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

Experience deploying an analysis facility for the Rubin Observatory's Legacy Survey of Space and Time (LSST) data

CHEP2023 Norfolk 8-12 May 2023

F. Hernandez, F. Jammes, Q. Le Boulc'h, **G. Mainetti**

8 May 2023

Outline

1. The Rubin Observatory and LSST
2. The analysis facility
3. The CC-IN2P3 infrastructure
4. Conclusions and Perspectives

The Rubin Observatory and LSST

1

Rubin Observatory and LSST

The Vera C. Rubin Observatory

- Cerro Pachón @ Chili (2647m asl)
- Main mirror **8.4m** ⌀
- **9.6 deg²** Field of View
- **3.2 G pixels** camera
- f/1.234 aperture

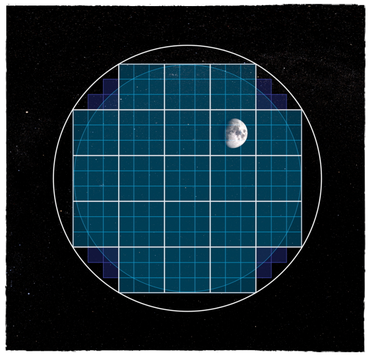

The Legacy Survey of Space and Time (LSST)

- Composition of the Universe: Dark Energy and Dark Matter
- The Solar System inventory
- The changing sky: variable and transient objects observation
- The Milky Way cartography

- 37 Billion objects (20B Galaxies, 17B Stars) catalog
- Each object observed more than 800 times
- Full visible sky survey in 3 nights
- One pointing each 38s, 10M galaxies per pointing

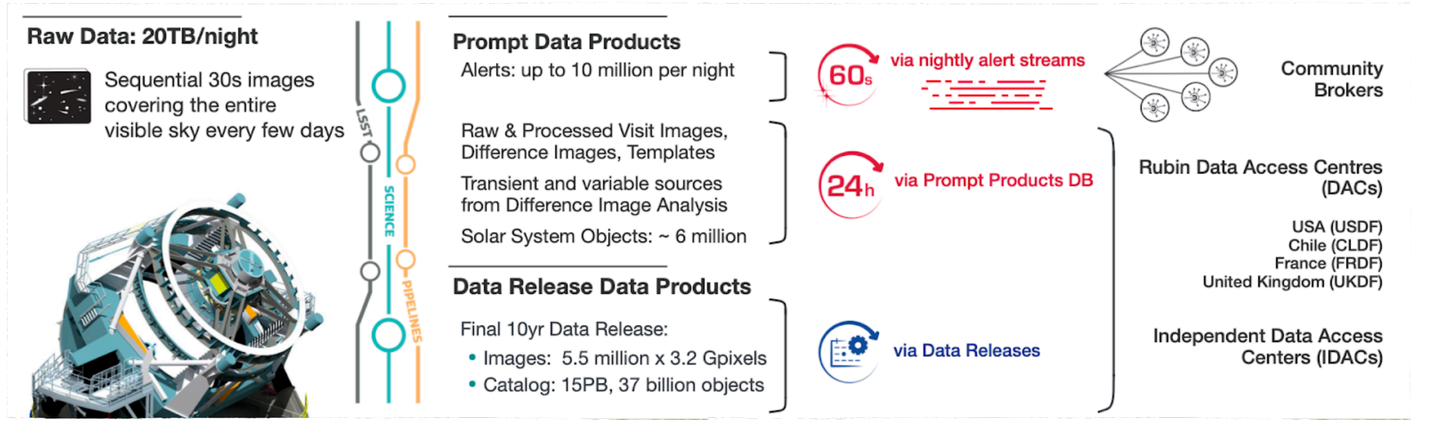

Images Source: Rubin Observatory

The analysis facility

2

The analysis facility

- Objectives:
  - Provide researchers with a platform to easily access/analyze survey data (images and catalogs)
  - Integrate it with the CC-IN2P3 (e.g authentication, $HOME and other file systems, …) for a smooth transition between environments
  - Deploy a scalable and resilient platform
- Technology:
  - **Kubernetes**
  - Open source development : https://github.com/lsst-sqre, https://www.lsst.io/
- Two main components:
  - **Qserv**: the astronomical catalog database
  - **Rubin Science Platform** (RSP): the interactive analysis platform

Qserv

2.1

Qserv

- Developed by SLAC with contributions from IN2P3
- Shared-nothing Massively Parallel Processing Relational Database
- Spherical partitioning with overlap, sciSQL (UDF)
- Shared scans (concurrent query load)
- Data Replication
- 100 % Open Source

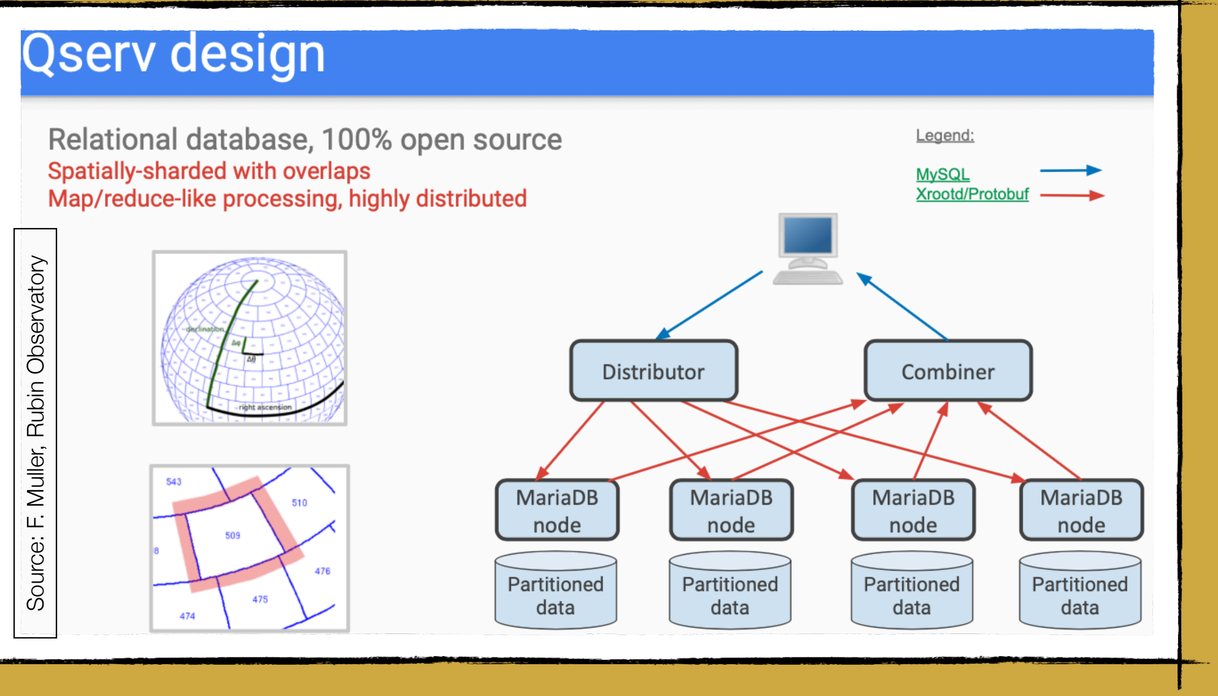


```
apiVersion: qserv.lsst.org/v1beta1
kind: Qserv
metadata:
  name: qserv
spec:
  queryService:
    type: NodePort
    nodePort: 30040
  storageClassName: "qserv-local-storage"
  storage: "100Gi"
  worker:
    replicas: 15
    replicationResources:
      limits:
        cpu: 36
  tolerations:
  - key: "dedicated"
    operator: "Equal"
    value: "qserv"
    effect: "NoSchedule"
```

