucio updates
- SHREK updates
- Production updates
 - First some bookkeeping. Proposed to collect results from official productions under rucio collections:
 - RHIC2023-sPHENIX-CALOR-EVENTS ... combined events
 - RHIC2023-sPHENIX-CALOR ... calorimeter DSTs
 - There will be only one such collection for each type of production in a given year.
 - Special containers (e.g. for testing, commissioning, etc...) would have additional information in their name, such as a build tag, database tag, the word “test” etc...
- Beneath the collections are datasets, with each dataset containing the output files from a given run



```
$ rucio list-content --short group.sphenix:RHIC2023-sPHENIX-CALOR-EVENTS
group.sphenix:group.sphenix.rn00022046-CALOR-ana379-20231009-173452_000_EventCombine_0
...
```
- Job submission builds an input dataset consisting of physical filelists. Name of the dataset is unique to the production for each run, e.g. **group.sphenix:rn00024316-CALOR-ana381-2023p001**. Encodes the run number, production type, software build tag, database tag.
 - IF we restrict ourselves to using the submission script...
 - THEN for a given run, only one instance of a workflow can be run with a given build / dbtag.
- Any attempt to submit same run with same build and dbtag will result in a warning issued to the user.
- Any reproduction *requires* a cleanup of the file catalog (and disk).
- Setup and ran calorimeter production over the 153-good-runlist.
 - 09 October 2023
 - These are currently mapped onto the official collections
 - Output filenames not ideal...
 - 13/14 October 2023
 - Revised naming conventions (DB tag replaces panda task ID in filename, calorimeter output is a DST_calor*.root file)
 - Most runs processed in ana.381. (Presuming ok?)
 - Runs 00021798,00020869,00022948,00021739,00023696,00021199,00023605,00023605 produced against build ana.380. (Reprocess?)
 - Replace official files. Create links.
- Soft links will be created under gpfs to provide for simple (ls) lookup of data files
- TODO:
 - Cleanup the Oct 9th production
 - Move Oct 13th datasets into the official collections
 - Create softlinks under gpfs
- sPHENIX QA for calorimeters... adding in additional (non gold) runs. Rucio filecatalog integration. ... consider items to ship with the job report for access through panda monitoring.
- AOB

Meeting Mon Oct 02 2023 1 PM EDT Room 1-224 [Meeting Link](#)

- PanDA updates

- Created new PQs for writing outputs to GPFS area, tested successfully on the test PanDA instance.
 - BNL_OSG_SPHENIX (lustre ...rse=BNL_PROD3)
 - BNL_OSG_SPHENIX_GPFS (gpfs rse=BNL_PROD2)
 - Expecting to switch next week
- Correction/clarification on users accessing PanDA monitor: for users to use PanDA monitor, they do need to register an account on PanDA IAM. Authentication is done through their home institute ID providers though. The first time a user tries to access the PanDA monitor, after the initial authentication step, the user will be prompted with an IAM registration page, just follow the page to register.
- Rucio updates
- Production updates
 - 164 jobs submitted, expecting to run all day
 - Running from the new ProdFlow github repository
 - Sasha had to update his credentials...
 - Clean directory... setup file assumes shrek is in subdirectory... pain point. Symlink used to solve.
 - (Added path issue to TODO list for shrek).
- Workflows
 - New repository for sPHENIX production workflows
 - <https://github.com/sPHENIX-Collaboration/ProdFlow>
 - Purpose is to collect and tag the workflow definitions
 - Fun4All macros to used from other github controlled directories
 - Development to be done on feature branches
 -
 - Naming workflows
 - No documented way to pass a user string for the workflow name via pchain (Wen noted that it can be passed as a task parameter when submitting via the python api...)
 - Refactored the calorimeter workflow to use dataset input rather than pfnlist.
 - <https://sphenix-panda.apps.rcf.bnl.gov/idds/wfprogress/>
 - Propose convention rnXXXXXXXXX-TYPE-PROD_TAG(-DATETIME)
 - User-defined string may be implemented in pchain soon-ish.
- Prompt processing
 - Logged into a worker node... I am able to run rucio commands (e.g. rucio ls, rucio list-file-replicas, etc...) for catalog queries... expect that registering a file will work fine. TBD.
 - Quick option: shell out and use python script
 - Best option: extend C++ API
- Rucio as a file catalog
 - File erasure / reuse
 - The sphnxpro account needed increased permissions to erase files
 - The magic command (one of three possible ways one can divine from the rucio python API documentation) is to “set_tombstone” on a file.
 - Apparently it is treated differently than other data identifiers (datasets and containers) which can be erased from the command line ‘rucio erase’... which just sets the lifetime flag in the metadata to 24h.
 - Once a tombstone is set, the file entry is gone from rucio within a minute or two.
 - Once erased, the filename can be reused and associated to a differently-check-summed file.
 - Erasing a dataset does not seem to effect the attached file DIDs (e.g. it is not a recursive operation)
 - Rucio is operating as a file catalog. Removal of files is the responsibility of the production team.
- AOB

Meeting Mon Sep 25 2023 1 PM EDT Room 1-224 [Meeting Link](#)

- Updates (PanDA, rucio, shrek/donkey/workflows...)
 -  donkey v2 test coverage
- Possible topics for round table discussion
- <https://sphenix-panda.apps.rcf.bnl.gov/>
-
-
- Conventions
 - Filename conventions
 - Where to stage raw data onto lustre
 - Where to write event builder data on lustre
 - Dataset, scope and/or container naming (raw data, intermediate jobs, final output)
 - metadata to fill in datasets / files
 - e.g. run_number...
- Where to run various scripts / how to pass signals
 - Scripting needs to accomplish
 - Raw Data transfer,
 - Create rucio dataset for each run
 - Register files with rucio
 - Close dataset
 - Dispatch appropriate workflows
- Prompt Processing of Event Builder Outputs

My understanding...

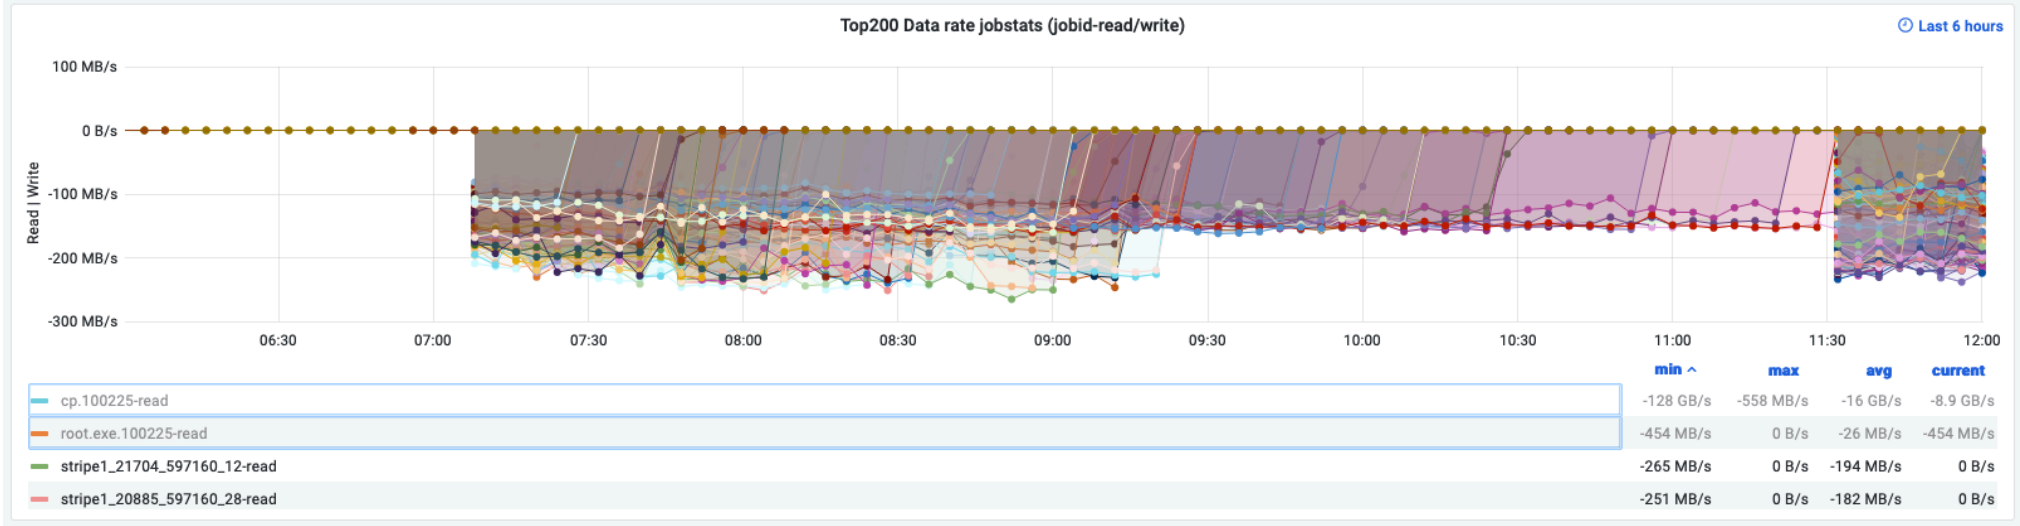
 - JEDI periodically monitors datasets / dispatches child tasks as data becomes available
 - pilot would be constrained to registering files at end of job...
 - could we use the rucio api to register the files as they are produced? i.e. from within Fun4All...
 - Example c++ wrapper for getting pfn from lfn ... <https://github.com/BNLNPPS/sancho>
 - Could / will extend to register files as well...
 - Any issue with making http requests from worker nodes?

- Alternatively, can rucio command line client be called from the worker nodes?
- AOB
- Verify we can rewrite *files* after an erase in rucio
- Workflow and/or task name should carry the run number and production type
 - Better traceability in the workflow monitor
- Shift focus to defining production workflows...
- Rescoping initial startup... production manager will kick off productions using a “good run list” with some periodicity.
 - Focus shifts to delivering workflows kicked off with shrek rather than (automated) donkey application
- Alm for NOv 15th demo
- Tadashi reported a few days ago: “Discussion about prompt processing in sPHENIX. Event builder jobs are long and produce many files. The idea is to start downstream jobs as soon as each file is produced while the EB job is running. Agreed to use JEDI’s capability to periodically check the contents of open datasets and generate jobs when new files are there. Jason and Chris will improve EB jobs to register files to Rucio while they are running.”
- Priority of tasks
 - Deliver sPHENIX workflows
 - MOve existing workflow definitions from shrek repo into the sPHENIX repo
 - PRovide assistance in creating the next / final set of shrek workflow definitions from sPHENIX macros
 - Optimize event builder / job creation using JEDI’s capabilities to montior datasets for new inputs
 - Automate submission of workflows on transfer
- Real time logging... cleared to turn this on for the full set of workflows.... (time to drink from the firehose...)

Meeting Mon Sep 18 2023 1 PM EDcT

- Face to Face meeting?
 - Next Monday same time? ~~Doodle poll to confirm.~~ Remote connection will be available.
- BNL_PROD / rucio storage elements
 - BNL_PROD, BNL_PROD3, ... panda usage ...
 - BNL_PROD and 3 are both setup as zfs volumes, same parameters
 - lustre migration we used prod2... either one can be used
 - don’t change anything continue to use 3.
 - histogram files will need their own area... want them under gpfs... need RSE for rucio.
 - direct posix access... how to cleanup? sPHENIX will handle cleaning up the temporary files.
 - Quantity of processing is one run. Once completed, cleanup is triggered. (Human QA at least from the start).
 - BNL_PROD2: /sphenix/data/data03/sphnxpro/rucio
- PanDA
 - Update of Prod instance postponed, more components have newer releases recently, will upgrade them altogether, after more testing
 - Ongoing discussion on getting more info about “lost heartbeat” jobs : pilot catches kill signal and save job logs to Rucio before exit, to be tested...
 - Optimization of event builder jobs : prompt triggering of secondary jobs that consume outputs from long running EB jobs. To collect more details and discuss with relevant experts.
 - Pilot registeres files to rucio once the user job has successfully completed
 - Request is to register during the execution of the job.
 - ... on kill signal, flush mesage to logstash?
 - logstash / elastic search retention time? 1wk?
 -
- Rucio
- SHREK / donkey application
 - donkey listen --dbfile messages.db --time 5m --user uname --pass pwd
 - All messages are saved in the “pending” list stored in the messages.db file
 - donkey dispatch --dbfile messages.db --rule raw-events --actor actors/dispatch.sh --regex 'r"(\w)+EVENTS-(\d+)"' --scope group.sphenix --event close
 - Messages which match the regex (and are in scope and match the specified event) are dispatched using the specified script.
 - Multiple rules may be defined
 - Any message that does not match is returned to the “pending” state, or “dropped” (e.g. we leave unmatched “closed” events in pending and move others to “dropped”.)
 - “listen” and “dispatch” have a reasonably comprehensive set of tests... *currently failing b/c I am implementing the command line parser above*

```
tests/test_donkey2_dispatch.py::test_donkey2_dispatch_default_rules FAILED
tests/test_donkey2_dispatch.py::test_donkey2_dispatch_matching_rules FAILED
tests/test_donkey2_dispatch.py::test_donkey_dispatch_class FAILED
tests/test_donkey2_dispatch.py::test_donkey_dispatch_class2 FAILED
tests/test_donkey2.py::test_donkey2_import_dataset PASSED
tests/test_donkey2.py::test_donkey2_init_dataset PASSED
tests/test_donkey2.py::test_donkey2_init_dataset_collection PASSED
tests/test_donkey2.py::test_donkey2_add_pop_dataset_collection PASSED
tests/test_donkey2.py::test_donkey2_update_dataset PASSED
tests/test_donkey2.py::test_donkey2_add_remove_dataset_collection PASSED
tests/test_donkey2.py::test_donkey2_add_to_unkown_list PASSED
tests/test_donkey2.py::test_donkey2_multiadd_dataset_collection PASSED
```
 - Next step will be to kick off the calorimeter analysis by “closing” a dataset. Then will add the capability to build/close dataset and process the calorimeter workflow.
- Lustre Scale test -
 - Chris provided the software - <https://chat.sdcc.bnl.gov/bnl/pl/9seqh3ipzpyf3mtb8j9zafmi9y>
 - Snapshot of initial tests - (~ 100 - 250 MB/s - stripe size = 1)



