

DAC Phase A review and Phase B plans

George Beckett

LSST:UK Project Manager

LSST:UK Phase A (Jul'15—Mar'19)

DAC Objectives

- become fully engaged in the LSST Data Management activities
- support the Phase A LUSC-DEV programme
- develop implementation plan for UK Data Access Centre, for Phase B and beyond

[\[Phase A Case for Support\]](#)

Stakeholders

- LSST:UK Consortium—including
 - Science Working Groups (and representation on Science Collaborations)
 - LSST Affiliate PIs and Junior Associates
 - Funded (Phase A and B) Science Centre activities
- LSST Project and LSST Corporation
- Other international contributors
 - In particular, France (IN2P3 and European DAC working group)
- IRIS (STFC-funded infrastructure for PPAN science)
- Other topical surveys
 - Euclid; GAIA; SKA; 4MOST; ...

Background

LSST Data

- Product categories
 - Prompt Products (formerly Level 1) – nightly alert stream and Prompt Products database
 - Data Release Products (formerly Level 2) – catalogues and processed images
 - User-generated Products (formerly Level 3) – created by community and incorporated into Project-curated data
 - (Raw data – not clear whether this will be available or useful)
- (At time of writing) Prompt Products database and Data Release Products only for data-rights holders (for two years)
 - Nightly alert stream is public via set of Community Brokers (to be selected)

[\[LSE-163: Data Products Definition\]](#)



LSST Data Product Categories



PROJECT

Prompt

Formerly "Level 1" data products

Real Time Difference Image Analysis (DIA)

- A stream of ~10 million time-domain events per night (Alerts), detected, characterized, and transmitted to event distribution networks with 60 seconds of shutter close.
- A catalog of orbits for ~6 million bodies in the Solar System

Data Release

Formerly "Level 2" data products

Reduced single-epoch & deep-to-added images, reprocessed DIA products

- A catalog of ~37 billion objects (20bn galaxies, 17bn stars), ~7 trillion observations ("sources"), and ~30 trillion measurements ("forced sources")
- Produced annually and accessible through online databases.

