Towards a sizing model for the Rubin/DESC analysis at CC-IN2P3

Fall 2023 – CC-IN2P3, Lyon

Recap of last episode: spring survey

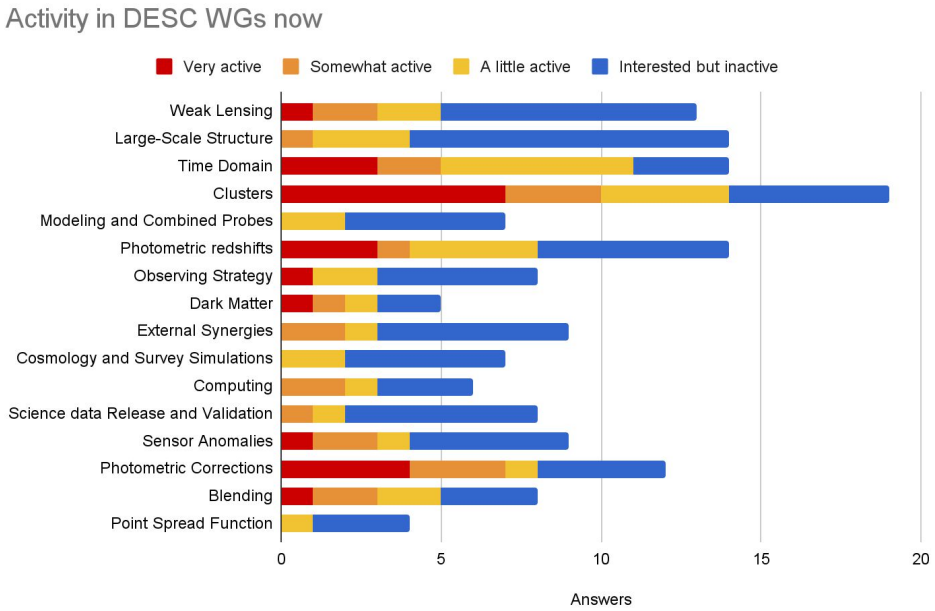

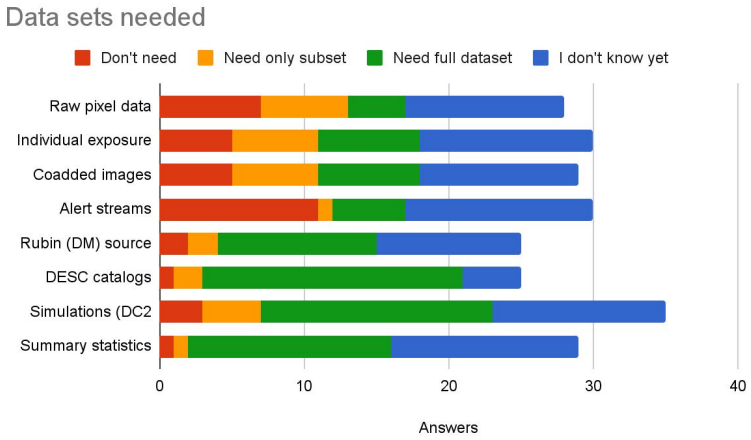

- "Interested but inactive" in various WGs
- Uncertainty in computing needs for "science" (data and cpu)

This otter doesn’t care about the computing requirements.

Why should you?

Context

What we know

- CC-IN2P3 will host the full dataset of raw images for archiving
- IN2P3 to fund 40% of Rubin Data Release Processing (DRP) at CC-IN2P3, to be merged at SLAC (USDF) into full release
- “Best effort” approach for (DESC) science analysis deemed insufficient at last “Evaluation annuelle projet” (EAP) to guarantee scientific return

What we need

- A *science-driven* proposal for a reasonable analysis model to guide CC allocation for LSST and infrastructure developments (software and hardware). *We need your input!*

This session

Goals

1. Start a discussion!
2. Identify critical components to complete “main” science goals (3x2+CL+SN)
    - Including simulations and intermediate products (i.e. not just final run)
    - Numbers needed are storage (PB), computing (Mcpu-h) and machine needs (GPUs? Large-memory? Infiniband? *vs* standard MPI nodes)
3. Identify other potential datasets/processing needed in the broader LSST-France community.