- AOB

Meeting Mon Sep 11 2023 1 PM EDT

- Calorimeter Production Update
 - Failed runs pass. Full workflow passes. “Heatbeat” issue resolved. Run # needs to be part of the output dataset name.
 - Heartbeat issue was jobs were written to local /tmp disk and overflowing (solution was to write output direct to lustre).
 - These types of issues are difficult to debug... required a lucky (educated?) guess
 - “pilot” is the job which runs (and user job runs w/in the pilot). So when the job is terminated by condor, pilot is terminated.
 - condor logs wiped,,, needs to be persisted. (work directory is removed / logs removed).
 - Condor “save on error”. Xin will interact w/ Chris H.
 - Revisiting naming conventions.
 - Pushed changes in the sphnxpro shreksys to github
 - Large number of runs (not in the golden run list) remain to be produced / updates on the macro side TBD.
- PanDA updates
 - Updates to DB schema and pilot wrapper scripts tested on the test instance, to be applied to the prod instance (this afternoon?)
- Countdown... 10 weeks to sPHENIX production dress rehearsal ...
 - Rucio Integration
 - Transfer from counting house to lustre
 - Register raw data with rucio / create and close datasets
 - python tool implemented to streamline process
 - needs work to automate the process
 - DonkeyApplication
 - Receives activeMQ message on dataset close
 - Launches arbitrary user scripts, e.g., to start a shrek job
 - Advanced prototype stage.
 - Was working, but currently throws exceptions.
 - Monolithic application which monitors active MQ in one thread, dispatches workflows in a second thread, and implements a command menu in the third.
 - Requires synchronization when accessing the shared cache of received messages.
 - No advantage to multithreading so...
 - Split into three (much simpler) applications which will be run sequentially...
 - donkey_listener.py [implemented / tested] – monitors active MQ and caches messages of interest into a lightweight DB (pickleDB). Straightforward to migrate to a real DB if/when needed.
 - donkey_show.py [implemented / tested] – tabulates the current set of messages in the DB
 - donkey_dispatch.py [expected to be ready this week] – examines the DB for messages that match a user-defined set of rules, and dispatches work to the corresponding scripts. Updates DB entry that the work has been dispatched.
 - Sancho
 - C++ interface to rucio’s REST api... but just the part which translates logical to physical filenames.
 - In principle full REST api could be exposed to C++ code in a straightforward way, if there were demand.
 - Relies on BOOST libraries for web interface, certain string manipulations, regex-es, json and ini file parsing, etc...
 - Lustre Testing / Scale
 - ... after QM ‘23
 - Chris will hand off macros to SDCC

- Rucio support for multiple copies... yes. On multiple RSEs. For “hot” files.
- Rucio updates?
- AOB

Meeting Mon Aug 28 2023 1 PM EDT

- Calorimeter production debugging
 - Access granted to worker nodes / tested
 - PanDA jobs not running in “expected” directory.
 - “pwd” reports /home/tmp/sphenix_od9GmgXI/PanDA_Pilot-2248570/workDir

- /home/condor/local/sdcc/execute/dir_296217/tmp/sphenix_od9GmgXI/PanDA_Pilot-2248570/workDir
 - pilot log issues warnings that it cannot find the /home/tmp/... directory
 - atlas queues work... so Xin will take a look.
- Lustre Testing
 - Resume after QM
 - Chris will pass code to sdcc
- Data Management / Integration rucio
 - Quick and dirty rucio access from c++
 - Returns PFN given the LFN
 - if its in dcache... prepend prefix
 - Uses a (beta, not actively developed) RESTful client api. (Would be better to use libcurl directly)
 - ...
- AOB

Meeting Mon Aug 21 2023 1 PM EDT

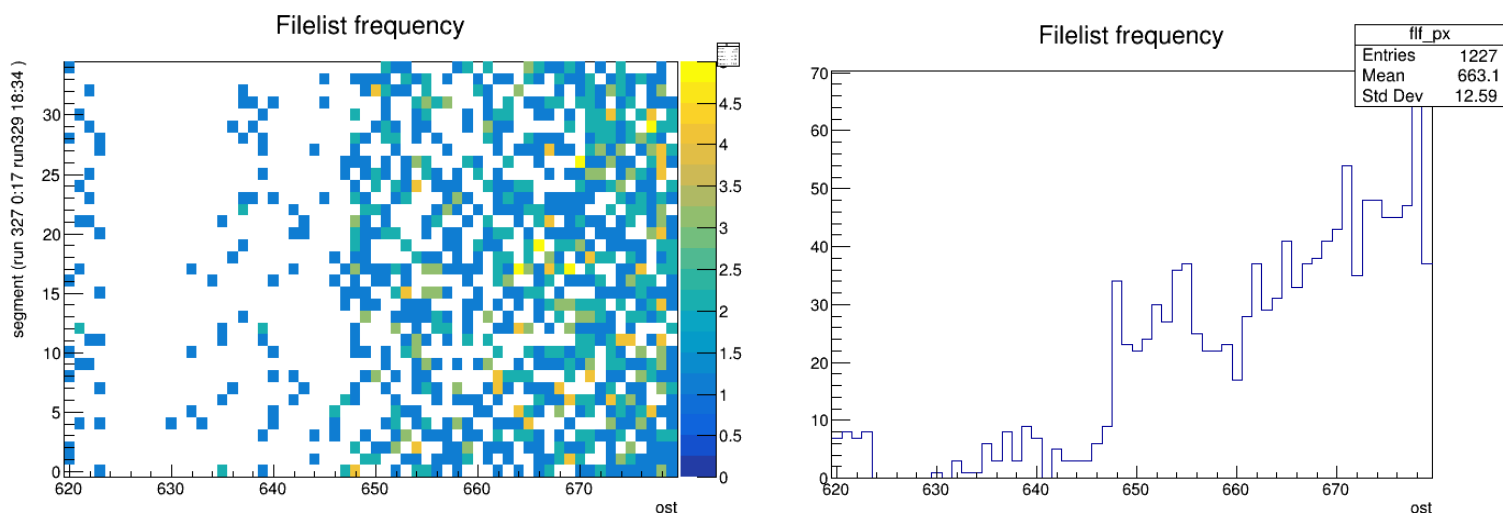
- Calorimeter Production Update
 - Llst of 154 “good” runs, 62 failed w/ lost heartbeat message
 - Partially successful Enterprise, beginning with panda ID 1701...
 - Chris peeking a log files... Issue w/ 2 or 3 osts?
 - Is it possible for pilot to have a wrapper to sleep after a crash is detected so that we have time to cleanup?
 - Harvester is terminating the jobs when pilot no longer returns a heartbeat
 - ...
- Data Management Discussion / Integration rucio
 - Develop a c++ interface layer to the functionality of rucio ... (“sancho” or “rocinante”)
 - Option 1: exec shell command and capture output ...
 - Option 2: boost.Python to integrate the rucio python API (rather... map the rucio app
 - Option 3: use the rucio RESTful api
 - Delete (rucio-erase) *may* allow to reuse a DID... i.e. it clears out the protected meta-data fields (eg checksum)
 - should cover the requirement of being able to reuse filenames
 - Needs testing
- AOB

Meeting Mon Aug 14 2023 1 PM EDT

- Important time scales. Site visit 1wk. Next run begins 20wks. (Dress rehearsal ideally 15 wks).
- New strategy Stress testing lustre
 - Lower overhead / direct submission of macro(s) to condor
- Calorimeter workflow / 1h limit / heartbeat error
 - Definitely a 1h time limit issue
 - Correlated with a (panda reported) increase in maximum PSS
 - Error msg indicates condor terminates the job
 - ... and *eats* the log files ...
 - ... smells like a memory leak
 - Will launch a job and post so the job can be monitored (and run interactively as well)
 - Does Pilot capture CE submitted to? What *local* condor job ID is.
- sPHENIX data management / cataloging / integration discussion
 - Naming conventions (scopes, filenames, etc...) adopt from phenix. Sasha should take the lead.
 - Feature requirement on Rucio... Able to overwrite files with duplicate filenames...
 - For *preparing* for production / testing / not standard procedure (but should not need to jump through many hoops either).
 - run few% of events for QA then run full sample, OR
 - production is started and issue discovered. Need to patch and restart.
 - Do not want to generate new tag on each iteration.
 - *Procedure exists* will be summarized at future meeting
 - Vincent: A procedure to reuse filenames and update metadata exists. This procedure was used by the ATLAS operator to fix file inconsistencies. Need to gather information regarding this procedure, and the outcomes will be presented at a future meeting.
- AOB

Meeting Mon Aug 7 2023 1 PM EDT

- Stress Tests
 - Disks on harvester filled due to large size of pilot logs.
 - Consequence of running stress tests with file lists (pfnList) rather than rucio datasets
 - Disk space increased (and limiting the size of the logs raised as a possibility). Cronjob modified, in addition to cleaning up pilot job files older than 3 days, it will also aggressively delete pilot stdout files if space usage reaches 90%.
 - Given end of run... future tests constrained by need to generate MC datasets for QM23 in sept. 2k cores may be reserved for testing.
 - Reworking the tests to
 - Read in data only (no real processing)
 - Small test on Friday 1k jobs... succeeded
 - Second small test failed... (understood... stray character in job description file...)
 - Balance the input file lists across lustre
 - Issue is we only have two runs...
 - Limit number of concurrent processes hitting each ost



- Monte Carlo production is constraining how the real data is being distributed. May need a separate rawdata pool.
- Production
 - Heartbeat?
 - 18th of August... production scheme demonstration. 11 days.
- AOB

Meeting Mon July 31 2023 1 PM EDT

- Stress Test
 - Submission scripts setup for *two tracking tests*:
 - reading raw event data from stripe1
 - reading raw event data from stripe5
 - TODO: same event file inputs
 - submit in 5k groups, up to 50k tasks for each test
 - expected duration of 1h at job0, 3h at jobA
 - Timeframe... tomorrow after current production jobs drain out... 10AM start time...
 - Restart of negotiator done / condor fine
 - Last Friday's issues due to decommissioned HW that condor thought was still active
 - PanDA IDs from task ids...
 - Plan is for two minitests today... full tests tomorrow @ 10am ish.
 - Event file distribution?
 - Send file list(s)
- PanDA updates
 - No major issues
- Calorimeter production ...
 - Anything needed to help migrate to PanDA?
 - Commissioning few runs / day
 - Naming conventions for DS and scopes
- AOB

Meeting Mon July 24 2023 1 PM EDT

- sPHENIX calorimeter workflows
 - Calorimeter workflow (almost) running under sphnxpro account
 - run_calorimeter.sh script submits the workflow
 - Builds input file list (pfnlist) from raw event files in a given directory for a given run number
 - Workflow in three sequential tasks:
 - EventCombine – event builder runs in single job, outputs dataset containing assembled event files
 - panda<taskID> is appended to the output filename to guarantee uniqueness
 - CalorimeterChain – task running the calorimeter reconstruction over the N files created by the event builder
 - panda<taskID> is appended to the output filename to guarantee uniqueness
 - MakeLinks – single job which creates a directory on lustre, and builds links between the sphenix-preferred filename and the panda<taskID> obfuscated name... e.g.

```
$ cd shreksys
$ ln -s workflows/rhic2023.AuAu200.calor/run_calorimeter.sh .

$ source run_calorimeter.sh 9245 \
    /sphenix/lustre01/sphnxpro/commissioning/emcal/beam \
    /sphenix/u/sphnxpro/shrek/run00009245 \
    NEVENTS

$ pbook show limit=3
```

| JediTaskID | ReqID | Status | Fin% | TaskName |
|------------|-------|--------|--------|---|
| 1008 | 1008 | done | 100.0% | group.sphenix.test.calor-calib-20230724-130657_002_MakeLinks |
| 1007 | 1007 | done | 100.0% | group.sphenix.test.calor-calib-20230724-130657_001_CalorChain |
| 1006 | 1006 | done | 100.0% | group.sphenix.test.calor-calib-20230724-130657_000_EventCombine |

```
$ rucio ls group.sphenix:*calor-calib* --filter type=DATASET
```

| SCOPE:NAME | [DID TYPE] |
|--|------------|
| group.sphenix:group.sphenix.test.calor-calib-20230724-130657_000_EventCombine_0.5598 | DIDType.DA |
| group.sphenix:group.sphenix.test.calor-calib-20230724-155057_000_EventCombine.log.5616 | DIDType.DA |
| group.sphenix:group.sphenix.test.calor-calib-20230724-155057_000_EventCombine_0.5617 | DIDType.DA |
| group.sphenix:group.sphenix.test.calor-calib-20230724-130657_002_MakeLinks_manifest.dat.5611 | DIDType.DA |

```
| group.sphenix:group.sphenix.test.calor-calib-20230724-130657_000_EventCombine.log.5597 | DIDType.DATASET |
| group.sphenix:group.sphenix.test.calor-calib-20230724-130657_001_CalorChain.log.5604 | DIDType.DATASET |
| group.sphenix:group.sphenix.test.calor-calib-20230724-130657_001_CalorChain_0.5605 | DIDType.DATASET |
| group.sphenix:group.sphenix.test.calor-calib-20230724-130657_002_MakeLinks.log.5610 | DIDType.DATASET |
+-----+-----+
$ cd /sphenix/u/sphnxpro/shrek/run00009245
$ ls -l
calor_emcal-00009245-0000.prdf ->
/direct/sphenix+lustre01/sphnxpro/zfs/rucio/group/sphenix/f0/56/calor_emcal-00009245-0000.prdf.panda10
...
calor_emcal-00009245-0023.prdf ->
/direct/sphenix+lustre01/sphnxpro/zfs/rucio/group/sphenix/b0/f4/calor_emcal-00009245-0023.prdf.panda1007
```

- Next steps... automation. tpc workflows? other?

- Lustre status / readiness for next stress test
 - Finished on Friday. Determine # needed for test. Few days notice. Plan to use commissioning data (which will be evenly distributed). Expect multiple processes accessing same data file(s).
 - Test anticipated... distribution of raw data...
 - MC data distribution (may be less significant).
 - Big files... stripe 1 vs stripe 5 ... → use old test files
 - Planning for July 31st for kickoff of next test.
- PanDA:
 - Prod instance updated with the fixes to support regex in output files, DB schema updates and several other config changes
 - Reformatting the HTCondor_Job_ID env variable inside pilot for better lustre I/O trace.
- AOB

Meeting Mon July 17 2023 1 PM EDT

- Lustre scaling tests / path forward
 - Last week’s scaling test showed less than desired performance reading files from lustre.
 - Only 9 servers participating in IO / jobs starved reading data.
 - Did not produce enough writes to stress the system / reproduce the issue with the first scaling test.
 - mdraid → zfs being migrated (need >> 9)
 - note: jobs are configured to run for 2h in job0 tracking, write data, run for ~6h in jobA tracking. Same or greater amount of IO expected from when we caused problems ~month ago.
- PanDA updates to support sPHENIX workflows / file naming conventions
 - Regular expressions implemented to define which files are added to a dataset from a given job
 - Single job can add an arbitrary number of files (compared w/ panda wildcard matching which adds a single tar file).
 - Responsibility for unique naming of files is on the user side.
 - Convention adopted is to append the panda task ID to the output files when a job finishes / strip it off when the next job in the workflow commences
 - The test instance should be working / confirm and should be able to run calorimeter commissioning workflows.
- AOB

Meeting Mon June 12 2023 1 PM EDT

- PanDA / pilot / shrek support for “regex” output datasets
 - Regex keyword introduced by Tadashi.
 - Files matching the supplied regular expression (python’s “re” module) will be added to DDM as specified (no mangling of names, no hiding beneath a tarball, etc...).
 - Implemented to satisfy requirements of sPHENIX I/O modules
- PanDA to be updated to save (use?) md5 checksum per sPHENIX request
- Discussion of Friday’s stress test
 - Tracking Workflow:
 - 400x120 events/job 40k jobs (job0 → jobA)
 - Submitted in 2x5k increments.
 - Simulated event + (random noise) “real” events.
 - Lustre issues
 - Able to read 160GB/s
 - Writing back to lustre... kaboom. 15 / 40 servers hung up.
 - job 0 memory intensive / job A slower but less memory intensive
 - Steady state 12GB/s incoming raw data ... but spikes in rate an issue...
 - expecting 70 GB/s steady state once the jobs are processing... +/- fluctuations
 - High loads noted by Chris H. (Significant # of threads... multithreaded tracking algo... # of threads need to be tuned...)
 - *Status – not yet ready to accept anything at scale*
 - If we have to... is there a way to throttle (PanDA restricts # and/or type of concurrent jobs).
 - Throttle at condor level (resources)? Using shrek? Need to decide (a) conditions and (b) responsibility for throttling
 - Direct write to lustre during the job? Would require change to pilot model? For testing we can simply output filelists through panda/pilot. And write direct to / read direct from lustre.
- Real time logging status
 - being tested on the test instance
- PanDA / pilot certificate
 - PanDA pilot cert expired over the weekend, a new cert was generated

- CEs mapping has been updated, testing ongoing, once successful, will update the prod instance.

