VERA C. RUBIN OBSERVATORY

Vera C. Rubin Observatory

Overview of Data Processing

fabio hernandez
CNRS / IN2P3 computing centre, Lyon, France

LOFAR meeting, ASTRON, Sep. 19th 2024

Vera Cooper Rubin
(1928 - 2016)

Legacy Survey of Space and Time

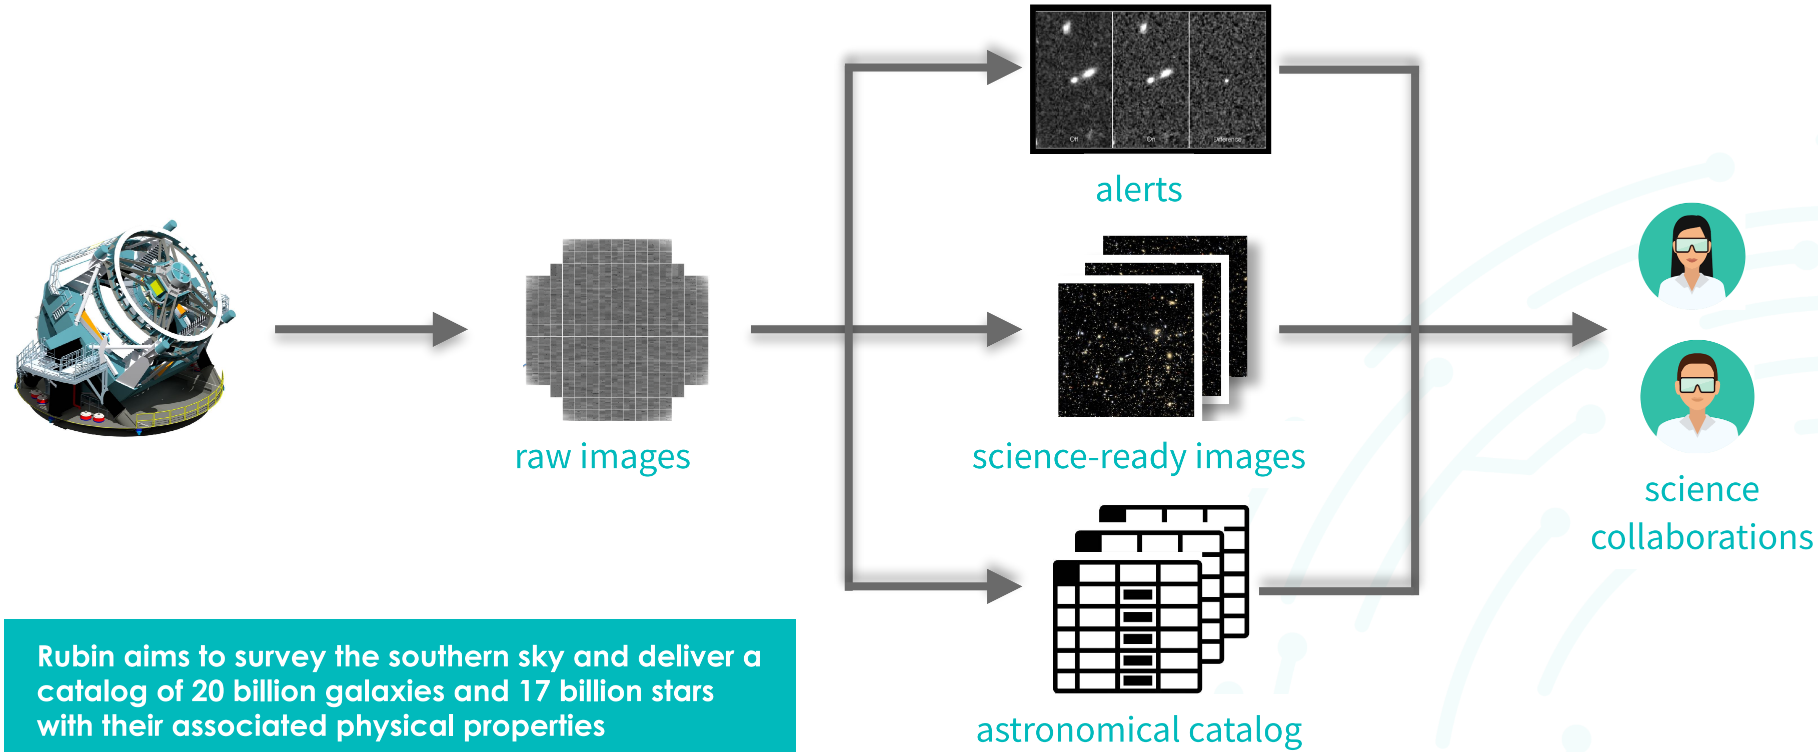

Rubin aims to survey the southern sky and deliver a catalog of 20 billion galaxies and 17 billion stars with their associated physical properties

LSST overview

- Principle of operations
  - 90% of the observing time of the telescope devoted to a deep-wide-fast survey
  - one complete visit of the southern hemisphere sky every 3-4 nights for 10 years, from late 2025
  - each patch of the observable sky to be visited about 800 times
  - 43% of the celestial sphere will be covered by this survey

- Science themes
  - determining the nature of dark energy and dark matter
  - taking an inventory of the solar system
  - exploring the transient optical sky
  - mapping the structure and evolution of the Milky Way

Observatory Overview

Observatory overview

SITE

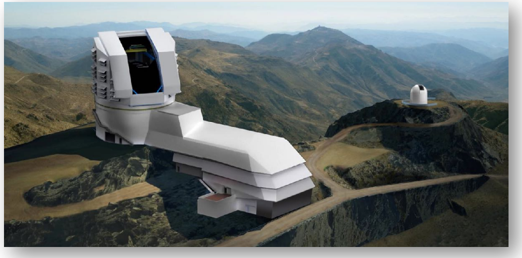

southern hemisphere | 2647m a.s.l. | stable air | clear sky | dark nights | good infrastructure

TELESCOPE

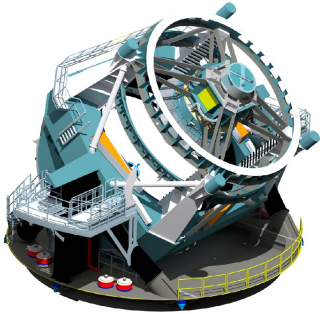

main mirror ⌀ 8.4 m (effective 6.4 m) | large aperture: f/1.234 | wide field of view | 350 ton | compact | to be repositioned about 3M times over 10 years of operations

CAMERA

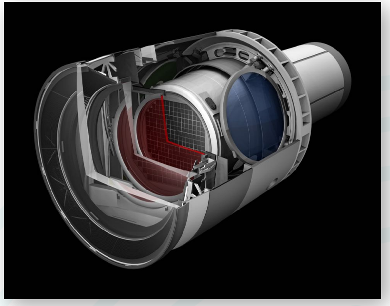

3.2 G pixels | ⌀ 1.65 m | 3.7 m long | 3 ton | 3 lenses | 3.5° field of view | 9.6 deg^2 | 6 filters *ugrizy* | 320-1050 nm

Source: LSST: from Science Drivers to Reference Design and Anticipated Data Products

