

**Centre de Calcul**  
de l'Institut National de Physique Nucléaire  
et de Physique des Particules

# Update on computing for LSST at FrDF

Rubin-LSST France - LAPP - Annecy

Fabio Hernandez, Quentin Le Boulc'h, Gabriele Mainetti

doc.lsst.eu

# Outline

1

DP2 Processing

2

Characterizing  
Memory  
Requirements

3

Road to DR1

4

Questions &  
Comments



# Pilot phase for DP2


- **Main goal:** Identify bugs in the software stack (v30) ahead of the official production run at USDF
  - Test the processing stages in advance of USDF operations
  - All discovered bugs / unexpected behavior reported via JIRA tickets
  - **Bug tracking → Resolution** Corrections were backported into v30
- Voluntarily limited to ~1000 visits (Cosmos and EDFs) to allow for rapid iterations

| Ticket                   | Stack       | Plots   |
|--------------------------|-------------|---|
| <a href="#">DM-53719</a> | v30.0.0.rc2 | <a href="https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_0_rc2/DM-53719">https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_0_rc2/DM-53719</a> |
| <a href="#">DM-53877</a> | v30.0.0     | <a href="https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_0/DM-53877">https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_0/DM-53877</a>         |
| <a href="#">DM-54061</a> | v30.0.1     | <a href="https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_1/DM-54061">https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_1/DM-54061</a>         |
| <a href="#">DM-54249</a> | v30.0.4     | <a href="https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_4/DM-54249">https://usdf-rsp.slac.stanford.edu/plot-navigator/collection/dp2_prep/LSSTCam/runs/DRP/v30_0_4/DM-54249</a>         |