- Updates
- AOB

Meeting Mon May 22 2023 1 PM EDT

- sPHENIX production workflow
 - Demo of the calorimeter, global and tracking reconstruction chains implemented
 - Single MC file used as input along side one “sequence” of raw event files.
 - Should be ready for an “end-to-end” test by Wednesday
 - Close dataset → panda submission → output results
 - Large scale stress test next priority
 - Where are intermediate files read from during a pchain workflow?
 - ... /home/temp/atlas_<blah>/PanDA_Pilot-<blah>/<zipped tarball containing my file>
- CE does not guarantee original clustering from harvester...
 - causes some inefficiency...
 - direct submission is the way to go (harvester on VM skipping CE loop). But...
 - 40 jobs / s starting performance currently ...
 - This is sufficient for sPHENIX currently.

- Updates
- AOB

Meeting Mon May 15 2023 1 PM EDT

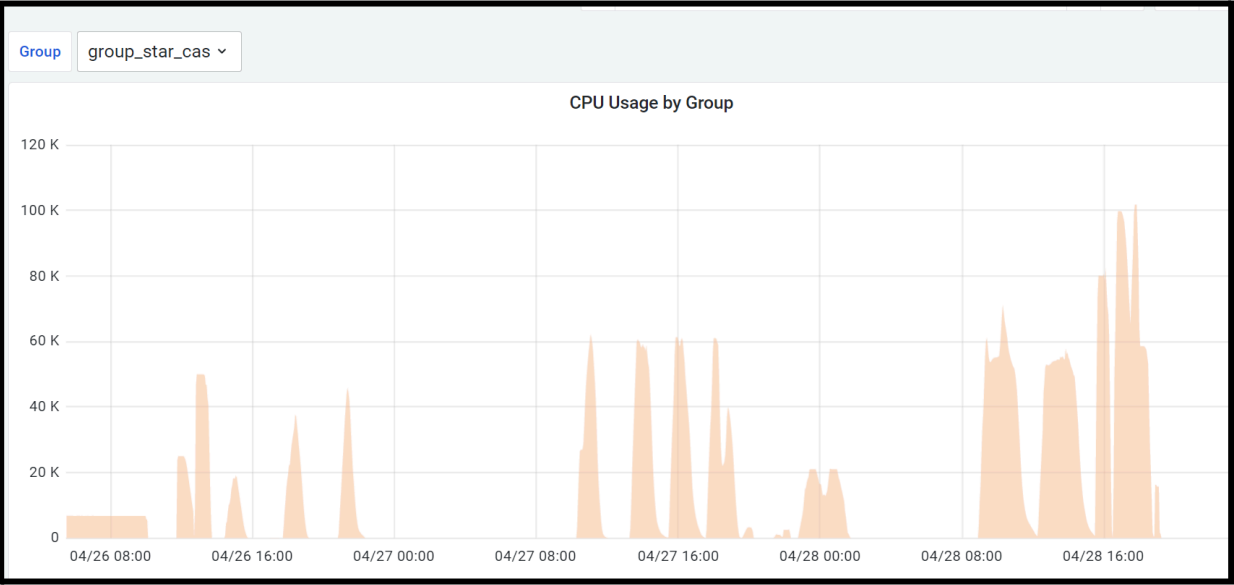
- Rucio Downtime tentatively scheduled for Wednesday morning 8AM-12PM
 - Upgrade to version 1.31...
- PanDA downtime to be coordinated w/ Rucio downtime
 - Renew token and host certificate
- Workflow status
 - Bash script in place to build “file lists” for task input:
 - (optional) MC hits file,
 - (optional) vertex file
 - List of Raw Event Files,
 - An event range.
 - PanDA tasks submitted with pfnList option... creates a job for each file list provided
 - pfnList is a prun feature. pchain can only make use of pfnlist IFF the filelist is present in the task’s payload. shrek will ensure that file list is available for pchain.
 - Test Calorimeter workflow in progress, Tracker workflow up next...
- Can PanDA submit in clusters?
 - Harvester currently does this... size=~400. Changes depending on available #
 - But subcluster broken at CE
 - Path forward... harvester on CE for local submission
- Updates
- AOB

Meeting Mon May 08 2023 1 PM EDT

- 100k Job Stress Tests
 - **Success** once job payload (including pilot) was *extremely light*. (b/c hardware was oversubscribed).
 - Could not run 4 pilots on one core...
 - Quantify resource load of the pilots?
- Workflow Status
 - Filelist scheme for job submission
 - sPHENIX processing will map event ranges onto single jobs
 - Each event range spans across multiple files in a run. And a file can be visited multiple times.
 - Initial task (processRun) in the workflow will generate a file list for each event range in the run
 - File lists become the output dataset of the processRun task, and feed in as input to the calorimeter and tracking chains
 - Prototype of the processRun task is implemented
 - Few tests... will likely need to build event range/file lists as a pre-processing step and either pass them directly to the jobs, or build them as a dataset
- Time scale
 - PanDA production instance... needs to have certificates renewed / short down time
 - 1-2 hours to renew / restart
 - Expected this week or next
 - Initial startup w/ raw condor
- Updates
- AOB

Meeting Mon May 01 2023 1 PM EDT


- 100k Job Stress Test
 - Started Wednesday 04/03
 - 25k->50k->75k->100k
 - Test jobs : 30~60m sleep jobs
 - Multiple tasks defined : 5k jobs per task → multiple processes generating jobs in PanDA/Jedi
 - 6~7 harvester pods, 5 CEs



- Several tunings done along the way, on both PanDA and CE/Condor sides
- Preliminary results :
 - Hit plateau around 60k while running the whole chain (PanDA/Harvester → CE → condor)
 - A lot of idle jobs on the CE :

```
[sphenx02] /usatlas/u/xzhao > date; condor_status -submitter | grep spce0[1,3,4,5,6]
Fri Apr 28 09:45:52 EDT 2023
group_sphenix.mdc2.sphenxpro@bnl.gov      spce01.sdcc.bnl.gov      16623      3373      0
group_sphenix.mdc2.sphenxpro@bnl.gov      spce03.sdcc.bnl.gov      6653      13460     0
group_sphenix.mdc2.sphenxpro@bnl.gov      spce04.sdcc.bnl.gov      7631      12075     0
group_sphenix.mdc2.sphenxpro@bnl.gov      spce05.sdcc.bnl.gov      9255      10795     0
group_sphenix.mdc2.sphenxpro@bnl.gov      spce06.sdcc.bnl.gov      14607     5515      0
[sphenx02] /usatlas/u/xzhao > date; oc get pod -o name | grep harvester | grep -v maria | xargs -i oc
dor/etc/condor_config; export PATH=$PATH:/data/condor/condor/usr/bin/; hostname; condor_q -nobatch |
Fri Apr 28 09:45:57 EDT 2023
n1-harvester-0
Total for query: 13953 jobs; 0 completed, 0 removed, 6415 idle, 7538 running, 0 held, 0 suspended
n2-harvester-0
Total for query: 13841 jobs; 0 completed, 0 removed, 6527 idle, 7314 running, 0 held, 0 suspended
n3-harvester-0
Total for query: 14083 jobs; 0 completed, 0 removed, 6390 idle, 7693 running, 0 held, 0 suspended
n4-harvester-0
Total for query: 13538 jobs; 0 completed, 0 removed, 5988 idle, 7550 running, 0 held, 0 suspended
n5-harvester-0
Total for query: 14440 jobs; 0 completed, 0 removed, 6318 idle, 8122 running, 0 held, 0 suspended
n6-harvester-0
Total for query: 15410 jobs; 1 completed, 0 removed, 7032 idle, 8377 running, 0 held, 0 suspended
n7-harvester-0
Total for query: 14731 jobs; 0 completed, 0 removed, 6733 idle, 7998 running, 0 held, 0 suspended
```

- For comparison
 - Local submission using phnxsub0X, reached 100k jobs
 - Harvester → single CE test : single CE can run 21k jobs
- Slow job status update between harvester condor queue and CE -> eventually catch up
 - One single CE test:
 - Around 04-27 23:56, submitted 21k jobs to spce04, from 6 harvester nodes
 - Around 04-28 00:18, 9.5k running jobs on the CE
 - Around 04-28 00:32, 21k running jobs on the CE
- Job failure rate : 1.2%



- One major failure : unable to fetch source tarball from panda server
 - Fixed in a new version from panda repo, deployed now on the production PanDA instance
- attemptNum = 1 in the test jobs
- Load on the PanDA DB ?
 - [Moderate](#) load, nothing special so far...
- Peak Auth rate 30Hz per CE.
-

- Production Workflow Status
 - Input datasets now working
 - <https://sphenix-panda-test.apps.rcf.bnl.gov/task/51/>
 - Single run, single sequence for each filestream
 - Rucio RSE = BNL_PROD
 - Files must be under /lustre01/sphenxpro/rucio/ ...
 - Only MD5 sum and GUID
 - Working on standing up a full example workflow
- Lustre Downtime schedule Until W May 3rd
 - [SDCC Announcements | Scientific Data and Computing Center](#)
 - major upgrade from version 2.12.8 to 2.15.2, starting Monday, May 1st at 9:00 AM and lasting until Wednesday, May 3rd. The system will be unavailable during this time, and users should plan accordingly.
 - Expected end time is 11:59 PM on 5/3/23.
- Updates
 - Pilot monitoring in Elastic Search:
 - Two logstash servers [configured](#) on splogstash.sdcc.bnl.gov:8080 :
 - http plugin
 - [Pipeline](#)
 - ES Index: sphenix_rucio-pilot_write
 - Available for integration

- AOB

Meeting Mon April 24 2023 1 PM EDT

- New meeting time
- 100k job stress test [...]
 - Plan is to submit jobs on Tuesday
 - 100k jobs w/ randomized sleep (20 to 40 minute, flat distribution)
 - Ramp slow 25k 50k 100k
 - Discussion of use (or not) of Atlas nodes... (should not be a dependency)
 - Memory used by pilot... need to ensure this is not an issue
 - Plan is now end of week
 - How many jobs to each CE... 5k/CE. (Scatters submission across the CEs).
 - Need ability to ensure that submitted load is uniform across all CEs, such that authentication is uniform
 - where does the 5k limit on the CE come from? ...
 - balance between buffering the idle jobs between harvester vs the CEs. (Idle jobs consume resources).
 - jobs in excess pileup in panda
 - ... workflow is setup ready to submit ...
- sPHENIX production workflow / Rucio datasets for PanDA
 - Setting up input dataset(s)
 - When attaching extant files to new datasets.
 - PanDA does not recognize files in datasets that lack the GUID
 - PanDA *also* does not recognize files in datasets that lack the Adler32 sum
 - Once set in rucio, you cannot reset the Adler32 checksum... (must create a separate DID which correctly sets the information)
 - GUID *must* be set when the DID is created... rucio will not accept a new meta-data value afterwards.
 - ... recommend sPHENIX provide Adler32 in addition to MD5 checksums
 - Raw “junk” event dataset for run # 255 has been created...
 - Simple workflow (print the input file locations) failed b/c of error on PanDA side ...
 - MOCK rse should not be used. Use BNL_PROD.
 - sPHENIX will provide MD5. No Adler.
- High / Low / Medium priority condor queues status? Changes made by SDCC on CEs, tested in test PanDA instance, prod instance to be updated.
- Updates
- AOB

Meeting Mon April 17 2023 9 AM EDT

- SHREK
 - Support added to declare processing types.
 - Configuration is in a yaml file / easily extensible ...
 - default: NA, testing: Ts, user: Us, calorimeter: Ca, calibration: Cb, Tracking: Tr
 - Pilot on production side exporting this into the global id
 - TODO finalize list / coordinate w/ SDCC on condor/lustre action
 - Added secondary dataset support to code which handles single-task (“prun”) submissions.
 - Data sets created for production workflow test
 - Main issue was that some fraction of jobs were taking >> longer than expected to complete. (Kill / resubmit)
 - rucio ls group.sphenix:mdc2.rawdata.junk.rhic2023-03-29
 - Contains all 50k files... may create DS’s for each “run”
 - rucio ls group.sphenix:mdc2.shijing_hepmc.fm_0_20.g4hits.run0006
 - First attempt to run from these datasets... (tasks 107, 108, 109...)
 - <https://sphenix-panda.apps.rcf.bnl.gov/task/?jeditaskid=107>
 - unknown input dataset "mdc2.rawdata.junk.rhic2023-03-29"
 - Does PanDA depend on either the Adler32 checksum or the GUID field in rucio?
 - Error in PanDA logs: `getFilesInDataset pid=801702`
`<datasetName=group.sphenix:mdc2.rawdata.junk.rhic2023-03-29> <class`
`'TypeError'> : 'NoneType' object is not subscriptable"`
- Preparation of 100k job scale test by the end of this week (2 new racks mid-late week)
 - based on ghost hardware / ghost job (sleep 30m)
 - separate IO / rucio test
- High / Low / Medium priority condor queues...?
 - ETA this week
- FluentD vs logstash:
 - Email thread with Xin&BigPanDamon team (Ye Shuwei, Tanya Korchuganova, Sasha Aleksandr) about fluentd for sPHENIX job live logging
 - Expansion of the support of real-time logging in debug mode to other experiments: <https://its.cern.ch/jira/browse/ATLASPANDA-803>
 - ATLAS jobs are using logstash (Elasticsearch + Kibana) and it needs the debug mode
 - Questions about emit rates, payload sizes and retention policies for sPHENIX
- Monitoring (what, how)?, prmon turned out to hang jobs occasionally in the sim production - its use is not an option
 - Need capability to identify / diagnose the <1% jobs that fail
 - panda monitor / elastic search / summary presentation? (SDCC supports elastic search..)
- Lustre upgrade between now and start-of-run
- +1wk delay cooldown
- Reopen discussion for meeting time/// doodle poll to be submitted. (Monday 10AM?)
- Updates
- AOB

