

DUNE Operations Meeting Week Of 2023-02-19

Call details are on <https://indico.fnal.gov/event/63079/>

General Information

- The purpose of the meeting is:
 - to report significant operational issues (i.e. issues which can or did degrade experiment or site operations) which are ongoing or were resolved after the previous meeting;
 - to announce or schedule interventions at sites or services;
 - To inform about recent or upcoming changes in the experiment activities or systems having a visible impact on sites;
 - to provide important news about the middleware;
 - to communicate any other information considered interesting for DUNE operations.
- Any topic requiring more than 5 minutes of discussion should be noted and then discussed at another venue
- Whenever a particular topic needs to be discussed at the operations meeting requiring information from sites or experiments, it is highly recommended to announce it by email to the dune-computing-sites list (at listserv.fnal.gov) to allow the meeting coordinators to make sure that the relevant parties have the time to collect the required information, or invite the right people at the meeting. It is requested that such topics be announced no later than COB the Friday before the next meeting."
- Relevant contacts for test beams, services, and sites below should replace the "NC" either with a status or with "NTR" (nothing to report)

Monday 19 Feb 2024

Attendance

S. Timm, A. McNab, J. Perry, E. Pennacchio

Production Calendar:

<https://calendar.google.com/calendar/u/0/r/month?cid=i7ljnj9ilk70hin7cr1l2e1kj4@group.calendar.google.com&pli=1>

Production graphs and slides

Doug is submitting workflows (HD)---> going on smoothly

Test workflow for VD ran last Friday (1558) : 1000 jobs, 117 failures related to Pandora validation file. Understood: jobs ran on one event only, for that event the validation tree was not created → jobs tagged as failed. If more events are processed the error disappears → can start also VD workflow submissions.

Significant number of held justin jobs over the weekend (see second link) stalled at submit Due to (we think) communication problems between justin and condor_schedd on the condor submit. Condor_schedd is known to have been oom-killed at least once, still investigating,

Also monitoring issues—not all stuff from global pool is getting to Rucio or Kibana
Also had to solve problem of missing sites in the monitoring, think this is fixed again now.
Need to extract throughput numbers out of justin for what's already been done.

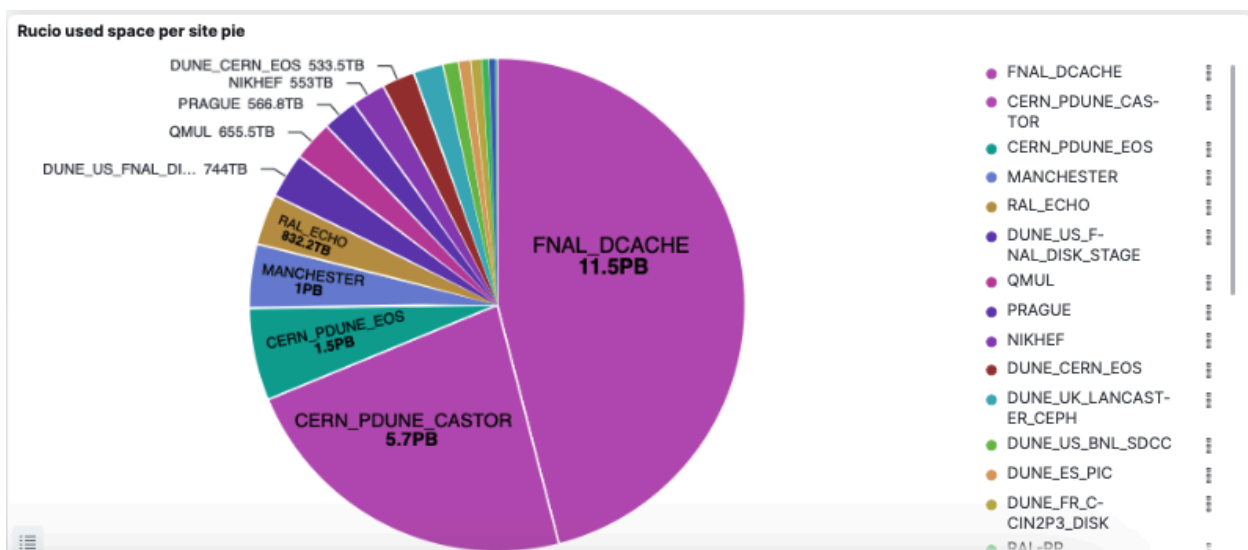
need to identify monitoring links (as discussed in slack):

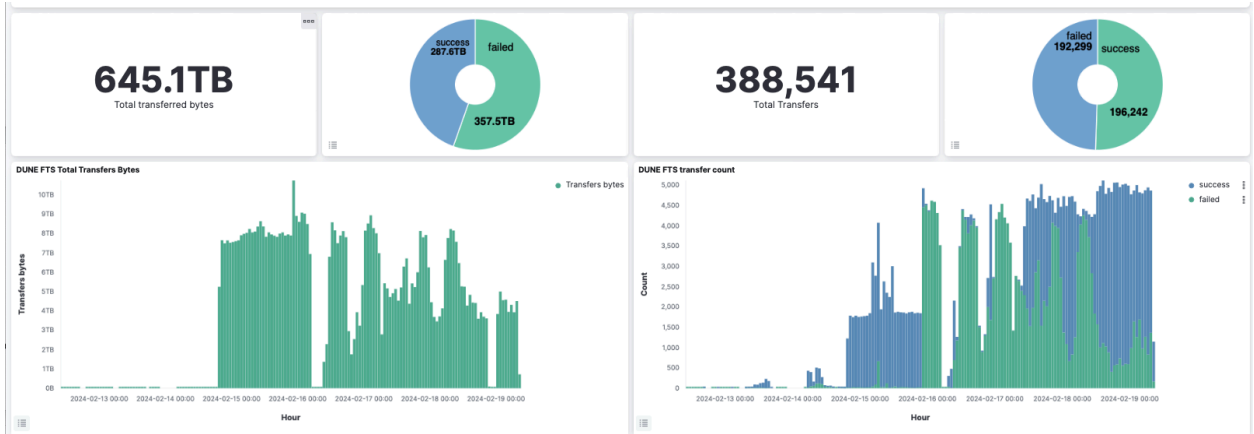

<https://fifemon.fnal.gov/monitor/d/000000053/experiment-batch-details?orgId=1&from=now-21d&to=now-5m&var-experiment=dune&viewPanel=1>