- Based on the Kubernetes operator-sdk framework
- 2 commands to deploy it on cloud or bare metal clusters
  - k apply -f manifest/operator.yaml
  - k apply -k manifest/<instance>

https://qserv.lsst.io

Deployment demo: https://is.gd/FK62Wa

Data Ingestion: Qserv Ingest

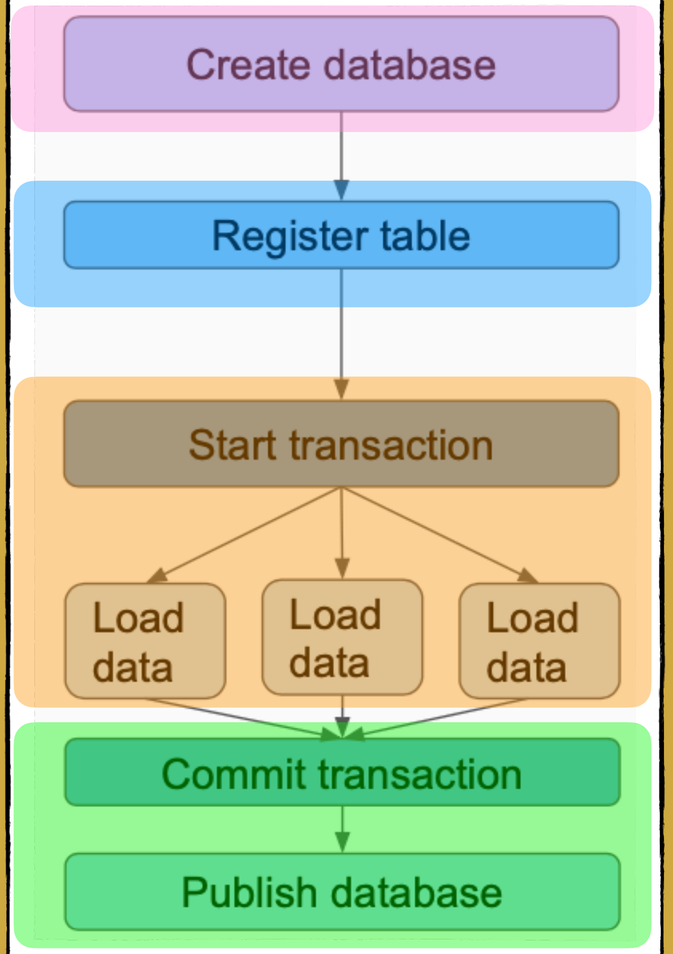

Tool used for data ingestion
Developed by F. Jammes (IN2P3)

argo

Argo Workflow

```
- name: main
  dag:
    tasks:
    - name: queue
      template: ingest-step
      arguments:
        parameters: [{name: script, value: load-queue.sh}]
    - name: register
      template: ingest-step
      arguments:
        parameters: [{name: script, value: register.sh}]
    - name: transactions
      template: transactions
      dependencies: [queue, register]
    - name: check-transactions
      template: ingest-step
      arguments:
        parameters: [{name: script, value: check-transactions.sh}]
      dependencies: [transactions]
    - name: publish
      template: ingest-step
      arguments:
        parameters: [{name: script, value: publish.sh}]
      dependencies: [check-transactions]
    - name: index-tables
      template: index-tables
      dependencies: [publish]
    - name: validate
      template: ingest-step
      arguments:
        parameters: [{name: script, value: validate.sh}]
      dependencies: [index-tables]
    - name: benchmark
      template: benchmark
      dependencies: [validate]
```

Implementation of a large-scale data loading algorithm:

2M files and ~40TB ingested in 5h

Gabriele Mainetti

https://github.com/lsst-dm/qserv-ingest

Rubin Science Platform

2.2

Rubin Science Platform

- What is it?
  - Web environment for interactive data analysis
- Why?
  - To provide astronomers with an all-in-one tool allowing easy and quick data access/analysis integrating:
    - catalogs and tables viewer (as TOPCAT, STILTS)
    - image viewer and analyzer (as DS9)
    - advanced analysis with LSST python stack via Jupyter
    - gateway to Qserv catalogs (Interoperability) for Virtual Observatory (VO) tools

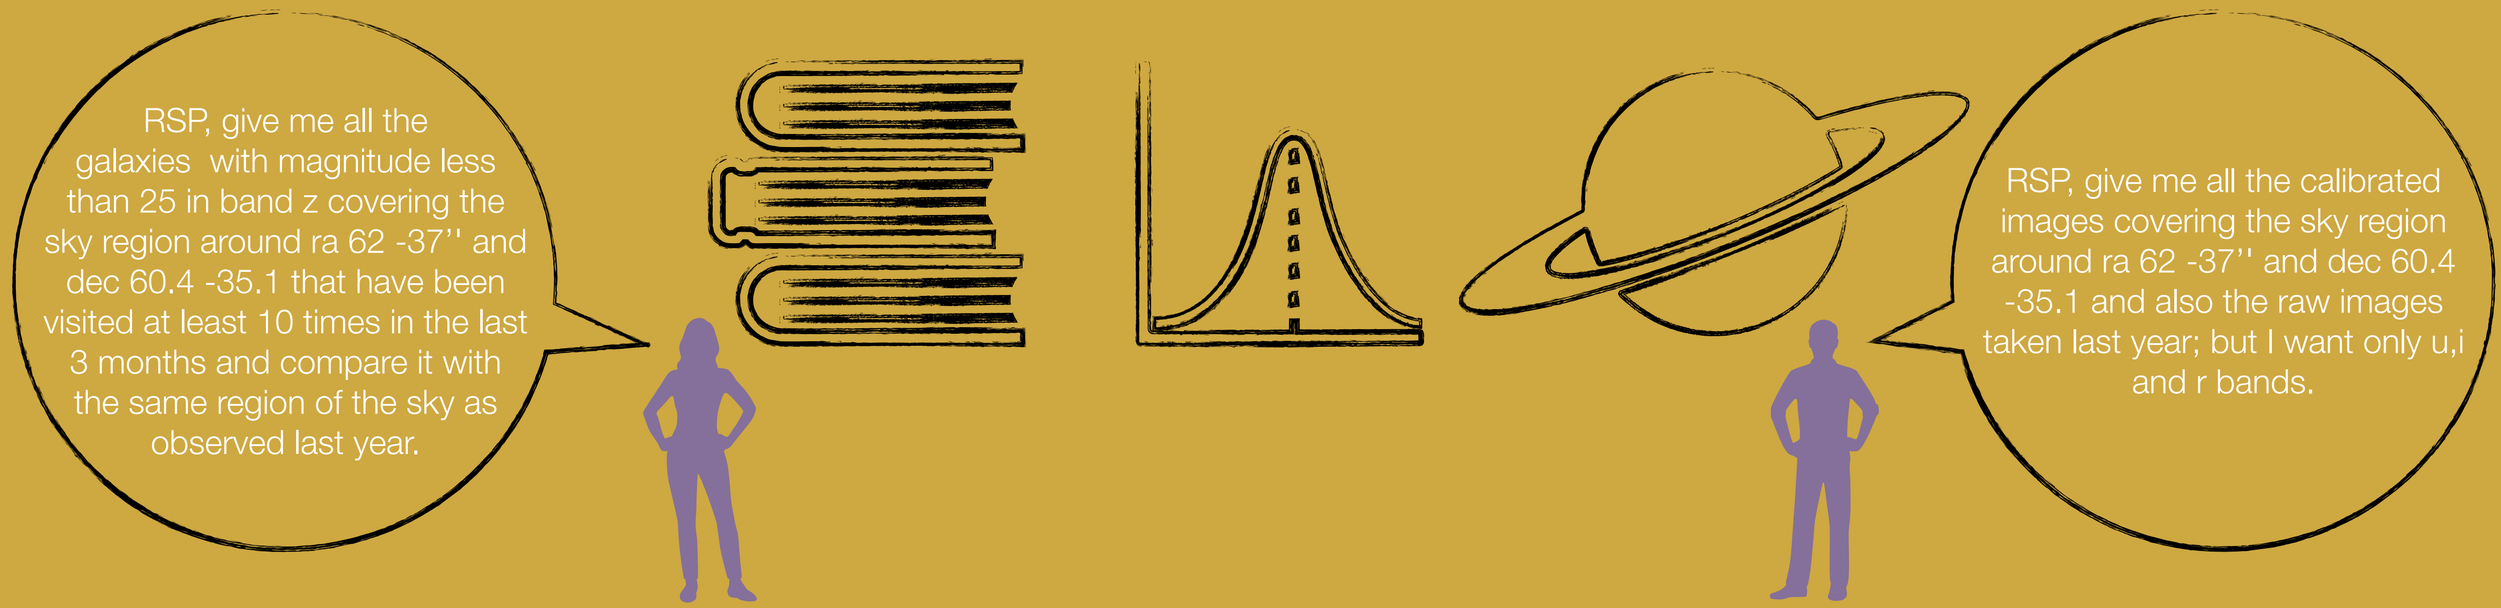

!!! It's not (yet?) ChatGPT, so the request must be translated into ADQL format !!!