Meeting Mon April 10 2023 9 AM EST

- Token lifetime 1yr for PANDA_AUTH_VO=sphenix.production.
 - shrek site.yaml as auth_vo
 - Chris, Sasha, Xin and Jason allowed. (POC to add additional users? A: Xin.)
- Retention time for intermediate datasets in a PanDA workflow? (PandDA setting?)
 - Users / when job is finished update lifetime / ... follow up with Tadashi et al. In Atlas it is prodsys.
 - Physical deletion? Procedures need clarified.
- Rucio namespace organization
 - Reminder... rucio has a *flat* directory structure
 - *scopes* used to partition the namespace / aid in data discovery
 - eg rhic2023.auau200.<datatype> {junk,online,production, ...} w/ runs mapped to datasets
 - specification from experiment
- sPHENIX production workflow
 - Rucio is now configured to use MD5 checksum rather than adler32. (Thanks Vincent)
 - Stalled on filling the input datasets
 - For MC need to calculate MD5 sum for each file (~30s @ 50k files...) → submit to condor
 - Expect jobs to take ~minute each. Significant fraction observed 2h+.
 - Have been chasing the wrong rabbit down the hole debugging this. Real issue is the python script which adds the file to the DS. Investigating.
 - Expecting to be in position to test full workflow (dataset close → automatic dispatch of production-like workflow) early this week.
- Lustre optimization / production type specification
 - -productionType flag in PanDA/prun passed through pilot as data["processingType"]
 - shrek syntax updated to support... will restrict to 2 characters... (e.g. Tr=tracking, Ca=calor, Cb=calib, ...)/// map at the level of shrek.
 - lustre job ID as env variable associates activity types to nodes... JOBID: condor ID, panda ID, workflow description limited to 31 characters.
- High / Low / Medium priority condor queues...
 - PanDA queues created (BNL_OSG_SPHENIX_HIPRI, BNL_OSG_SPHENIX_MDPRI), waiting for changes to the router on the CE
 - (shrek supports by specifying in either site.yaml or command line option at submission time)
- Fluentd
 - SDCC (Vincent) asks for details about the setup, collecting info from relevant experts.
- Updates
- AOB

Meeting Mon April 03 2023 9 AM EST

- New meeting time
 - sPHENIX Production Workflow
 - Transfer Script coordination copy files / dataset close
 - [Scalable Job Tracking — HTCondor Manual 10.3.1 documentation](#) [\[example\]](#)
 - Mock production workflow (sim + raw data)
 - Calorimeter workflow (almost) ready... need to fill input datasets
- ```
farquaad add-file --pfn /sphenix/lustre01/sphnxpro/mdc2/shijing_hepmc/fm_0_20/g4hits/run0006/*
--dataset mdc2.shijing_hepmc.fm_0_20.g4hits.run0006
--containers /TEST-CONTAINER-2023/mdc2.shijing_hepmc.fm_0_20.g4hits
--simulate
```
- Adler32 checksum required for each file... **50k files @ 20s per operation 10 CPU days**
    - Is a checksum required? What functionality missing if set = 0?
    - Yes.
    - sPHENIX *computes* MD5 checksum would prefer that,, in DB for g4hits
      - Psql access for the md5sum of files from the sim production:  
psql -h sphnxdbmaster.rcf.bnl.gov FileCatalog -c "select md5 from files where lfn = 'DST\_CALO\_G4HIT\_pythia8\_PhotonJet\_sHijing\_0\_20fm\_50kHz\_bkg\_0\_20fm-0000000006-01617.ro  
ot'"
    - Eddie noted in chat md5 checksum is provided as an option when adding replicas...
      - [http://rucio.cern.ch/documentation/client\\_api/replicaclient/#add\\_replica](http://rucio.cern.ch/documentation/client_api/replicaclient/#add_replica)
  - Namespace organization
    - [sPHENIX naming convention](#)
    - Multiple types of data...
    - TYPE (junk)/YEAR-SPECIES/.../filename (note that *containers* do not *behave* as filesystem directories)
      - RHIC-2023-AuAu/JUNK/
      - RHIC-2023-AuAu/REAL/
  - PanDA status – config issues over weekend / debugging
    - Multistep / monitoring issue now fixed
    - Config changes ... multiple queues in progress some failures
    - some instability today
    - Need to pass an env variable (31 char) that specifies the stage of data production
      - Lustre optimization based on workflow / task details
  - Updates
  - Open / Pending / Critical items
  - AOB

Meeting Wed March 29 2023 10 AM EST

- [Doodle poll for new meeting time](#)



- sPHENIX Production Workflow
  - [sPHENIX production workflow slide from Chris 20230328](#)
- Namespace organization
  - [sPHENIX naming convention](#) [open file?](#)
- PanDA
  - Broken link on PanDA monitor for multiple-step workflow is fixed
  - Making some config changes to PanDA (disable local postgres DB in yaml file, adding more PQs etc), to be tested.
  - New schema update to PanDA DB, tested in test instance.
- SHREK
  - Task Management (aka how to kill another production manager's job)
    - On submission of task, save the user's panda token in submission area / pbook-wrapper script uses token in submission directory for authentication / anyone with execute priv in submission directory may now monitor / terminate jobs production tasks.
  - Data Production Pipeline
    - Transfer to lustre → **create / fill dataset** → **close dataset** → **Donkey launches workflow on PanDA**
    - Script implemented to create / fill and close dataset. Small tests done. Test at scale TBD.
      - Currently done by hand. Automated DS creation / closing on transfer to be done.
      - Donkey tested already at appropriate scale (but needs integration test)
  - Tests of a commissioning-like workflow
    - Last week Tadashi fixed issue blocking re-submission of tasks
    - Ran simple prun task under shrek, generating 5 jobs.
    - Resubmitted task with nJobs=10. 5 new jobs processed & files added to original output dataset
    - Added a "file tag" command line argument which creates env variable in job script enabling users to append a grep-able string to their output files.
- Updates
- Open / Pending / Critical items
- AOB

## Meeting Wed March 22 2023 10 AM EST

- Meeting time
  - NPPS group meeting will be bumping us from this timeslot. Doodle poll will be circulated for new time.
- Stress Test 03/21/2023
  - Problems in PanDA job submission observed from the last stress test solved, saturated 30k slots for several hours.
    - Previous issue w/ harvester not seen (interference).
    - 15k ramp up in 30min... will be faster under typical use.
      - Speed matter of panda setup and queue configuration
    - Stripe size = 5 works well. <https://chat.sdcc.bnl.gov/api/v4/files/pq61o4fsnbffif7buygsny94we/preview>
    - Tested pushing data in while event building
    - Next step... file transfer to rucio w/ work dispatch to panda through donkey/shrek
      - Quantity of work is 1 run ~1h of data ~10k files
      - Further discussion on namespace / organization of files into datasets
      - [sPHENIX naming convention](#)
- Multi-step workflow in PanDA stops working after the recent upgrade, under investigation. (Possible rucio issue).
- sPHENIX data processing readiness / path forward
  - Milestones from last week
    - Signup instructions w/ panda IAM
    - Instructions for using prun etc...
      -
    - Data → rucio
      - test of automated production process (aka donkey)
      - Address namespace issues during testing (rucio dataset... unique filenames... multiple users...)
    - Raw data production data is in place | intermediate production move data to worker node... (confirm that data is moved to the worker node)... and in appropriate numbers of files.
      - *Not yet confirmed. Added capability in shrek to pass (or not) the forceStage flag to prun / pchain. Note that in my (Jason) previous tests, the `–forceStage` flag *did not result in input files being copied into the working directory on the worker node. Is there a rucio / panda configuration that needs to be set for this?**
    - Admin job killing
      - Tested using my (Jason) credentials to kill Xin's tasks... failed.
      - Users submitting under sphnxpro will obtain a token, which will be saved in some directory. Said token can be used by the sphnxpro account to cancel jobs from other uses. Confirmed yesterday. Sufficient?
- Status on FluentD service?
- Review of action items
- AOB

FILE TRANSFER → kick off workflow

## Meeting Wed March 15 2023 10 AM EST

- Happy Ides
- Stress Tests?
  - Several changes made in PanDA , harvester pods are more isolated with their own IDs.

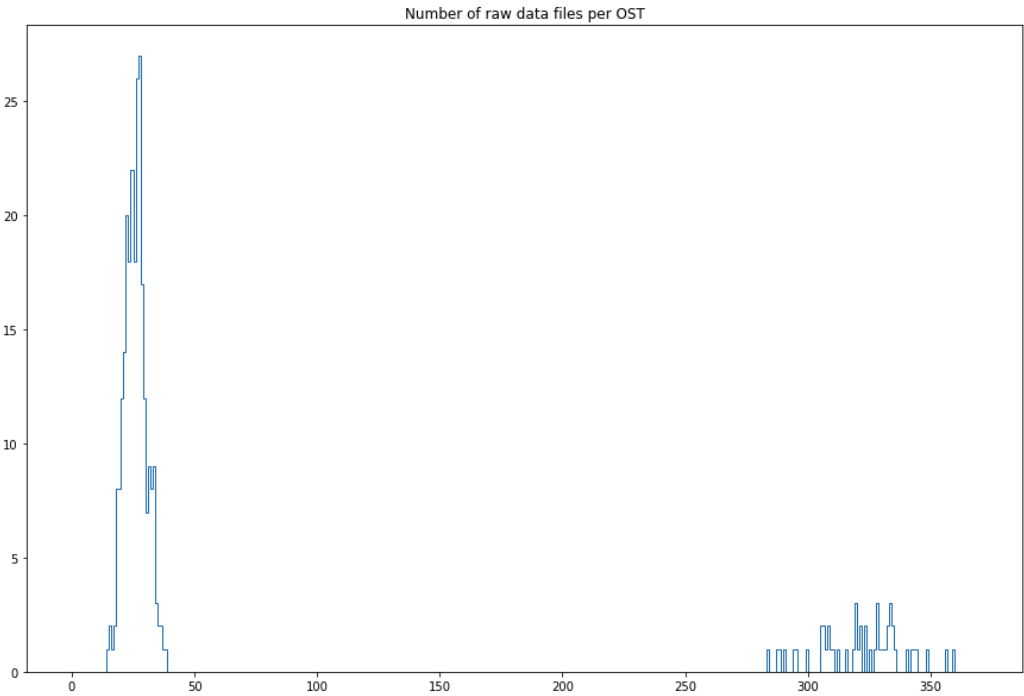
- Follow up stress test TBD. Production instance to be upgraded (DB schema change required, wipe and reinitialize.) ~2 days. Expect 1h downtime.
- Shadow condor universe? Still in progress. (Stress test after).
- 30GB/s hammering of lustre... concentrated onto 5 servers... issue solved... better distribution... 100+GB/s achieved. May want to split further.
- New token job submission to CE... job submission works, next how to have harvester update access token.
- sPHENIX data processing readiness / path forward
  - Commissioning phase.
    - May 1st. ... 12 weeks of commissioning
    - Data to be pushed to lustre, cataloged in rucio, processed using both user code and sPHENIX libraries. prun and/or shrek task submission. No automation.
    - Many users, private codes ...
    - Once data files closed, push to SDCC... users kickoff jobs.
    - Anyone has to be able to run jobs @ SDCC. Register w/ PanDA IAM. (Need ~ continuous support to add users ). Use SDCC credentials
  - Calibration and QA.
    - TBD. Soon afterwards (~2wk?)
    - Data to be pushed to lustre, cataloged in rucio, processed using sPHENIX libraries. Automated workflow submission. Simple prun and/or shrek tasks.
  - Steady state production. Physics operations ~end of June.
    - TBD.
    - Data to be pushed to lustre, cataloged in rucio, processed using sPHENIX libraries. Automated workflow submission. Full reconstruction workflow.
    - Production manager responsible for data processing
  - Review of open action items
- AOB:
  - [Mattermost channel for operational support for sPHENIX](#)

MILESTONES:


- Signup instructions w/ panda IAM
- Instructions for using prun etc...
- Data → rucio
  - test of automated prodction process (aka donkey)
  - Address namespace issues during testing (rucio dataset... unique filenames... multiple users...)
- Raw data production data is in place | intermediate production move data to worker node... (confirm that data is moved to the worker node)... and in appropriate numbers of files.
- Admin job killing

Meeting Wed March 08 2023 10 AM EST

- Review of action items
- Stress Tests
  - Raw data file likely not spread uniformly across the OST's
  - Unbalanced job assignment among harvester pods, lead to high load on some pods, under investigation.
- PanDA new schema changes and feature release (close a dataset to signal no more files to add), tested on test PanDA instance successfully, production instance to be done.



- 
- Here is the number of raw input files per OSS (The oldest nodes have the most files)



- 
- PanDA DB: stress test results
  - [summary document](#)
  - Currently working with Jason S. on Lustre-based DB backups (access is tricky)
- SHREK
  - Donkey application now classifies messages by the event type in the header...
    - rucio create dataset
    - rucio close dataset
    - rucio (re)open dataset
    - ... work can be dispatched on one or more of those messages
    - If needed, we can extend to rucio containers as well
  - Tested usage of <https://git.racf.bnl.gov/gitea/> as repository for submission artefacts.
    - e.g. create submission directory, initialize a local git repo, set a BNL repo as the upstream instance
    - Will try to configure usage of ssh so that constant authentication not required on a push to the repo.
- 
- AOB

Meeting Wed March 01 2023 10 AM EST

- Review of action items and open issues
  - SHREK profiling / latency
    - 4.3s to dispatch a (simple) task.
    - 0.3s to prep scripts, support files, directory.
    - 3s for PanDA to submit.
    - 1s to git / add / commit artefacts
- File persistence in Rucio / distortion correction use case
- Group production scope / new panda client
- PanDA DB on VMs: connection to services
  - DB setup and integration with OKD-based pods is now complete.
    - Schema installation, pg\_cron, pg\_partman - debugged and fixed several issues, everything works as expected now.
    - Short test jobs submitted by Xin succeeded: prun, pchain work fine;
    - ToDo:
      - Larger test to see if multiple jobs succeed
      - Read/Write splitter test
      - Short stress test, ~full farm occupancy, ~a minute per job
      - Long stress test, ~full farm occupancy, several hours
    - Rucio VM-based db setup - need coordination with Vincent and Matt here
      - ORM-based client, no cron, no stored procedures
- AOB
  - shadow condor universe

Meeting Wed February 22 2023 10 AM EST

- PanDA@OKD
  - Stress test
    - Multiple CEs speed up pilot submission.
    - Longer delay between bulk job creation in PanDA observed, to be followed up.