# DP2 HIPS

Cosmos



EDFS

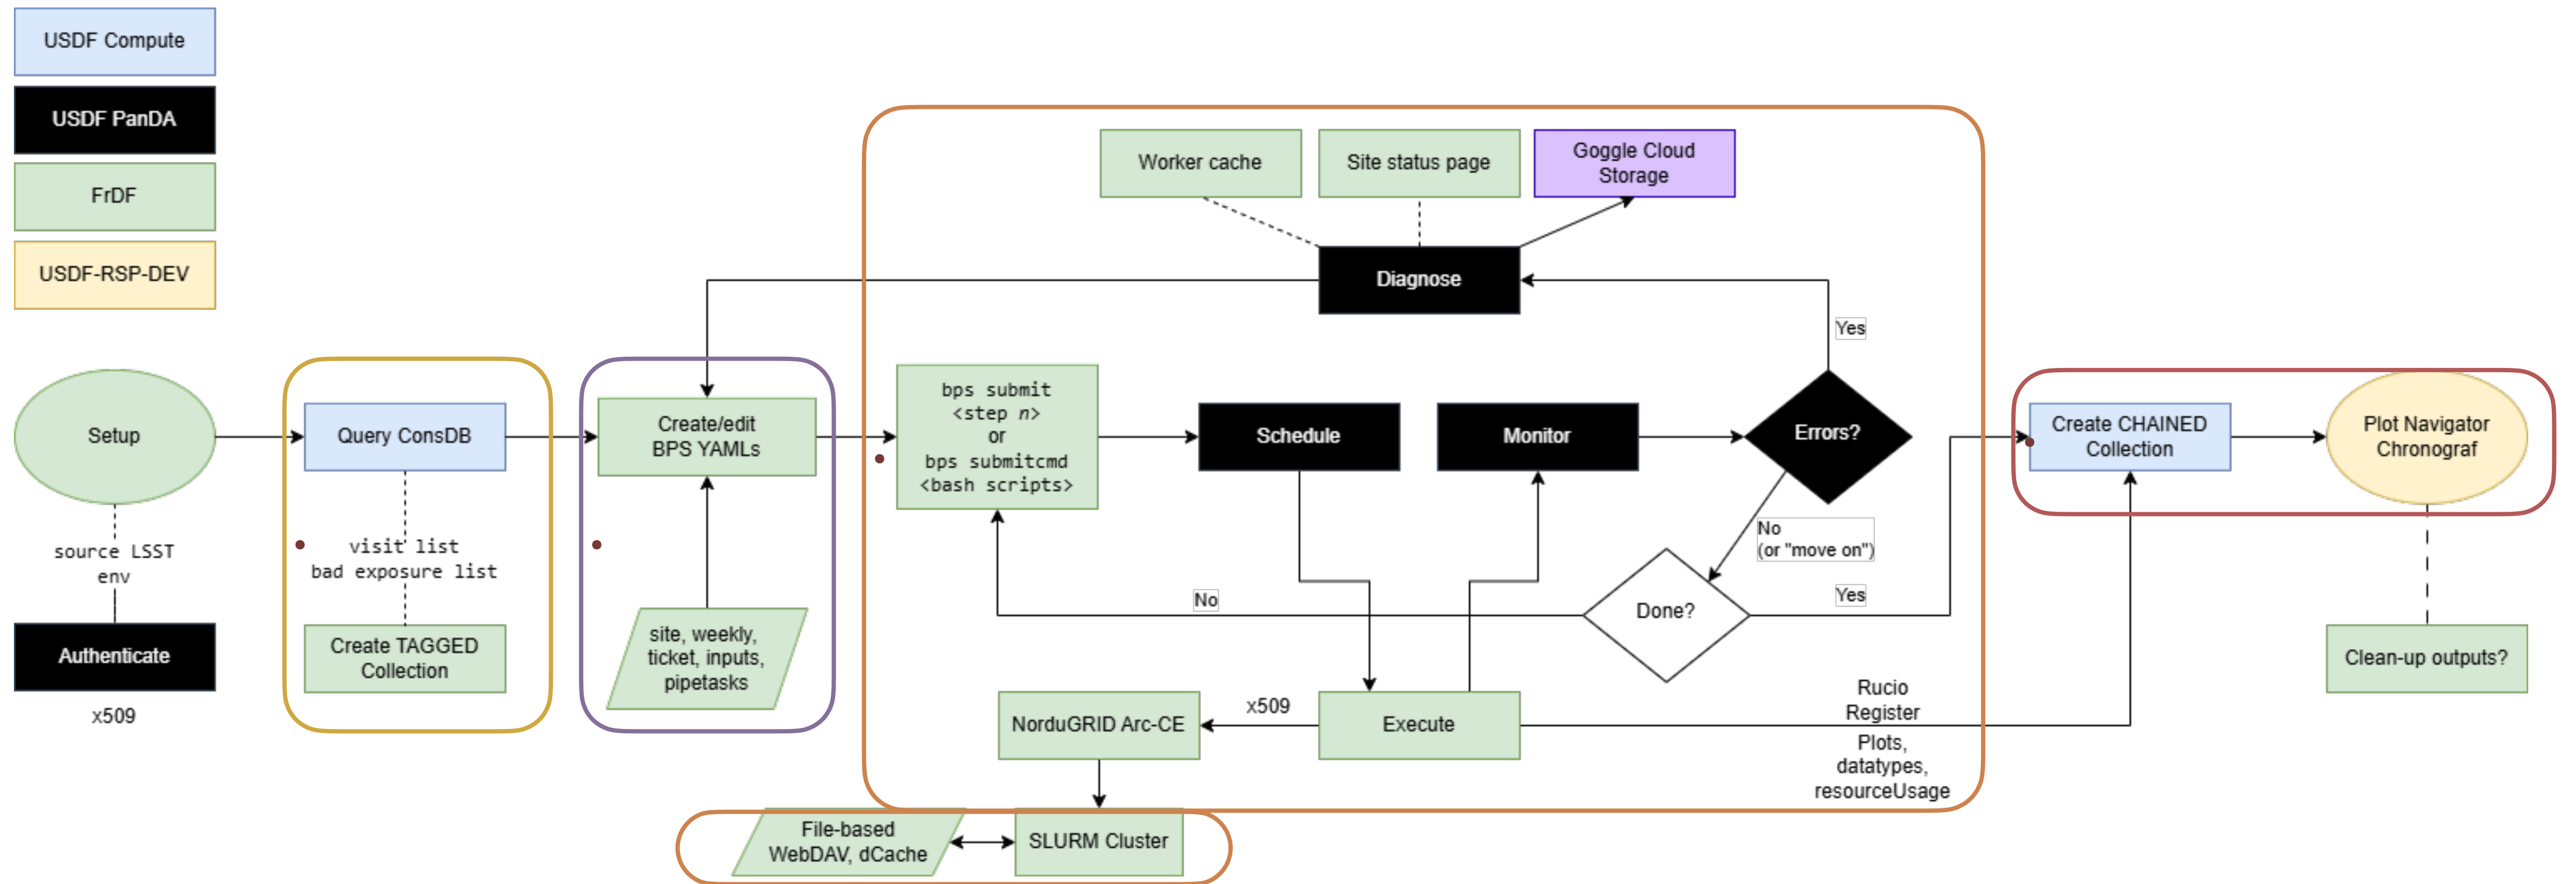


# Processing how-to #1



- We are using our **dp2\_prep** butler repository that contains all ingredients needed for DP2 (raws, calibs,...):
  - <https://rubinobs.atlassian.net/browse/DM-52520>
  - <https://rubinobs.atlassian.net/browse/DM-52680>
- But the repository often needs to be updated before running a campaign:
  - New calibrations to ingest, update of pretrained models, new orbits catalog...
  - Create corresponding tagged collections containing the list of visits/raws to be processed

# Processing how-to #2



# Processing how-to #3

## Execution of DP2 pilot run at FrDF using v30.0.4

**Key details**

Description **Unsaved changes**

- Campaign starting from stage 1 with v30.0.4.rc1, then switching to v30.0.4 starting from stage 3
- Butler repository: dp2\_prep (davs://ccdavrubint.in2p3.fr:2880/pnfs/in2p3.fr/lsst/butler/dp2\_prep). sasquatch disabled.
  - Output collection: LSSTCam/zuns/DRP/v30\_0\_4/DM-54249
- Skymap: lsst\_cells\_v2
- Input collection: LSSTCam/zaw/DM-53877, LSSTCam/calib/DM-53135, LSSTCam/calib\_refcats, skymaps, pretrained\_models, LSSTCam/calib\_fgmcml, mpcorb, sorchache with LSSTCam/zaw/DM-53877 (1067 visits) created from the following collections:
  - LSSTCam/zaw/EDFS/DM-53877: 350 visits
  - LSSTCam/zaw/COSMOS/DM-53877: 717 visits
- Data query: "instrument='LSSTCam' and skymap='lsst\_cells\_v2'"
- For stage 3 and step 4c, a tract constraint is applied: tract > 9800 OR tract in (2398, 2561, 2562, 2563, 2882, 2727, 2728, 2729, 2730, 2731, 2732, 2902, 2983, 2984, 2985, 2238)
- Pipeline definition YAML file: \${DRP\_PIPE\_DIR}/pipelines/LSSTCam/DRP.yaml
- Submission directory: /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249 (FRDF)

The input collection includes three datasets:

- LSSTCam/zuns/DRP/v30\_0\_4/DM-53877 output collection from DM-53877: Execution of DP2 stage2+ pilot run at FrDF using v30.0.0 [IN REVIEW](#)
- LSSTCam/zaw/COSMOS/DM-53877 and LSSTCam/zaw/EDFS/DM-53877, generated from the consDB query up to 6 January 2026 with these constraints:

```
1. LSSTCam/zuns/DRP/v30_0_4/DM-53877 output collection from DM-53877: Execution of DP2 stage2+ pilot run at FrDF using v30.0.0 IN REVIEW
2. LSSTCam/zaw/COSMOS/DM-53877 and LSSTCam/zaw/EDFS/DM-53877, generated from the consDB query up to 6 January 2026 with these constraints:
```

```
cosmos = ("science_program IN ('BLOCK-365', 'BLOCK-407', 'BLOCK-408', 'BLOCK-416', 'BLOCK-417') "
          " AND vignette='NO' AND day_obs >= 20250424 AND day_obs <= 20260106 AND img_type = 'science' AND "
```

```
includeConfigs:
- ${CTRL_BPS_PANDA_DIR}/config/bps_panda_df.yaml
- ${CTRL_BPS_PANDA_DIR}/config/bps_remote.yaml
- ${DRP_PIPE_DIR}/apps/clusterlog/LSSTCam/DRP-clustering.yaml
- ${DRP_PIPE_DIR}/bps/resources/LSSTCam/DRP.yaml
- ${DRP_PIPE_DIR}/bps/caching/LSSTCam/frdf/DRP-caching.yaml
edfs =
- /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249/DRP-memory.yaml
- /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249/DRP-caching.yaml
- /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249/DRP-clustering.yaml
LSST_VERSION: v30.0.4.rc1
pipelineYaml: "${DRP_PIPE_DIR}/pipelines/LSSTCam/DRP.yaml#step1-single-visit-detectors,step1b-single-visit-visits"
payload:
runInit: true
payloadName: lsstcam_drp
butlerConfig: dp2_prep
inCollection: LSSTCam/zaw/DM-53877,LSSTCam/calib/DM-53135,LSSTCam/calib_refcats,skymaps,pretrained_models,LSSTCam/calib_fgmcml,mpcorb,sorchache
output: LSSTCam/runs/DRP/v30_0_4/DM-54249
dataQuery: "instrument='LSSTCam' and skymap='lsst_cells_v2'"
numberOfRetries: 3
computeSite: CC-IN2P3
requestMemory: 4000
memoryMultiplier: 2
extraGraphOptions: "--dataset-query-constraint raw -c parameters:sasquatch_dataset_identifier=LSSTCam/DRP -c parameters:sasquatch_timestamp_version=explicit_timestamp:20260225T090452"
remoteBuild:
requestMemory: 500000
pipetask:
pipetaskInit:
requestMemory: 64000
numberOfRetries: 3
finalJob:
command: pipetask report --butlerConfig --fileDistributionEndPoint --graphFile --force-v2 --full-output-filename /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249/reports/pipetask_report_output_step1a_1b.json
extraAggregateOptions: --output /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249/reports/qg_step1a_1b.qg --worker-log-dir /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249/logs/step1a_1b/
requestCpus: 32
requestMemory: 500000
numberOfRetries: 10
extraRunQuantumOptions: "--no-raise-on-partial-outputs"
custom_lsst_setup: >
```

There are no p...  
The visits have

| Name                            | Last commit                                   | Last update  |
|---------------------------------|---|--------------|
| configs                         | removed < 4GB drp memory and fix yaml         | 4 weeks ago  |
| templates                       | fix ruclio call in resource usage move script | 1 week ago   |
| README.md                       | update readme                                 | 4 months ago |
| create_scripts.sh               | fix version in move data                      | 2 months ago |
| create_scripts_resourceUsage.sh | Fix resourceusage scripts                     | 3 months ago |

### Template Processing Script

**Overview**

This script automates the generation of configuration files from templates by replacing placeholders with user-provided or default values. It's designed for ease of use both in batch and interactive modes, making it adaptable for various workflow scenarios.