COMMUNITY

User Generated

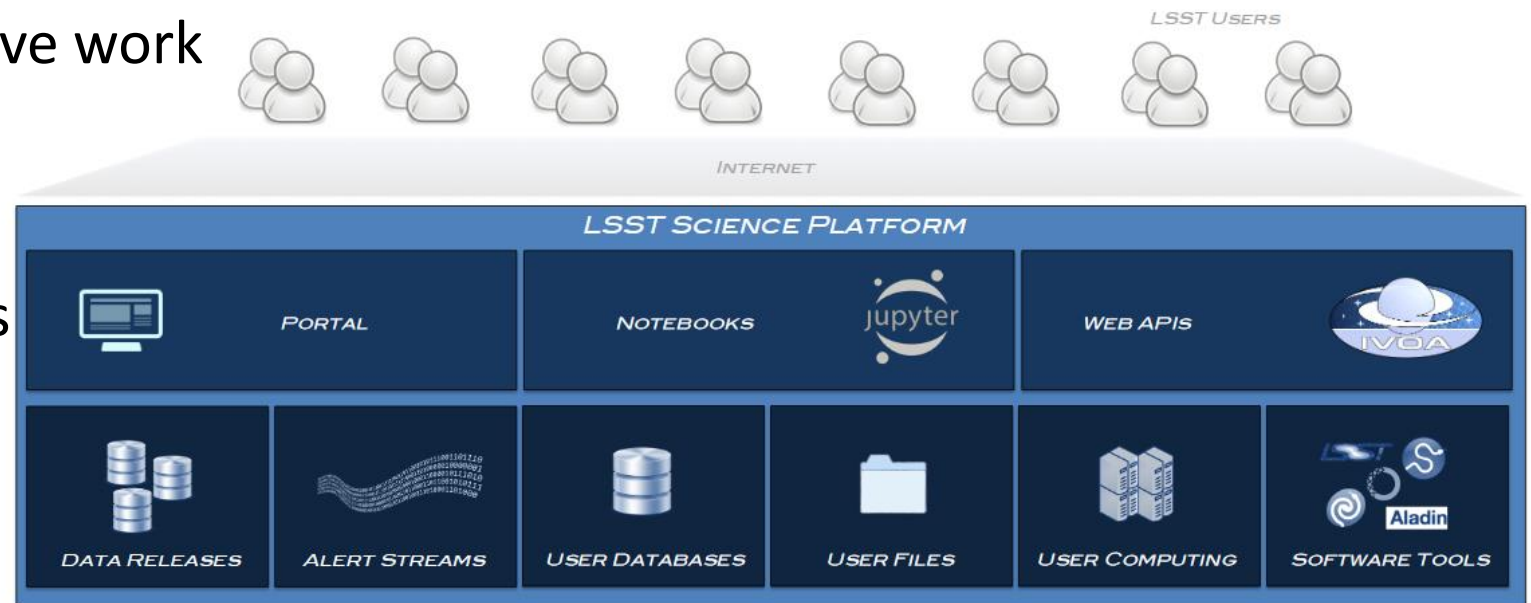
Formerly "Level 3" data products

User-produced added-value data products

- Deep KBO/NEO, variable star classifications, shear maps, etc
- Enabled by services and computing resources at the LSST Data Access Centers (DACs) and via the LSST Science Platform

LSST Science Platform

- Substantial LSST data archive
 - Estimated 500 PB (over 10 years)
 - Apply next-to-data approach to analysis
 - Promote collaborative work
- Three interfaces
 - Web portal
 - (Jupyter) Notebooks
 - Web API



LSST Science Platform (back-end)

- Databases
 - Prompt Product and Data Release Product catalogues (via SQL / ADQL)
 - Plus User databases
 - File services
 - Image Server
 - User file-space
 - Batch computing
 - Route to more substantial computing resources (HPC/ cluster/ grid)
-
- Science Payloads
 - Processing of raw image data into Data Products
 - Bulk distribution
 - Transmission of data to partners such as LSST EPO and LSST:UK

LSST DM Software Stack

- Software status
 - All software will be Open Source
 - DM Roadmap tracks milestones and metrics (available on request)
- Distribution mechanism to be finalised
 - Docker container images
 - CERN VMFS also in use
 - (Source code)
- No software planned for User-generated Products
 - though may benefit from functionality of pipeline software

UK DAC Overview

LSST:UK DAC

- Plan to follow LSST DM team closely
 - LSST DM effort order of magnitude greater than UK DAC team
 - DM Software is Open Source and portable
 - Simplifies integration of UK DAC into international DAC network
- Differences
 - UK aims to run Community Broker
 - UK DAC hosted on shared infrastructure (IRIS)
 - UK DAC will enable synergies with other surveys – GAIA; IR (UKIDSS, Euclid); Spectroscopy; SKA; ...
 - UK DAC needs to ingest data from processing centres in US/ France promptly

Usage categories

Possible DAC user classification	Competence	Level of engagement	Workload sophistication	Infrastructure requirements	Software	Effort to engage	Level of Collaboration	Number of such users
Guest	Assume low	Low--sporadic single-visit sessions lasting few minutes	Low--may be limited to running prescribed queries through web interface	Low--satisfied by resources of web portal or interactive webpage. No persistence.	Query Tool (e.g. ODBC/ Firefly)	None -- expect to be able to engage without a priori effort. However, will need to check has data rights for restricted-access data	Assume none, as no session persistence	O(100s) acting as individuals
Power User	Good	Semi-persistent sessions lasting a few days. Repeat visitor	Interactive scripting -- e.g. Jupyter notebook	Modest -- potentially engaging few cores of compute and utilising persistent storage for few weeks	Scripting env w/ libraries -- Python, NumPy, AstroPy, ...	Lightweight--may be required to sign-up and confirm identity	Modest--may wish to share notebooks, results with small number of individuals	O(100s) acting as individuals or informal collaborations of couple of people
Science Team	Advanced	Potentially run experiments over months/ years	High--bring along own software, potentially require workflow management/ batch system or bespoke access to infrastructure	High --potentially large amount of persistent storage and large allocations of compute time. May produce data to be published to community	Bring along own software to be ported to platform, plus own data	Modest--may be required to apply for time, identify team of users, and assume responsibility for authorising/ training individual members. May be awarded resources over fixed period of time -- e.g. 1PB and 20M core hours for 1 year	Potentially high-- science team will likely wish to share code, data, files. Plus, outputs of project may need to be published/ curated for much wider consumption	~10 teams consisting of people with different roles