Outline

1. Report of estimations and ongoing discussion between LSST-France scientific coordinators and CC staff (lots of work done by Dominique)
2. Discussion with everyone!

From Rubin data to static cosmology

A worked example for 3x2+CL, DESC work includes:

1. Generating and validating catalogs of source/lens/cluster/photoz
2. Validating models and analysis choices on simulations
3. Validating and running cosmological inference

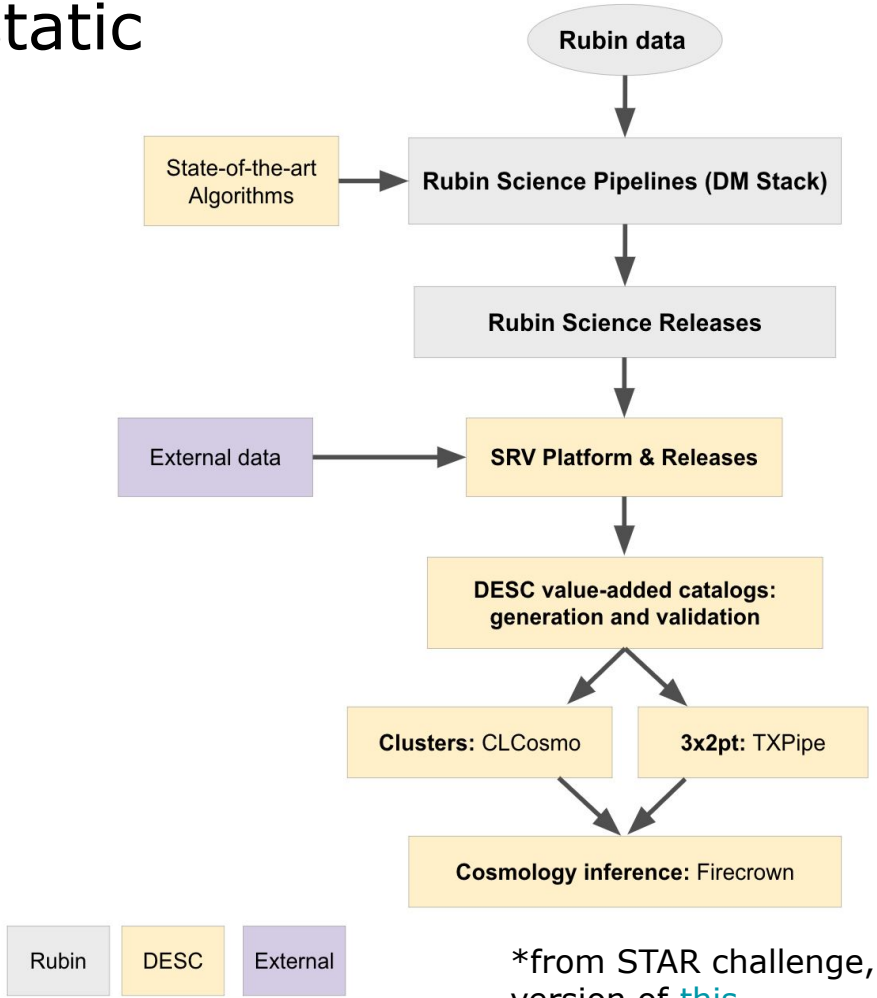

*from STAR challenge, simplified version of this.