- Database built by Dmitry on the external PanDA DB server, testing ongoing, also investigating some other errors after switching to a newer idds image.
- Setting up a test instance of PanDA@OKD, need a hostcert for the test hostnames from SDCC
  - (Independent test and production versions)
- Group production scope for job outputs tested, works for prun but not pchain, to be followed up
- Pilot tarball now from sphenix cvmfs, works fine for BNL PQ.
- SHREK / Donkey demo
  - Identified bottlenecks in submission and removed. Tested... able to dispatch ~20 jobs in approximately 1 min. Assuming ~10-100 runs to be pushed out in one shot every ~24 hours... should be able to dispatch jobs to PanDA in 1-10 minutes.
  - Tested over the weekend. Ran for 72 hours, during which time Xin pushed several tests. Messages were logged and acknowledged. No actions were setup for his data sets. They did not interfere with my test DSs.
  - Message monitoring and work dispatch run in separate threads, with a python cmd.Cmd command loop providing user the ability to inspect current set of pending messages, define new job dispatch rules in flight.
- Populate Action items

## Meeting Wed February 15 2023 10 AM EST

- Stress testing
  - After some config changes on gridmanager on harvester pods, pilots were submitted to CE much faster, but the overall step (harvester condor – CE) still slow.
  - Post-mortem ongoing (tracing some aborted pilots from the panda logs)
  - Next step to use multiple CEs
    - Fast path – local submission? Focus on production.
    - CH asked for 3 CEs... sometime this week
    - minimum acceptable job startup rate?
- PanDA DB update:
  - VMs are configured, DB servers are installed, DB schema for PanDA and IDDS (as provided by Xin) is imported into sp-pandadb01/02/03/04.
  - Xin is configuring the deployment of a new PanDA service at OKD that will use the new DB backend;
    - To capture the load of Panda and validate the deployment, having scaling tests with it is needed. (Vincent)
  - Note: PanDA server database contains stored procedures => CPU-heavy service, less scalable than expected. ~~Master DB server VM may benefit from resource sharing during the job launch hours (=spiky load) [ upd: Jason S. thinks it would not }~~
- Setting up a test instance of PanDA@OKD, for testing new DB and token etc ...
- SHREK
  - New DonkeyApplication being implemented using stomp / ActiveMQ
    - Need to take some care in design b/c stomp.py is inherently multi-threaded
    - Two implementations planned.
      - fire-and-forget – program sees message, process and dispatch to panda... simple daemon running in the background. Easy, no race conditions to consider b/c everything happens in one happy thread
        - Currently able to receive messages. Dispatch to PanDA through shrek requires ~10 LOC. Should be able to test by lunch.
        - Needs to drink from the firehose... single thread may not suffice... shrek has bottlenecks as well...
      - command loop – production managers can run, suspend, inspect progress, dispatch jobs manually, etc... multithreaded, so have to consider thread safety.
        - Lower priority
- AOB

## Meeting Wed February 8 2023 10 AM EST

- BNL Instance Stress Test #1 Takeaways
  - Ran 14400 jobs @ 1k ev/job w/ avg duration 12h (*n.b. 14390 @ 1k ev, first 10 @ 10 ev*) **[Test #1]**
    - *Expectation: reach full utilization of 14400 nodes w/in reasonable time scale*
    - Maxed farm utilization at 8500 slots, slow ramp up over ~1h, hit a few plateaus along the way
      - sPHENIX simulations (geometry) occupy too much memory... limit of 4GB / job (average?)
      - Overwhelmed mem/swap of the single CE (compute endpoint?) in use  
<https://chat.sdcc.bnl.gov/npps/pl/spamso35xinijpxikgod1wx3dr>
      - plot of memory weighted slots -  
[https://monitoring.sdcc.bnl.gov/grafana/d/000000039/shared-pool?orgId=1&from=1675231200000&to=1675835999000&var-Group=group\\_sphenix&viewPanel=26](https://monitoring.sdcc.bnl.gov/grafana/d/000000039/shared-pool?orgId=1&from=1675231200000&to=1675835999000&var-Group=group_sphenix&viewPanel=26)
      - CH will add memory to the CE ~today
  - Long duration test w/ 24h duration broke in scouting jobs **[Test #1a]**
  - Short duration test ~30 min duration **[Test #0]**
    - Xin identified (fixed) configuration limiting to 3k, increased space allocated for condor job logs,
    - Ramp up ... harvester parameters tweaked... shorten cycle time and limit that each harvester can handle
  - Next stress test to see how fast we can ramp up
  - Window of opportunity w/ lustre down... (but need to have another FS configured for panda/rucio)
    - Action item: add one+ RSEs: BNL\_PROD2 **[done]**
      - Additional output path (make yourself a subdir): /sphenix/data/data03/sphnxpro/rucio/
    - Action item: PanDA nap time (submit 15k sleep 10min) try to observe the true ramp up time.
- PanDA DB update:
  - PostgreSQL v14 servers installed and configured at sp-pandadb01|02|03|04
    - 2TB storage available per VM for DB needs



- ~30 TB storage for backups is not available (need to check with Jason S.)
- Replication: async streaming Master-Replica replication setup has been created
  - sp-pandadb01 - read-write Master
  - sp-pandadb02|03 - read-only Replicas
  - sp-pandadb04 - read-only backup Replica
- Scalability:
  - based on the information provided by Tadashi, PanDA cannot split reads and writes (it does immediate reads after writes) unless there is a misunderstanding of some sort. This should effectively prevent horizontal scalability via read-only replicas, and would rely primarily on the master node to do all the work. [need to see if the hardware is able to perform well for Run 1]
  - PgPoolII (read/write splitter mode) - I am going to test it just to see if PanDA server really fails to use the replicated setup or not (assuming splitter and DNS RR enabled). If it does not, we may still get some benefit from partial read/write splitter. If things go bad, we can easily revert to basic M-R setup (no major changes needed).
- TODO (this week):
  - set up DB access accounts (Dmitry, Xin)
  - upload DB structures/content (Xin)
  - do basic tests with/without PgPool (Dmitry, Xin, Jason?)
  - request a DNS RR for sp-pandadb02 and sp-pandadb03 => sp-pandadb
- TODO (next week):
  - enable backups, monitoring
  - switch from the OKD-based panda db to the VM-based panda db and re-test again
- ActiveMQ status:
  - Activemq notification has been enabled on the Rucio@sphenix instance
  - ActiveMQ server: spmq01.sdcc.bnl.gov:61613
  - Destination for Rucio (dataset) events: /queue/events
  - Action: user:pwd consumer account for ActiveMQ will be communicated [done]
- ...
- Need IAM account / rucio scope defined for sphnxpro. (?)
  - IAM server for sphenix ?
  - sphenix production scope being tested
- SHREK
  - Moved to BNLNPPS github: <https://github.com/BNLNPPS/shrek>
  - [Quicker-start document](#) written
    - Help pages for accounts / authentication tbd
  - no-scout option implemented, logging and std output more useful (e.g. panda output both logged and echoed to stdout)
- Rucio namespace / scope / naming convention for sPHENIX files
  - [https://rucio.cern.ch/documentation/file\\_dataset\\_container](https://rucio.cern.ch/documentation/file_dataset_container)
  - Files, datasets and containers are uniquely identified over all time. This means that a data identifier, once used, can never be reused to refer to anything else at all, not even if the data it referred to has been deleted from the system.
  - [https://docs.google.com/document/d/1sZrMnawDf2P8RJ6xu5k4P07m\\_fIQ5R0qRn-zwltmOLc/edit](https://docs.google.com/document/d/1sZrMnawDf2P8RJ6xu5k4P07m_fIQ5R0qRn-zwltmOLc/edit)
  - Each file / dataset must be unique (otherwise rucio throws an error)
  - recommendation – (scope → contains activity and time)
  - Cf. [https://docs.google.com/document/d/1sZrMnawDf2P8RJ6xu5k4P07m\\_fIQ5R0qRn-zwltmOLc/](https://docs.google.com/document/d/1sZrMnawDf2P8RJ6xu5k4P07m_fIQ5R0qRn-zwltmOLc/)
- Live job logging
  - Priority?
  - Need SDCC to setup a Fluentd data collector, to which pilots will publish live logs from WNs.
- AOB

## Meeting Wed February 1 2023 10 AM EST

- Friday's NPPS meeting will focus on sPHENIX: where are we in NPPS deliverables to sPHENIX, what does the timeline look like, where are gaps in our present commitments, where can we/should we provide further help. This meeting has a scope that includes SDCC related activities as well. It would be good to have the picture of SDCC related matters from this meeting to inform Friday's discussion.
- Rucio file catalog:
  - Request to enable activeMQ notifications - Follow-up on ActiveMQ server@sdcc, timeline, etc (Vincent)
- SHREK Update
- BNL Instance Update
  - Running... CERN cric dependence eliminated. Please test.
  - Monitoring : <https://sphenix-panda.apps.rcf.bnl.gov/> (behind the BNL campus firewall)
- PanDA DB status: VMs were provided last Friday. Configuration/testing is ongoing and PanDA people will be contacted for their inputs. Timescale: 2-3 weeks
  - Pgpool is needed and should be added on the requirement document - Xin
- Documentation (where to point the sPHENIX production manager to run the show?)
  - Move to sPHENIX wiki
  - Improve SHREK failure output / job monitoring link should reflect site.
- How to move production to the sphnxpro account
  - Need different authentication strategy for sphnxpro account (can't be manually logging in every week).. attach authentication to the sphnxpro (unix) account
  - Productions should be attached to the sphnxpro scope regardless of who
  - Production roles can be defined w/in IAM

Meeting Wed January 11 2023 10 AM EST

- Wider meeting sPHENIX/SDCC/NPPS discussion on DB infrastructure
- Clearpath forward to establish a production PanDA DB at SDCC ASAP
  - Panda DB for sPHENIX - requirements
  - ... update in ~ 1 month. Feb 8th.
  - ... need to be able to test ~ March.
- SHREK
  - Enabled prun-based submission for simple tasks
  - Reworking the job archiving to be branch-based rather than tag based
  - Tested on BNL instance
    - Simple workflows run
    - Issue with charm simulation not executing the “finish” task... TBD

Meeting Wed December 14 2022 10 AM EDT

- Wider meeting sPHENIX/SDCC/NPPS discussion on DB infrastructure
- Authentication periodicity... connect to the account
- CRIC@BNL
- dress rehearsal status - discuss panda, rucio roles
  - raw data files are to be in rucio
  - panda runs with rucio as data source
- offsite computing for sPHENIX
  - start with simu on OSG, can leverage EIC experimentation with PanDA+Rucio on OSG
- Data taking / production target ~March (simulations of course can be sooner)
- Followup meeting ~January 11th
- SHREK
  - deployed on sphnxpro
  - Quick Start Guide for building workflows and submitting jobs through PanDA
  - First successful user test (thanks Kolja)
  - TODO: initial tests w/ BNL PanDA instance

Meeting Wed November 22 2022 10 AM EDT

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- Meeting cadence going forward 2 weeks
- PanDA @ BNL
  - PanDA@OKD : test jobs run successfully. Users can get to the job output via Rucio and condor job log through web server.
    - All services installed
    - CRIC to be installed locally
  - For scalability :
    - SDCC to provide a reliable and scalable DB service. Stress test will be conducted against it later.
  - Reinstate sPHENIX/SDCC/NPPS discussion...
- Full dress rehearsal upcoming
  - Files produced in counting house to HPSS to ... ruico
    - Hook into PanDA
- SHREK update
  - ...
- EIC piggybacking
  - Timescale for PanDa EIC... ASAP...
- RHIC S&T review: sPHENIX needs more computing than is currently budgeted! So need to use offsite resources...
- External sites? Makes sense for simulation.
- sPHENIX focus is on commissioning
  - Workflow run on single sample repeatedly...

Meeting Wed September 27 2022 10 AM EDT

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- PanDA @ BNL
  - PanDA is being deployed on the OKD cluster
    - Xin is learning the full stack on OKD
    - Taking Tadashi’s test instance and deploying it...
    - PanDA is “running”... sorting through issues to get it properly “working”
    - Xin will be at CERN next week and can directly interact with experts
- SHREK Update
  - Implemented a “dispatcher” script (donkey)
    - Monitors rucio for new datasets matching some criteria defined on the dataset’s meta-data
    - For example, when a given run is completed, a corresponding dataset is uploaded to rucio and closed, an analysis job can be dispatched
    - One or more “actors” can be assigned to dispatch multiple “workflows” on a single
  - Discussion
    - Main idea is to enable automated processing of datasets... but also
    - Useful for commissioning sPHENIX...

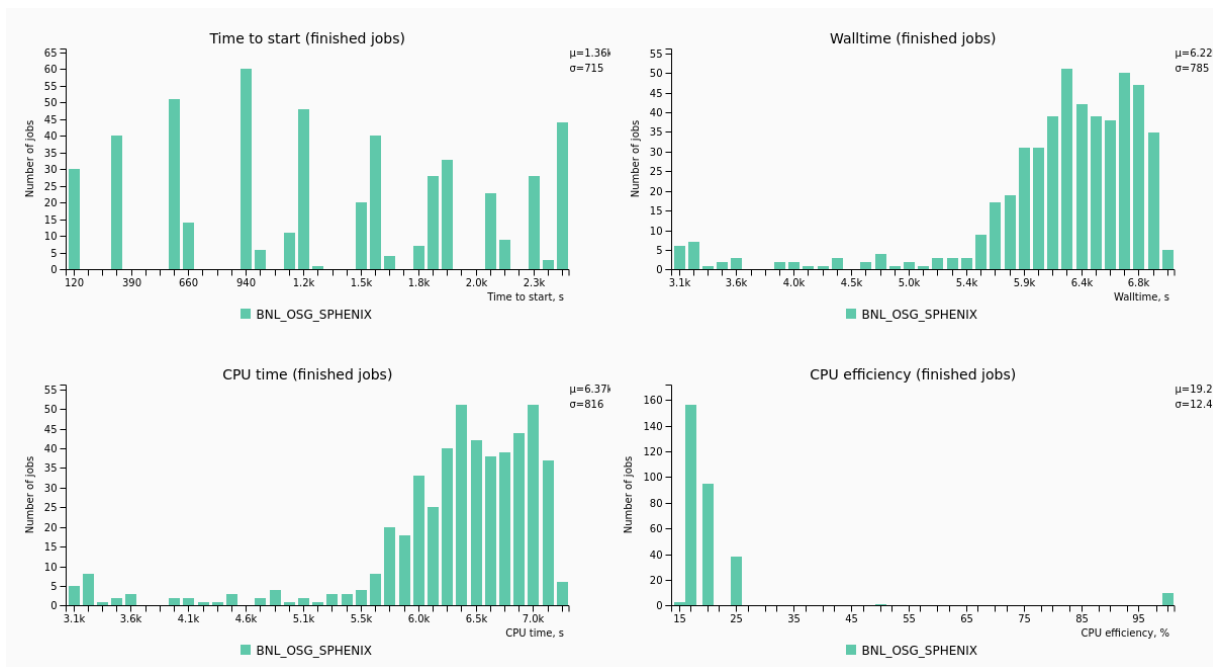
- Users can define tasks / workflows to be automatically dispatched when runs are pushed out to Rucio
- Could also be used in cases where software stability makes pchain painful to use (i.e. crash somewhere along the chain aborting the workflow).
- Discussion about how the sPHENIX event / daq file model maps onto PanDA processing of datasets / files.
  - Events are distributed (in sequence) across multiple files
  - Solution discussed is that PanDA should launch jobs based on *lists* of input files, rather than the input files themselves
  - Will also need to ensure that files are processed *in place*... i.e. **not** copied to the worker node, but rather read in over the lustre filesystem