**Usage**

```
./create_scripts.sh #in interactive mode
./create_scripts.sh [WEEKLY] [TICKET] [BUTLER] [INPUT_COLLECTION] [OUTPUT_COLLECTION] [NUMBER_OF_GROUPS]
```

| Name                                  | Last commit   |
|---------------------------------------|---|
| README.md                             | added option to set log names to the campaign ticket... |
| WARNING                               | commit with changes in automatize script                |
| execute_campaign_improved.sh          | commit with changes in automatize script                |
| execute_campaign_improved_complete.sh | commit with changes in automatize script                |
| sbatch_scri                           | added option to set log names to the campaign ticket... |

### Campaign Execution Scripts

This repository contains two shell scripts to automate the execution of a multi-step data processing campaign using BPS (Big Pipeline System) and PanDA workflow management.

**Files**

```
execute_campaign_improved.sh
```

A Bash script that:

- Executes a series of BPS steps in order.
- Submits jobs to PanDA via bps submit.
- Monitors job status until completion or failure.
- Supports resuming from a specific step (optional).

**Usage:**

```
./execute_campaign_improved.sh <working_directory> <start_step>
```

## Step 1

Step1a\_1b

```
bps submit /pbs/throng/lsst/groups/drp/LSSTCam/DM-54249/step1/lsstcam_drp_v30.0.4.rc1
```

Submitted  
Completed

| Panda Workflow | Note | Status   | URL  |
|----------------|------|----------|--|
| 15982          |      | Finished | https://usdf-panda-bigmon.slac.stanford.edu:8443/tasks/15982 |

Comments:

Report

```
bps report --id 15982
```

| X         | STATE | KS    | ID                | OPERATOR | PROJECT          | CAMPAIGN | PAYLOAD |
|-----------|-------|-------|-------------------|----------|------------------|----------|---------|
| SUCCEEDED | 100   | 15982 | Gabriele Mainetti |          | LSSTCam_runs_DRP |          |         |

Path: None  
Global job id: None

|       | UNKNOWN | MISFIT | UNREADY | READY | PENDING | RUNNING | DELETED | HELD | SUCCEEDED | F |
|-------|---------|--------|---------|-------|---------|---------|---------|------|-----------|---|
| TOTAL | 0       | 0      | 0       | 0     | 0       | 0       | 0       | 0    | 10000     |   |

Execution of DP2 pilot run at FrDF us...