Rubin Science Platform

templates	Clean up the sqlproxy service
Chart.yaml	Fix typo in science-platform Chart.yaml
README.md	Updated missed values to reflect new naming
values-base.yaml	Delete obstap service
values-ccin2p3.yaml	activate datalinker
values-idfdev.yaml	adjusted naming to be more generic
values-idfint.yaml	Update values-idfint.yaml
values-idfprod.yaml	Delete obstap service

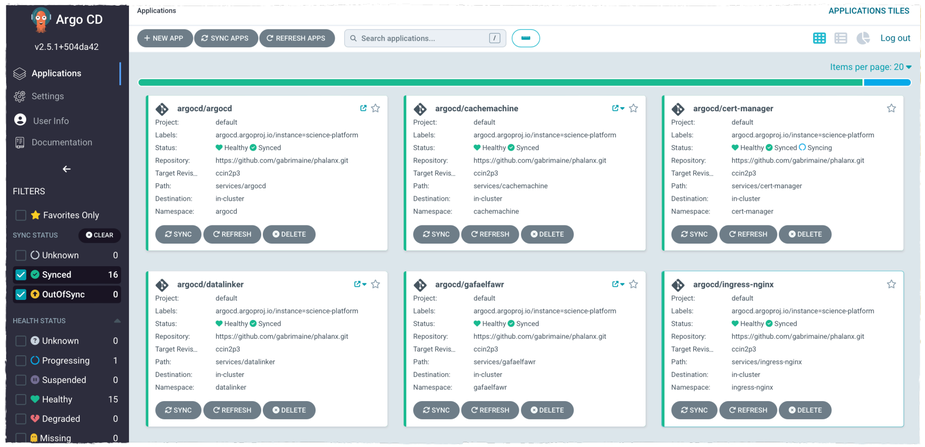

- Developed mainly by the Rubin SQuaRE Team
- Configured via Helm Chart (https://phalanx.lsst.io)
  - One config per data facility and one config per application
- Deployed via ArgoCD
- 4 Core applications:
  - **argocd** for deployment orchestration
  - **cert-manager** for certificates management
  - **Ingress-nginx** for traffic routing
  - **vault-secrets-operator** for secret management
- Authentication managed via in-house **gafaelfawr** IDM application (supporting OpenId, LDAP, github, …)
- Many others applications available

VERA C. RUBIN OBSERVATORY

Portal Notebooks APIs Documentation Support Community

Home page

mainetti

Rubin Science Platform

Portal

Discover data in the browser

Learn more about the portal.

Notebooks

Process and analyze LSST data with Jupyter notebooks in the cloud

APIs

Learn how to programatically access data with Virtual Observatory interfaces

RSP TAP Search | External Images | External Catalogs | Add Chart | Upload

TAP Searches

1. Select TAP Service — Using LSST RSP https://data-dev.lsst.eu/api/tap - Replace...

2. Select Query Type — Single Table (UI assisted) | Edit ADQL (advanced) | Image Search (ObsTAP)

3. Select Table — Table Collection (Schema): **dp02_dc2_catalogs** — 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: **dp02_dc2_catalogs.Object** — Properties of the astronomical objects detected and measured on the deep coadded images.

4. Enter Constraints

Spatial — *no target found*

Longitude Column : coord_ra
Latitude Column: coord_dec
Shape Type: Cone
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°

Temporal

Output Column Selection and Constraints

column_name char	constraints char	unit char	ucd char	description char	datatype char	arraysize char	utype char	xtype char	princi int
coord_dec		deg	pos.eq.dec;meta.main	Fiducial ICRS Declination of centroid used	double				
coord_ra		deg	pos.eq.ra;meta.main	Fiducial ICRS Right Ascension of centroid	double				
deblend_nChild				Number of children this object has (defau	int				
deblend_skipped				Deblender skipped this source	boolean				
detect_fromBlend				This source is deblended from a parent w	boolean				
detect_isDeblendedModelSource				True if source has no children and is in th	boolean				
detect_isDeblendedSource				True if source has no children and is in th	boolean				
detect_isIsolated				This source is not a part of a blend.	boolean				
detect_isPatchInner				True if source is in the inner region of a c	boolean				
detect_isPrimary				True if source has no children and is in th	boolean				
detect_isTractInner				True if source is in the inner region of a c	boolean				
footprintArea		pixel		Number of pixels in the sources detection	int				
g_ap03Flux		nJy		Flux within 3.0-pixel aperture. Forced on	double				
g_ap03Flux_flag				General Failure Flag. Forced on g-band.	boolean				
g_ap03FluxErr		nJy		Flux uncertainty within 3.0-pixel aperture	double				
g_ap06Flux		nJy		Flux within 6.0-pixel aperture. Forced on	double				
g_ap06Flux_flag				General Failure Flag. Forced on g-band.	boolean				
g_ap06FluxErr		nJy		Flux uncertainty within 6.0-pixel aperture	double				
g_ap09Flux		nJy		Flux within 9.0-pixel aperture. Forced on	double				
g_ap09Flux_flag				General Failure Flag. Forced on g-band.	boolean				
g_ap09FluxErr		nJy		Flux uncertainty within 9.0-pixel aperture	double				
g_ap12Flux		nJy		Flux within 12.0-pixel aperture. Forced on	double				
g_ap12Flux_flag				General Failure Flag. Forced on g-band.	boolean				
g_ap12FluxErr		nJy		Flux uncertainty within 12.0-pixel apertur	double				
g_ap17Flux		nJy		Flux within 17.0-pixel aperture. Forced on	double				
g_ap17Flux_flag				General Failure Flag. Forced on g-band.	boolean				
g_ap17FluxErr		nJy		Flux uncertainty within 17.0-pixel apertur	double				
g_ap25Flux		nJy		Flux within 25.0-pixel aperture. Forced on	double				
g_ap25Flux_flag				General Failure Flag. Forced on g-band.	boolean				
g_ap25FluxErr		nJy		Flux uncertainty within 25.0-pixel apertur	double				
g_ap35Flux		nJy		Flux within 35.0-pixel aperture. Forced on	double				
g_ap35Flux_flag				General Failure Flag. Forced on g-band.	boolean				
g_ap35FluxErr		nJy		Flux uncertainty within 35.0-pixel apertur	double				
g_ap50Flux		nJy		Flux within 50.0-pixel aperture. Forced on	double				

Search | Cancel | Row Limit: 50000 | Populate and edit ADQL

Qserv data access

RSP TAP Search | External Images | External Catalogs | Add Chart | Upload

TAP Searches

1. Select TAP Service — Using LSST RSP https://data-dev.lsst.eu/api/tap - Replace...

2. Select Query Type — Single Table (UI assisted) | Edit ADQL (advanced) | Image Search (ObsTAP)

3. Advanced ADQL — ADQL edits below will not be reflected in **Single Table** view

Schema Browser
Schema -> Table -> Column
- dp02_dc2_catalogs
- ivoa
- dp01_dc2_catalogs
- tap_schema
- uws

ADQL Query:

```
SELECT objectId, coord_ra, coord_dec, detect_isPrimary,
scisql_nanojanskyToAbMag(g_cModelFlux) as gmag, scisql_nanojanskyToAb
scisql_nanojanskyToAbMag(r_cModelFlux) as rmag,
scisql_nanojanskyToAbMagSigma(g_cModelFlux, g_cModelFluxErr) as gmag_
scisql_nanojanskyToAbMagSigma(i_cModelFlux, i_cModelFluxErr) as imag_
scisql_nanojanskyToAbMagSigma(r_cModelFlux, r_cModelFluxErr) as rmag_
FROM dp02_dc2_catalogs.Object
WHERE CONTAINS (POINT('ICRS', coord_ra, coord_dec), CIRCLE('ICRS', 62
AND detect_isPrimary = 1
```

ADQL to submit to the selected TAP service

Type ADQL text; you can use the Schema Browser on the left to insert table and column name