UK DAC Activities

Software Survey


- Introduction to DM stack (pipeline and analysis software)
 - and set up information for tutorial-style notebooks
- LSST Simulation software
 - Metrics Analysis Framework for evaluation survey schedules
 - ... and tutorial-style notebooks
- Packaged up as Docker image
 - For use with observing-strategy databases
 - Used for Lasair notebook interface
 - Used by DESC team for PSF image simulations




Lasair—Pilot Broker for LSST Alert Stream

- Operating broker for Zwicky Transient Facility
 - Treat as pre-cursor for LSST Alert Stream
- Ingest nightly alerts into relational database
 - Filter, annotate, and follow-up
- Publish via web interface
 - Includes standard queries, plus cross-match to other surveys
 - Watch lists and individual alerts
- Pilot analysis platform using Jupyter Notebooks
- Plus, early experiments with stream-based and NoSQL alternatives for database

lasair.roe.ac.uk

[LUSC-A-08]




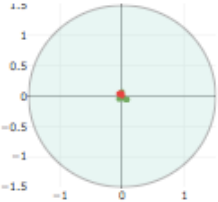




Login | Signup

- [Home](#)
- [Query Objects](#)
- [Cone Search](#)
- [Coverage](#)
- [Watchlists](#)
- [Jupyter](#)
- [Ingestion Status](#)
- [Release Notes](#)
- [Team and Contact](#)

Object ZTF18adbntwo





- Object has 10 candidates, at mean position:
 - (RA, Dec) = (53.575547, 1.082069)
 - (RA, Dec) = (03:34:18.131, 01:04:55.450)
 - (l, b) = (183.875397, -41.746524)
- Classified as SN at distance 8.67 arcsec.
- The transient is possibly associated with 1237666302167613543/2MASXJ0334; a B=16.04 mag galaxy found in the SDSS/NED/GLADE/2MASS catalogues. It's located 7.72 S, 4.86 E (8.1 Kpc) from the galaxy centre. A host z=0.048 implies a transient $M = -17.28$.
- Information on this webpage also [available as JSON](#).
- Conesearch Links (at 5 arcsec): | [Simbad](#) | [NED](#) | [Transient Name Server](#)

Comments


Lasair Bot	Jan. 17, 2019, 3:21 p.m.	In TNS as SN2018les at 0.1 arcsec, discovered 2018-12-30 05:17:04 (MJD 58482.00) by ZTF, ATLAS
------------	--------------------------	--

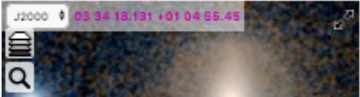
You must be signed in to post comments.

Crossmatches

rank	ID	Catalog	Type	Separation
1	1237666302167613543/2MASXJ0334	SDSS/NED/GLADE/2MASS	galaxy	8.67

AladinLite

Image layer is PanSTARRS DR1; use the layers icon to change it . You can also overlay PanSTARRS and/or Gaia DR2 catalog.



Jupyter Notebook Platform

- Suite of custom Jupyter notebooks
 - set up with astronomy/ LSST software tools and database access
- Deployed on cloud platform (OpenStack/ Kubernetes)
- Workstation-class resources
 - W/ access to catalogues and image repositories
 - Plan for scale-out to HPC via batch system

```
figtitle += ' = '+objectdict[i]['iau_name']
if 'class' in objectdict[i]:
    figtitle += ' ('+objectdict[i]['class']+)'
plt.title(figtitle)


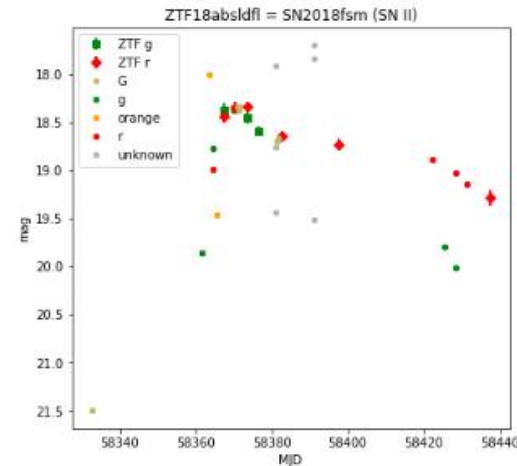
# show cutout
if 'cutout' in objectdict[i]:
    if os.path.exists(objectdict[i]['cutout']):
        plt.subplot(122)
        img = mpimg.imread(objectdict[i]['cutout'])
        plt.imshow(img)

print('https://lasair.roe.ac.uk/object/%s/' % i)
plt.show()

n += 1

print('Done.')
```