○ AOB

## Meeting Wed August 24 2022 10 AM EDT

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- sPHENIX workflows
  - Macros now moved to the ana.315 software release
    - Issue with file naming (?) required reproduction from pass1 (?)
    - Had to disable (again) the file management in the job macros

### ■ Scaled successfully to 100 events / job



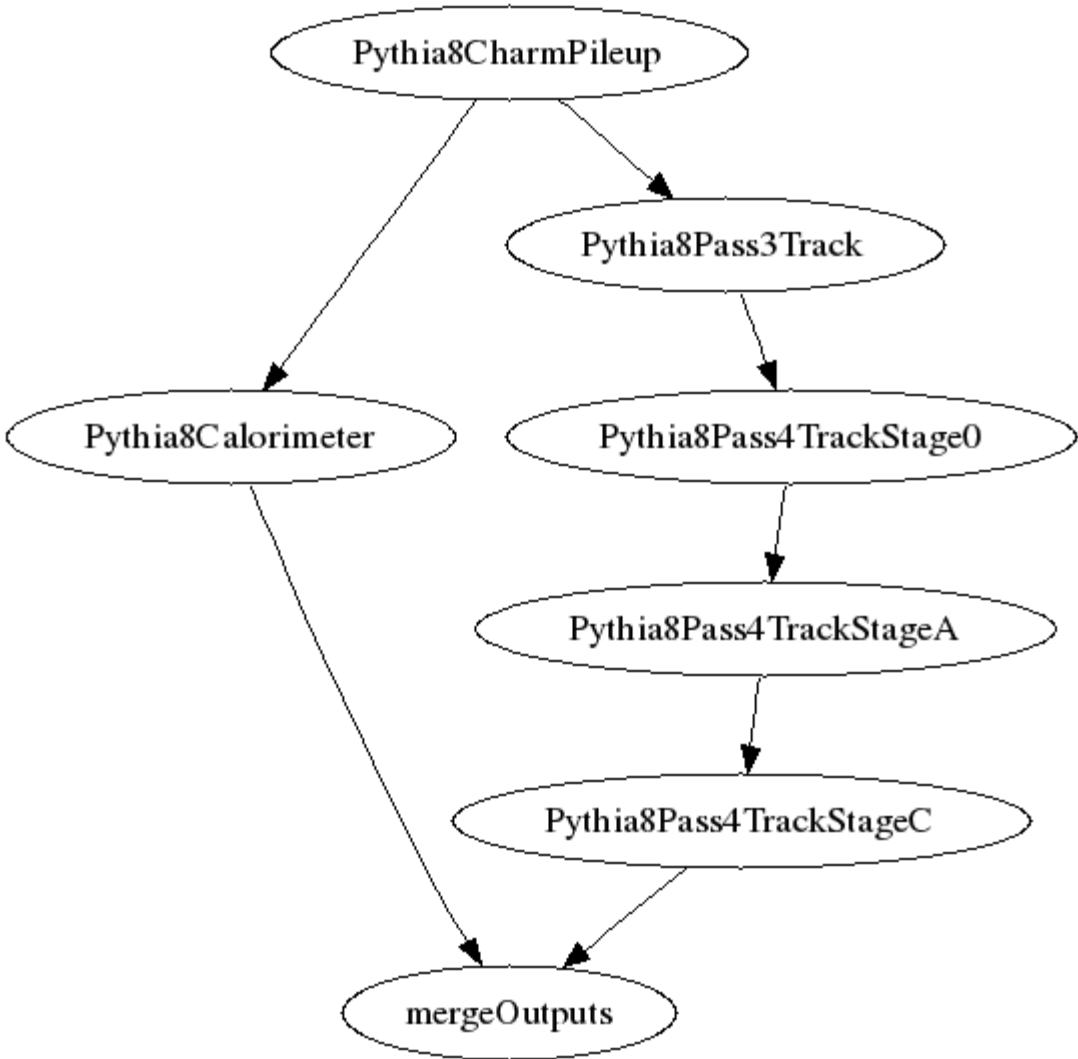
- Fails at 200 events / job
  - 2 hour walltime limit ...
  - [Discussion on mattermost...](#) possibly a harvester setting? TBD
  - ... 2h limit here ... [https://panda-doma.cern.ch/site/BNL\\_OSG\\_SPHENIX/](https://panda-doma.cern.ch/site/BNL_OSG_SPHENIX/)
- Discussion
  - Panda queue config on cric 2h needs to be changed. Xin has permission now... has increased wall time limit to 48h
    - Note... the limit is injected as a parameter to the request to the batch system
  - 48h is reasonable... increase if/when needed
  - Discussed whether this should (can?) be a user specified value... i.e. defaults to cric setting and user value can take precedence... tabled.
  - cric is currently running at CERN. Plan to move to BNL.
  - Scaling test today should (hopefully) succeed... will know about 2h after the meeting ends...
  - Plan for full test production
    - Aim for 1M events to test (1k jobs x 1k events)
    - Some coordination w/ Chris would be useful at this point b/c:
      - Keeping the farm nice and toasty warm with jobs...
      - Everything is running under sphnxpro credentials... so priority
      - Current “ana.315” macros are not actually producing tracks...
  - PanDA jobs relative priority? PanDA jobs on shared pool? Dedicated hardware?
    - Only goes to sPHENIX hardware
  - All jobs (Chris’s, Jason’s, ...) running under the sphnxpro account... conflicts in priority
    - Will discuss w/ Chris H about setting PanDA jobs @ higher priority

- Question to Kolja about piggy-backing off of this work...
  - Planning to, but just at the initial stages of trying to get things to run and see the jobs post to the panda monitoring

## Meeting Wed August 17 2022 10 AM EDT

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- Rucio / Panda Integration
  - Use following to iterate on rucio download / rucio upload issues...
    - <https://chat.sdcc.bnl.gov/bnl/channels/sdcc-support-for-sphenix>

- Panda @ BNL
  - OKD access instructions recieved
- sPHENIX workflows
  - Stability issues
    - Able to scale to 500 jobs by breaking into two steps (1) generate events, (2) pileup merge and reconstruction
    - Scaling to > ~20 events / job runs into a 2h wall clock limit.
      - –maxWalltime option appears not to be accepted
  - Updating from the mdc2.8 build to the ana.315 build of sPHENIX codes...
    - Currently debugging why pileup jobs are not unpacking the background files...



- 
- Splits track reconstruction into several additional steps, allowing (in the future) distortion corrections to be computed and applied.
  - Distortion corrections will (probably) require usage of conditional processing in PanDA and possibly looping. TBD.
  - Some care will need to be taken w/r to how the distortion corrections are calculated b/c this could become a bottleneck.
  - → run by run
  - → compartmentalize further ?

Meeting Wed August 03 2022 10 AM EDT

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- PanDA BNL update
  - Tadashi and Xin discussed asking SDCC for a production cluster on OKD / when Xin is back from vacation
    - Currently using a test cluster
- SHREK / sPHENIX workflow update
  - HF Charm ... scaling up the jobs
    - Scaling 100 events @ 1 charm job (5 bg jobs)
      - [Success](#)
    - Scaling 10 events @ 10 charm jobs (50 bg jobs)
      - [Success](#)
      - [Failure](#)
        - Some jobs creating truncated output in the minbias simulation stage... suggests issue with code crashing and confirmed with logfile...
        - I cannot reproduce the issue offline
        - The pileup task appears to launch prior to the input becoming available? Resulting in workflow crashing?
  - Scaling 1 event @ 100 charm jobs (500 bg jobs)
    - 12-13 input files fail to produce output in the final tracking pass
    - This prevents the final merge job from executing
    - This prevents the full workflow from completing
    - There is a big-red-finish button that, when pressed, asks me to log in with my (nonexistent) CERN credentials



- In principle should be able to use pbook here? But pbook does not show any information about any processes...
  - Inspecting the datasets in rucio... ~87% of the jobs completed with good data
- Summary:
  - Workflows appear somewhat fragile... if there is any instability in the user code there will be problems in pushing the workflow through
  - Plan to add feature to shrek to allow each step in a workflow to be executed “by hand” (prun each step by hand)
  - In an ideal world... could submit based on prior results in a workflow

■ Rucio

- Able to access “files” created by PanDA using sphnxpro creds
- Able to erase “files”... but ... not sure why this is still hanging around...

```
jwebb2@sphnx01 [Tue Aug 02]$ rucio stat
user.jwebb2:user.jwebb2.TestPileup-5c495990-d164-11ec-9093-00163e101048_000_Pythia8CharmSimulation.log
account: iddsv1
bytes: None
expired_at: 2022-07-29 18:56:48
length: None
monotonic: False
name: user.jwebb2.TestPileup-5c495990-d164-11ec-9093-00163e101048_000_Pythia8CharmSimulation.log
open: True
scope: user.jwebb2
type: CONTAINER
```

- Unable to upload / download files
  - note: inability to upload not a show stopped b/c I can always add a “copy data from disk and create output dataset” job.
- There is an operational issue with pchain workflows and rucio.
  - All intermediate datasets that are created during the PanDA run are persistent
  - There are duplicate files propagated through the jobs
    - *final merge step takes the calo and tracking output and creates a single dataset*
  - Unless rucio is just creating references under the hood, in which case nevermind...

- Chris will update the HF charm macros to use the latest tracking passes
- Jason will integrate into new workflow
- Discussed where distortion corrections fit into the picture ...
  - will need to implement seperate steps for each tracking pass so that distortion corrections can be inserted in the future
  -

Meeting Wed July 27 2022 10 AM EDT


- meeting link
    - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
  - SHREK / sPHENIX workflow update
    - Scaling simulation to multiple jobs
      - Successfully ran 2x charm : 10x background => 2x pileup : reco  
<https://panda-doma.cern.ch/tasks/?jeditaskid=131708|131716|131721|131728|131707|131720|131719>  
 (trick seemed to be that nJobs needed to be set on the pileup task)
      - Scale to 10x : 50x ? error reported on pileup task  
<https://panda-doma.cern.ch/tasks/?jeditaskid=131727|131726|131725|131724|131722|131723>  
 (possibly due to typo in yaml file... resubmitting)
    - First crack at using rucio...
      - [https://rucio.cern.ch/documentation/setting\\_up\\_the\\_rucio\\_client](https://rucio.cern.ch/documentation/setting_up_the_rucio_client)
        - ☒ ~~pip install rucio-clients~~
        - ☒ ~~pip install --upgrade rucio-clients~~
        - ☐ rucio whoami
      - rucio hits a stack trace...
 

```
RuntimeError: Could not load Rucio configuration file. Rucio looked in the following paths for a configuration file, in order:
/opt/rucio/etc/rucio.cfg
```
- setup script under cvmfs... /cvmfs/sphenix.sdcc.bnl.gov/rucio-clients
- pp 1k events @ 24h 2-4GB filesize

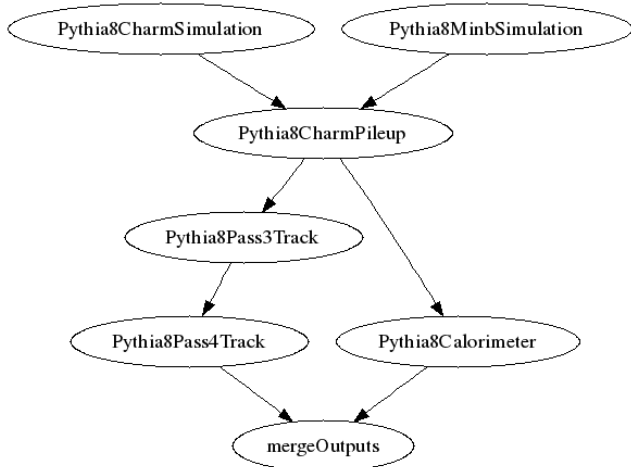
- PanDA @ BNL update on Friday
- Tracking update @ sPHENIX... code may be crashing ...
  - Updates to workflow to be passed to Jason
    - @ new release / build

Wed July 20 2022 10 AM EDT

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- SHREK / sPHENIX workflow Update
  - Pileup task is now running
    - But only able to process 1 event per input job!?
    - Longer jobs (5 events) are being terminated by the condor. (?)
    - Trouble ticket to be submitted...
  - [Calorimeter reco task running](#) based on pileup inputs



- Working to get 1st pass of tracking running
- Will need to add a final “merge” step to collect the outputs



- <https://panda-doma.cern.ch/tasks/?jeditaskid=131247|131246|131239|131238|131248|131244|131250>

- Once scalability addressed + rucio access / integration... can begin using in a production mode
- Jobs require 1.5 GB, ... up to 12 GB
  - Need to be able to specify memory requirements @ job submission
  - PanDA needs to interact w/ Condor

- Reverse proxy is up... should be able to interact w/ monitoring from with the BNL network.  
<https://sphenix-panda.apps.rcf.bnl.gov:443>

**Wed July 13 2022 10 AM EDT**

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- SHREK update
  - HF charm pileup jobs still having issues
    - Signal and BG jobs both complete successfully (according to PanDA)
    - Pileup job fails (reason unknown)
    - Confirmed this is not an issue w/ PanDA (all inputs avail, unpacked... and simple workflows succeed)
    - Confirmed this in not an environment issue [?] (same environment across all jobs, same environment vs interactive)
    - Ran signal and BG by hand interactively ... ROOT files are truncated...
    - Updated codes from main repo / same result
    - Ran using sPHENIX production scripts... same result

possibly b/c disk full? mdc2.8 preferred library.

- PanDA @ BNL
  - all components running at bnl okd (not monitor)
  - fully commissioned / pruned via rcas
  - outputs → bnl rucio
  - pchain under testing (goal this week)
  - panda monitoring...

switchover following scaling tests... end user will need to update environ variables... but otherwise transparent.