Inputs

To Do list after execution

To Do list after execution

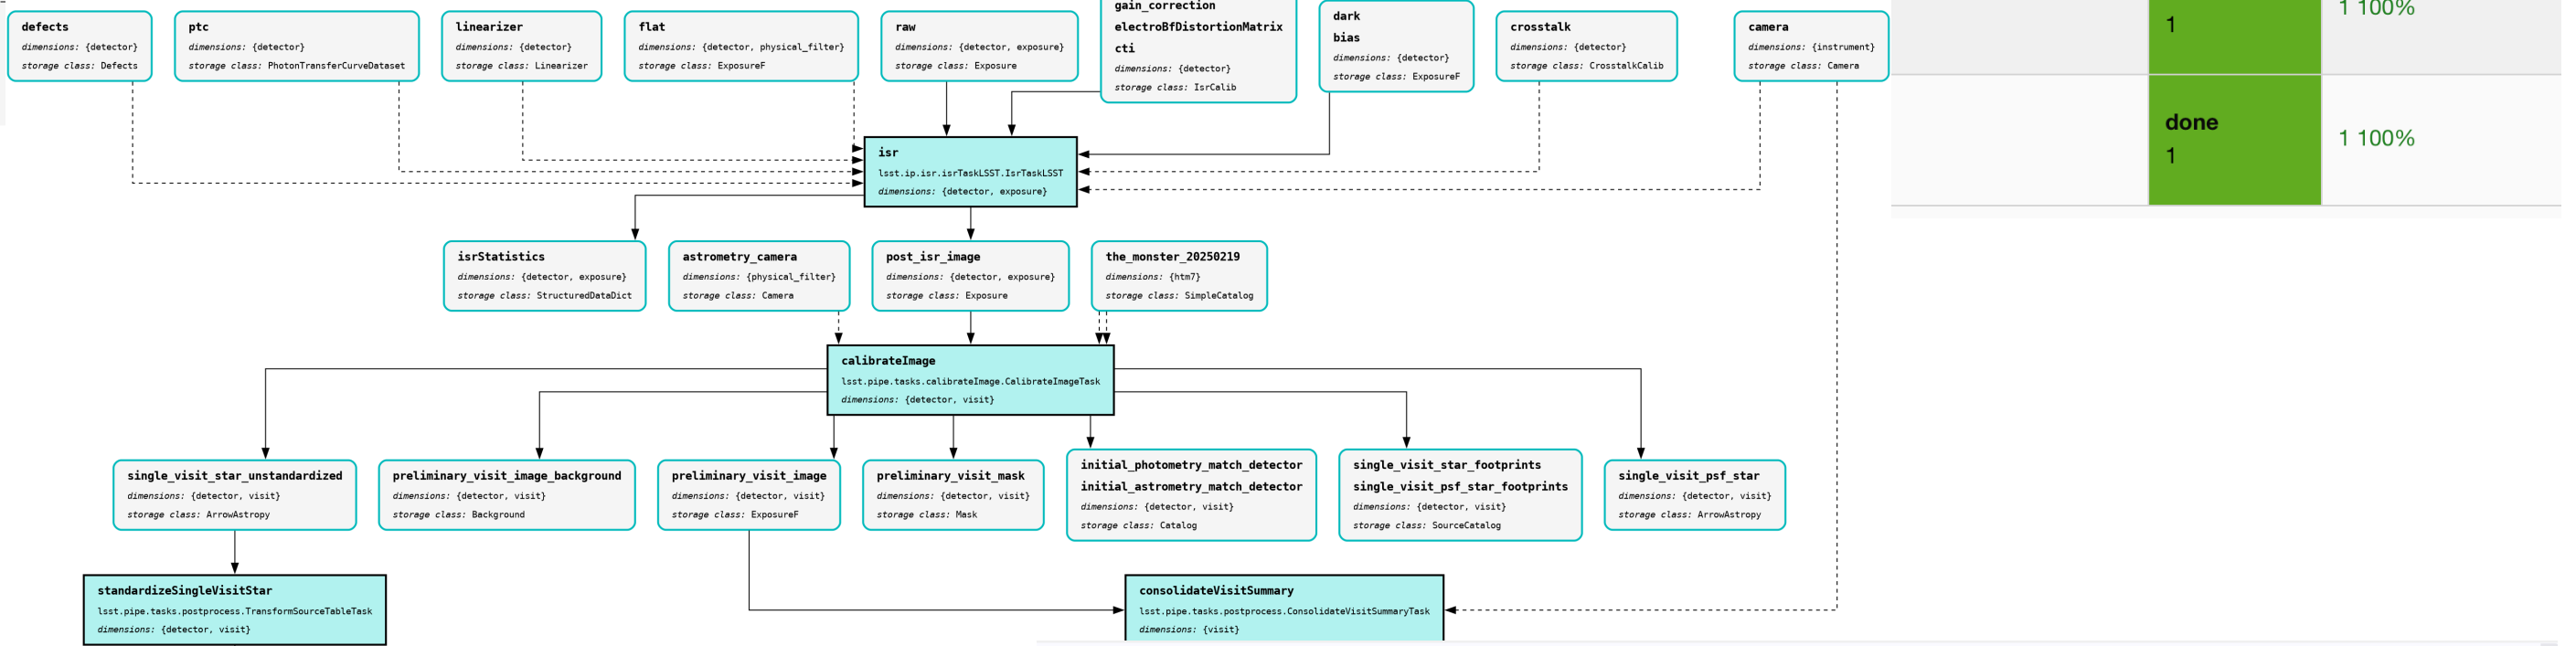
Executions

- Step 1
  - Step1a\_1b
  - Step1c\_1d
- Step 2
  - Step2a\_2c
  - Step 2d
  - Step 2e\_2f
- Step 3
  - Step 3a WAPPS
  - Step 3a COARDS
  - Step 3a GROUPS
  - Step 3b
- Step 4
  - Step 4a
  - Step 4b
  - Step 4c
  - Step 4d

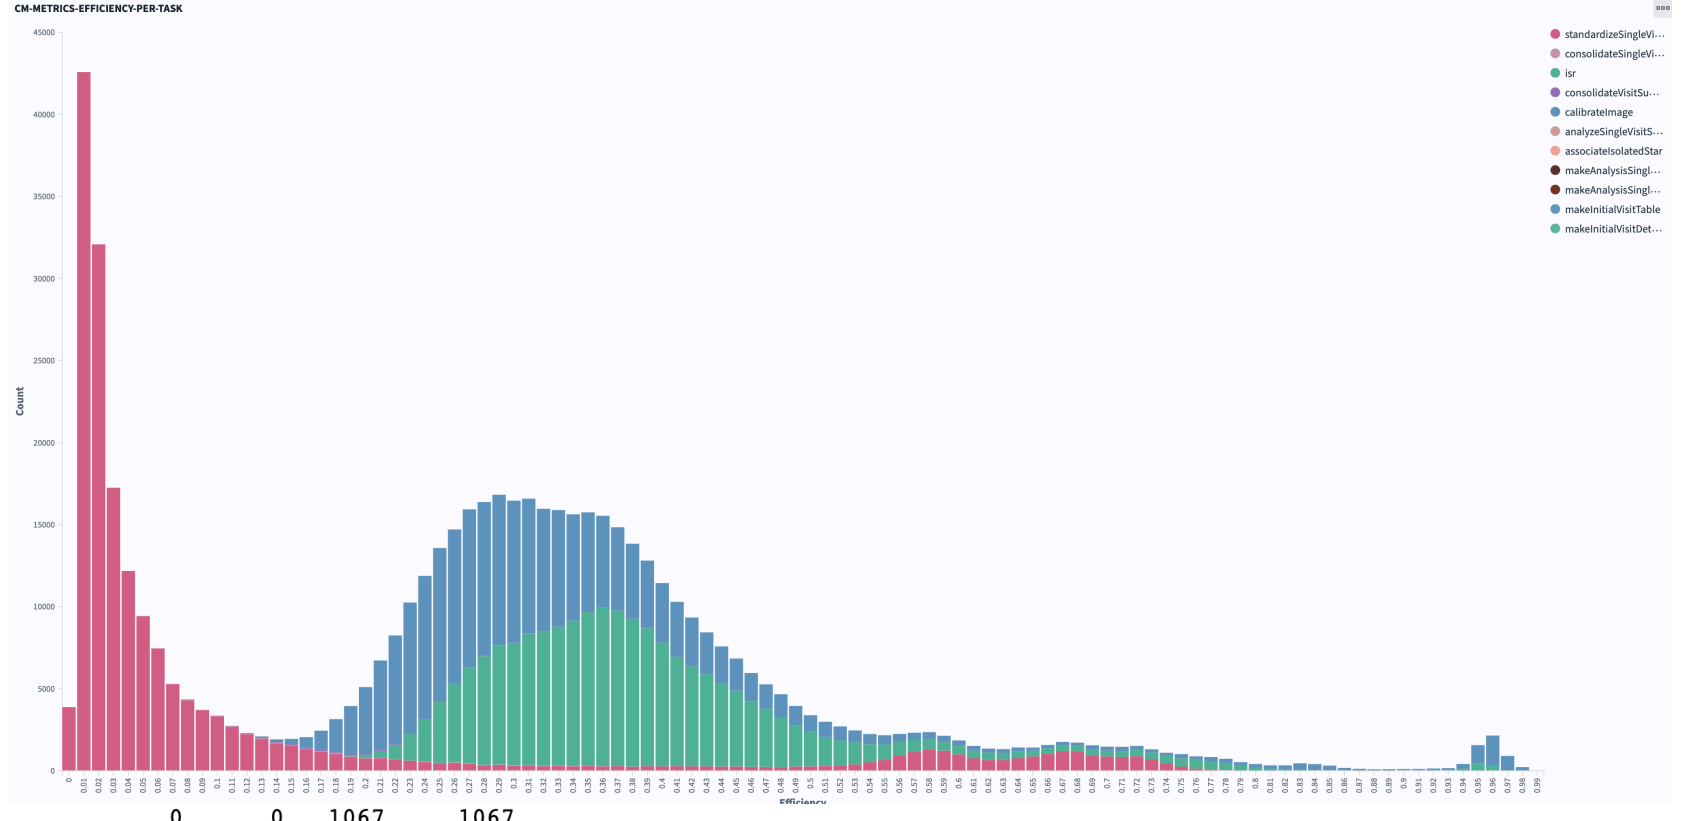
Collapse all  
Back to top  
Go to bottom