<https://lasair.roe.ac.uk/object/ZTF18absldf1/>



Characterisation of dark energy

- Scale-out experiment using grid
- Galaxy classification pilot
 - Shape determination for 30,000 galaxies, over 3 wave bands, from Dark-energy Survey (30 Terabytes)
 - Embarrassingly parallel – suitable for grid (GridPP)
 - (LSST implies 1,000×data to analyse)


[[LUSC-A-01 \(case study\)](#)]



Credit: Joe Zuntz/ Manchester

Classification of Supernovae

- Interest in Type 1a S/N
 - Separating Type Ia from others (Ib, II, ...) not straightforward
- For LSST, classification based on light-curve alone
- Team ran survey simulation
 - testing how schedule and cadence of survey affects ability to classify S/N
 - Use machine learning to enable large-scale S/N classification
- Deployed on IRIS grid resources
 - fast evaluation of new survey strategies



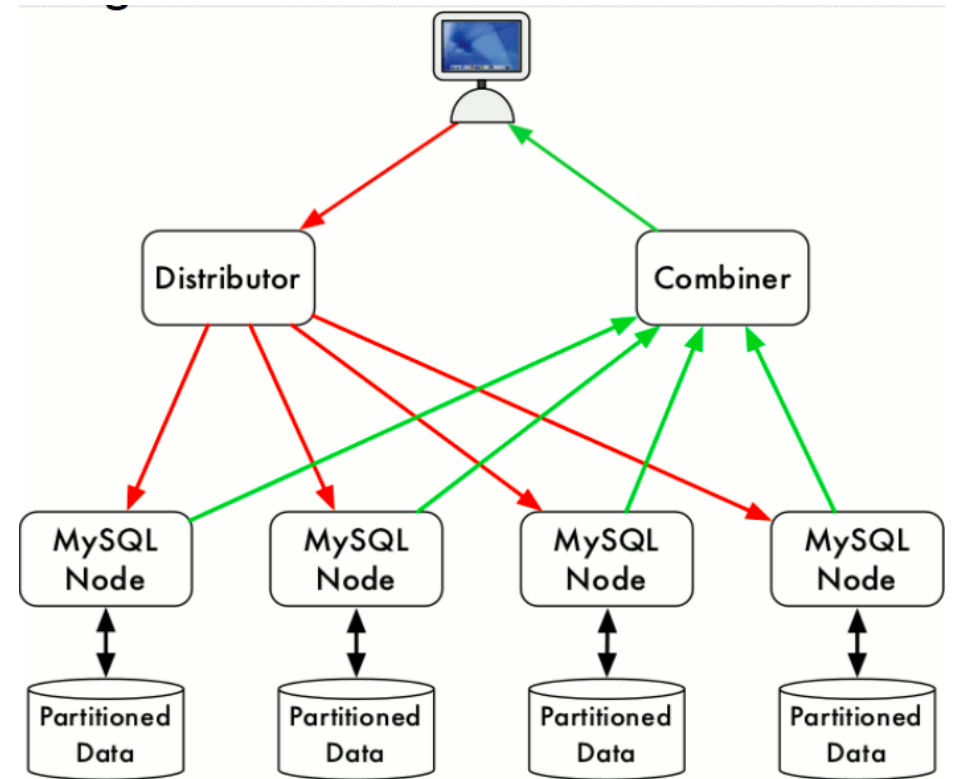
Credit: Natasha Karpenka, Mark Sullivan/ Soton, Michelle Lochner, Hiranya Peiris, Jason McEwen, Ofer Lahav/ UCL

Dark Energy Science Collaboration

- Optimisation of LSST simulator (ImSim)
 - Multicore parallelism/ distributed-memory
- Contribution to Data Challenges
 - Develop/ automate simulation workflow—using HPC and Grid (LHC/ WLCG)
- Undertaking data-transfer experiments w/ IN2P3
 - Ingest Instance Catalogues from NERSC
 - Ingest outputs (catalogues) into UK DAC
 - Understand process for receiving Data Release Products