Insert fully-qualified column names (recommended for table joins)

Popular Functions

```
TOP n  -- Limit the results to n number of records
ORDER BY [ASC/DESC] -- Used for sorting
POINT('<coordinate system>', RIGHT_ASCENSION, DECLINATION)
CIRCLE('<coordinate system>', RIGHT_ASCENSION_CENTER, DECLINATION_CENTER, RADIUS)
BOX('<coordinate system>', RIGHT_ASCENSION_CENTER, DECLINATION_CENTER, WIDTH, HEIGHT)
POLYGON('<coordinate system>', POINT1, POINT2, POINT3...)
DISTANCE(POINT1, POINT2)
CONTAINS(REGION1, REGION2)
INTERSECTS(REGION1, REGION2)
```

Sample Queries

Query the object table to get positions and composite model magnitudes and their errors in three filters using a CONE search to define a region on the sky. Filter on deblended sources with i-band magnitudes brighter than 25 mag.

```
SELECT objectId, coord_ra, coord_dec, detect_isPrimary,
scisql_nanojanskyToAbMag(g_cModelFlux) as gmag, scisql_nanojanskyToAbMag(i_cModelFlux) as imag,
scisql_nanojanskyToAbMag(r_cModelFlux) as rmag,
scisql_nanojanskyToAbMagSigma(g_cModelFlux, g_cModelFluxErr) as gmag_err,
scisql_nanojanskyToAbMagSigma(i_cModelFlux, i_cModelFluxErr) as imag_err,
scisql_nanojanskyToAbMagSigma(r_cModelFlux, r_cModelFluxErr) as rmag_err
```

RSP TAP Search | External Images | External Catalogs | Add Chart | Upload | Background Monitor | Logout

Table Search

IRSA Catalogs | VO SCS Search | NED

Select Proje — Gaia, WISE, 2MASS, Spitzer, Planck, Herschel, COSMOS, ZTF, PTF, IRAS, MSX, AKARI, Bolocam, USNO, Simulated, Composite_Catalogs, Contributed_Data_Sets, DENIS

Select Catal

Gaia Sou

Coordinates or Object Nam — 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)

Search Method: Cone

Radius: 10 arcseconds
Valid range between: 1° and 3600°

External catalog access

Table Selecti

	constraints	description	units	indx	dbtype	tableflg	sel
solutio		Solution Identifier	null	gn	number(20.0)	2	y
designation		Unique source designation (unique across all Data Releases)	null	n	varchar2(32)	2	y
source_id		Unique source identifier (unique within a particular Data Release)	null	y	number(20.0)	2	y
random_index		Random index used to select subsets	null	y	number(20.0)	2	y
ref_epoch		Reference epoch	yr	n	binary_double	2	y
ra		Right ascension	deg	y	binary_double	2	y
ra_error		Standard error of right ascension	mas	n	binary_double	2	y
dec		Declination	deg	y	binary_double	2	y
dec_error		Standard error of declination	mas	n	binary_double	2	y
parallax		Parallax	mas	y	binary_double	2	y
parallax_error		Standard error of parallax	mas	y	binary_double	2	y
parallax_over_error		Parallax divided by its error	null	n	binary_float	2	y
pm		Total proper motion	mas/yr	n	binary_float	2	y
pmra		Proper motion in right ascension direction	mas/yr	y	binary_double	2	y
pmra_error		Standard error of proper motion in right ascension direction	mas/yr	n	binary_double	2	y
pmdec		Proper motion in declination direction	mas/yr	y	binary_double	2	y
pmdec_error		Standard error of proper motion in declination direction	mas/yr	n	binary_double	2	y
ra_dec_corr		Correlation between right ascension and declination	null	n	binary_float	2	y
ra_parallax_corr		Correlation between right ascension and parallax	null	n	binary_float	2	y
ra_pmra_corr		Correlation between right ascension and proper motion in right ascension	null	n	binary_float	2	y
ra_pmdec_corr		Correlation between right ascension and proper motion in declination	null	n	binary_float	2	y
dec_parallax_corr		Correlation between declination and parallax	null	n	binary_float	2	y
dec_pmra_corr		Correlation between declination and proper motion in right ascension	null	n	binary_float	2	y
dec_pmdec_corr		Correlation between declination and proper motion in declination	null	n	binary_float	2	y

Additional constraints (SQL)

Ex: w3snr>7 and (w2mpro-w3mpro)>1.5 and ra>102.3 and ra<112.3 and dec<-5.5 and dec> -15.5
(source_id_ref = '1861p075_ac51-002577')
The format for date type is yyyy-mm-dd

Search | Cancel

RSP TAP Search | External Images | External Catalogs | Add Chart | Upload | Background Monitor | Logout

Image Search

1. Choose Image Type — View FITS Images | Create 3-Color Composite | View HiPS Images

2. Select Image Source — Search | Use my image | URL

3. Select Target — Coordinates or Object Nam — 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)
Cutout size (leave blank for full size): 500 arcseconds
Valid range between: 1° and 3600°

4. Select Data Set

Filter By: •Clear Filters

Selection: •Clear Selections •Expand All •Collapse All

▼ MISSION:
SDSS (1)
DSS (1)
2MASS (8)
WISE (4)
AKARI (1)
ZTF (1)
more

▼ PROJECT TYPE:
all-sky (16)
extragalactic (50)
compilation (9)
galactic (17)

▼ BAND:
X-ray (1)
UV (3)
optical (14)
near-IR (12)
mid-IR (46)
far-IR (46)
more

- SDSS: Sloan Digital Sky Survey
- DSS: Digitized Sky Survey
- 2MASS All-Sky
- 2MASS Six Degree Mosaics
- 2MASS LGA: Large Galaxy Atlas
- 2MASS Full Survey
- 2MASS 6X Catalog (Images)
- 2MASS Full 6X
- 2MASS 6X Lockman Hole (LH) Ancillary Data Atlas
- 2MASS Calibration
- WISE AllWISE Atlas
- WISE AllSky Atlas
- WISE unWISE co-adds
- z0MOS: The z=0 Multiwavelength Galaxy Synthesis
- AKARI Far-Infrared All-sky Maps
- ZTF Reference Image
- PTF Reference Image
- MSX: Midcourse Space Experiment
- SEIP: Spitzer Enhanced Imaging Products
- Abell1763 Data
- C2D: From Molecular Cores to Planet-Forming Disks
- CLASH: Cluster Lensing And Supernova survey with Hubble
- Cygnus-X: A Spitzer Legacy Survey of the Cygnus-X Complex
- DeepDrill: Spitzer Survey of Deep Drilling Fields
- DUSTINGS: Dust in Nearby Galaxies with Spitzer
- ELFLock: The Eureka Lawrence Berkeley National Laboratory (LBNL) Far-Infrared Lockman Hole Map

Search | Cancel

External images access

Table and image visualization

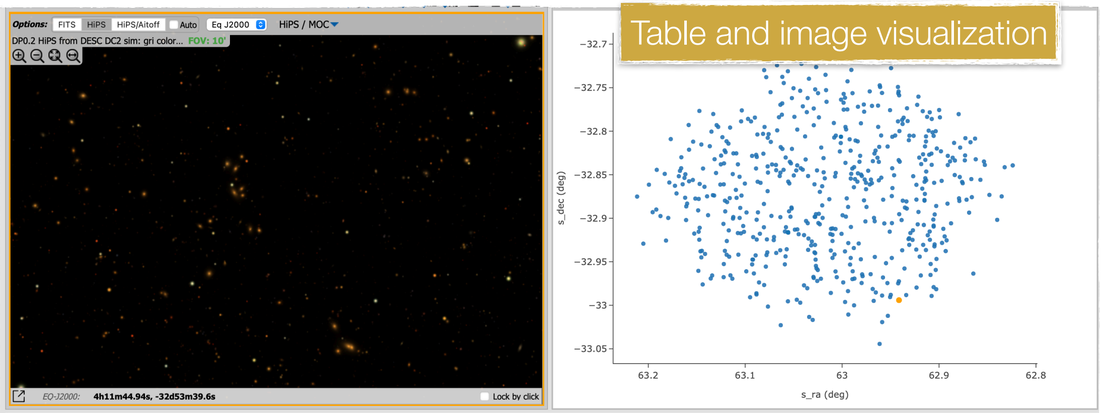

ivoa.ObsCore - data.lsst.cl... 1 of 17 (1 - 100 of 1,620)

dataproduct_type	dataproduct_subtype	calib_level	lsst_band	em_min (m)	em_max (m)	lsst_tract	lsst_patch	lsst_filter	lsst_visit	lsst_detector	lsst_ccdvisitid	t_exptime (s)	t_min (d)
char	char	int	char	double	double	long	long	char	long	long	long	double	double
image	lsst.raw	1	i	6.91e-7	8.18e-7			i_sim_1.4		167		30	60330.11023461111
image	lsst.raw	1	z	8.18e-7	9.22e-7			z_sim_1.4		23		30	60753.02893761111
image	lsst.raw	1	r	5.52e-7	6.91e-7			r_sim_1.4		152		30	60942.25857761111
image	lsst.raw	1	r	5.52e-7	6.91e-7			r_sim_1.4		122		30	60890.3697256111
image	lsst.raw	1	z	8.18e-7	9.22e-7			z_sim_1.4		102		30	61041.21133361111