Image: Google Earth

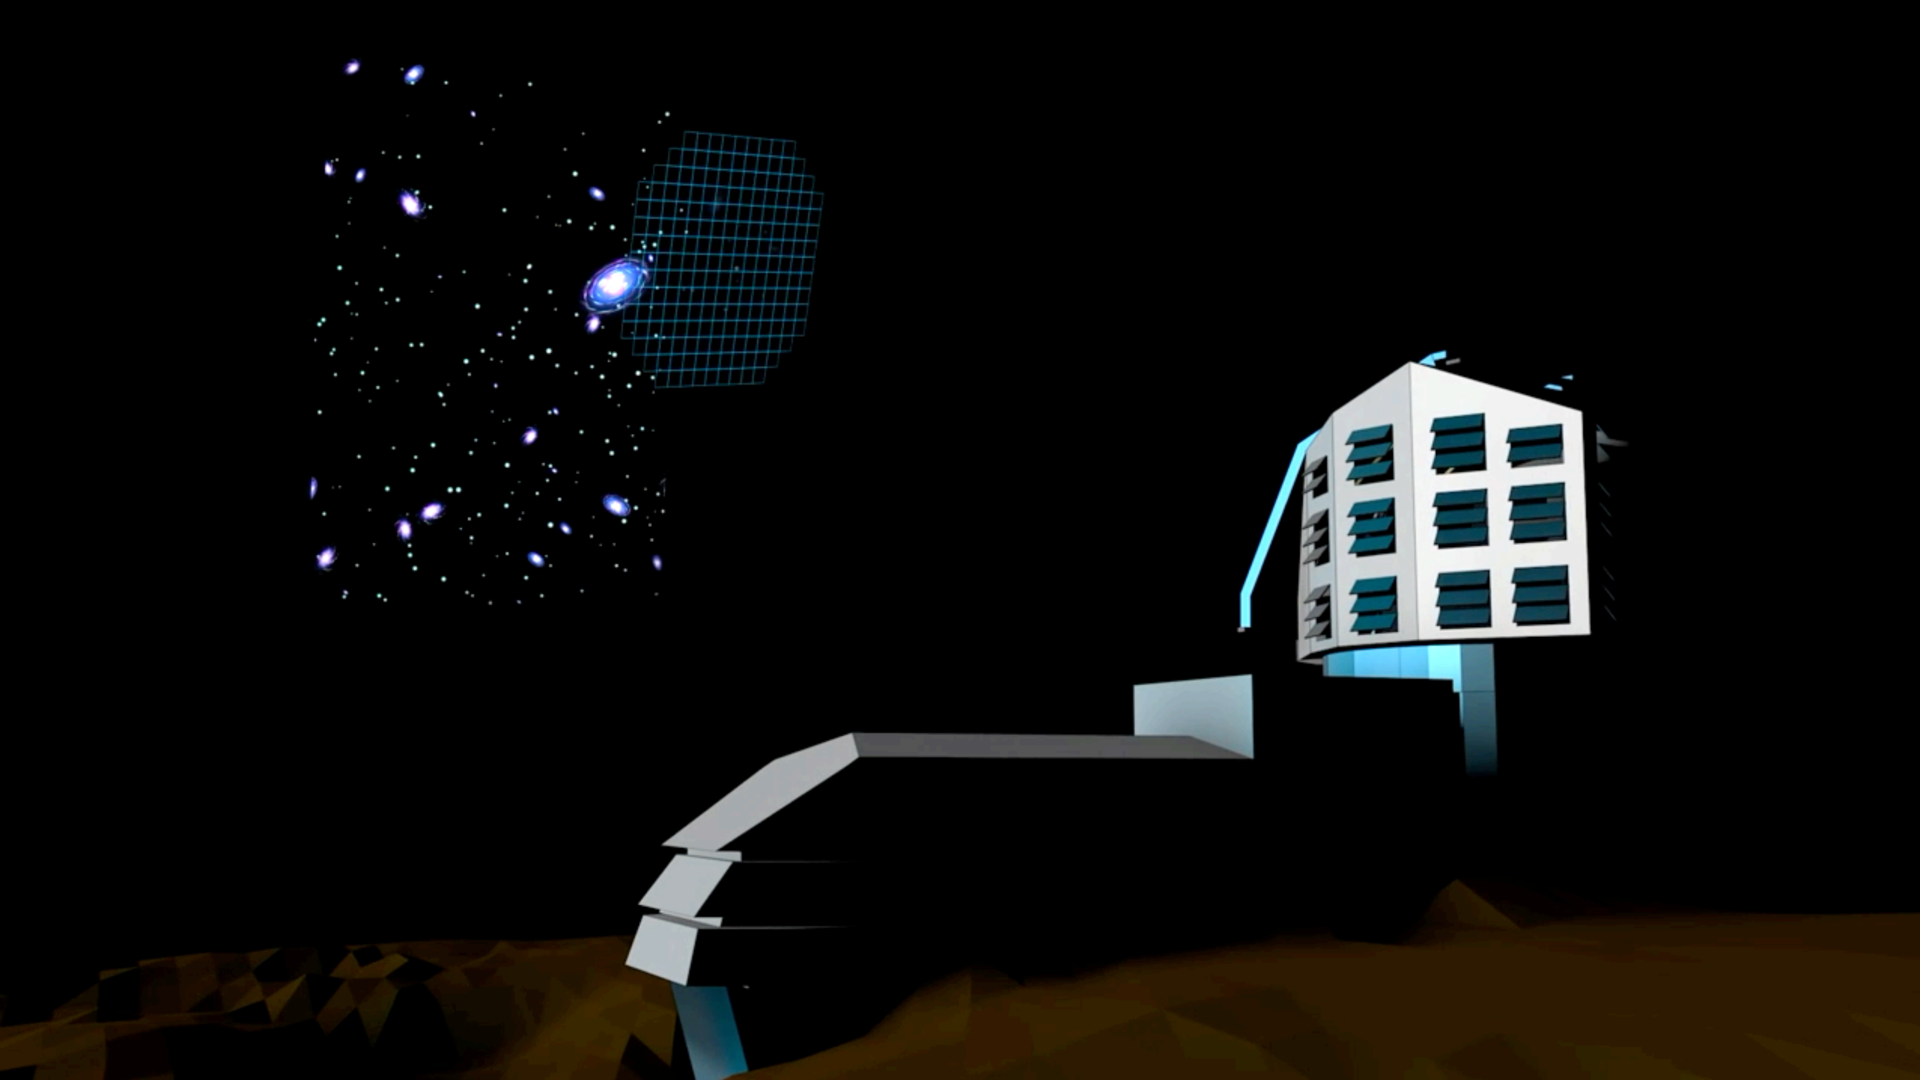


Image credit: Rubin Observatory


Credit: Rubin Obs/NSF/AURA/H. Stockebrand

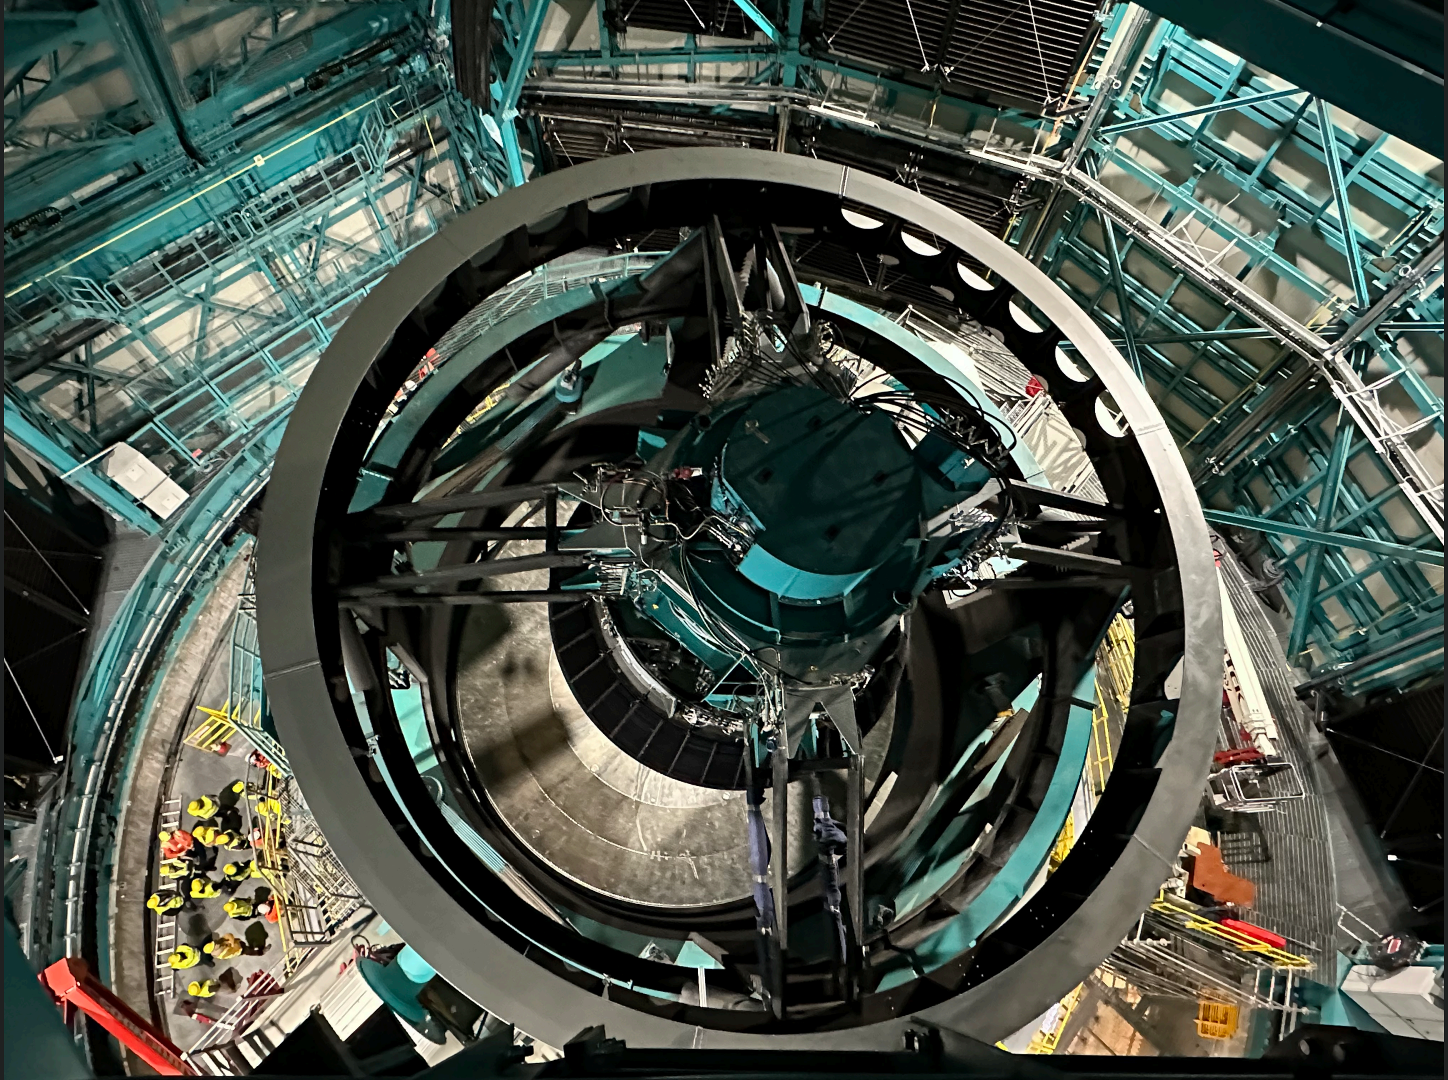

Image credit: RubinObs/NSF/AURA/A. Alexov

Raw data

6.4 GB per exposure (compressed)

2000 science + 500 calibration images per night

16 TB per night

300 nights per year, ~5 PB per year

Aggregated data over 10 years of operations*, including derived data

image collection: ~6M exposures

final catalog database: 15 PB

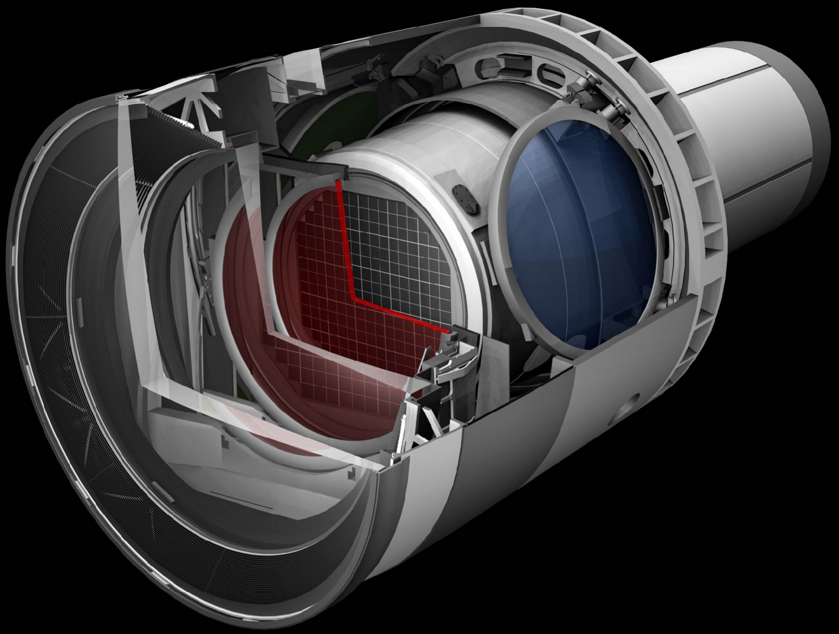

* source: [LSST Key Numbers]