- Chris – workflow split 2GB/core vs 4GB/core ... PanDA will need to send jobs to appropriate queues.
- AOB

**Wed June 22 2022 10 AM EDT**

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- SHREK / sPHENIX workflows update
  - Integration with version control for documentation
    - Job summaries <https://github.com/klendathu2k/sPHENIX-test-production/tree/test#readme>
  - [Pythia8 charm simulation](#)
    - Both charm signal and minbias backgrounds are being produced

```

$ ls /sphenix/u/sphnxpro/shrek/*1234567890*.root -l
/sphenix/u/sphnxpro/shrek/G4Hits-Charm-1234567890-00001.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00001.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00002.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00003.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00004.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00005.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00006.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00007.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00008.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00009.root
/sphenix/u/sphnxpro/shrek/G4Hits-ppminbias-1234567890-00010.root

```

- Pileup stage is crashing / reason TBD
  - cron jobs / regression tests to be implemented under sphnxpro account
  - global tag from DB to be added...
- rucio startup document?
  - [https://rucio.cern.ch/documentation/using\\_the\\_client/](https://rucio.cern.ch/documentation/using_the_client/)
- status of rucio / panda integration @ bnl?
- AOB

**Wed June 15 2022 10AM EDT**

- meeting link
  - <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWFEFBKzYwRXQ5U3NXU1dBZz09>
- PanDA @BNL update - Tadashi (will miss due to ATLAS S&C week)
  - Harvester and Condor up and running on BNL OKD to submit the pilot to BNL CE
  - Configured OKD JEDI and PanDA server to interact with BNL Rucio
  - The entire task/job sequences successfully tested with OKD PanDA server, JEDI, database, and Harvester
  - iDDS will be the next, while Eddie is working on containerization and helm for BigPanda monitoring
  - Will use CERN CRIC for a while as the development status is unclear due to Russia, but that should not be a big issue since it is just an information source
- PanDA Pilot update - Paul
  - Pilot development and testing
    - In ATLAS, the Pilot is tested initially with a standard test job that uses input data replicated on a all SEs -> Easy and quick to test Pilot everywhere
    - When a dev Pilot is stable, it is moved to the Release Candidate (RC) test framework (essentially a special tarball is created and different RC test jobs are submitted by HammerCloud)
    - Note: the RC testing is mainly useful when there are many sites/queues involved, like in ATLAS. If only a couple of queues are available, manual testing basically works just as well
    - Need to identify person that can help with testing new Pilot releases, and ideally who can do some development as well
      - Paul Nilsson is available for developments needed for sPHENIX (up to 10% on monthly average) but has no current access for testing on non-ATLAS resources
        - Recent developments for sPHENIX: dedicated Pilot plug-in where all code is placed relevant for sPHENIX; mv copy tool (in testing)
      - Doug Benjamin is interested in/willing to help with both testing and some development
  - Upcoming development
    - Dynamic selection of copy tool
      - If lustre is not available on current computing node, mv cannot be used and Pilot should switch to xrdcp instead
      - Consult e.g. condor class ad and make decision dynamically – both copy tools should be specified/listed in queuedata.json
- SHREK / sPHENIX workflows update
  - Integration with version control for documentation / reproducibility
    - <https://github.com/klendathu2k/sPHENIX-test-production/tree/test>
    - test, devel and production branches for (1) system testing, (2) workflow development, (3) production documentation
    - TODO: provide link to monitoring, archive input files, update the README.md file to provide a summary of productions, post png image of the workflow graph,...
  - Simple workflows are succeeding
    - <https://panda-doma.cern.ch/tasks/?jeditaskid=121474|121448|121449>
    - confirms that secondary inputs are handled correctly
  - Heavy flavor charm workflow
    - <https://github.com/klendathu2k/sPHENIX-test-production/tree/test/sP22a-hfcharm>
    - 1 HF signal job as primary input
    - 10 minbias background jobs as secondary input
    - Pileup job
    - The minbias jobs are failing in an odd way (failed inputs) leading to the pileup job failing. Debugging...
      - <https://panda-doma.cern.ch/tasks/?jeditaskid=121472|121446|121447>
- 
- Other business

**Wed May 25 2022 10am**

- <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWFEFBKzYwRXQ5U3NXU1dBZz09>
- sPHENIX review: outcomes, news since - Chris
  - Review went well
  - Full report expected next week
- real time logging options in PanDA - Shuwei:
  - <https://drive.google.com/file/d/1XCMMpVXliiY27pJk2W1-k6BI9SXLgOnc/view?usp=sharing>
  - Sdcc has elk servers
  - Needs to be set up in pilot (not Shrek)?
  - Doug is looking into this
- SHREK update - Jason
  - PanDa pilot does not find input files for pileup (known issue - will be solved by Paul)

- PanDA@BNL update - Tadashi (arrives 10:30)
  - DB's set up in OKD
  - Harvester needs to be ported to OKD (Doug)
  - Cric needs to be done (need to get OKD compatible container - in discussions with the CERN cric development team already)
  - iDDS needs to be done (what is needed here?)
  - Current configuration of OKD servers - (4 nodes 40 Logical cores (2xXeon Silver 4210) - 128 GB RAM)

Wed April 27 2022 10am

- <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- general sPHENIX
  - prmon - issues with event level info (of most interest to sPHENIX, summing of network interfaces, granularity of memory usage info, etc.)
    - perhaps enlist Shuwei (working with prmon for Rubin) to incorporate/feed back some improvements
- sPHENIX PanDA, Shrek status

Wed April 6 2022 1pm

- <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- [SHREK status](#)

Wed Mar 23 2022 1pm

- <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>
- Unfortunately the announcement mail got stuck somewhere, me and Jason had a good nuts and bolts discussion
- [PanDA Update](#)
- Attendees please help correct/complete the notes!

Wed Mar 9 2022 1pm

- sPHENIX update
  - Will have to solve X509 -> token issue
- PanDA update
  - [sPHENIX production sys](#)
  - Next step: focus on pp production with full downstream simu/reco workflow, and data flow (getting files back)
    - address how to add Rucio registration? Jobs write directly to shared file space. So PanDA knows instance locations as well as LFNs. PanDA does the necessary registration.
- Rucio update
  - The rucio service / catalog is available outside BNL and configured for PanDA(account and permission). Tadashi confirmed that the panda server@cern can reach it with the provided rucio configuration file.
  - Data storage location
    - RSE: BNL\_PROD
    - -----
    - protocol: posix file://
    - path: /sphenix/lustre01/sphnxpro/rucio/
    - Protocol: https (s3 minio)
    - hostname: sphenixs3.rcf.bnl.gov
    - port: 9000
    - path: /storage
- Target for next meeting in 2 weeks
  - be ready to switch simu production to PanDA. Would be simpler than what Chris is currently doing.
- OKD PanDA update? Doug is working on it. Hardware is there. Need an update from Doug.
  - (Doug) - PanDA DB server installed - issue with configuration. One of the configuration script fails. Will need to work Tadashi to debug issue. Waiting on OKD instance to be fully deploy before proceeding w/ PanDA
  - This proceeds in parallel with current running which uses CERN DOMA instance.
- EIC update
  - Momentous news!

6. Recommendations:

The panel unanimously recommends ECCE as Detector 1. The proto-collaboration is urged to openly accept additional collaborators and quickly consolidate its design so that the Project Detector can advance to CD2/3a in a timely way.

The panel supports the case for a second EIC detector, however, given the current funding and available resources, the committee finds that a decision on Detector 2 should be delayed until the resources and schedule for the Project detector (Detector 1) are more fully realized.

- 
- ECCE has close ties with sPHENIX - ECCE jobs are a lot like (simpler) sPHENIX jobs
  - Join ECCE S&C and begin shadowing PanDA / Rucio
  - Interaction / Integration of ATHENA S&C TBD

Wed Feb 16 2022 1pm NEW TIME in NPPS zoom room


<https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>

- Attendees please help correct/complete the notes!
- Brief update from Jason on PanDA
  - PanDA jobs timing out, put in ticket, getting the needed attention
    - Chris - sounds similar to problems I've seen
  - delving into doc describing chaining jobs (pchain is the new/preferred way)
- Focus on Rucio integration
  - sPHENIX Rucio instance up and working, ready for Chris' use
  - first use case of file catalog demonstrated
  - draw distinction between what Chris and Jason are doing, Jason needs to work with Rucio within PanDA workflows
    - server
      - we probably need PanDA core (Tadashi, Fernando) for this, telling PanDA to talk to BNL sPHENIX Rucio



- pilot/worker node
    - client available in cvmfs
    - don't need data mover or Rucio at pilot/worker level for now. The payload handles ingress/egress of data.
    - Rucio for now is limited to file catalog, the payload interacts with Rucio to load to file catalog
    - so Jason doesn't have to worry about Rucio for the moment (except for it appearing in the payload scripts, and authentication)
- sPHENIX has unique filenames which helps with Rucio implementation. Discussing the scopes that are needed.
  - sPHENIX might modify their naming convention
- Rucio passes to sPHENIX full filepath including protocol, directly usable
  - filename/path can be user defined, or algorithmically generated (and not very human readable). sPHENIX probably prefers the latter (as does ATLAS, except for T0 export; both can coexist).
  - sPHENIX discourages use of ls to find files. Data discovery should be done through file catalog.
- have to sort out Rucio client authentication (which will happen in the pilot)
- Rucio folks will create a 'panda account' for PanDA access to Rucio
- on PanDA server matters, the contact (at least to start with) is Tadashi
- access to the Rucio web UI from outside BNL is needed from the beginning (PanDA server access from CERN, client access to the file catalog from all over)
- sPHENIX interested to use Rucio to store metadata (beyond file catalog)? Maxim is thinking about this. sPHENIX hasn't discussed its metadata plans yet (FAIR etc.)
- EIC
  - could make sense to 'follow along' with Rucio for EIC
  - Chris prefers collaboration, more people working gets better results
  - agreement it makes sense to pursue both at the same
  - good question for Tadashi et al: have an independent EIC PanDA instance from the beginning or down the road (a bit)? sPHENIX will want an sPHENIX-only production instance.
- Please what are the doc links!
  - <https://panda-wms.readthedocs.io/en/latest/>
  - <https://rucio.cern.ch/documentation/>
- Next meeting Wed Mar 9 1pm. Should be news on OKD then too.

**Wed Feb 9 2022 10am** in NPPS zoom room <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>

- some (e.g. Torre) have to leave at 30min for a [GPU/HPC meeting with US ATLAS content](#)
- PanDA update from Jason
  -  Nemo Job Submission System
- What Jason is showing looks good.
- First step and target: replace first stage of Chris' production system, evgen.
  - Rucio comes in here for the outputs. Current prod script adds files to current sPHENIX file catalog so this is a good place to introduce Rucio.
- Next step: DAG defined workflows, which brings in iDDS/pchain
- When a question comes up, use Mattermost distributed software channel, and don't be shy!
- Let's please use this distributed software channel for data management matters (ie Rucio integration) as well.
- Next meeting in about a week, probably in a different timeslot.
- Tony will look for a time that is compatible with Vincent attending our meetings.

**Wed Feb 2 2022 10am** in NPPS zoom room <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>

- nice recommendation from the DOE in the just-concluded US ATLAS review... sPHENIX, EIC and Rubin were mentioned explicitly... we're now mandated by DOE to propagate PanDA and Rucio :-)

■ US-ATLAS should continue exploring ways of leveraging synergistic efforts beyond the Energy Frontier into other scientific communities, for example in deploying Rucio and PanDA. Provide updates to the funding agencies annually.

- 
- update on MDC2, needs - Chris
  - still problems writing from shared pool to Lustre. Following up.
  - met with Vincent on Friday, invited him to this meeting. Rucio concepts map well onto what we want in sPHENIX (plus some goodies).
    - need to establish where/how propagation of file info to Rucio happens.
    - metadata capabilities of Rucio - Maxim interested in this
    - Rucio instance 'up and running' but needs a week or two to be brought to a usable state
- PanDA/Prodsys update - Jason et al
  - Jobs submitted previously (by me) have been failing.
    - They take ~1 hour to submit...
    - Additional 3 hours for Panda to recognize jobs fail / then panda retries / fails again...
    - Reasons are not really clear to me based on the PanDA web pages for the jobs
  - Obtaining a token...
    - Able to authenticate with SDCC credentials
      - Followed link provided when job submitted / follow instructions to authenticate and link to (new) PanDA account
    - Submitted jobs...
      - ERROR : Not good. ERROR : token-based authentication failed on the server side with invalid member in sphenix
  - Jason will
    - apply for sPHENIX SDCC account
    - ask to be added to sPHENIX for purposes of PanDA recognizing him as sPHENIX (via ticket)
    - put in a ticket about (what sounds like) a pilot provisioning issue
    - soon set up github repo :-)
- Rucio integration discussion - Chris, Vincent et al

- Alexei - very strong Rucio team in ATLAS, part of their job is promoting Rucio outside ATLAS (also cf. new DOE review outcome above), probably worth discussion. And some are explicitly ‘beyond ATLAS’ people (DOMA).
- Chris - Vincent very willing to work with us which is great (and he is a long-time member of the team Alexei is mentioning, as is Cedric who is willing and interested to help also)
- Harvester, authentication, OKD etc update - Tony et al  
(Notes by Doug)
  - VM’s requested for Harvester (after getting details from Fa-Hui Lin (ATLAS)). Working out details for SDCC HTCondor schedd’s (with Farm team lead by Chris H, Xin is primary HTCondor admin for SDCC), VM’s should be available very soon no precise time scale yet (but very very soon < 1 week)
  - Received firm delivery time scale for production OKD hardware. Will have them available to sPHENIX by (Tony’s correction) April.
  - sPHENIX storage server head nodes now have a delivery date and tracking information. We in SDCC are working through the details on deployment to provide an accurate schedule.
  - All DB’s testing of sPHENIX Buffer Boxes (3) and ATLAS HPSS and Lustre is complete.
- EIC plans and synergy - Kolja et al
  - need EIC account, or sPHENIX account?
  - Torre - +1 for starting soon to jump the hurdles of getting things set up to work with EIC identity
  - Chris - should work for EIC in parallel with sPHENIX
  - Kolja will steal Jason’s setup once he has things going and go from there in a week or so
  - Kolja will make sure both ECCE and ATHENA are aware of what we’re doing in prototyping PanDA for EIC
- Next meeting in same slot in 1 week

**Wed Jan 26 2022 10am** in NPPS zoom room <https://bnl.zoomgov.com/j/16157150845?pwd=NXNqTi9ZWEFBKzYwRXQ5U3NXU1dBZz09>

- Organizational discussion - Chris et al
  - 52 weeks to data!
  - Chris and Jason have successfully submitted sPHENIX simu PanDA jobs
  - simu jobs as the first real world usage. Straightforward in their data handling and workflow.
  - Chris - “I think all the pieces are in place”
  - authentication - Doug
    - Chris has a token, used also by Jason. Chris was told “yes” others can use it, as a “production identity” (as ATLAS does too)
    - or, Jason et al should get their own keys?
    - assume Harvester will submit condor jobs, handles authentication itself
  - Harvester is SDCC territory, Doug will carry that to SDCC to be clear
- Tasks/priorities
  - list below is accurate
- Status, open questions/issues
  - PanDA running - Jason
    - Jason running at/from SDCC, updated instructions below
    - successfully submitting jobs as below, with sPHENIX input files
    - submitting tasks as well, using the panda monitor task page
    - how to retrieve outputs? Dir name is passed down as a parameter. Generated on local worker disk and then copied out.
    - next: more sophisticated MC example to run
    - moving to a more production campaign context. Rucio comes in when?
      - Chris: Rucio only as a file catalog at first, not for data movement.
    - Chris - will show you how to how to run interactively so PanDA doesn’t need to be always in the dev cycle
  - monitoring status - Amol
    - feedback
      - Need as function of event number. Not time.
      - Needs to be fast. Need to do binning, not show 20k points etc. Need to average over N data points to reduce number of plotted values and to reduce the jitter.
      - Full spec of what’s on the page has to be in the url. No POST, no javascript hiding configuration. (ie the present behavior has to be changed)
  - Rucio: how to organize technical discussions between sPHENIX, NPPS, SDCC on Rucio? We have a designated SDCC contact, Vincent (CERN based). Use these meetings? Intimately connected to panda/prodsys. Yes we agree we’ll use these meetings, starting next week
  - processing/storage resource issues?
  - EIC? use these meetings to advance PanDA/Rucio for EIC (big overlap with sPHENIX) or keep purely sPHENIX focused?
    - Yes, the more the merrier
  - **authentication. one production token OK?**
- Next meeting
  - broaden scope to Rucio, invite (at least) Vincent, meet in one week (perhaps exceptionally, 2 week cadence when things stabilize?)
  - in 1 week, tentatively in same slot:
    - Rucio integration discussion, updates on PanDA, output handling, monitoring, clarifying authentication

## Task/priority list

- sPHENIX production jobs running via PanDA at SDCC
- monitoring integration
  - Need as function of event number. Not just time.
- job choreography. Role for iDDS/pchain in supporting DAGs. Event service for parallel populating multi-core (covered by iDDS)?
  - topic for next meeting?