Jupyter platform

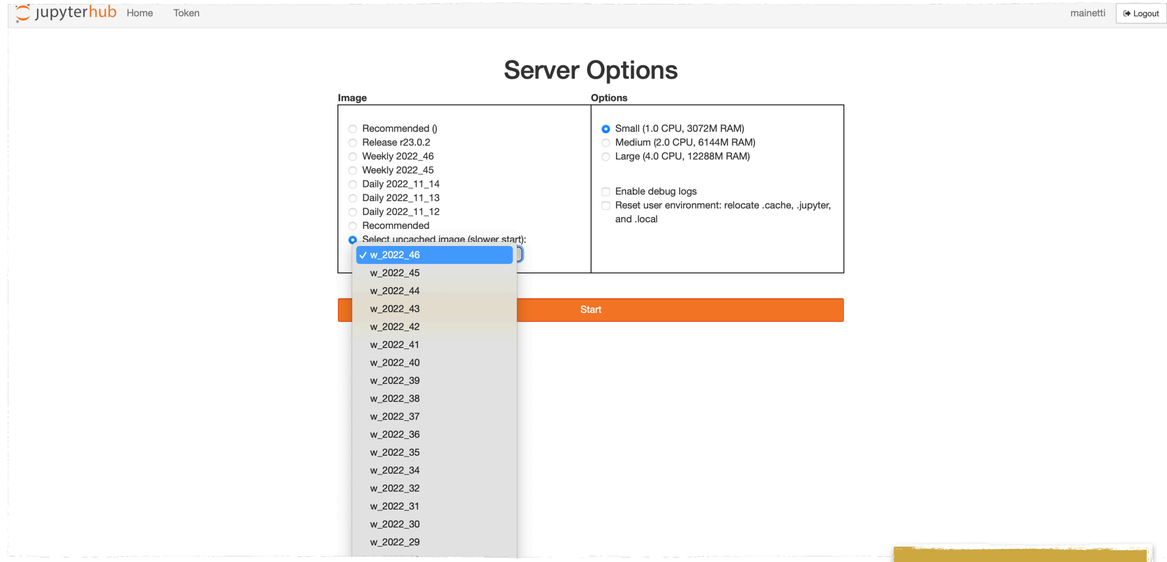

Chart creation

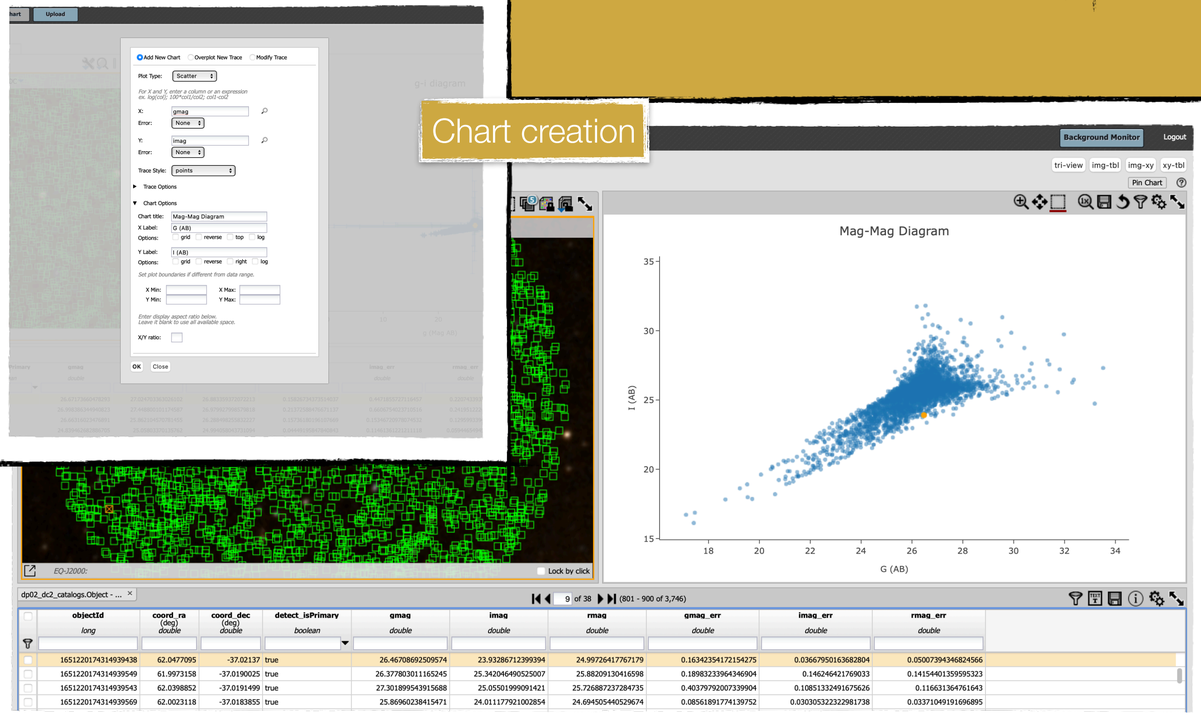

dp02_dc2_catalogs.Object - ... 9 of 38 (801 - 900 of 3,746)

objectId	coord_ra (deg)	coord_dec (deg)	detect_isPrimary	qmag	imag	rmag	qmag_err	imag_err	rmag_err
long	double	double	boolean	double	double	double	double	double	double
1651220174314939438	62.0477095	-37.02137	true	26.46708692509574	23.932867123993394	24.99726417767179	0.16342354172154275	0.03667950163682804	0.05007394346824566
1651220174314939549	61.9973158	-37.0190025	true	26.377803011165245	25.342046490525007	25.88209130416598	0.18983233964346904	0.14624642176903	0.14154401359595323
1651220174314939543	62.0398852	-37.0191499	true	27.301899543915688	25.055019990914	25.726887237284735	0.40379792007339904	0.10851332491675626	0.116631364761643
1651220174314939569	62.0023118	-37.0183855	true	25.86960238415471	24.011177921002854	24.694505440529674	0.08561891774139752	0.030305322322981738	0.03371049191696895

The CC-IN2P3 infrastructure

3

CC-IN2P3 role and resources

- CC-IN2P3 is the French Data Facility (**FrDF**, with USDF, UKDF)
- Annual processing of 40% of the cumulated data
- Long term storage of selected subsets of the data releases
- Eventually several hundreds Pb stored on disc/tapes, 15 Pb for the astronomical catalog of the last data release