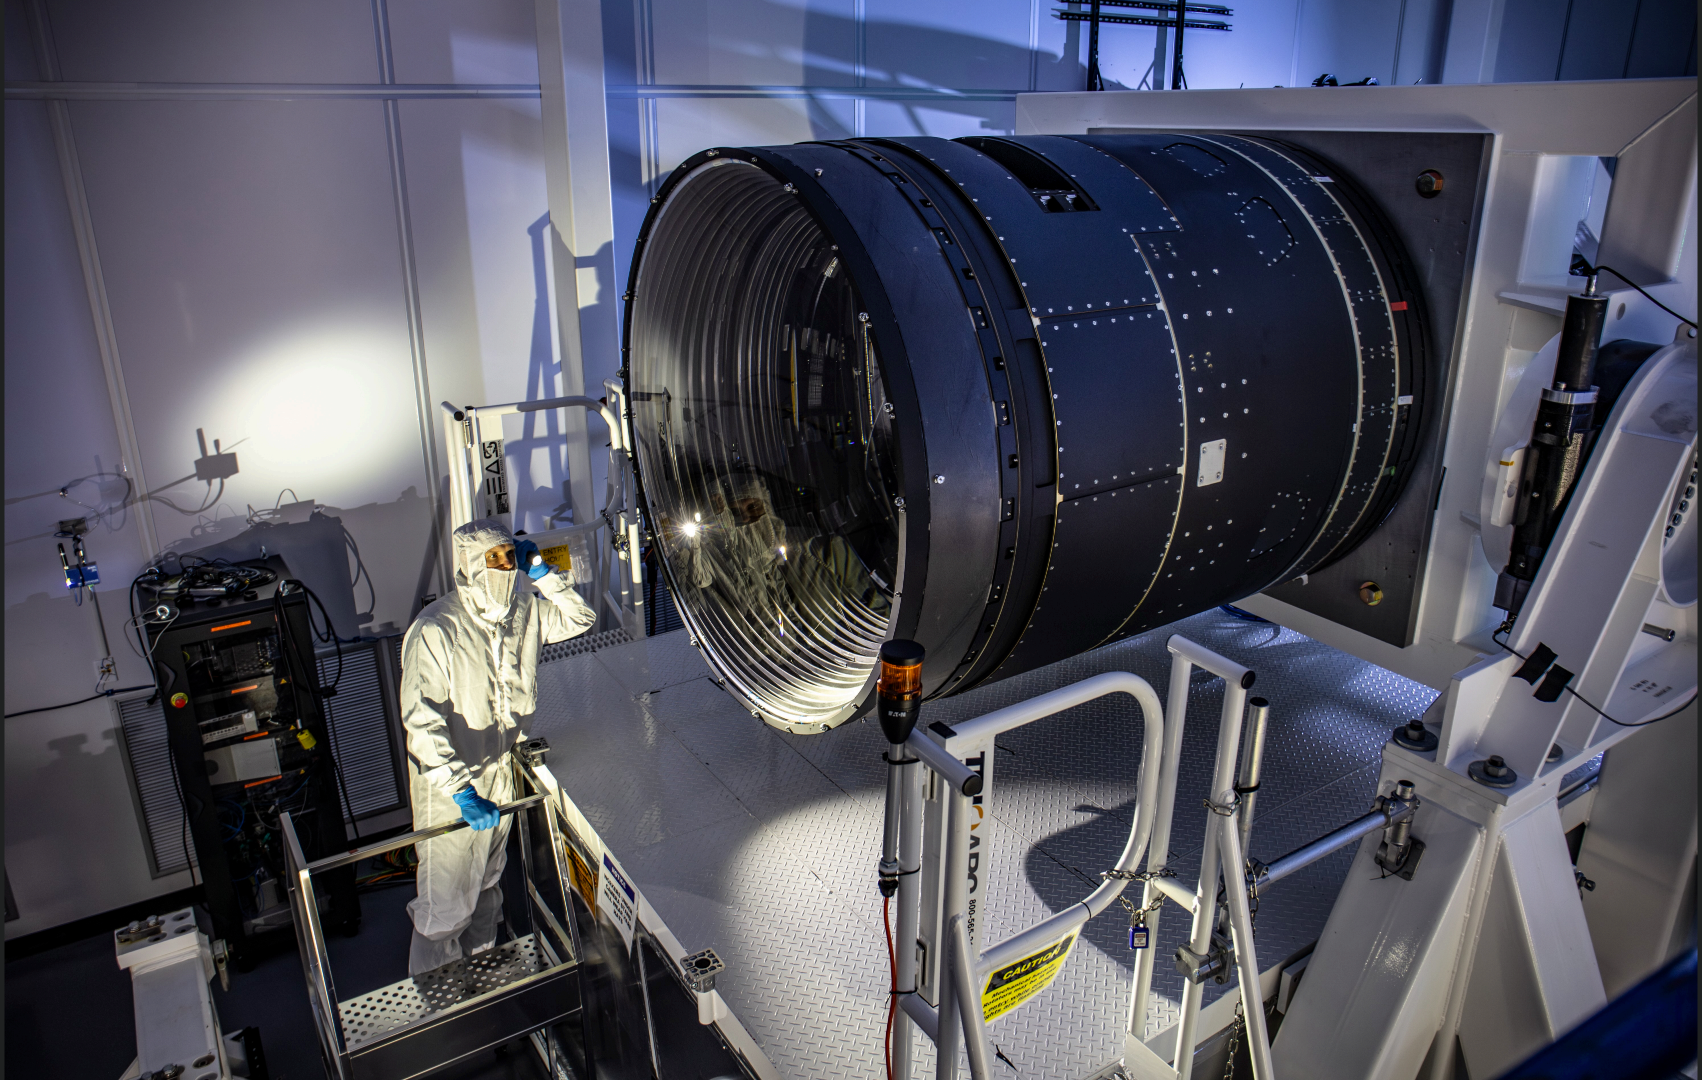

Image credit: Jacqueline Ramseyer Orrell / SLAC National Accelerator Laboratory

Cumulative data volume

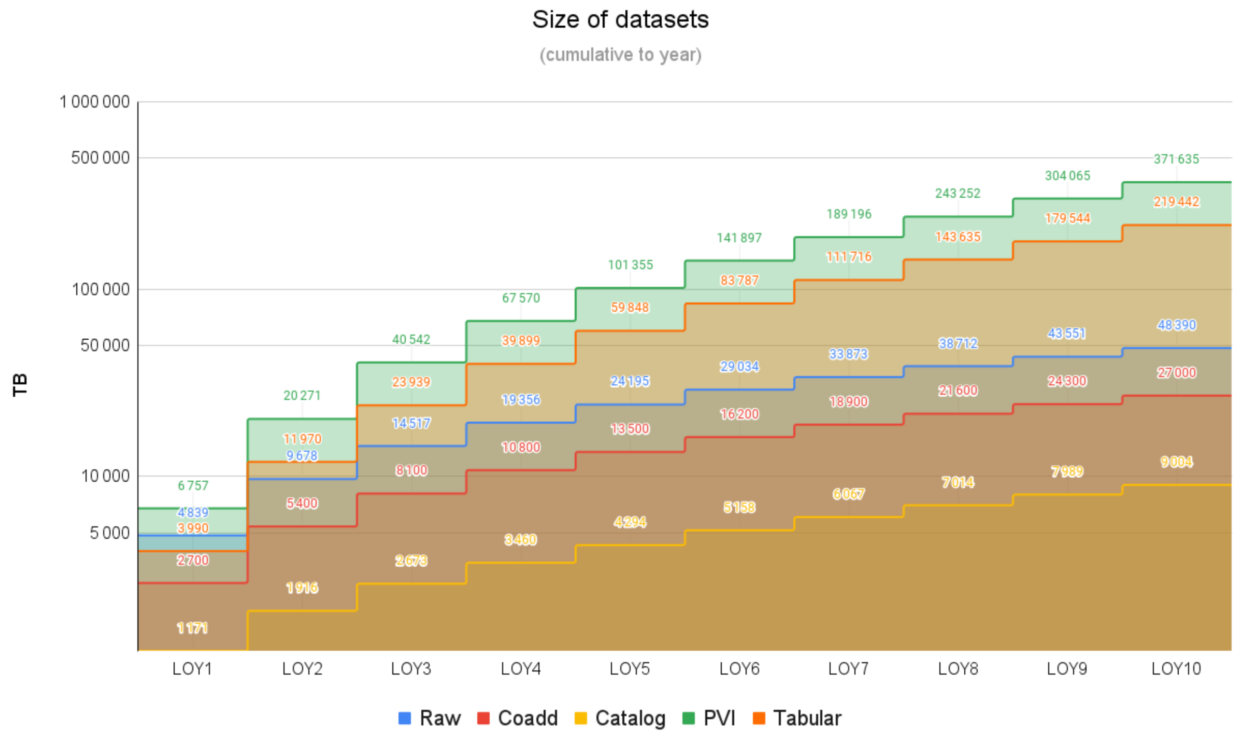

~0.5 EB of data by the end of the survey by 2035

raw image data (~50 PB)

Data Processing and Data Products

Rubin data processing and data products

VERA C. RUBIN OBSERVATORY

Raw Data: 20TB/night

Sequential 30s images covering the entire visible sky every few days

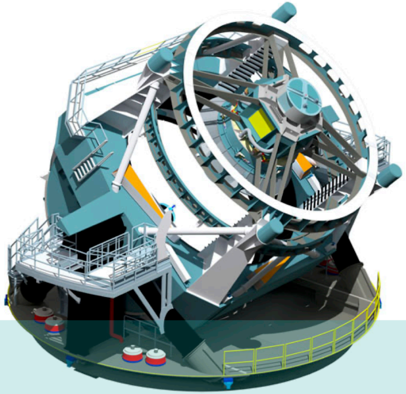

LSST SCIENCE PIPELINES

Prompt Data Products

- Alerts incl. science, template and difference image cutouts — 60s via Alert Streams
- Catalogs of detections incl. difference images, transient, variable & solar system sources — 24h via Prompt Products
- Raw & processed visit images (PVIs), difference images — 80h via Image Services

Data Release Data Products

Final 10yr Data Release:

- Images: 5.5 million x 3.2 Gpixels
- Catalog: 15PB, 37 billion objects

via Data Releases

Community Brokers

Rubin Data Access Centres (DACs)

USA (USDF)
Chile (CLDF)
France (FRDF)
United Kingdom (UKDF)

Independent Data Access Centers (IDACs)

Access to proprietary data and the Science Platform require Rubin data rights

Rubin Science Platform

Provides access to LSST Data Products and services for all science users and project staff.