## 5 tasks, sorted by jeditaskid-desc

| ID    | Parent | Workflow ID | Task name  | TaskType/ProcessingType | Campaign Group | User              | Errors | Task status | Input files   |
|-------|--------|-------------|--|-------------------------|----------------|-------------------|--------|-------------|---|
|       |        |             |  |                         |                |                   |        | Nfiles      | Most <span style="color:red">●</span> Nfinish % Nfail % |
| 59251 | 15982  |             | LSSTCam_runs_DRP_v30_0_4_DM-54249_20260225T092325Z_04_finalJob_01_15982_65551      | test                    | Rubin_Merge    | Gabriele Mainetti | Errors | done 1      | 1 100%  |
| 59249 | 15982  |             | LSSTCam_runs_DRP_v30_0_4_DM-54249_20260225T092325Z_03_step1b_visits_01_15982_65549 | test                    |                | Gabriele Mainetti | Errors | done 1067   | 1067 100%   |
| 59248 | 15982  |             | LSSTCam_runs_DRP_v30_0_4_DM-54249_20260225T092325Z_02_step1detector_01_15982_65548 | test                    |                | Gabriele Mainetti | Errors | done 9834   | 9834 100%   |
|       |        |             | LSSTCam_runs_DRP_v30_0_4_DM-54249_20260225T092325Z_01_pipetaskinit_01_15982_65547  |                         |                |                   |        | done 1      | 1 100%  |
|       |        |             |  |                         |                |                   |        | done 1      | 1 100%  |



|                                     |        |   |      |
|-------------------------------------|--------|---|------|
| isr_metadata                        | 190734 | 0 | 0    |
| isr_log                             | 190734 | 0 | 0    |
| preliminary_visit_mask              | 186747 | 0 | 3987 |
| preliminary_visit_image_background  | 190698 | 0 | 36   |
| single_visit_star_footprints        | 186889 | 0 | 3845 |
| initial_photometry_match_detector   | 186747 | 0 | 3987 |
| single_visit_psf_star               | 187092 | 0 | 3642 |
| initial_astrometry_match_detector   | 186889 | 0 | 3845 |
| preliminary_visit_image             | 190698 | 0 | 36   |
| single_visit_psf_star_footprints    | 187464 | 0 | 3270 |
| single_visit_star_unstandardized    | 186887 | 0 | 3847 |
| calibrateImage_metadata             | 190734 | 0 | 0    |
| calibrateImage_log                  | 190734 | 0 | 0    |
| single_visit_star_detector          | 186887 | 0 | 3847 |
| standardizeSingleVisitStar_metadata | 190734 | 0 | 0    |
| standardizeSingleVisitStar_log      | 190734 | 0 | 0    |
| visit_geometry                      | 1059   | 0 | 8    |
| preliminary_visit_summary           | 1067   | 0 | 0    |
| consolidateVisitSummary_metadata    | 1067   | 0 | 0    |
| consolidateVisitSummary_log         | 1067   | 0 | 0    |
| single_visit_star                   | 1059   | 8 | 0    |
| consolidateSingleVisitStar_metadata | 1067   | 0 | 0    |
| consolidateSingleVisitStar_log      | 1067   | 0 | 0    |