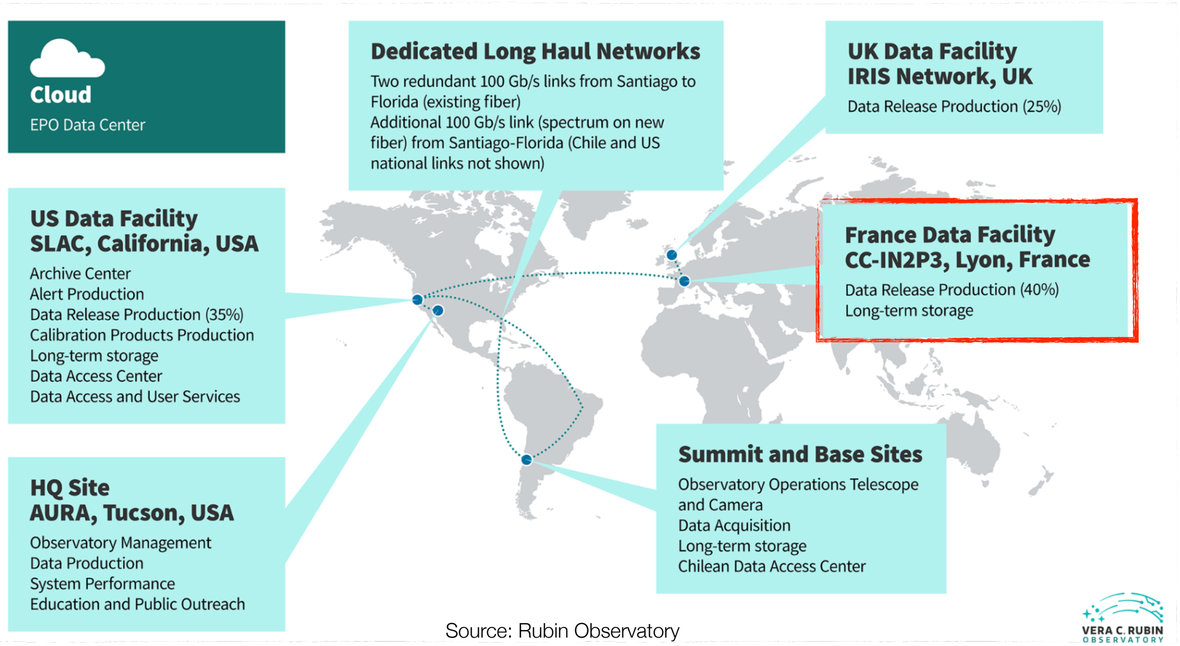

Source: Rubin Observatory

- Bare-metal cluster dedicated to Qserv and RSP:
  - 25 worker-nodes
  - 5 DELL PowerEdge R440, 20 DELL PowerEdge R540
  - 3 worker-nodes as K8S control plane
  - 17 worker-nodes dedicated to Qserv (via taint)
  - RSP worker-nodes :
    - RAM: 256 GB
    - Local storage: 50 TB
- OpenStack cluster as test-bench:
  - 8 VM
  - 3 VM as K8S control plane
  - 4 VM dedicated to Qserv
  - 1 VM dedicated to the RSP and other
- 4 data transfer nodes to expose Qserv data via Caddy web server

Conclusions and Perspectives

4

Conclusions and Perspectives

- Rubin-LSST challenge: the amount of data and the computing resources needed to process it
  - “Overview of the distributed image processing infrastructure to produce the Legacy Survey of Space and Time (LSST)” (Track 1, Thursday, 11:30am, F. Hernandez)
  - “The Rubin Observatory’s Legacy Survey of Space and Time DP0.2 processing campaign at CC-IN2P3” (Track 4, Thursday, 2:15pm, F. Hernandez)
- Scalable, shared and resilient database to deal with the astronomical catalog deployed
- Scalable analysis platform to access and analyze data deployed

- Collect the RSP user feedback to improve it with new functionalities
- Complete the RSP integration with the CC-IN2P3 environment
- Scaling up the amount of data in QServ (~55T at the moment)

Backup

B

RSP Current Status

- Deployed on the production cluster
  - 17 applications activated
  - Simulated DP02 and DP01 catalogs available
  - CC-IN2P3 environment accessible from the notebook platform
  - /sps/lsst not yet accessible
  - Images not yet available (SPS needed to expose the butler)
- Development mainly US:
  - Use of $ Google services (e.g. GCS)
  - Efforts have been needed to adapt some RSP applications to the CC-IN2P3 environment, but modification proposals well-accepted and integrated

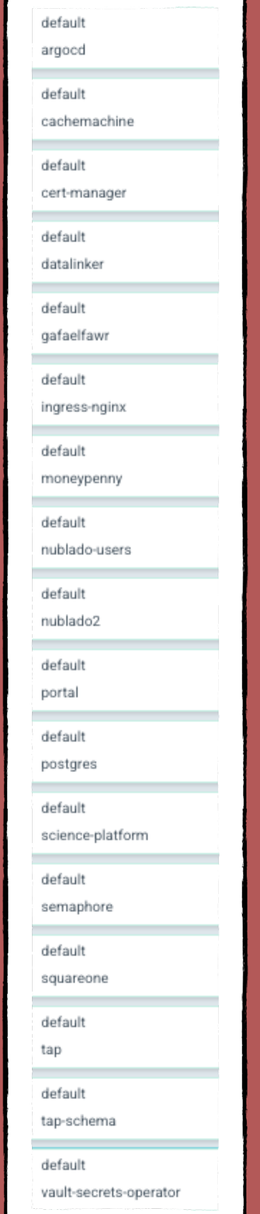

Qserv Current Status

Catalog	Taille (To)	# Lignes (Milliards)
idf-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

			Data [GB]																		
			in unique chunks										in all replicas								
	#chunks		chunks			overlaps			regular				chunks			overlaps			regular			
Database	unique	replicas	data	index	Σ	data	index	Σ	data	index	Σ	Σ	data	index	Σ	data	index	Σ	data	index	Σ	Σ
cosmoDC2_v1_1_4_image	1730	1744	3569.4	69.4	3638.7	41.9	<0.1	41.9	0.0	0.0	0.0	3680.7	3569.4	69.4	3638.7	41.9	<0.1	41.9	0.0	0.0	0.0	3680.7
dp01_dc2_catalogs	1398	1412	915.3	58.9	974.2	114.3	<0.1	114.3	0.0	0.0	0.0	1088.5	915.3	58.9	974.2	114.3	<0.1	114.3	0.0	0.0	0.0	1088.5
dp02_dc2_catalogs	1478	1492	31746.3	2737.4	34483.7	2138.3	<0.1	2138.3	0.0	0.0	0.0	36622.0	31746.3	2737.4	34483.7	2138.3	<0.1	2138.3	0.0	0.0	0.0	36622.0
skysim5000_v1_1_1_parquet	18738	18752	13171.2	261.5	13432.7	157.9	<0.1	158.0	0.0	0.0	0.0	13590.7	13171.2	261.5	13432.7	157.9	<0.1	158.0	0.0	0.0	0.0	13590.7
Total [TB for data]	**23344**	**23400**	**49.4**	**3.1**	**52.5**	**2.5**	**<0.1**	**2.5**	**0.0**	**0.0**	**0.0**	**55.0**	**49.4**	**3.1**	**52.5**	**2.5**	**<0.1**	**2.5**	**0.0**	**0.0**	**0.0**	**55.0**

Gafaelfawr

- Provides authentication and identity management services.
- It also manages the tokens
- OpenID Connect compatible
- @CC-IN2P3 : Keycloak+LDAP