RUBIN SCIENCE PLATFORM

PORTAL | NOTEBOOKS | WEB APIS

DATA RELEASES | ALERT FILTERING SERVICE | USER DATABASES | USER FILES | USER COMPUTING | SOFTWARE TOOLS

Credit: Leanne Guy

Cloud

EPO Data Center

Dedicated Long Haul Networks

Two redundant 100 Gb/s links from Santiago to Florida (existing fiber)
Additional 100 Gb/s link (spectrum on new fiber) from Santiago-Florida (Chile and US national links not shown)

UK Data Facility
IRIS Network, UK

Data Release Production (25%)

US Data Facility
SLAC, California, USA

Archive Center
Alert Production
Data Release Production (35%)
Calibration Products Production
Long-term storage
Data Access Center
Data Access and User Services

France Data Facility
CC-IN2P3, Lyon, France

Data Release Production (40%)
Long-term storage

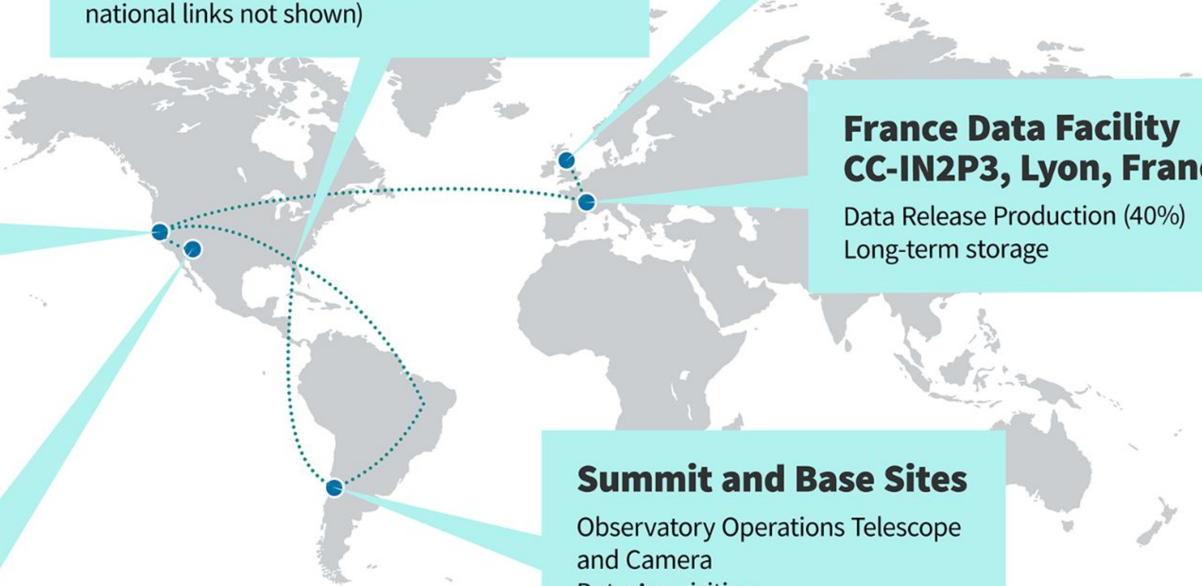

Summit and Base Sites

Observatory Operations Telescope and Camera
Data Acquisition
Long-term storage
Chilean Data Access Center

HQ Site
AURA, Tucson, USA

Observatory Management
Data Production
System Performance
Education and Public Outreach

Rubin Long Haul Network

- More than 100,000 kms of fiber (46% underwater) crossing 6 countries
- 16 points of presence along the primary path
- 15 engineers from several research and education organizations

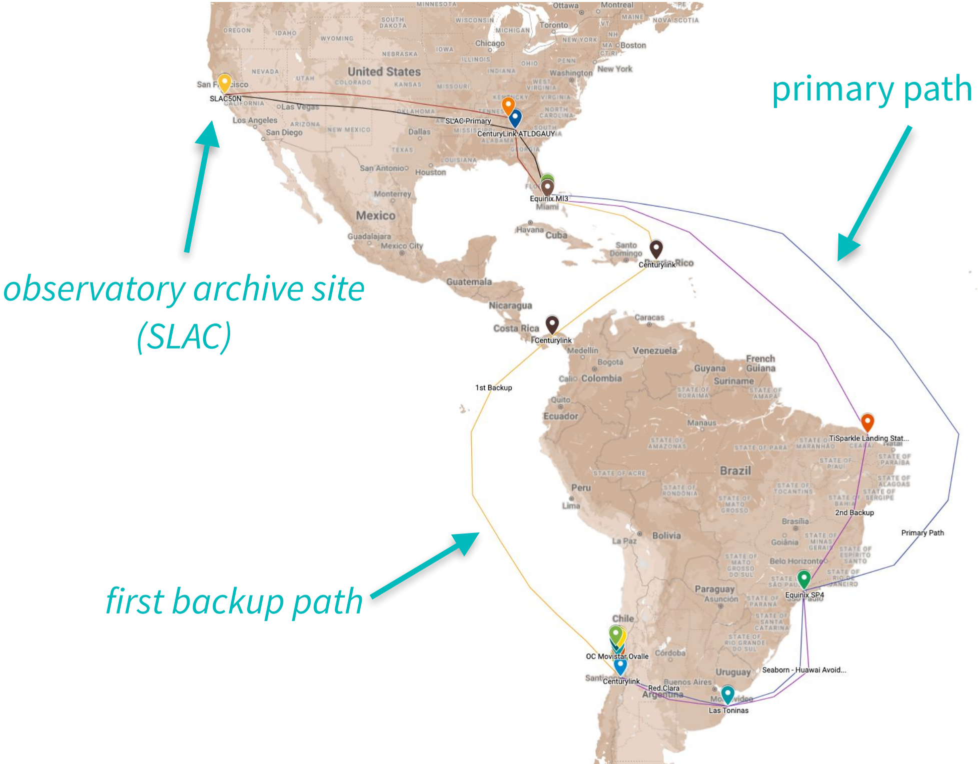

Adapted from C. Silva, Rubin Observatory

Alert Production

- Difference Image Analysis performed on science images to detect varying sources (position or flux)
  - an alert is generated, recorded and emitted for public consumption
  - alert packet composed of data about the source (coordinates, photometry) and 30x30 pixels image cutout
- Alert processing performed at the US data facility and alert stream sent to community-developed alert brokers
  - within 60 seconds of image readout
  - about 10 million alert packets emitted per night, ~82 KB each
  - by the 10th year of the survey, size of the alert database to be about 2.2 PB
- Seven community brokers will receive the Rubin alert stream, ingest it, filter it, augment it (e.g. classification) and redistribute it publicly
  - ALERCE, AMPEL, ANTARES, Babamul, Fink, Lasair, Pitt-Google

More information: Rubin Alert Brokers

Fink (2019 –)

Operating 24/7 since 2019, serving 100+ unique users per day (scientists & follow-up facilities).

- Real-time components (million+ event/night)
- Event database (~1B entries)

Processing the ZTF alert stream since 2019

- 210 million ZTF alerts received
- ⅔ is classified: 50% galactic, 15% Solar System, few% extra-galactic
- Coupled to GCN: Fermi, Swift, Icecube, LVK, …

Community-driven: scientists bring building bricks

- 60+ members, 15+ scientific topics covered

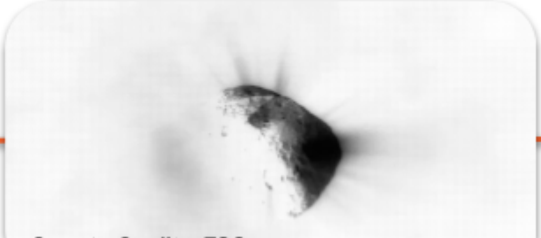
Comet - *Credits: ESO*


Active Galactic nucleus

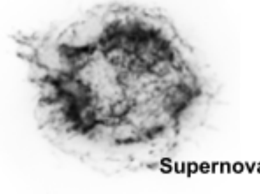
Supernova

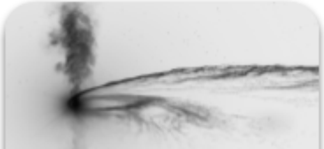
Tidal disruptive event

Source: J. Peloton (Fink)
More information: https://fink-broker.org

LSST Science Pipelines

LSST Science Pipelines

- Major processing steps
  - single-frame processing
  - calibration
  - image coaddition
  - coadd processing
  - catalog production

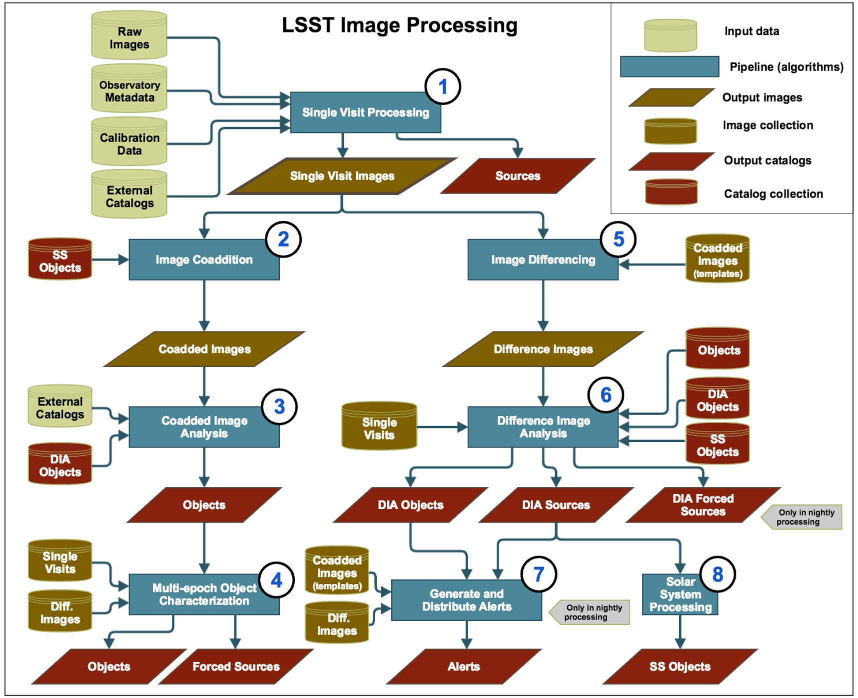

Details available in:

- An Overview of the LSST Image Processing Pipelines, by J. Bosch et al.
- Rubin Observatory Data Products Definition Document, by M. Jurić et al.