# Characterizing Memory Requirements

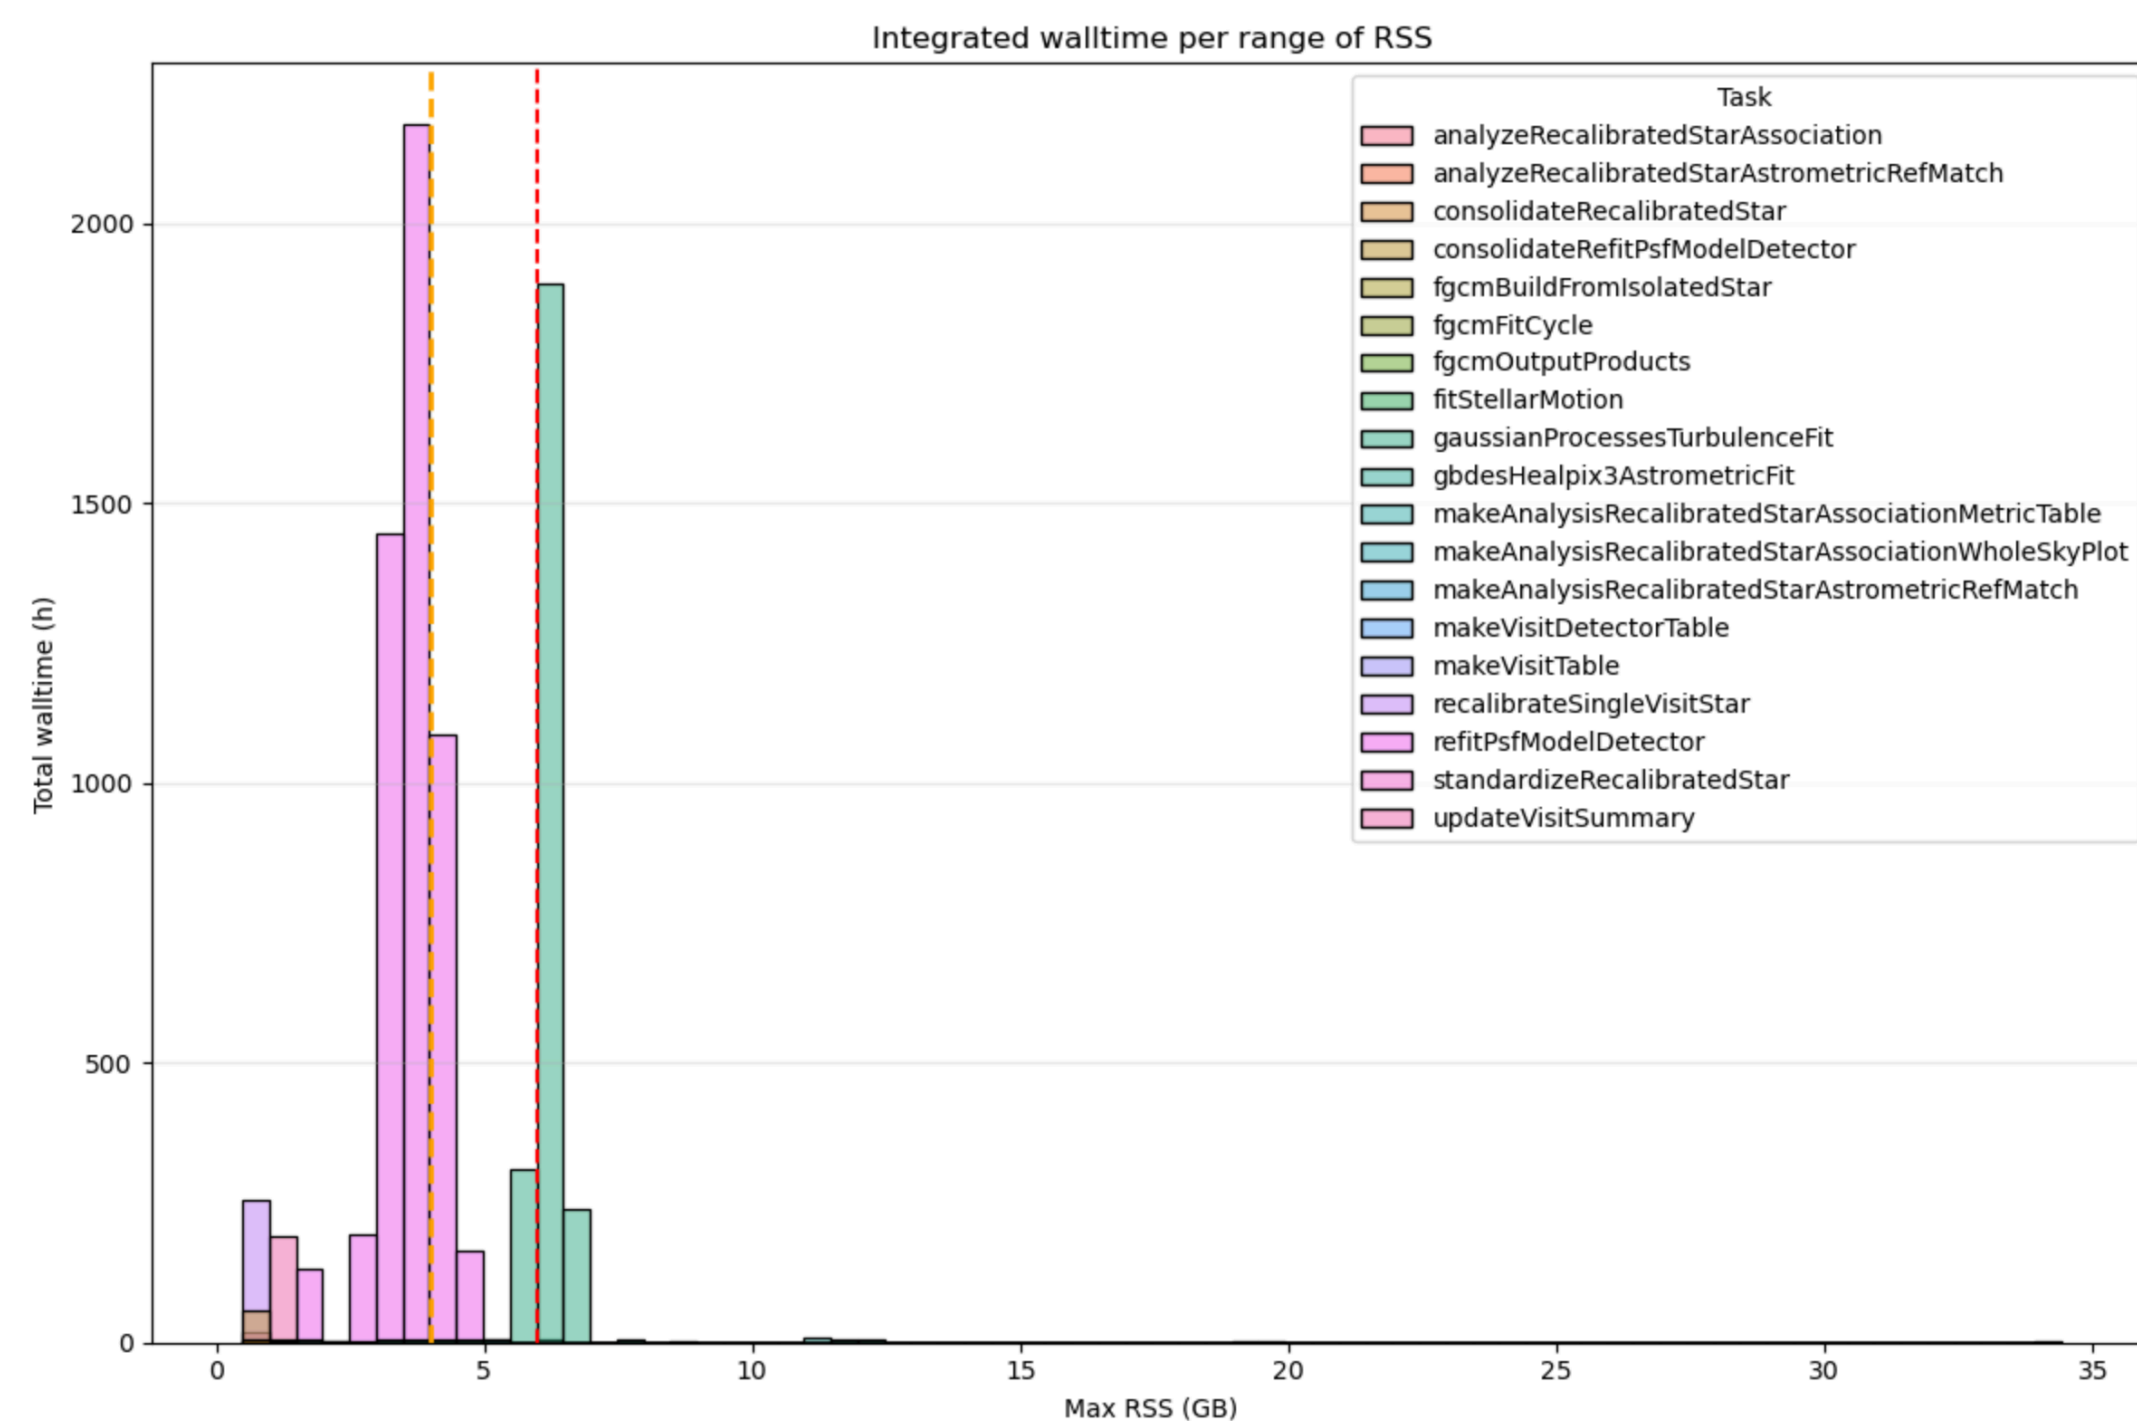
**2**

# Characterizing Memory Requirements

- Goal: determine whether we can reduce memory per core to 6 GB
- Current worker-node configuration:
  - lsst partition: 10 GB per core
  - htc partition: 4 GB per core
- Pilot phase campaign DM-54249 (COSMOS + EDFs fields) metrics produced by the pipetask extracted from the butler and analyzed stage by stage, pipetask by pipetask, quanta by quanta
- Tech note available : <https://dmtn-334.lsst.io/>