Evaluation of Qserv

- Candidate database for LSST catalogues
- Maintain deployment of Qserv on UK testbed
 - Track roadmap and functional developments
 - Engaging w/ Community site re. progress and issues
- Ingest non-LSST data (UKIDSS and SDSS)
 - Grow understanding of ingest process (ready for DR)
 - Optimise for UK environment (IRIS)
 - Enable experiments w/ multi-wavelength analysis
- Develop benchmark suite of science-realistic queries
 - Assess Qserv functionality vs. UK science requirements
 - Develop strategy for cross-matching other catalogues



Cloud deployment of Qserv

- Containerised deployment, using cloud-like storage
 - Set up of Qserv databases on Ceph cluster
 - Head node and worker node on OpenStack/ Kubernetes
- Automated deployment and catalogue partitioning
 - Rancher/ Magnum for deployment
 - Scripted chunking and ingest
- Ceph experiments
 - CephFS vs Block Storage
 - Ceph hardware optimisations: SSD-hosted journals; RAID vs. replication; interconnect
- Balance performance vs. ease of administration

International Context

International DAC Network

- Project/ Corp set up working group to evaluate network of DACs (2018Q4)
 - New opportunities for science with LSST data
 - More computing infrastructure for user-generated analysis
 - Better scope to integrate user-generated products into D/R
 - Greater availability and resilience of DAC services
- Introduces technical challenges
 - Distribution of Data Releases
 - Increased work and complexity
 - Risk to enforcement of data-access policy
- Activity stalled due to ODF discussions
 - Hope to produce recommendation report in May

Collaboration and Support

Collaboration Tools

- LSST:UK Confluence site [<https://lsst-uk.atlassian.net/wiki>]
 - Shared repository for all of UK collaboration
 - Access Phase B prop.; Science Requirements; Project Management Plan; etc.
 - Spaces for Science Working Group; DAC and DEV teams; work w/ collab's
 - Confluence also used by Project (though not shared accounts)
 - Contact george.beckett@ed.ac.uk for account
- Mailing lists
 - LUSC-ANNOUNCE – (low-traffic) posts interesting to Consortium
 - Many other lists (SWG, science-collaborations, Exec Group, ...)
 - Hosted on JISCMail—see [guidelines](#) on Confluence

Collaboration Tools, cont'd

- Software repo on GitHub [<https://github.com/lstt-uk>]
 - Hosts various DAC outputs (Qserv benchmarks, Lasair Notebooks, ...)
 - In process of restructuring, ready for wider use in Phase B
 - More info soon
- LSST Virtual Organisation
 - Authorisation group for various LSST (mostly grid) activities
 - Requires user to hold valid certificate (from UK CA)
- Video conferencing
 - Consortium has licence for GotoMeeting (single-session, 10 participants)
 - Speak to George.beckett@ed.ac.uk for access

Pool Travel Fund

- Pool Travel Fund
 - Substantial funds to support UK participation in strategically important LSST meetings/ workshops
 - Available to anyone in LSST:UK (not holding direct travel funds)
 - Information on upcoming meetings on [Confluence](#)
 - Focused calls on LUSC-ANNOUNCE plus responsive-mode mechanism
 - See <https://lsst-uk.atlassian.net/wiki/display/HOME/LSST%3AUK+Pool+Travel+Fund>

Future Plans

Plans for Phase B

- Deploy DAC (technology preview) for early experiments
 - Possible for Commissioning (assuming UK secures role)
- Achieve LSST Community Broker status
 - Develop Lasair as production service for ZTF
- Finalise infrastructure provision on IRIS
 - Shared-access cloud infrastructure
- Mature plans for community-wide User-generated Products
- Contribute to International DAC Network

Thank-you

Extra Slides

Stream-based approach

- Prototype infrastructure for low-latency event stream processing
 - Inc. support for end-user analysis components and workflows
- Distributed architecture designed to scale to data rate and volume
- Trialling “big data” stream processing technologies (SMACK)
 - Spark
 - Mesos
 - Akka
 - Cassandra
 - Kafka