From Rubin data to static cosmology: storage

- DESC static catalogs (~ year 10)

	Lens	Cluster	Shear	photo-z
Nbr. of rows	3.10^8	3.10^5	3.10^9 * 5 versions	3.10^9 * 5 versions
Nbr of columns	300 floats	300 floats	300 floats	200 bytes
Total Size	660 GB	0.3 GB	18 TB	3 TB

  - Several versions needed, so O(100s) TB
  - In comparison, DES DR2 = 3.1 TB

- DESC data vectors
  - 2pt functions + n(z)'s + covariance + metadata (eg, masks) represent < 1TB

From Rubin data to static cosmology: computing

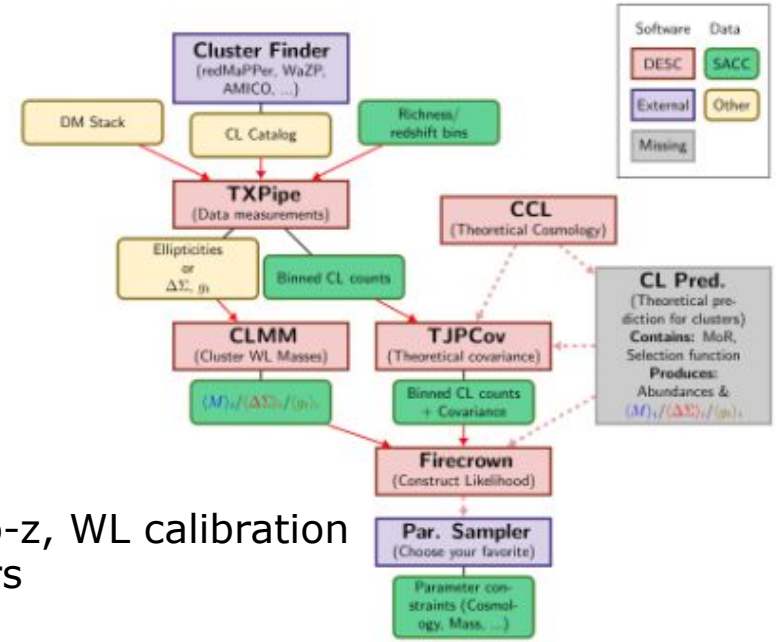

- DESC computing model at NERSC: O(25M) CPU core hours + O(10k) GPUs
- Estimation of DESC computing needs:
  a. *Image simulations* for WL calibration
     - For DES Y3, O(1M) core hours.
  b. *Synthetic source injection*
     - Re-processing of raw images for photo-z, WL calibration and LSS systematics, O(1M) core hours
  c. Catalogs to *summary statistics*
     - TXPipe run ~ O(10k) core hours, several runs to optimal configuration.
  d. Model validation and *inference*
     - For DES Y3 3x2, ~1000 MCMC chains, each O(10k) core hours; CL integrals?
- So overall, **DESC 3x2+CL ~ O(10M) core hours**

From Rubin data to SN cosmology

- Storage
  - Actual cosmology data vectors are not large (few TB).
  - What about transient catalogs (DIASRC, DIAOBJ) from DRP?
  - We probably will need to go back to pixel level to develop tools
    - Scene modeling, uniformity (uber)-calibration, maybe even instrument removal (detrending)
    - Will need image data (raw*, calexp, calibration[?])
    - Will need catalogs (refcats and rubin star catalogs)
  - Hope is to then push these tools to DM to run in next DR.
  - One solution would be to train on DDF (~7% of images)
    - ~500TB/yr for calexp (or own calibrated data) x nb_of_version

* Should already be on disk at CC.
? Are they already at CC in current plan?

From Rubin data to SN cosmology

- Computing (mostly scaled from ZTF)
  - Scene modelling will need something like 1 core-week per SN on small machines (~10GB) → ~O(10M clockhours) for one 10yr prod.
  - Ubercalibration would need **hugemem** machines
  - Cosmo: w and forward fs8 will probably need **GPUs** (jax)
  - Maximum likelihood fs8 would need **hugemem** machines
  - Detrending (ISR) is massively // (GPUs or many cpus)

* Should already be on disk at CC.
? Are they already at CC in current plan?