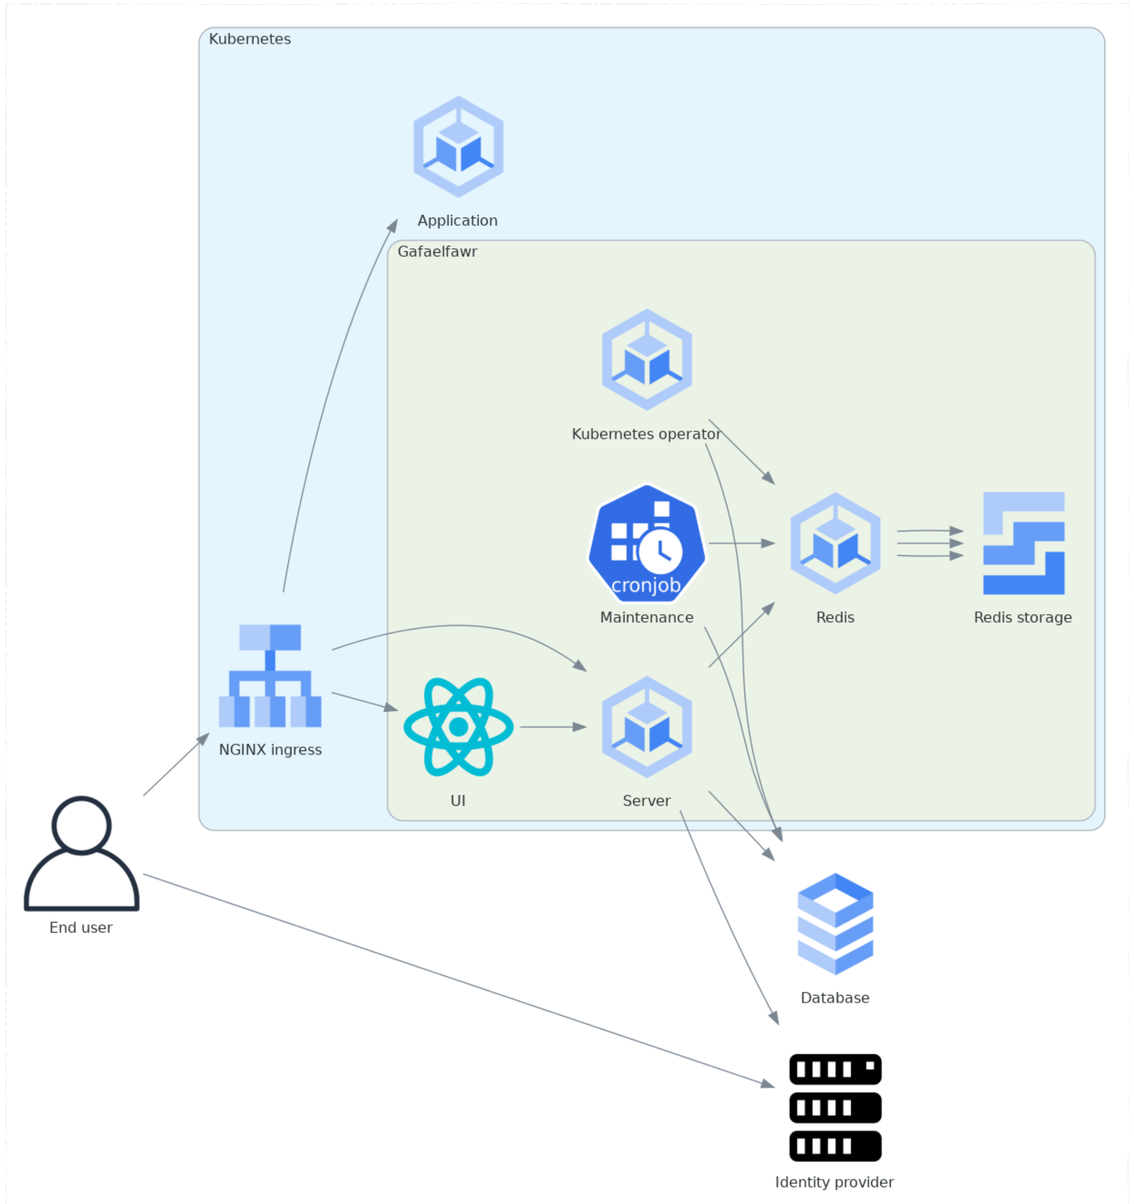

Gafaelfawr is named for Glewlwyd Gafaelfawr, the knight who challenges King Arthur in *Pa gur yv y porthaur?* and, in later stories, is a member of his court and acts as gatekeeper. Gafaelfawr is pronounced (very roughly) gah-VILE-vahwr. (If you speak Welsh and can provide a better pronunciation guide, please open an issue!)