# Characterizing Memory Requirements

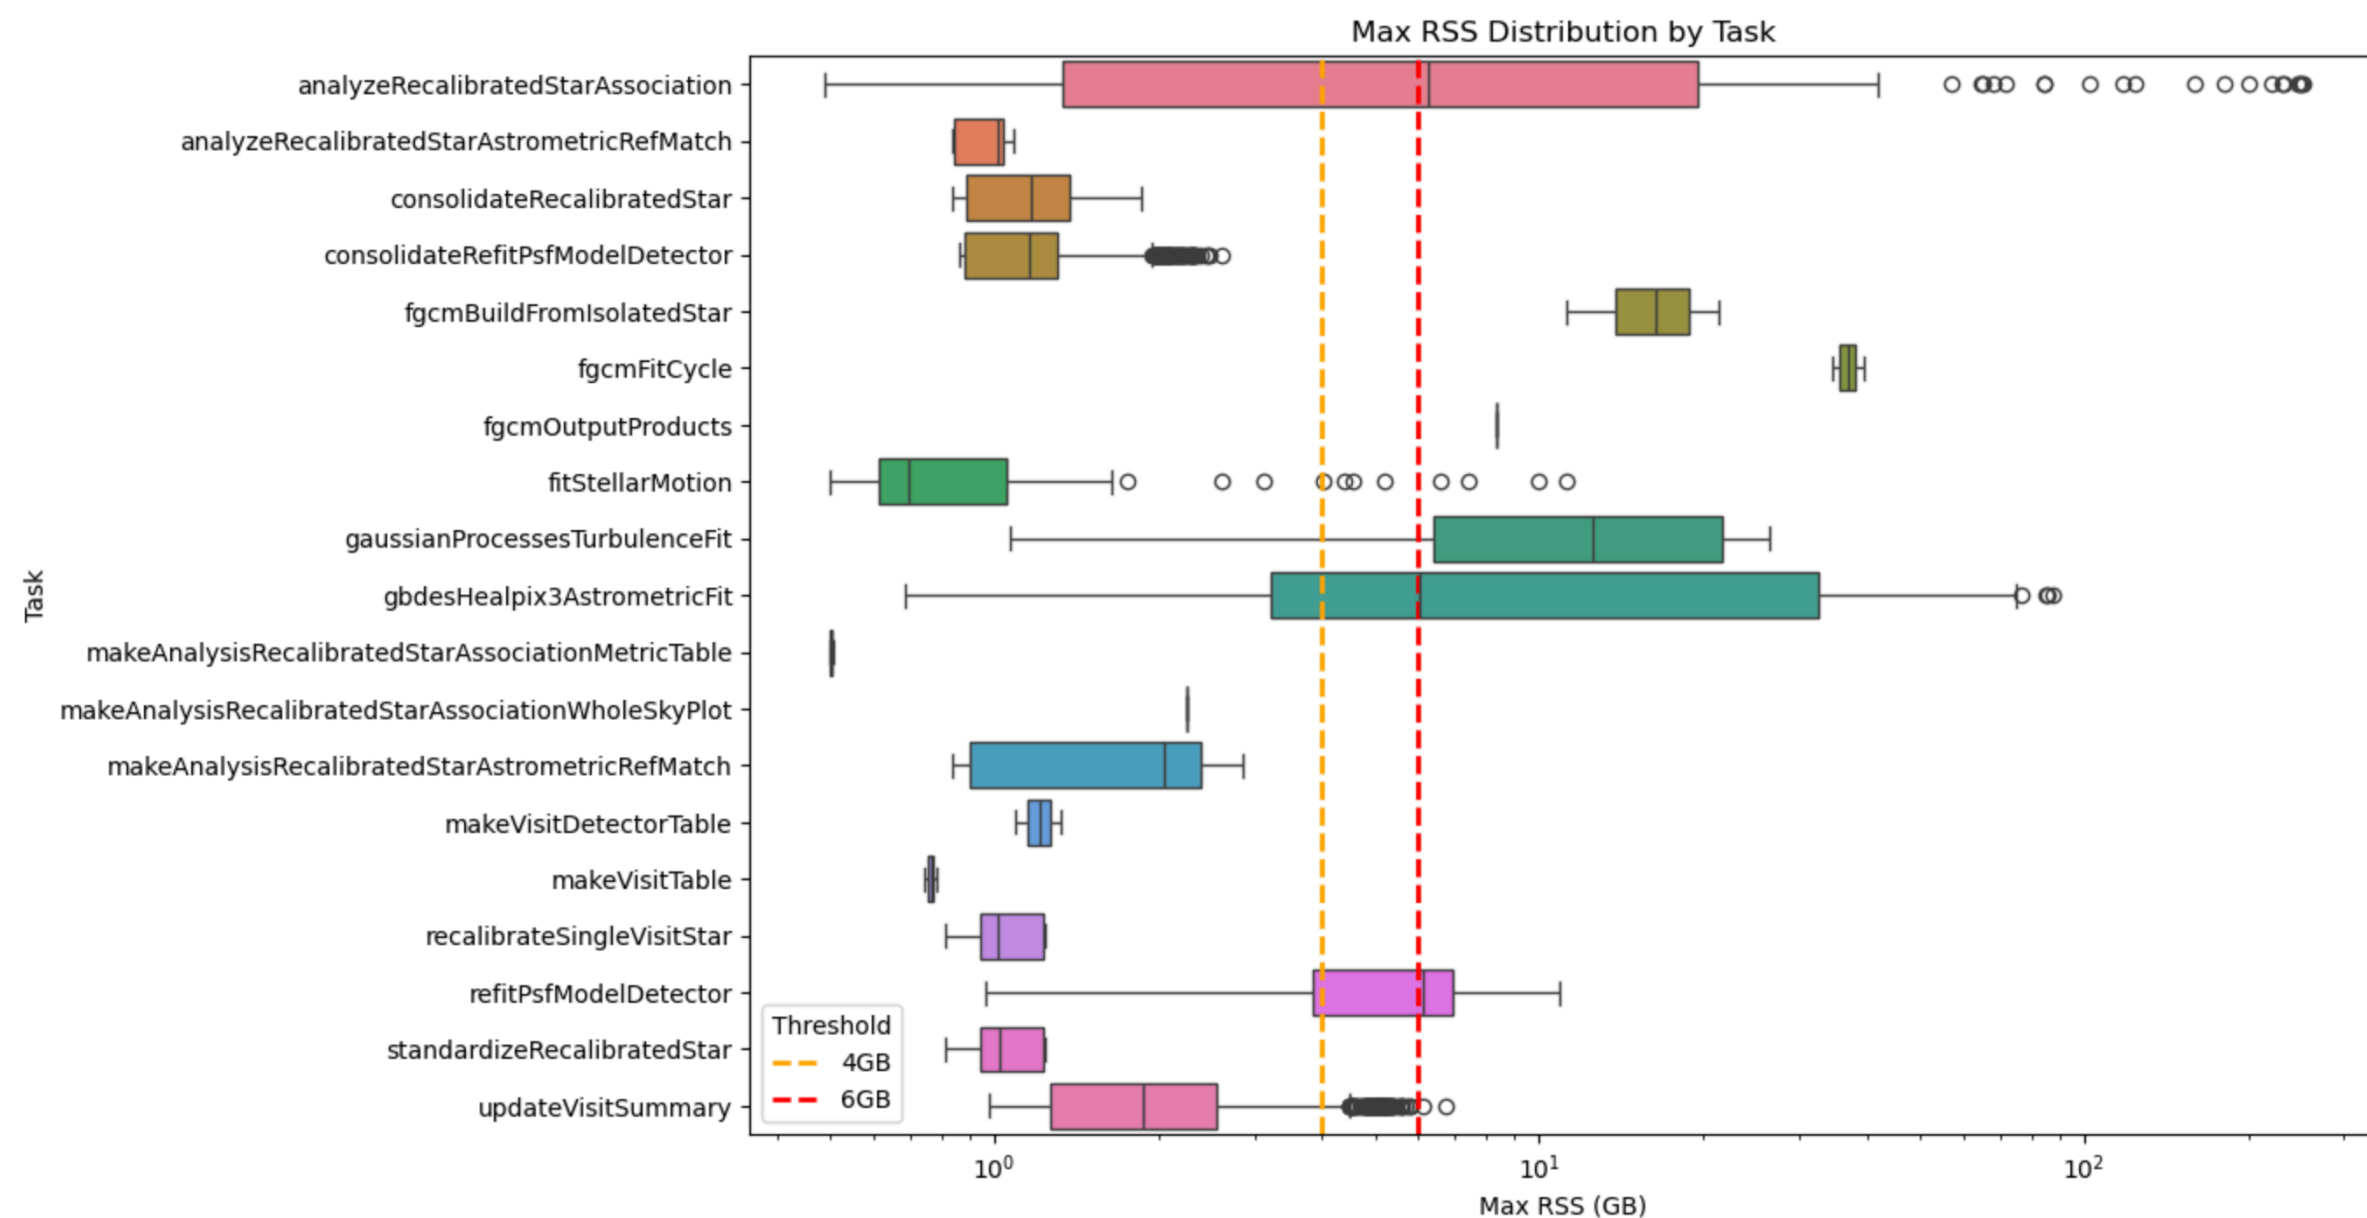
- 99.9%+ of quanta across all stages use <6GB RAM
- In Stage 2, 99.67% of quanta stay below 6GB, but a single pipetask (*gaussianProcessesTurbulenceFit*) accounts for  $\approx 25\%$  of the wall-time