LSST Science Pipelines (cont.)

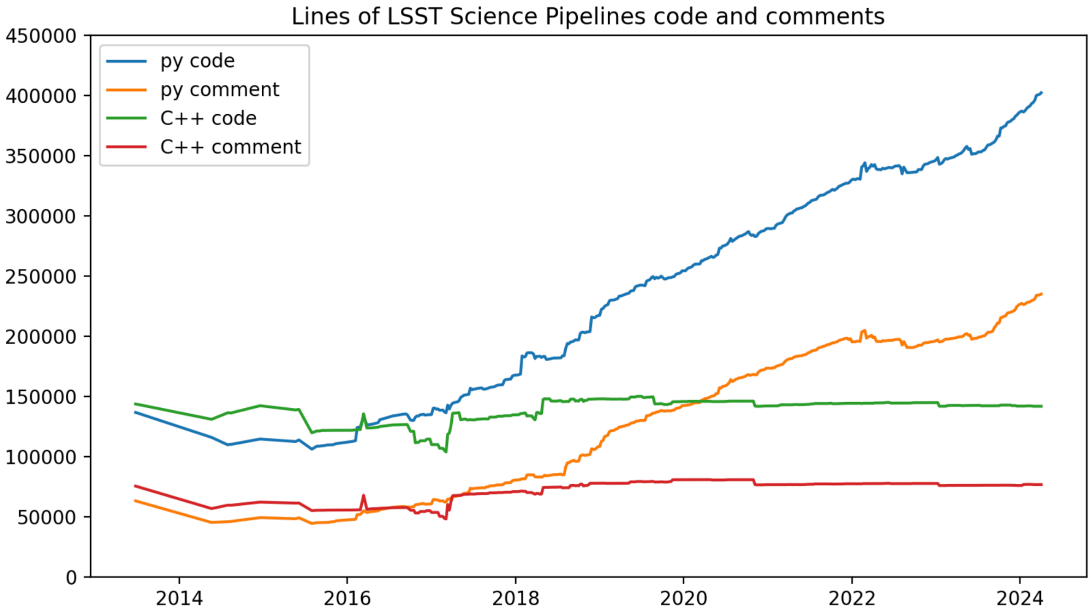

Lower layer written in C++ for performance

Upper layer in Python for convenience and expressivity

Open source: https://github.com/lsst

Very good documentation: https://pipelines.lsst.io

Source: T. Jenness, Rubin Observatory

LSST Science Pipelines (cont.)

- We use CernVM-FS for distributing the LSST science pipelines to the three Rubin data facilities and beyond
  - personal computers, other sites (NERSC, OpenScience Grid, etc.)
  - weekly and stable releases
- Two packaging formats
  - conda-based environment
  - Apptainer container images
- 800+ third party packages in the conda environment + 100 LSST packages
- Base operating system: currently transitioning from CentOS 7 to AlmaLinux 9
  - execution of the pipelines work on several flavors of Linux (both `x86` and `aarch64`) as well as on macOS (both `x86` and `arm64`)

- More information: https://sw.lsst.eu

Rubin Data Butler

- Middleware which abstracts the data access details from the developers of the image processing science algorithms
  - exposes *in-memory* Python objects: hides the details of where the data are being read or written, what file format is used, what data transport protocol is used, etc.
  - users can locate data using common astronomical concepts (e.g. patch of the sky, physical filter, detector id, exposure id, observation identifier, etc.)
- Main components:
  - ***registry***: catalog of files organized conceptually and associated with astronomical concepts; not concerned with where or how data is persisted
  - ***datastore***: responsible for serialization (deserialization) of Python objects to (from) a storage system
- Registry is backed by a relational database
  - typically PostgreSQL for production, sqlite for small amounts of data and for testing
- Datastore is where files are actually persisted
  - supported protocols: S3, Google Cloud Storage, webDAV and POSIX
  - implements a caching system using local storage (memory or disk)
- Butler exposes both a CLI and a Python API

Details: The Vera C. Rubin Observatory Data Butler and Pipeline Execution System by T. Jenness et al.

Data Release Processing

Data release processing

- Data release processing: **annual campaign of reprocessing all the raw images collected** since the beginning of the survey
  - data products: *science-ready images* and *astronomical catalog*
- To be performed at three Rubin data facilities
  - US data facility (SLAC National Accelerator Laboratory, CA, USA) — 35%
  - UK data facility (IRIS and GridPP, UK) — 25%
  - French data facility (CC-IN2P3, Lyon, FR) — 40%
- US data facility is the observatory archive site: stores and serves a full copy of raw and published data products

- Connectivity among those facilities provided by ESnet (transatlantic segment from/to SLAC), GEANT (within Europe), JANET (UK) and RENATER (FR)

Distributed image processing

- A *Rubin data facility* is typically composed of a compute element and one or more Data Butler repositories
- Compute element
  - exposes the site's batch farm to the workflow executor
  - typically composed of Nordugrid's ARC CE and Slurm
- Data Butler
  - registry database: PostgreSQL
  - datastore: Weka (S3) at US facility, dCache (webDAV) at French facility, XrootD (webDAV) at UK facility
- Rucio storage elements exposing the Butler data stores
  - for intermediate and final data products exchange between the facilities

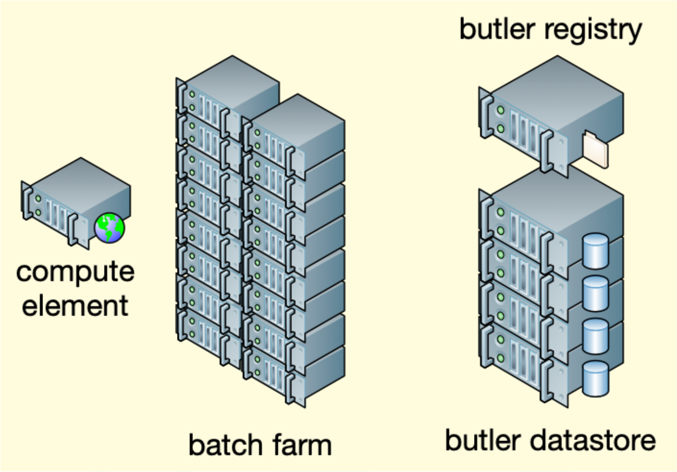

Distributed image processing (cont.)

- Rubin's Batch Processing Service (BPS)
  - generates the workflow to be executed at each facility: a directed acyclic graph of independent units of work
  - takes into account data dependencies and data location
- PanDA
  - creates pilot jobs and coordinates the execution of the workflow
  - each job executes one or several science algorithms over a set of input data, stores output data in the Butler repository local to the facility

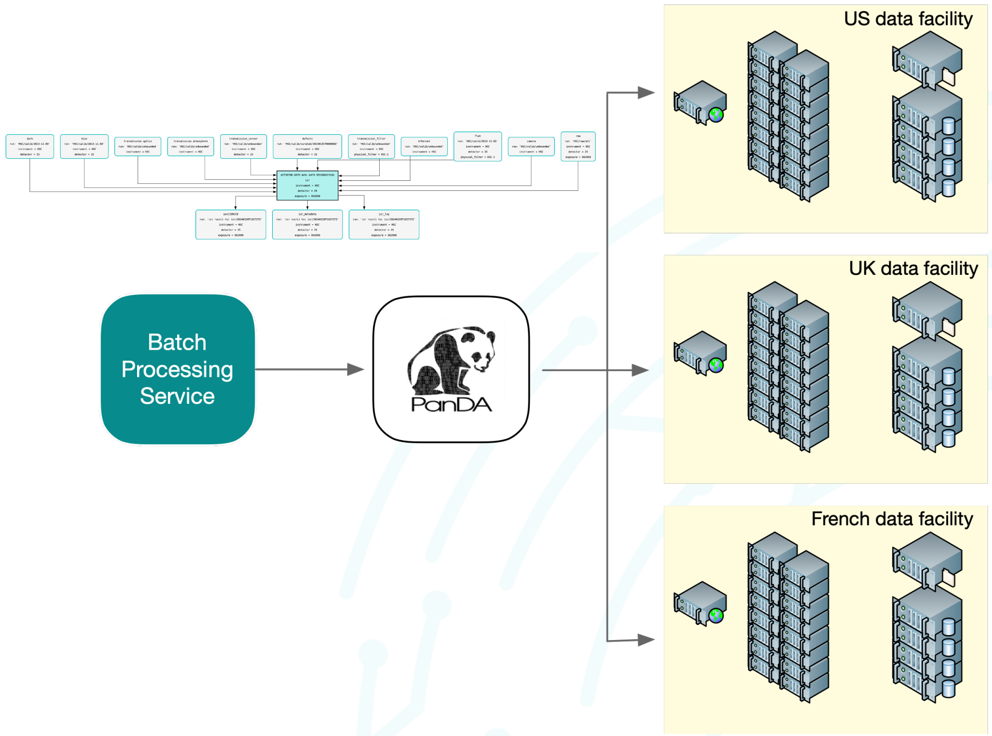

Inter-site data replication

- Butler repositories only record information about data local to each data facility
- To compile a data release we need to transfer final (and some intermediate) data products to the archive site at the US facility and ingest them into the local Butler repository
- Data replication is driven by Cern's Rucio & FTS
  - relevant Butler datasets to be replicated among the facilities are registered into Rucio and replication rules configured
  - FTS executes the transfers on behalf of Rucio
- Upon arrival of files at each facility, a daemon ingests the replicated files into the local Butler repo
  - file arrival is detected by listening to Rucio's Hermes events emitted through Kafka
- Identification of what files need to be replicated and registration into local Butler is performed by Rubin-specific tools