Discussion: data sets

-----DRP-----

Data set	Size (Y1 → Y10)	Availability at CC-IN2P3
Raw images	5 PB → 50 PB	Full (archiving) via Butler
Individual exposures (calexps)	6 PB → 67 PB	On demand for pre-identified fields (SN/transients) ?
Co-added images (coadds)	2.7 PB	Full via Butler
Alert streams	1TB/night → 3 PB	Full (guaranteed for FINK)
Rubin DRP catalogs	1 PB → 9 PB*	??
DESC catalogs + metadata (masks) + data vectors	< 1 PB	Full via DESC tools (GCR, TXPipe)?
Extra: simulations, precursor, spectro	??	**Sims** ~ 150 TB (DC2+skySim5000+buzzard) **Reference catalogs** ~ 2 TB (Gaia, SDSS, Panstarrs)

* direct access to parquet files *without* Qserv – otherwise, 3 PB → 27 PB

Discussion: computing

- Rubin DRP computing at CC-IN2P3
  - Allocation 2023 ~ 48 M core hours / year
  - Full DRP = 40M → 400M cpu hours (https://dmtn-135.lsst.io/, table 25), 40% to be done at CC
- DESC 3x2+CL+SN analysis at CC-IN2P3
  - Participation to computationally expensive systematics calibration efforts (e.g. SSI, image sims for WLSS/CL) ~ O(few 1M) cpu hours?
  - Validation efforts + final runs ~ also O(few 1M) cpu hours?
  - DESC analysis to be distributed across facilities, including NERSC, which means moving data products around: scp/globus tools needed.

Discussion: what else?

- What about your analysis?
  - What other data products do you need? What are your major computing needs?
  - Need an Rubin Science Platform (RSP) at CC for the broader community? Need Qserv?
- Proposal for DESC + other science analysis separated from proposal for DRP ?
  - Note: eventually, proposal in Mcpu-hours to be converted to € for IN2P3 (cost of cpu/storage decreases over 10 years…)

References

1. Requirements for analysis and references therein
    a. DESC Computing Model Update 2023
    b. Summary of (DESC) Data Management Principles
    c. LSST DESC Science Overview Document