Qserv Ingest

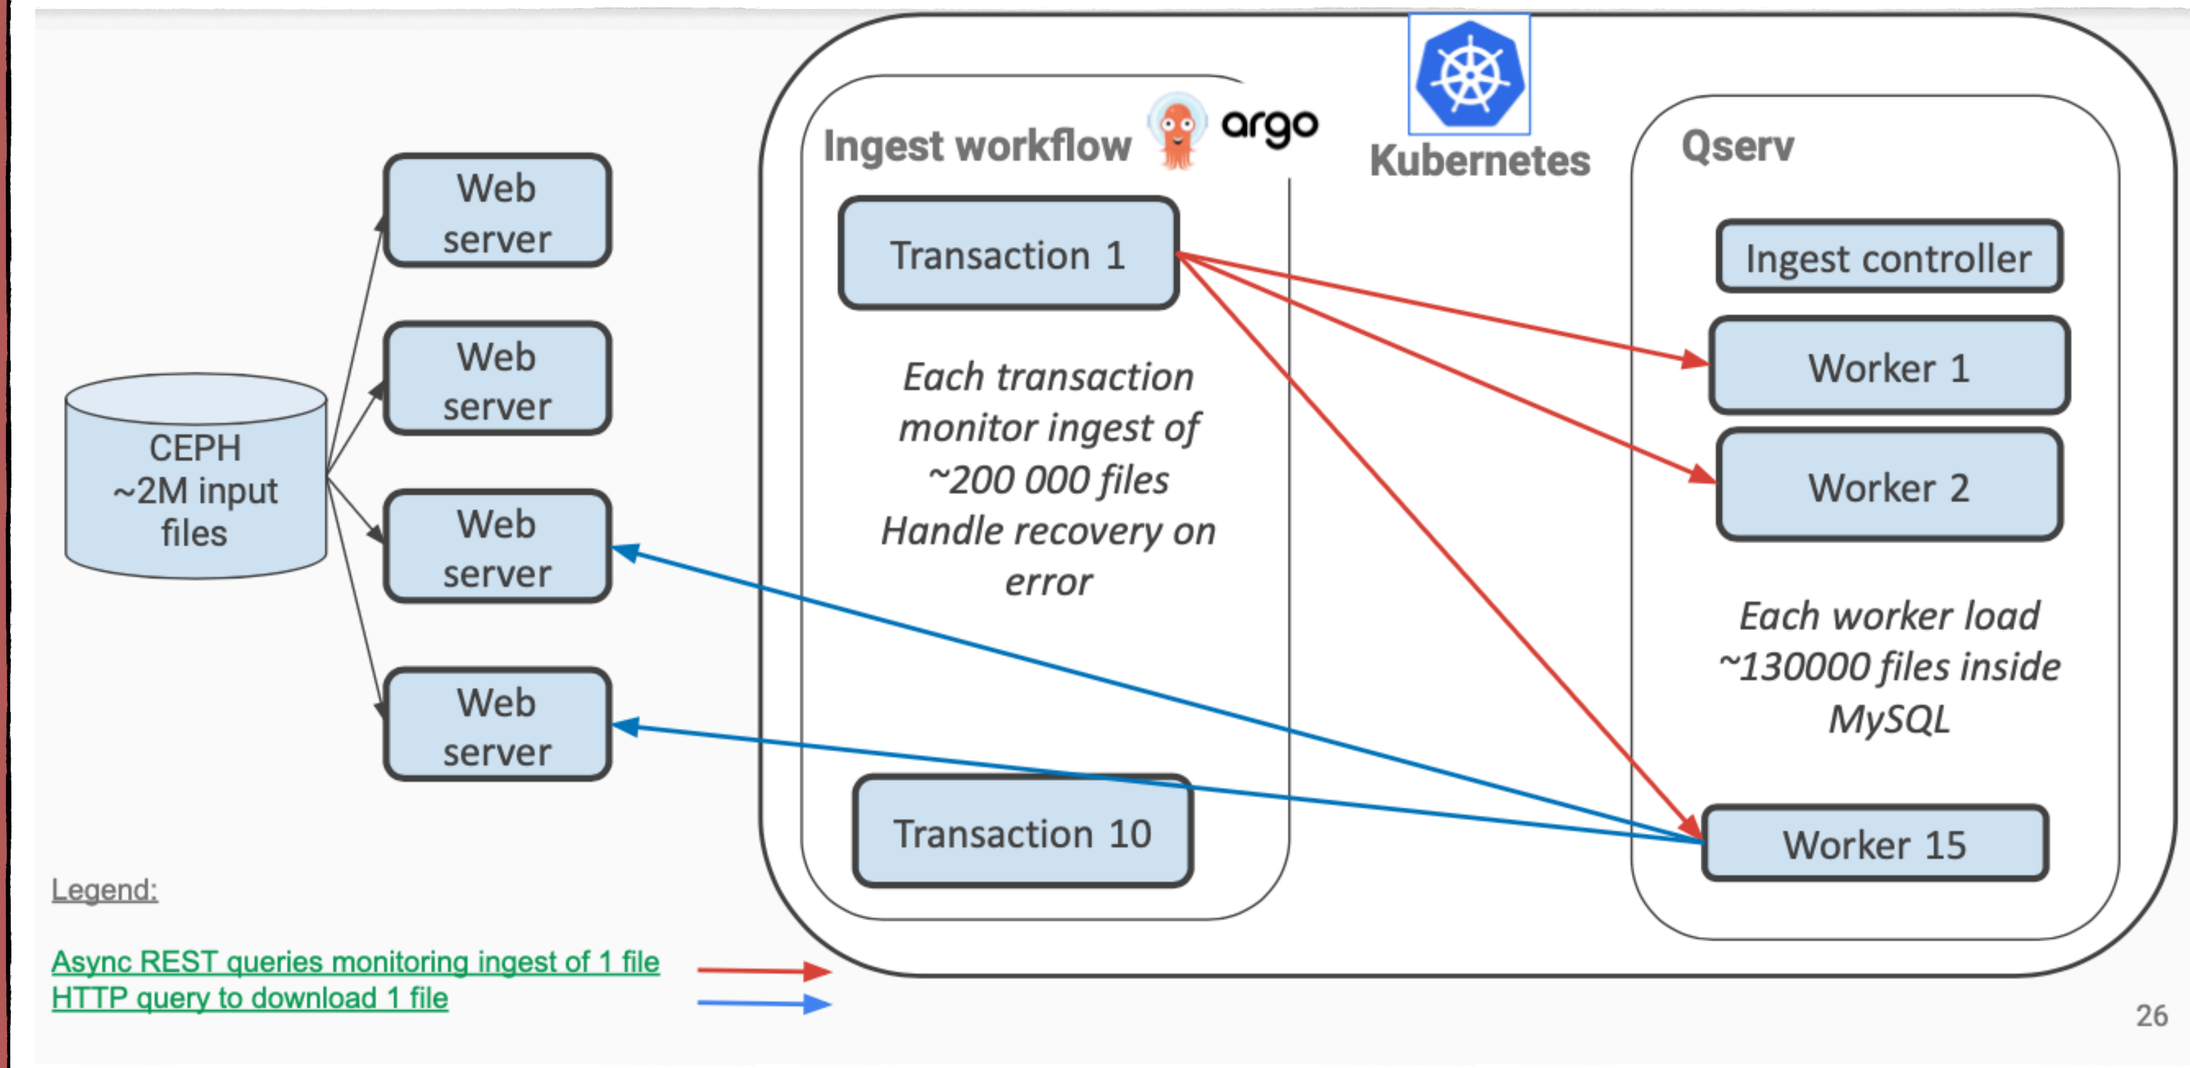

Source: F. Jammes

Rubin FoV

15x

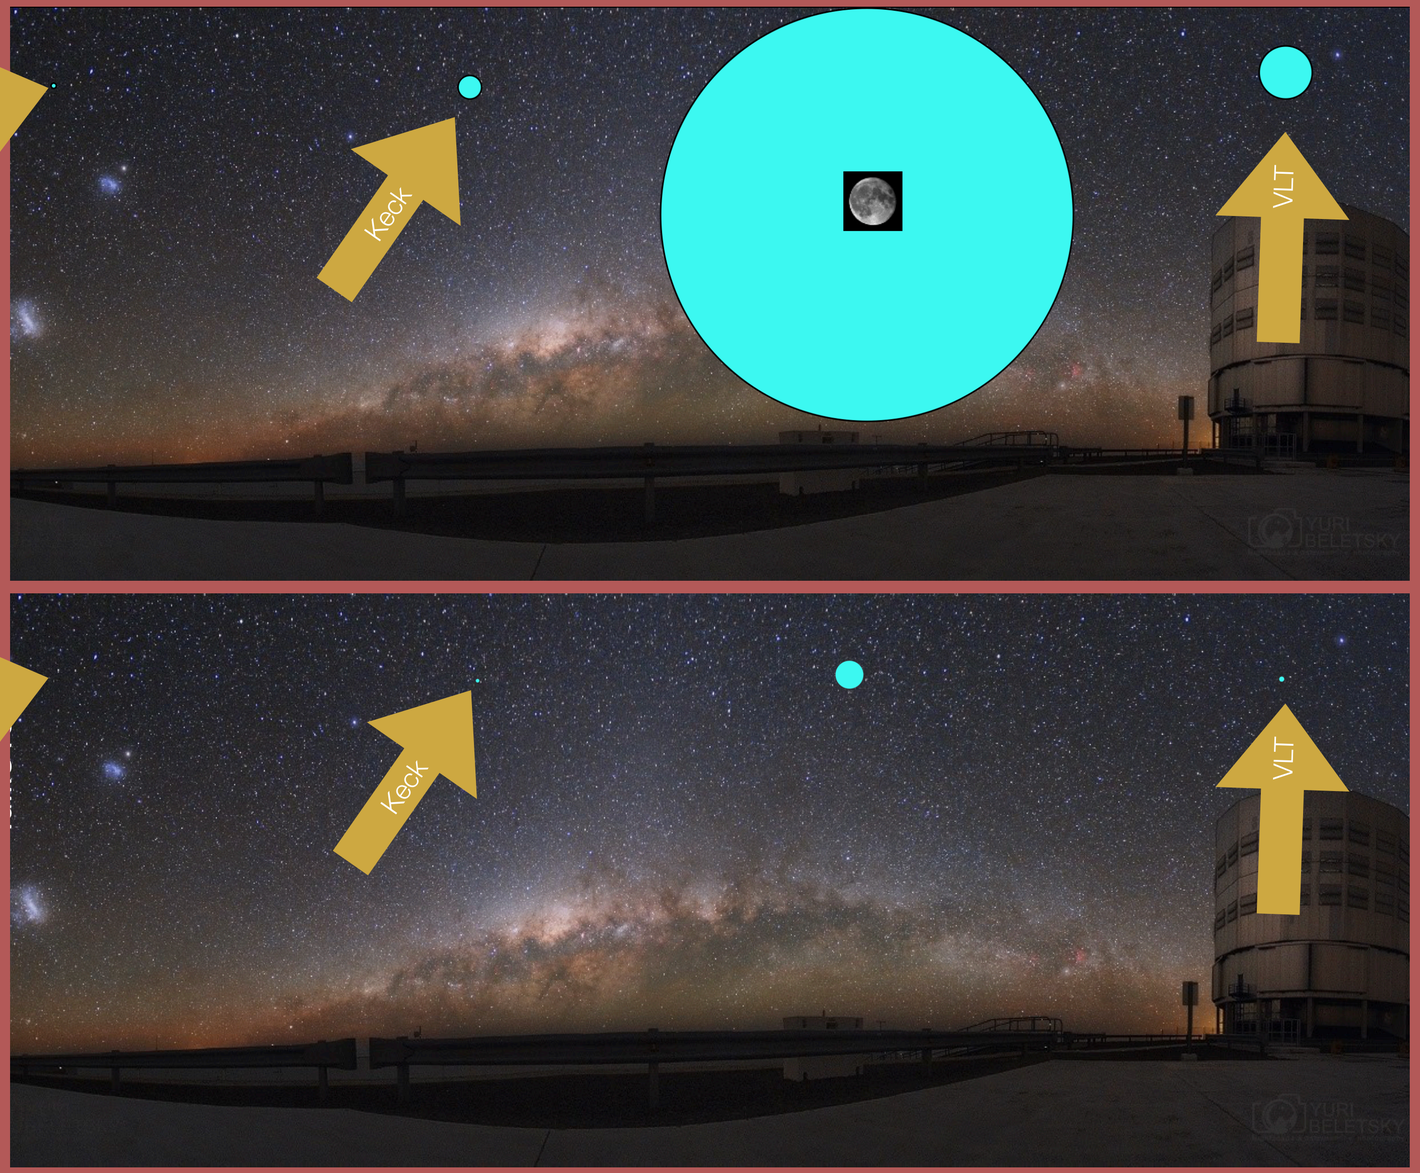

Source: P. Antilogus