| Stage   | Quanta    | N. quanta > 6GB | % quanta > 6GB | Total Walltime (h) | Walltime (h) > 6GB | % wall-time > 6GB |
|---------|-----------|-----------------|----------------|--------------------|--------------------|-------------------|
| Stage 1 | 574 472   | 36              | 0.006          | 10 049.86          | 3.28               | 0.03              |
| Stage 2 | 567 886   | 1 916           | 0.337          | 8 267.89           | 2 077.15           | 25.12             |
| Stage 3 | 887 683   | 637             | 0.072          | 25 973.91          | 123.53             | 0.476             |
| Stage 4 | 2 593 783 | 1 104           | 0.043          | 52 124.19          | 2 184.912          | 4.192             |

# Characterizing Memory Requirements

- Dense field SV225 (DM-54372) – special case:
  - Stage 2 memory demand increases dramatically:  $\approx 27\%$  of quanta exceed 6GB, consuming  $\approx 82\%$  of the total wall-time.
  - Peak memory can reach 250GB (*analyzeRecalibratedStarAssociation*) and 450GB (*analyzeSourceAssociation* in Stage4).



Stage 2

| Stage   | Quanta    | N. quanta > 6GB | % quanta > 6GB | Total Walltime (h) | Walltime (h) > 6GB | % wall-time > 6GB |
|---------|-----------|-----------------|----------------|--------------------|--------------------|-------------------|
| Stage 1 | 1 095 651 | 1 880           | 0.172          | 23 142.842         | 67.186             | 0.29              |
| Stage 2 | 1 071 583 | 288 912         | 26.961         | 60 984.435         | 50 123.035         | 82.19             |
| Stage 3 | 2 182 659 | 1 617           | 0.074          | 76 739.012         | 1 073.987          | 1.4               |
| Stage 4 | 6 425 437 | 42 119          | 0.656          | 178 482.496        | 15 385.953         | 8.62              |

# Characterizing Memory Requirements

- Reducing memory allocation to 6 GB per core is feasible (for the majority of pipetasks across all stages)
- Memory configuration for future purchase of compute nodes will take into account these results
- DRP memory configuration has been adapted to reflect the results
- For dense field we configured the BPS multiplier factor to deal with memory needs
- We continue monitoring memory because software updates can introduce significant changes, e.g. DM-54633



# Replication of LSSTCam raw exposures

176 154 exposures replicated up to May 12, 638 TB

Number of file transferred correctly

**41137**

No error

Number of errors during the transfer

**9**

10027 - Error code

**1**


10001 - Error code

Data transferred

**68TB**

Total transfer

Data Transfer Status Timeline



- All datasets from the official **dp1** butler repository at USDF have been registered in Rucio and copied to CC-IN2P3
- They have been ingested in a similar **dp1** butler repository (see technical note: <https://rtn-103.lsst.io>)
- Accessible by all CC-IN2P3 **lsst** group members, using the same mechanism as for the main repo: <https://doc.lsst.eu/tutorial/community-butlers.html>
- Feel free to have a look and give us some feedback.
- We are currently investigating the possibility to ingest additional data that are not part of the official dp1 release. **Let us know what data are required for your analysis.**

```
[(lsst-scipipe-13.0.0-exact) 15:30 leboulch@ccacc001 ~> butler query-datasets dp1 deep_coadd
py.warnings WARNING: /cvmfs/sw.lsst.eu/almalinux-x86_64/lsst_distrib/w_2026_17/conda/envs/lsst-scipipe-13.0.0-exac
actions specified. The --collections argument will become mandatory after v28.
  for table in script.QueryDatasets(butler=butler, **kwargs).getTables():
```

| type       | run                              | id                                   | band | skymap        | tract | patch |
|------------|----------------------------------|--------------------------------------|------|---------------|-------|-------|
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | af9c7765-8295-42f3-8368-0e6569d12c8a | g    | lsst_cells_v1 | 453   | 11    |
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | 27bebfdc-dfe9-4067-a56c-0128d285b9c0 | i    | lsst_cells_v1 | 453   | 11    |
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | 4798a2cc-7735-4dc2-b6b1-93fa4192b454 | r    | lsst_cells_v1 | 453   | 11    |
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | d7125f39-010a-4ea1-a117-01630a8ac98f | y    | lsst_cells_v1 | 453   | 11    |
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | 57a62e73-3106-419f-a456-b4998ef593cb | g    | lsst_cells_v1 | 453   | 12    |
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | 41f58b23-1d4d-4581-86c4-91c40f2bb04f | i    | lsst_cells_v1 | 453   | 12    |
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | 2c6273b7-1b15-42a5-ae4-e0ff1cb22453  | r    | lsst_cells_v1 | 453   | 12    |
| deep_coadd | LSSTComCam/runs/DRP/DP1/DM-51335 | 5fb56945-e8f8-4e92-8457-f53e8c19eab  | y    | lsst_cells_v1 | 453   | 12    |

# Other updates

- We continue test campaign using weekly release, e.g. DM-54601, DM-54760
  - To test the stack
  - To consolidate our infrastructure
- We are testing Rucio to move back to USDF data produced at FrDF (DM-54867)
- A tool to simplify the butler access has been developed and deployed on cca machines, see <https://doc.lsst.eu/tutorial/indigoiam.html>
- Equipment purchase for the 2027 needs ongoing: high increase in price
  - CPU allocation in the Slurm farm: 12k CPU threads
  - storage under /sps/lsst managed by CephFS: 8.5 PB
  - image storage and data products, managed by dCache: 31PB now, 36 PB in 2027

- Qserv actually is configured to provide access to the following databases:

| Catalog            | Size (Tb) | # Lines (Billions) |
|--------------------|-----------|--------------------|
| idf-dp0.2-catalog  | 36.6      | 139                |
| frdf-dp0.2-catalog | 36.6      | 139                |
| dp01_dc2_catalogs  | 1.1       | 1.7                |
| skysim5000_v1.1.1  | 13.6      | 20.5               |
| cosmoDC2_v1.1.4    | 3.7       | 5.5                |

- The machines behind Qserv are at the end of life and we cannot replace them
- We cannot provide enough resources to expose the DR data on Qserv
- **Qserv will be discontinued in few months**
- RSP instance redeployment is in progress
- We want also building tools for direct catalog access from Parquet files — but **we need use cases** to test and and help to develop a solution.

Questions & Comments