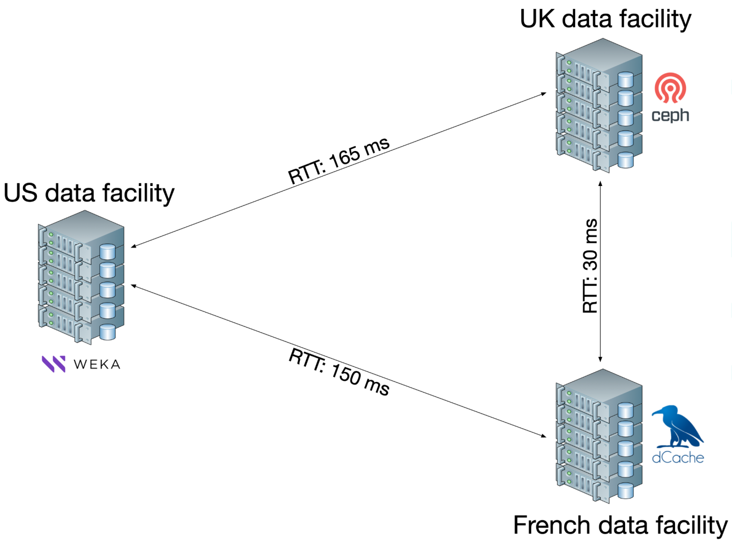

Data replication over high-latency network links

Projected data transfer rates

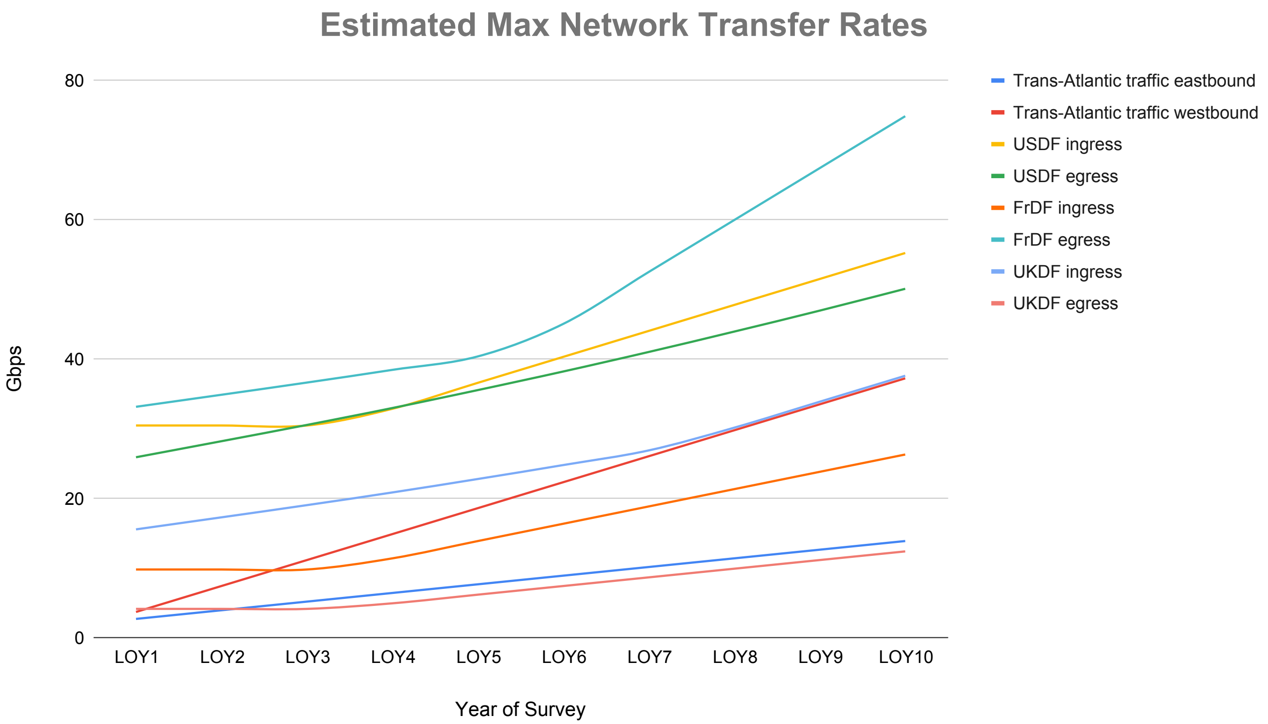

These estimations make some assumptions that we may need to revisit as we learn how data reprocessing will proceed in real-life conditions

Adapted from R. Dubois, Rubin Observatory

Data Products

Astronomical catalog database

- Read-only relational database
  - contains the physical properties of the celestial objects and light sources detected by processing the science images
  - exposes a subset of SQL to scientists
  - repopulated every year with the contents of the latest data release
- Few number of *very tall* tables

table	rows	columns	storage
ForcedSource	50 T	10s	2 PB
Source	9 T	100s	5 PB
ObjectExtra	1.5 T	1000s	1.2 PB
Object	47 B	1000s	100 TB

Adapted from F. Mueller, Rubin Observatory

Details about the columns: Data Preview 0.2 Schema Browser

Astronomical catalog database (cont.)

- Data spatially partitioned: catalog contents physically distributed over a set of independent database servers
- Qserv — packaged as a Kubernetes application

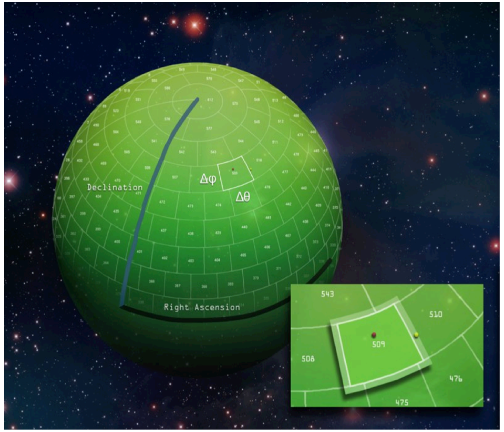

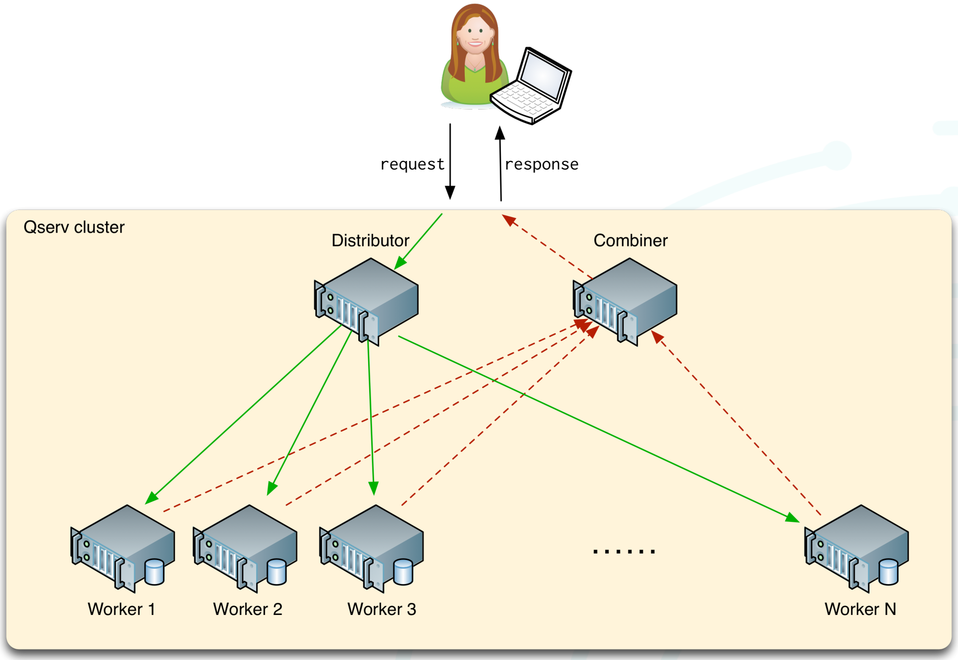

Details: https://qserv.lsst.io

Rubin Science Platform

https://data.lsst.cloud

- Integrated web-based environment for interactive scientific exploration of image and catalog data
  - both GUI and programmatic interfaces
  - execution of Python notebooks
  - visualization of images, navigation, tabular data, graphics
  - interfaced with LSST own catalog as well as external catalogs (via IVOA interfaces)
- Architected as a set of cooperating services orchestrated by Kubernetes

- Click here for a short video (8 min) demonstrating the science platform in action, prepared by G. Mainetti