- optimizing file size for HPSS storage. So files are in sequence but start/end dynamic. Managing ‘filling’ these files is a choreography question.
- Rucio integration stands high on the priority list. Is it needed/desirable for SDCC running? Yes. As a file catalog if nothing else. SDCC instance is working, should start using it.
- once SDCC PanDA production running is working, try out running in Google
- down the road... NERSC, but takes sPHENIX action/attention
- BNL PanDA instance: may have the OKD hardware in Feb. We (PanDA team) said May as target to have it up.

# Setting up and working with PanDA

## Install PanDA client (on laptop or any machine):

**[virtual environment]**

Your system should have bzip2, wget and glibc installed. If not (and you have root privilege...)  
\$ yum install bzip2 wget glibc.i686

*On 32bit systems:*  
\$ wget <https://repo.anaconda.com/miniconda/Miniconda3-latest-Linux-x86.sh>  
\$ chmod +x ./Miniconda3-latest-Linux-x86.sh  
\$ ./Miniconda3-latest-Linux-x86.sh  
\$ source ~/.bashrc

*On 64bit systems:*  
\$ wget [https://repo.anaconda.com/miniconda/Miniconda3-latest-Linux-x86\\_64.sh](https://repo.anaconda.com/miniconda/Miniconda3-latest-Linux-x86_64.sh)  
\$ chmod +x ./Miniconda3-latest-Linux-x86\_64.sh  
\$ ./Miniconda3-latest-Linux-x86\_64.sh  
\$ source ~/.bashrc

- Notes:*
- Before running the Miniconda3 script, be sure to unset your PYTHONPATH environment variable (if it was set).
  - When you run the Miniconda3 script, you will be prompted for a path into which the environment will be installed. By default it will install under \$HOME/miniconda3. The rest of this document assumes the default location.
  - If you'd prefer that conda's base environment not be activated on startup, set the auto\_activate\_base parameter to false:  
\$ conda config --set auto\_activate\_base false
  - TODO: How to activate it if the above is set to false...

**[mandatory]**  
pip install panda-client

- Notes:*
- This installs the panda client (software and configuration files) under \$HOME/miniconda3. YOu may choose an alternate path. (If you do, modify the path to miniconda accordingly in the instructions below...)

## How to initialize proper environment for PanDA client:

**[optional, point local client to the DOMA Panda server]**  
export PANDA\_URL\_SSL=https://ai-ids-01.cern.ch:25443/server/panda  
export PANDA\_URL=http://ai-ids-01.cern.ch:25080/server/panda

**[optional, select an auth method if other then VOMS/Grid certificate + ATLAS VO is used]**  
export PANDA\_AUTH=oidc  
export PANDA\_VERIFY\_HOST=off  
export PANDA\_AUTH\_VO=sphenix

**[mandatory]**  
source \$HOME/miniconda3/etc/panda/panda\_setup.sh

export PANDA\_URL\_SSL=https://pandaserver-doma.cern.ch:25443/server/panda  
export PANDA\_URL=http://pandaserver-doma.cern.ch:25080/server/panda  
*#export PANDA\_URL\_SSL=https://ai-ids-01.cern.ch:25443/server/panda*  
*#export PANDA\_URL=http://ai-ids-01.cern.ch:25080/server/panda*  
export PANDA\_AUTH=oidc  
export PANDA\_VERIFY\_HOST=off  
export PANDA\_AUTH\_VO=sphenix

source \$HOME/miniconda3/etc/panda/panda\_setup.sh

*Note:*

The file \$HOME/miniconda3/etc/panda/panda\_setup.sh should resemble...

```
export PATH=/home/pinkenbu/.local/bin:$PATH
export PYTHONPATH=/home/pinkenbu/.local/lib/python2.7/site-packages${PYTHONPATH:+:$PYTHONPATH}
export PANDA_CONFIG_ROOT=~/.pathena
export PANDA_SYS=/home/pinkenbu/.local
export PANDA_PYTHONPATH=/home/pinkenbu/.local/lib/python2.7/site-packages
```

## How to prepare a submission script:

Depending on the scale and nature of the planned computation work a user may choose to run each job as an individual payload or unify few or thouthands jobs into a single or few tasks. PanDA supports handling a number of jobs united into tasks as a single object which makes it easy to monitor and control them (e.g. <https://panda-doma.cern.ch/task/423/> ). When an individual job in a task is submitted into the computation backend, PanDA can provide an index



as a job parameter. This index could be translated into a random number seed, events index, etc to make unique, repeatable initial conditions for each job in a task.  
In the script below the input

```
import sys
from pandaclient import Client

inFileList = ['000000'] # How is the inFileList element made available in the running job? (%IN ?)

taskParamMap = {}

taskParamMap['vo'] = 'wlcg'
taskParamMap['site'] = 'BNL_OSG_SPHENIX'
taskParamMap['workingGroup'] = 'sphenix'

taskParamMap['nFilesPerJob'] = 1
taskParamMap['nFiles'] = len(inFileList)
taskParamMap['noInput'] = True
taskParamMap['pfnList'] = inFileList

taskParamMap['taskName'] = 'test sPHENIX submission'
taskParamMap['userName'] = 'A user Name'
taskParamMap['processingType'] = 'step1'
taskParamMap['prodSourceLabel'] = 'test'
taskParamMap['taskType'] = 'test'

taskParamMap['transPath'] = 'https://atlpn.web.cern.ch/atlpn/bash-c'
taskParamMap['taskPriority'] = 900
taskParamMap['architecture'] = ''
taskParamMap['transUses'] = ''
taskParamMap['transHome'] = None
#taskParamMap['coreCount'] = 1
#taskParamMap['ramCount'] = 4000

taskParamMap['skipScout'] = True
taskParamMap['cloud'] = 'US'
taskParamMap['jobParameters'] = [
 {'type':'constant',
 'value':"""source /cvmfs/sphenix.sdcc.bnl.gov/gcc-8.3/opt/sphenix/core/bin/sphenix_setup.sh -n;
git clone https://github.com/sPHENIX-Collaboration/macros.git;
... I think this will echo the inFileList element ...
echo %IN
cd macros/detectors/sPHENIX;
echo 'here comes your environment';
printenv;
echo running root.exe -q -b Fun4All_G4_sPHENIX.C\\\(10\\\)
root.exe -q -b Fun4All_G4_sPHENIX.C\\\(10\\\)
echo "script done";
""""
 },
]

print(Client.insertTaskParams(taskParamMap,verbose=True))
```

If one needs to submit multiple jobs, the inFileList should be updated to contain as many indexes as needed and an additional parameters provided to the jobParameters entry, e.g.:

```
...
inFileList = [('0000'+str(i)) for i in range(10,20)]
...

taskParamMap['jobParameters'] = [
....
root.exe -q -b Fun4All_G4_sPHENIX.C\\\(IN/L\\\)
```

## How to submit:

python ./your\_submission\_script

Then a user gets a temporary link as this:  
[https://panda-iam-doma.cern.ch/device?user\\_code=AZ5VYC](https://panda-iam-doma.cern.ch/device?user_code=AZ5VYC)  
Open it in the browser and follow instructions. (Requires authentication, see below). At the end a user gets message like:  
(0, '(I01\nL460L\ntp0\n.)'  
(0, (True, 460))

Where 460 is the id of created task which could be monitored.

NOTE: If you do not have a previously cached token, you will be prompted to go to a web page to authenticate...

INFO : Please go to https://panda-iam-doma.cern.ch/device?user\_code=XXXXXX and sign in. Waiting until authentication is completed  
INFO : Ready to get ID token?  
[y/n]

Note: The task IDs should probably be captured and associated to the job in some persistent form, so that one can launch / manage / monitor multiple productions and not lose track of their status...

### Authentication

Following the link from submission will take you to a sign in page w/ Panda-doma. You have two options: “Your ID provider” and “Create an Account”. Selecting “Your ID provider” takes you to a page where you can authenticate with BNL SDCC credentials (and others). It required setting up 2FA (using google authenticator, or other such app). Once authenticated, I needed to link to a (new) Panda-doma account. You can request one through this dialog, providing your preferred username, email account, and justification for the request. After submitting the request you’ll be asked to confirm (an email will be sent to your linked email account).

### How to monitor:

Monitoring is provided here: <https://panda-doma.cern.ch/>  
URLs usually encode parameters to shortcut access to the information needed. E.g.: <https://panda-doma.cern.ch/task/460/>  
... where the task number is the number returned by the script above.

NOTE: Will ask for an authentication method (CERN credentials, gmail creds (I think), or github credentials). I was able to give my github account, and it worked w/out issue.

### How to kill a task:

```
[podolsky@localhost ~]$ python
>>> from pandatools import Client
>>> Client.killTask(466)
(0, (True, 'command=kill is registered. will be executed in a few minutes'))
```

### How to retry a task:

```
[podolsky@localhost ~]$ python
>>> from pandatools import Client
>>> Client.retryTask(466)
(0, (True, 'command=kill is registered. will be executed in a few minutes'))
```

### How to retrieve output datasets:

???????

### How to find out why jobs failed:

???????

## PanDA Rucio integration

- Jan 20 2022 SDCC liaison meeting: Rucio service for sPHENIX is up and running at SDCC
- Jan 27 MDC2 meeting: contradicted the above. Stated there that there is no service at present, can be brought up in next weeks if needed, to be clarified in discussion with Vincent.
- Contacts to work with us on integrating Rucio are Vincent and Matt, Vincent the lead contact

----

## Action Items / Issues

- ☒ Lifetime of authenticated tokens, are they long enough? if not extremely long is renewal automated?
- ☒ ~~Is there a remaining issue here~~
- ☐ Do we have the role management we need, ie production role? Does everyone who wants to run jobs need a token?

- ☐ How do we fulfill the need for user-level rights, e.g. production rights (as VOMS is used in ATLAS)?
- ☐ where are such rights checked and the user informed? e.g. ‘in Shrek’? Or Shrek invokes something provided by PanDA?
  - ☐ By default, shrek handshakes with the panda server via pandaclient api “hello()” function. Should trigger authentication if required.
- ☐ Open for all sPHENIX users? (Suitable for complex workflows / advanced users)
- ☐ **Is there a remaining issue here**
- ☐ High / Low / Medium priority condor queues... (--site option... in prun)
  - ☐ Direct end-user use of prun by multiple people (e.g. different subdetector people)
  - ☐ There is need for a **quick-turnaround PanDA queue** (or condor config) for use by human beings submitting prun jobs and wanting quick results
  - ☐ High prior queue in progress...
- ☐ sPHENIX data transfer script needs to become Rucio aware [80% there...]
- ☐ storage of artifacts in git should be to a database (ie persistent git service like the BNL git service). Not github, not appropriate. Not local disk which is not compatible with distributed computing.
  - ☐ Rethink repo sizing
  - ☒ ~~Action: switch from local git to using BNL git service.~~
    - ☐ Test Authenticate via access token
    - ☐ TODO update quickstart instructions for setting up the submission area
    - ☐ TODO: Need to ensure git push does not required auth (should be able to achieve this by pushing to BNL git over ssh... TBD)
- ☐ How to handle tons of transient files with a life time of days weeks (the actual distortion correction machinery)?
  - ☒ ~~plan is to record them in rucio, for their lifetime (which is no more than a ~week)~~
    - ☒ ~~in principle, these are intermediate files w/in a run’s reconstruction workflow~~
  - ☐ Does rucio file deletion really mean permanent deletion, or just hiding? ie does the DB grow due to ‘deleted’ files?
    - ☐ has implications for the recording of high rate (~15ms) TPC calib files
    - ☐ **clarify rucio behavior**
      - ☐ Default is soft delete. Hard delete (db entry) can be enabled / attached to a “transient scope”. Vincent will discuss further with Chris.
    - ☐ if rucio actually does save them in perpetuity in the DB, reducing the file count has benefit. Could e.g. zip the files as part of the workflow with ~second granularity, gain an order of magnitude or two on file count. Rucio understands zipped files, so does ROOT.
  - ☐ **PanDA can (should?) mark transient files as not needed any more, at which point they can be archived (to tape) and removed from disk. Dispatch blocks (datasets) have limited lifetimes.**
- ☐ we agreed a while back that near real time job log access (as implemented for Rubin) is something sPHENIX should have. Anything happened on this?
  - ☐ Fluentd service @ SDCC needed in order to activate near real time job logs for sPHENIX... Xin can setup on OKD.
- ☐ Shadow universe for condor... ETA 2wks to setup. [2 weeks ago... Xin and/or Chris will ping] Blocks scale tests. ETA 2 weeks ~mid april.
- ☐ Need scheme for dataset/file/job management compatible with commissioning now and ~~production later~~
  - ☐ eg RHIC-2023-AuAu200 / RUN-NUMBER-12345
- ☐ Specify space requirements on worker node? Pilot double checks space before launching on the node. prun option.
- ☐ Add tp sphenix gitea the panda queueconfig files, which are used to deploy panda services in OKD. Xin will need access to the sphenix gitea repo.
- ☐ Monitoring (what, how)?, prmon turned out to hang jobs occasionally in the sim production - its use is not an option.

## Closed Action Items

- ☒ ~~what are the limiters on the ~20 jobs/min dispatch rate and are they reducible. The latency may not be significant for daily prod job submission but it could be relevant for quick turnaround submissions in other contexts.~~
  - ☒ ~~task submission latency reported as ~4sec on Mar 1 in which case we are OK~~
- ☒ ~~is rucio being populated with transient (semaphore?) files for the sake of using rucio notification messaging? Not what rucio is for. Is there a better way? (others free to correct/improve this item)~~
  - ☒ ~~Depends on the design of the process which builds the dataset? (All semaphores required to create the DS vs all semaphores to close the DS).~~
  - ☒ ~~need a mechanism to inform shrek that a run is completely available (transferred).~~
  - ☒ ~~**Use notification from Rucio of dataset closure when fully populated/transferred? Would be a good solution**~~
  - ☒ ~~Discussion w/ Chris. The buffer boxes will enter into DB when they have transferred all files to lustre. This signals that a dataset can be created, populated and closed. Shrek (donkey) can submit work when the dataset is closed.~~
- ☒ ~~Submission, particularly of production jobs which are always in bulk, should be done in bulk, so that overheads are amortized over a large number of submitted jobs.~~
  - ☒ ~~It is being done in bulk. Reported ~4s latency is for a task submission.~~