

# ComCam Data processing with PanDA

5 Dec 2024

Brian Yanny

Fermilab and the Vera Rubin Observatory Survey




8m diameter Mirror (M1)  
With 5m diameter hole for M3→  
6.4m effective diameter.

Cerro Pachon, Andes, Chile  
2200 m elevation



ComCam is a 9 CCD subset (144Mp) of the full 189 CCD LSSTCam (3Gpixel) (~5%).



LSSTCam:

The ComCam survey is taking place from Oct 24 - Dec 12, 2024 – a six week commissioning survey (Com = Commissioning) with the full mirrors and telescope in place.

Purpose: To help commission the telescope, optics and electronics, as well as the production software systems.

Scale: About 2000 9 CCD images in 6 filters (ugrizy).

Output: The data will be processed and reprocessed in early 2025 and a release to the public is planned about June 2025.

Data Processing: PanDA (and some HT-Condor)

On the USDF cluster (16K cores at SLAC)

On the UKDF cluster (8K cores at Lancaster, UK)

On the FRDF cluster (8K cores at IN2P3, France)

## Types of Processing using PanDA:

### Cumulative Data release processing:

Weekly tests of latest LSST software stack on all 'good' ComCam data taken to date. Currently done all at USDF.

### Production Processing: (Jan-May 2025):

Continual reprocessing with updated software stack

At FRDF, UKDF, USDF planned resulting in:

“Data Preview 1” public release of all CamCom data in Jun 2025.

# Data Release Processing PanDA and processing setup:

Single Instance of PanDA setup at SLAC, backed by Postgresql 'in the cloud' (CNPG) Database.

PanDA has connections to 16K + 8K + 8K core compute clusters at USDF, FRDF and UKDF. Each site has several PB disk store, either POSIX disk or S3: storage or both.

High speed network also in place between sites – past use by HEP community.

Special note: There is an 'embargo' period on the ComCam data where raw and processed images must remain isolated from the outside world. This embargo rack is at USDF/SLAC and this is why all initial processing is done at USDF only.

(SLAC cpus can be used for processing, but input/output ends up on the embargo rack S3: storage only). [This is also why no images shown here yet].

# Data Release Processing (DRP) steps 1-7:

Step 1: Instrument signature reduction (flat field, bias, gain correction) runs on each ccd in 2 minutes

Step 2: photometric and astrometric calibration vs. reference catalogs (hours)

Step 3: coaddition of all exposures (visits) that target the same piece of sky (tract) measure photometry of all detected object in all visits vs. coadd.

Steps 4-5-6: difference imaging vs. templates – derive time sequences of objects flux (light curves), find objects on the sky which vary significantly in flux or are moving (near Earth asteroids) vs. the background sky. [There is a separate pipeline which also does this quickly, separate from DRP].

Step 7: consolidate statistics of observed footprint and determine S/N of sky.

# PanDA IDDS wfprogress dashboard with ComCam weekly cumulative processing steps (1-7)

username (12) Wen Guan (57) Antonia Villarreal (11) Zhaoyu Yang (113) Jen Adelman-mccarthy (177) Brian Yanny (70) Peter Love (54) Huan Lin (81) Mikolaj Kowalik (29) Dave Mckay (27) Toby Jennings (5) Edward Karavakis (7) Michelle Gower (2)

## Requests:

Show 10 entries

Search: 20241101\_20241120

request id	username	workflow status	graph	workflow name	created on (UTC)	total tasks	tasks	transform type	total files	released files	unreleased files	finished files	failed files
9596	Huan Lin	Finished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T224720Z	2024-11-22 22:47:34.746049	8	Finished(8)	Workflow	13	13	0	100%	-
9595	Huan Lin	Finished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T220731Z	2024-11-22 22:08:06.933942	4	Finished(4)	Workflow	962	962	0	100%	-
9591	Huan Lin	SubFinished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T203237Z	2024-11-22 20:35:30.972595	9	Finished(7) SubFinished(2)	Workflow	6097	6085	12	99.98%	0.02%
9589	Huan Lin	Finished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T172644Z	2024-11-22 17:30:55.106268	7	Finished(7)	Workflow	32020	32020	0	100%	-
9588	Huan Lin	Finished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_hips_20241122T171726Z	2024-11-22 17:18:02.098660	3	Finished(3)	Workflow	3	3	0	100%	-
9584	Huan Lin	Finished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_hips_20241122T162631Z	2024-11-22 16:27:02.147819	3	Finished(3)	Workflow	3015	3015	0	100%	-
9583	Huan Lin	SubFinished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T151549Z	2024-11-22 15:16:58.079541	13	Finished(12) SubFinished(1)	Workflow	454	454	0	99.8%	0.2%
9582	Huan Lin	SubFinished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T031047Z	2024-11-22 03:13:58.321324	12	Finished(6) SubFinished(6)	Workflow	28635	28628	7	99.997%	0.003%
9581	Huan Lin	Finished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T022737Z	2024-11-22 02:27:49.070129	4	Finished(4)	Workflow	4	4	0	100%	-
9580	Huan Lin	Finished	plot	LSSTComCam_runs_DRP_20241101_20241120_w_2024_47_DM-47746_20241122T012757Z	2024-11-22 01:30:04.007578	6	Finished(6)	Workflow	5736	5736	0	100%	-