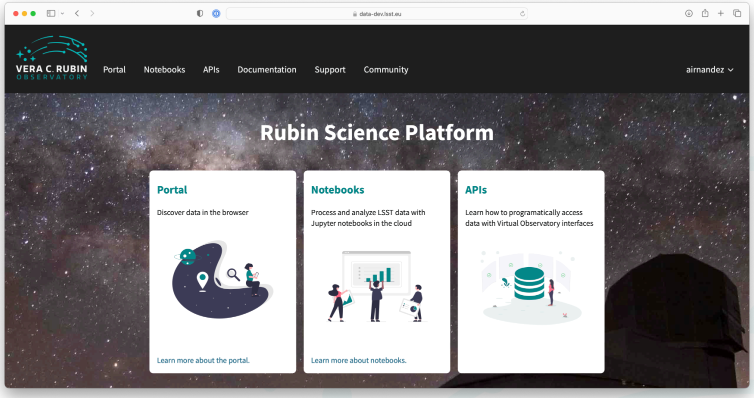

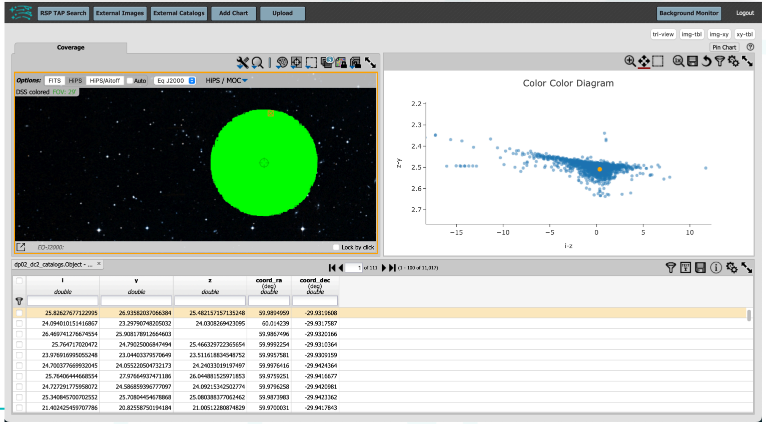

Further details: LSST Science Platform Vision Document data.lsst.cloud

Results | DP0.2 Images | DP0.2 Catalogs | DP0.3 Catalogs | Upload

Background Monitor | Logout

LSST DP0.2 DC2 Tables

Use Image Search (ObsTAP)

Table Collection (Schema): dp02_dc2_catalogs (tables: 11)
Data Preview 0.2 contains the image and catalog products of the Rubin Science Pipelines v23 processing of the DESC Data Challenge 2 simul...
Table Collection (Schema) count: 5

Tables: dp02_dc2_catalogs.Object
Properties of the astronomical objects detected and measured on the deep coadded images.
Table count: 11

View: UI assisted | Edit ADQL

Enter Constraints

Spatial — no target found

Shape Type: Cone Shape / Polygon Shape

Coordinates or Object Name | Try NED then Simbad

Examples: '62, -37' '60.4 -35.1' '4h11m59s -32d51m59s equ j2000' '239.2 -47.6 gal'
'NGC 1532' (NB: DC2 is a simulated sky, so names are not useful)

Radius

10 arcseconds

Valid range between: 1" and 360000"

Position Columns: coord_ra, coord_dec (from the selected table on the right)

Temporal

Object ID Search

Output Column Selection and Constraints

Reset Column Selections & Constraints

Name (char)	constraints (char)	unit (char)	ucd (char)	description (char)	datatype (char)	arraysize (char)	utype (char)	xtype (char)	principal (int)	size (int)	column_index (int)	indexed (int)	std (int)
coord_dec		deg	pos.eq.dec;meta.main	Fiducial ICRS Declination of centroid u	double				0			0	0
coord_ra		deg	pos.eq.ra;meta.main	Fiducial ICRS Right Ascension of centro	double				0			0	0
deblend_nChild				Number of children this object has (de	int				0			0	0
deblend_skipped				Deblender skipped this source	boolean				0			0	0
detect_fromBlend				This source is deblended from a paren	boolean				0			0	0
detect_isDeblende				True if source has no children and is in	boolean				0			0	0
detect_isDeblende				True if source has no children and is in	boolean				0			0	0
detect_isIsolated				This source is not a part of a blend.	boolean				0			0	0
detect_isPatchInn				True if source is in the inner region of a	boolean				0			0	0
detect_isPrimary				True if source has no children and is in	boolean				0			0	0
detect_isTractInne				True if source is in the inner region of a	boolean				0			0	0
footprintArea		pixel		Number of pixels in the source's detect	int				0			0	0
g_ap03Flux		nJy		Flux within 3.0-pixel aperture. Forced	double				0			0	0
g_ap03Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap03FluxErr		nJy		Flux uncertainty within 3.0-pixel apert	double				0			0	0
g_ap06Flux		nJy		Flux within 6.0-pixel aperture. Forced	double				0			0	0
g_ap06Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap06FluxErr		nJy		Flux uncertainty within 6.0-pixel apert	double				0			0	0
g_ap09Flux		nJy		Flux within 9.0-pixel aperture. Forced	double				0			0	0
g_ap09Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap09FluxErr		nJy		Flux uncertainty within 9.0-pixel apert	double				0			0	0
g_ap12Flux		nJy		Flux within 12.0-pixel aperture. Forced	double				0			0	0
g_ap12Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap12FluxErr		nJy		Flux uncertainty within 12.0-pixel aper	double				0			0	0
g_ap17Flux		nJy		Flux within 17.0-pixel aperture. Forced	double				0			0	0
g_ap17Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap17FluxErr		nJy		Flux uncertainty within 17.0-pixel aper	double				0			0	0
g_ap25Flux		nJy		Flux within 25.0-pixel aperture. Forced	double				0			0	0
g_ap25Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap25FluxErr		nJy		Flux uncertainty within 25.0-pixel aper	double				0			0	0
g_ap35Flux		nJy		Flux within 35.0-pixel aperture. Forced	double				0			0	0
g_ap35Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap35FluxErr		nJy		Flux uncertainty within 35.0-pixel aper	double				0			0	0
g_ap50Flux		nJy		Flux within 50.0-pixel aperture. Forced	double				0			0	0
g_ap50Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap50FluxErr		nJy		Flux uncertainty within 50.0-pixel aper	double				0			0	0
g_ap70Flux		nJy		Flux within 70.0-pixel aperture. Forced	double				0			0	0
g_ap70Flux_flag				General Failure Flag. Forced on g-band	boolean				0			0	0
g_ap70FluxErr		nJy		Flux uncertainty within 70.0-pixel aper	double				0			0	0
g_apFlux_flag				General Failure Flag. Measured on g-ba	boolean				0			0	0

Search | Row Limit: 50000 | Populate and edit ADQL

/ notebooks / tutorial-notebooks /

Name	Modified
data	2 minutes ago
DP02_09_Custom_Coadds	2 minutes ago
DP02_11_User_Packages	2 minutes ago
DP02_01_Introduccion_a_DP0...	2 minutes ago
DP02_01_Introduction_to_DP0...	2 minutes ago
DP02_02a_Introduction_to_TA...	2 minutes ago
DP02_02b_Catalog_Queries_wi...	2 minutes ago
DP02_02c_Image_Queries_with...	2 minutes ago
DP02_03a_Image_Display_and...	2 minutes ago
DP02_03b_Image_Display_with...	2 minutes ago
DP02_03c_Big_deepCoadd_Cu...	2 minutes ago
DP02_04a_Introduction_to_the...	2 minutes ago
DP02_04b_Intermediate_Butler...	2 minutes ago
DP02_05_Source_Detection_a...	2 minutes ago
DP02_06a_Interactive_Image_...	2 minutes ago
DP02_06b_Interactive_Catalog...	2 minutes ago
DP02_07a_DiaObject_Samples...	2 minutes ago
DP02_07b_Variable_Star_Light...	2 minutes ago
DP02_08_Truth_Tables.ipynb	2 minutes ago
DP02_10_Deblender_Data_Pro...	2 minutes ago
DP02_12a_PSF_Data_Products...	2 minutes ago
DP02_12b_PSF_Science_Demo...	2 minutes ago
DP02_13a_Image_Cutout_SciD...	2 minutes ago
DP02_14_Injecting_Synthetic_...	2 minutes ago
DP02_15_Survey_Property_Ma...	2 minutes ago
DP03_01_Introduccion_a_DP0...	2 minutes ago
DP03_01_Introduction_to_DP0...	2 minutes ago
DP03_02_Main_Belt_Asteroids...	2 minutes ago
DP03_03_Trans-Neptunian_Obj...	2 minutes ago
DP03_04a_Introduction_to_Ph...	2 minutes ago
DP03_04b_Advanced_Phase_C...	2 minutes ago
DP03_05_Near-Earth_Objects.i...	2 minutes ago
DP03_06_User_Uploaded_Tabl...	2 minutes ago
LICENSE	2 minutes ago
mobu.yaml	2 minutes ago
NOTICE	2 minutes ago
README.md	2 minutes ago

4. Extract cutout images

When a closer look at images in the vicinity of a star or galaxy is needed, visualizing a small cutout instead of the full image is more effective.

4.1. A cutout from a `calexp`

In a science use case, the coordinates of the object of interest may be known already, but for this example choose pixels near an extended source in the image above.

[21]:
```
wcs = calexp.getWcs()
radec = wcs.pixelToSky(2250, 700)
ra, dec = radec.getRa().asDegrees(), radec.getDec().asDegrees()
print(ra, dec)
```
Last executed at 2024-09-18 11:45:25 in 10ms

```
53.09004683168595 -33.97759584632161
```

Use the same visit and detector as were used for the `dataId` in Section 3.

> **Notice:** In order to make a cutout of a `calexp`, the RA, Dec, visit, *and* detector must all be specified (recall that visit uniquely identifies band, so band does not need to be explicitly identified). This is different from making a cutout with a `deepCoadd` (Section 4.2 below), which only requires RA, Dec, and band. This is because there are multiple `calexp` images for a given coordinate and band, but only one `deepCoadd` per band.

[22]:
```
visit = 192350
detector = 175
```
Last executed at 2024-09-18 11:45:25 in 9ms

Use the `cutout_calexp` function defined in Section 1 to extract a 301x301 pixel cutout.

> **Optional:** print the help documentation for the function.

[23]:
```
# help(cutout_calexp)
```
Last executed at 2024-09-18 11:45:25 in 8ms

[24]:
```
my_cutout_calexp = cutout_calexp(butler, ra, dec, visit, detector, cutoutSideLength=301)
```
Last executed at 2024-09-18 11:45:25 in 400ms

Display the cutout from the `calexp`.