<https://lsvip.fnal.gov/monitor/d/000000053/experiment-batch-details?orgId=1&var-experiment=dune&var-pool=dune-global&var-pool=fifebatch&viewPanel=13>

links to monitor resources usage (CPU, memory) ?

Data management graphs and slides





Announcements of Significant Operational Issues:

NFS->CEPH THIS WEDNESDAY 21 FEB: /dune/app will be *UNAVAILABLE* all day.

Running Expt. Round Table

NP02: NC

NP04: NC

ArgonCube 2x2: NC

TOAD: NC

Iceberg: NC

DUNE Services Round Table:

Data Management: 2 major outstanding issues with Rucio—one a config change to re-establish srm<->davs interoperability (without that can't copy anything off of FNAL_DCACHE). Also need to patch the DeclaD to be able to declare metacat datasets as well as rucio datasets. Until this is done can't do any more production using poms/jobsub or ingest new raw data. Also requested doubling of space of DUNE_US_FNAL_DISK_STAGE to 2 PB

Authentication: Trying to push CERN on VOMS-Admin issue.

Global Pool: See above—with out-of-memory issues on justin-prod-sched01

Monitoring: off-site dune global pool stuff is not coming to Kibana. And dune GP sched not reporting to GRACC..

Workflow (justIN):

justIN version 00.13.06

The central justIN services and the justin commands in cvmfs are now at version 00.13.06. If you use the justin command, please make sure you are using that version with justin version . Earlier versions will produce SSL errors as the certificates have been updated and a hostname changed. Recreate your environment and run setup justin .

This version has some urgent new features and various fixes:

- There are downloadable CSV and JSON files with stats for the processed input files and uploaded output files of workflows/stages. See this page for details: https://justin.dune.hep.ac.uk/docs/download_statistics.md
- --lifetime-days is required if creating a stage which has output to a Rucio dataset which does not yet exist (which justIN will create for you, with that lifetime)
- When creating a dataset, justIN now uses the RSE Expression DUNE_US_FNAL_DISK_STAGE - if you want to avoid this you need to precreate the datasets with your preferred rule(s).

- justin-test-jobscript now faithfully creates the environment seen by jobs again, including the new variables \$JUSTIN_PROCESSORS, \$JUSTIN_RSS_MB and \$JUSTIN_WALL_SECONDS
- An experimental feature has been added to the wrapper jobs to help identify partial failures during the outputting phase. If a job has multiple output files and some succeed but then some fail to upload, justIN will retry processing in another job. This can lead to confusing file states in MetaCat. Now output files are declared to MetaCat with dune.output_status = 'recorded'. As each output is successfully uploaded, its dune.output_status is changed to 'uploaded'. Once all of the output files for that job have been uploaded and the job has reported its final success to justIN, then the job updates the dune.output_status of all its outputs to 'confirmed'.

Ken Herner—wondering how reprocessing would work with such a setup, or any kind of multi-stage workflow (e.g. reco1 going to reco2)

Steve: Rucio upload has a multiple file mode where if any of the scheduled transfers fails, Rucio will undo them all.

AWT/ETF:

AWT: <https://justin.dune.hep.ac.uk/dashboard/?method=awt-results>

AWT benefits from the above upgrade to 00.13.06.

DUNE_US_FNAL_DISK_STAGE is now used for outputs. Weekly datasets for upload test files have been created with expirations a week after the end of each week for the rest of the year.

Still need to understand lack of CNAF storage tests.

ETF:

https://etf-dune-prod.cern.ch/etf/check_mk/index.py?start_url=%2Fetf%2Fcheck_mk%2Fview.py%3Fview_name%3Dallhosts

Database Servers: NC

Networking: NC

Sites Round Table

BR-CBPF: NC

CA-Victoria: NC

CA-SFU: NC

CERN: outstanding ticket on EOS

CH-BERN: NC

CZ-FZU(Prague): NC

ES-PIC+CIEMAT: NC

FR-CCIN2P3: NC

IN-TIFR: NC

IT-CNAF: NC

NL-NIKHEF/SurfSARA: NC

UK-GRID-PP:

Edinburgh: NC

Manchester: NTR

RAL-Tier1: NC

RAL-PP: NC

QMUL: NC

Liverpool: NC

Durham: Ticket still open [Glideins got in to durham. But maybe didn't match.](#)

Oxford: NC

Bristol: NC

Sheffield: NC

Imperial: NC

Brunel: NC

Lancaster: NC

Other:

US-BNL: NC

US-FNAL:

Enstore/dCache: Reboot of Dcache servers on Feb 21 to pick up a bug fix for stashcache

Ceph migration: This Wednesday!!!

Jobsub_Lite: New release coming this week. 1.6.1

Batch Clusters: NC

Computing Access: NC

Wiki/Documentation: NC

Alma9 conversion: Coming sooner than we're ready...DUNE-specific Spack tutorial on Mar 1.

US-HEPCloud (NERSC): Still outstanding issue with GPU detection

US-OSG-Opportunistic: NC

Review of github issues <https://github.com/DUNE/dist-comp/issues>