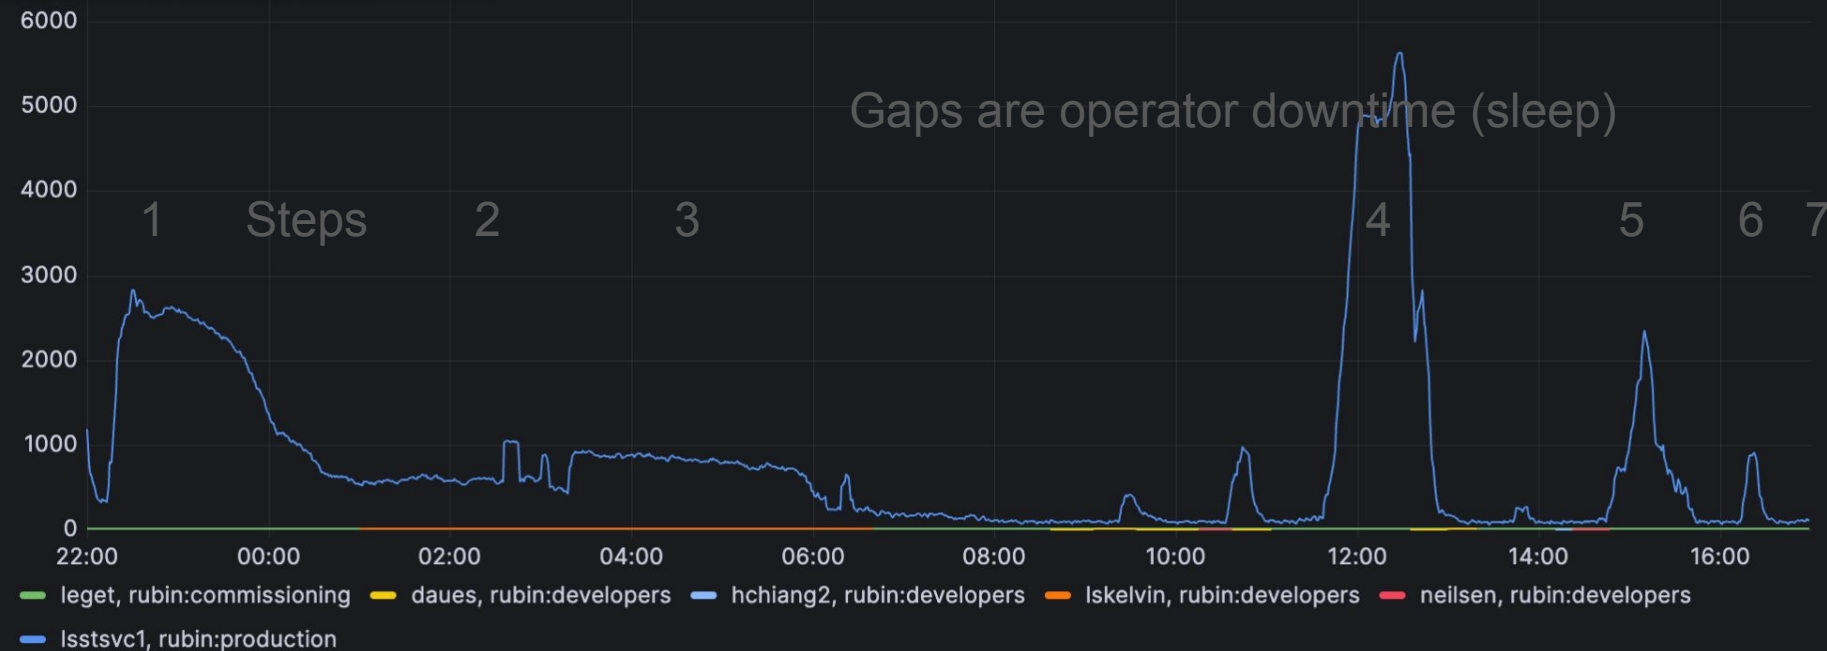
Showing 1 to 10 of 14 entries (filtered from 633 total entries)



Cumulative DRP processing with PanDA currently takes about 1 day step1-7 using 10-6000 cores at USDF.

S3DF Running Cores by User for Rubin Accounts


S3DF Running Cores by User for Rubin Accounts



# PanDA core usage Ramp up and Ramp down is quite smooth

This is a big advance over a couple of years ago when we did the HSC PDR2 test processing campaign:

S3DF Running Cores by User



# Still looking for ways to 'even out the core usage' for efficiency.

S3DF Running Cores by User for Rubin Accounts

S3DF Running Cores by User for Rubin Accounts



# Efficiencies and Scaling

Much of this is intrinsic to how Rubin processing works, with Step1,4,5 having millions of small short duration (2 minute) jobs where step2,3 can have many fewer (hundreds) of long duration (12 hour) jobs.

Need to figure out a way to 'keep all the cores busy' and still handle the wide

Difference in number of jobs vs. run time for the different steps.

(This is not a PanDA issue, but one we are trying to figure out).

One thing we are doing is using 'clustering' to group many short jobs into running as one job in PanDA. Another thing we are doing is 'file zipping' to group together, say 189 CCDs of data together into one file for network transmission with Rucio/FTS-3.

# Conclusions

PanDA working well for early repeat processing of Rubin ComCam dataset!

- Less than 1 day turnaround for testing science code changes. Uses 6000 cores efficiently, smooth ramp up and ramp down.
- Multisite processing (3 sites currently) coming early next year, works in testing
- Looking to scale up – full camera is 21x the data rate. Full annual DRP will take 200 days at 3 sites – 32K cores, 50PB (between the three sites), network.
- Need to work on ‘leveling off’ core usage so that many short jobs and fewer very long jobs can be run using the same compute system efficiently.
- Also need to package data into clusters and zip files for more efficient transmission.