[25]:
```
fig, ax = plt.subplots()
display = afwDisplay.Display(frame=fig)
display.scale('asinh', 'zscale')
display.mtv(my_cutout_calexp.image)
plt.show()
remove_figure(fig)
```
Last executed at 2024-09-18 11:45:26 in 627ms

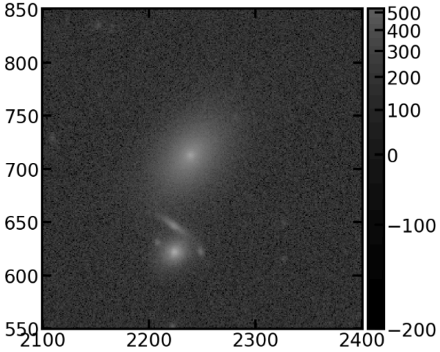

> Figure 6: A cutout from a `calexp` image, displayed in grayscale, with scale bar at right and axes in pixels. The three brightest objects are galaxies.

[26]:
```
del wcs, radec, ra, dec, visit, detector, my_cutout_calexp
```
Last executed at 2024-09-18 11:45:26 in 4ms

4.2. A cutout from a `deepCoadd` image

Data Access Centers

Data Access Centers (DACs)

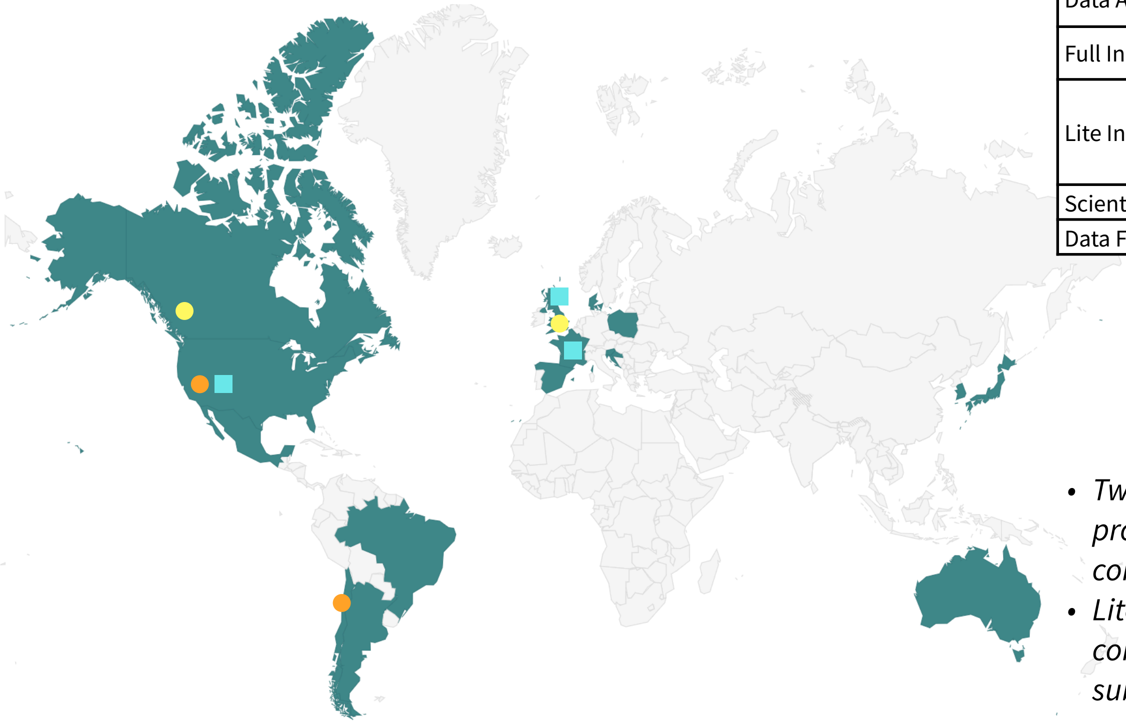

Data Access Center	● (orange)	Chile, US
Full Independent DAC	● (yellow)	UK
Lite Independent DAC		Argentina, Australia, Brazil, Canada, Denmark, Japan (x2), Mexico, Poland, Slovenia, South Korea, Spain
Scientific Processing Center		Croatia
Data Facilities	■ (cyan)	France, UK, US

- *Two Data Access Centers managed by the Rubin project to serve released data to scientific communities*
- *Lite Independent Data Access centers are in-kind contributions. They intend to store and serve a subset of data*

Technical Documentation

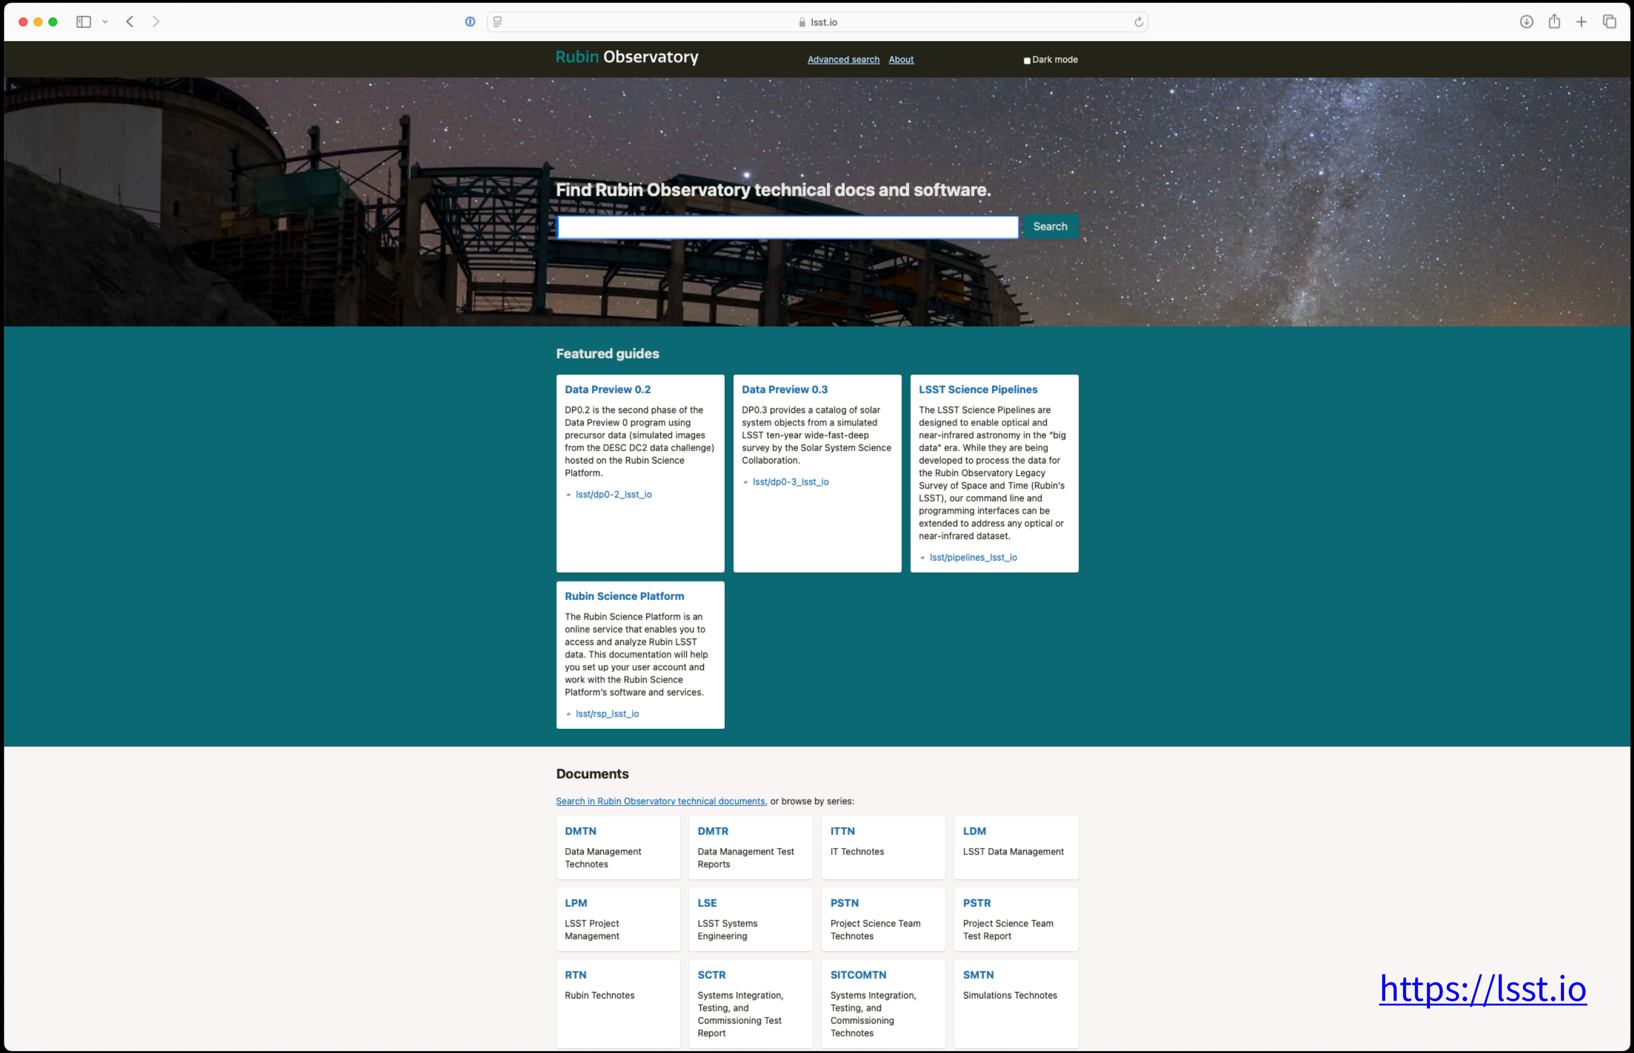

https://lsst.io

Summary

Summary

- Rubin is currently in an advanced phase of preparation of its distributed data processing infrastructure
  - preview of data products are available: https://dp0-2.lsst.io
  - still a lot of work to do
- Eager to re-use and integrate processing models, practices and tools proven effective for other science projects with Rubin-specific tools
  - Rucio, FTS, CernVM FS, PanDA
  - IVOA standards
- Many details still to be sorted out and adapted to real-life processing conditions

Questions & Comments