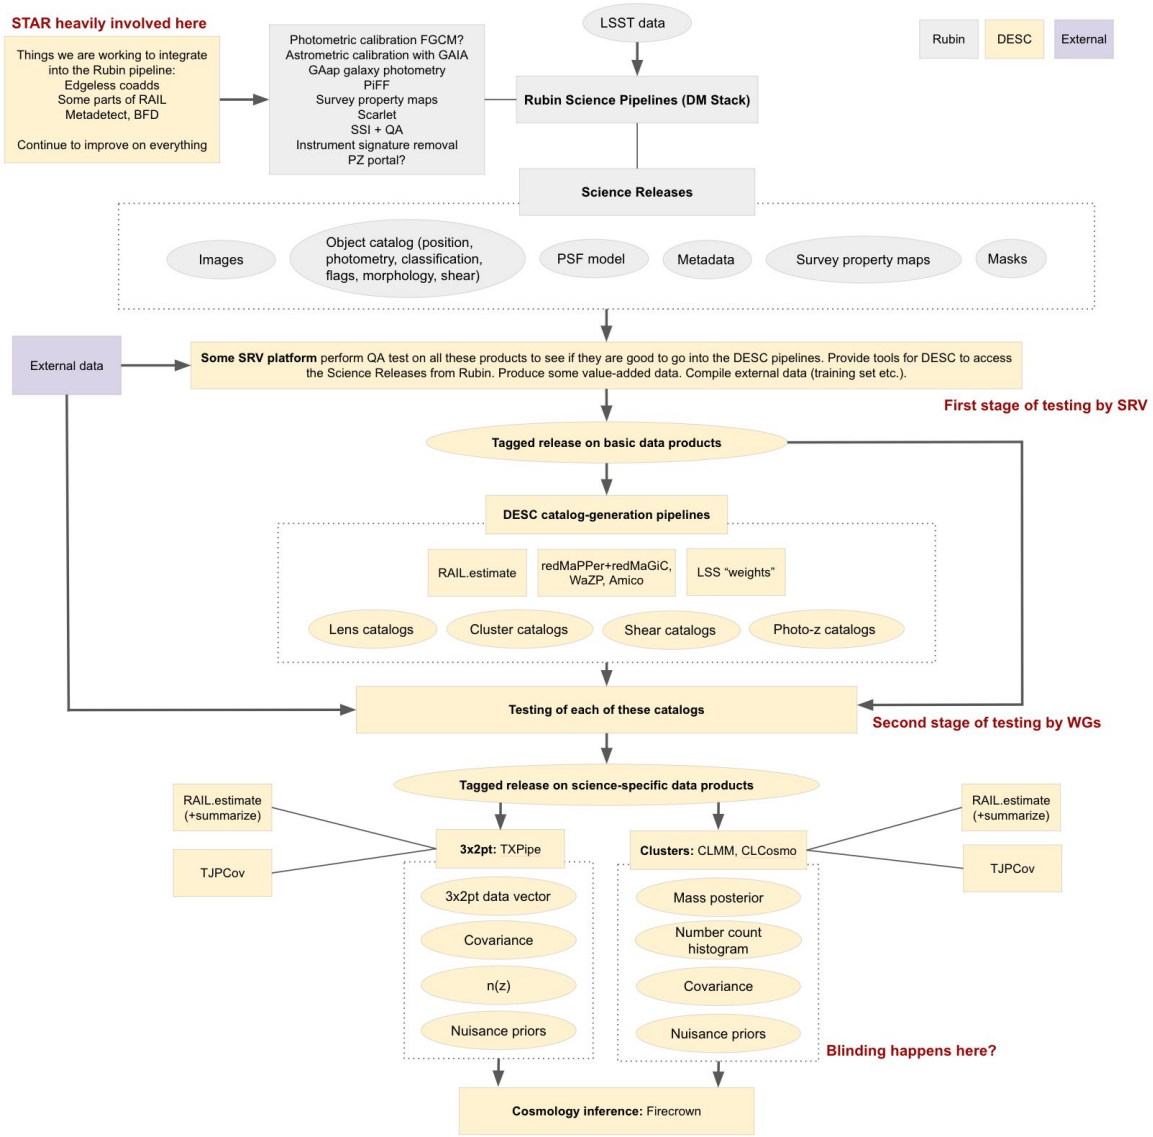
STAR heavily involved here
Things we are working to integrate into the Rubin pipeline:
Edgeless coadds
Some parts of RAIL
Metadetect, BFD
Continue to improve on everything
Photometric calibration FGCM?
Astrometric calibration with GAIA
GAap galaxy photometry
PiFF
Survey property maps
Scarlet
SSI + QA
Instrument signature removal
PZ portal?
LSST data
Rubin Science Pipelines (DM Stack)
Rubin
DESC
External
Science Releases
Images
Object catalog (position, photometry, classification, flags, morphology, shear)
PSF model
Metadata
Survey property maps
Masks
External data
Some SRV platform perform QA test on all these products to see if they are good to go into the DESC pipelines. Provide tools for DESC to access the Science Releases from Rubin. Produce some value-added data. Compile external data (training set etc.).
First stage of testing by SRV
Tagged release on basic data products
DESC catalog-generation pipelines
RAIL.estimate
redMaPPer+redMaGiC, WaZP, Amico
LSS "weights"
Lens catalogs
Cluster catalogs
Shear catalogs
Photo-z catalogs
Testing of each of these catalogs
Second stage of testing by WGs
Tagged release on science-specific data products
RAIL.estimate (+summarize)
TJPCov
3x2pt: TXPipe
3x2pt data vector
Covariance
n(z)
Nuisance priors
Clusters: CLMM, CLCosmo
Mass posterior
Number count histogram
Covariance
Nuisance priors
RAIL.estimate (+summarize)
TJPCov
Blinding happens here?
Cosmology inference: Firecrown