

---

---

# Running astronomy workflows on GridPP

Rohini Jbshi  
University of Manchester

---

---

# Aim: Distribute astronomy workflows on existing infrastructure

- Existing infrastructure?
- Can it be done?
- Do astronomy compute models place restrictions on underlying hardware, or any of the middleware?

# skatelescope.eu VO on GridPP

Jobs run at

- Manchester
- SURFSara
- Imperial College
- Glasgow
- Birmingham
- Bristol
- Cambridge
- Liverpool
- Oxford
- Queen Mary

Storage available at

- Manchester
- SURFSara
- Imperial College
- Cambridge
- Queen Mary

## GridPP DIRAC SAM Tests for skatelescope.eu

This monitoring is still experimental, and doesn't take into account whether sites are properly registered in DIRAC, or are [in downtime](#).

The [GridPP DIRAC Portal](#) is very useful for looking at sites and jobs in detail. For VM-based sites there is a [table](#) in the [VM](#).

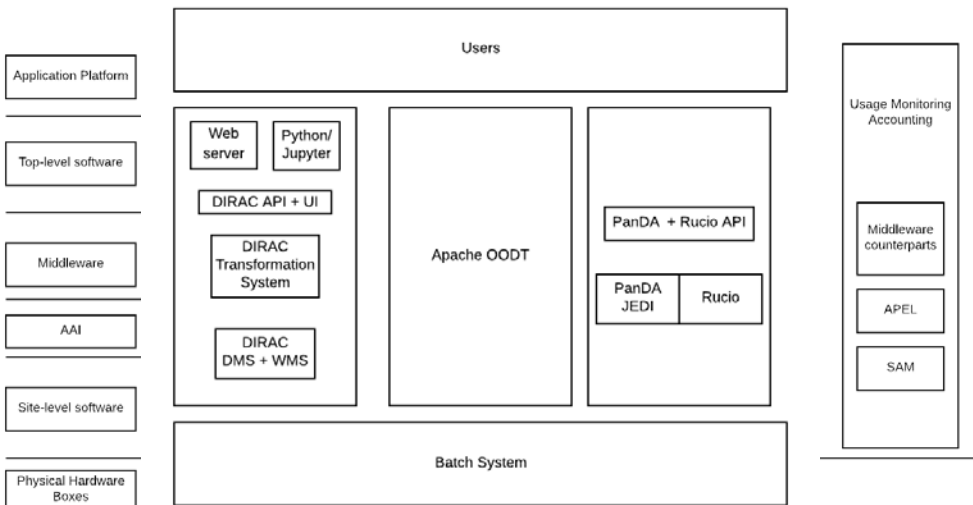
### Vac / Cloud sites

DIRAC site	Last job run	Wait time	Last submitted
CLOUD_CERN-PROD.ch	Sun 17 Jun 2018 11:11:21	60m	4h
***_RAL-LCG2.uk			
***_UK0-LT2-Brunel.uk			
***_UK0-LT2-QMUL.uk			
VAC_UK0-LT2-RHUL.uk			4h
VAC_UK0-LT2-UCL-HEP.uk			4h
CLOUD_UK0-GridPP-Cloud-IC.uk	Sun 1 Jul 2018 4:18:41	8m	4h
***_UK0-NORTHGRID-LANCS-HEP.uk			
VAC_UK0-NORTHGRID-LIN-HEP.uk			4h
VAC_UK0-NORTHGRID-MAN-HEP.uk	Tue 14 Aug 2018 10:54:33	42m	4h
***_UK0-NORTHGRID-SHEF-HEP.uk			
***_UK0-SCOTGRID-DURHAM.uk			
***_UK0-SCOTGRID-ECDF.uk			
VAC_UK0-SCOTGRID-GLASGOW.uk	Tue 14 Aug 2018 10:42:14	30m	4h
VAC_UK0-SOUTHGRID-BHAM-HEP.uk	Tue 14 Aug 2018 10:31:44	19m	4h
***_UK0-SOUTHGRID-BRIS-HEP.uk			
VAC_UK0-SOUTHGRID-CAM-HEP.uk	Tue 14 Aug 2018 10:31:16	19m	4h
***_UK0-SOUTHGRID-RALPP.uk			
***_UK0-SOUTHGRID-SUSX.uk			

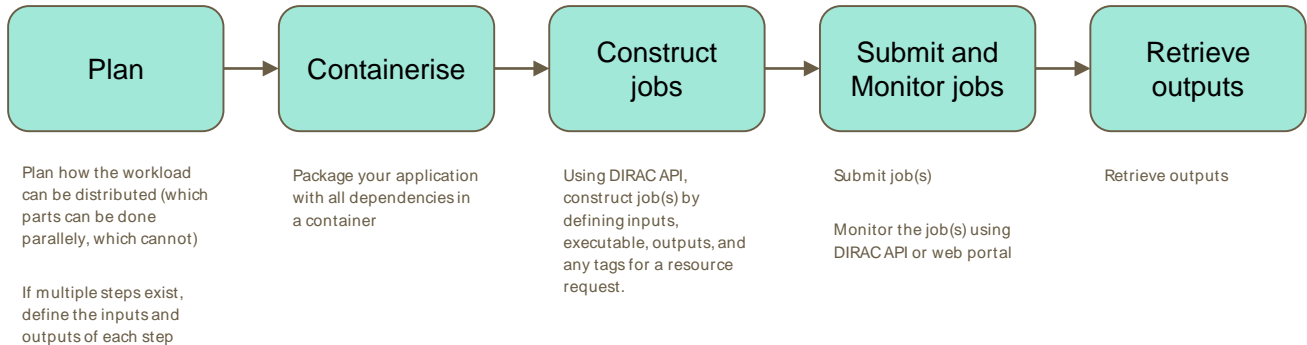
### CREAM / ARC / HTCondorCE sites

DIRAC site	Last job run	Wait time	Last submitted
LOG_RAL-LCG2.uk	Tue 14 Aug 2018 10:16:20	4m	4h
LOG_UK0-LT2-Brunel.uk			4h
LOG_UK0-LT2-IC-HEP.uk	Tue 14 Aug 2018 1:30:59	3h	4h
LOG_UK0-LT2-QMUL.uk	Tue 14 Aug 2018 10:15:27	3m	4h
LOG_UK0-LT2-RHUL.uk			4h
LOG_UK0-NORTHGRID-LANCS-HEP.uk			4h
LOG_UK0-NORTHGRID-LIN-HEP.uk	Tue 14 Aug 2018 10:16:29	4m	4h
LOG_UK0-NORTHGRID-MAN-HEP.uk	Tue 14 Aug 2018 10:15:58	4h	4h
LOG_UK0-NORTHGRID-SHEF-HEP.uk			4h
LOG_UK0-SCOTGRID-DURHAM.uk			4h
LOG_UK0-SCOTGRID-ECDF.uk			4h
LOG_UK0-SCOTGRID-GLASGOW.uk			4h
LOG_UK0-SOUTHGRID-BRIS-HEP.uk			4h
LOG_UK0-SOUTHGRID-CAM-HEP.uk			4h
LOG_UK0-SOUTHGRID-OX-HEP.uk	Tue 14 Aug 2018 10:16:48	4m	4h
LOG_UK0-SOUTHGRID-RALPP.uk			4h
LOG_UK0-SOUTHGRID-SUSX.uk			4h

# Software stack options



# Porting a workflow to the Grid

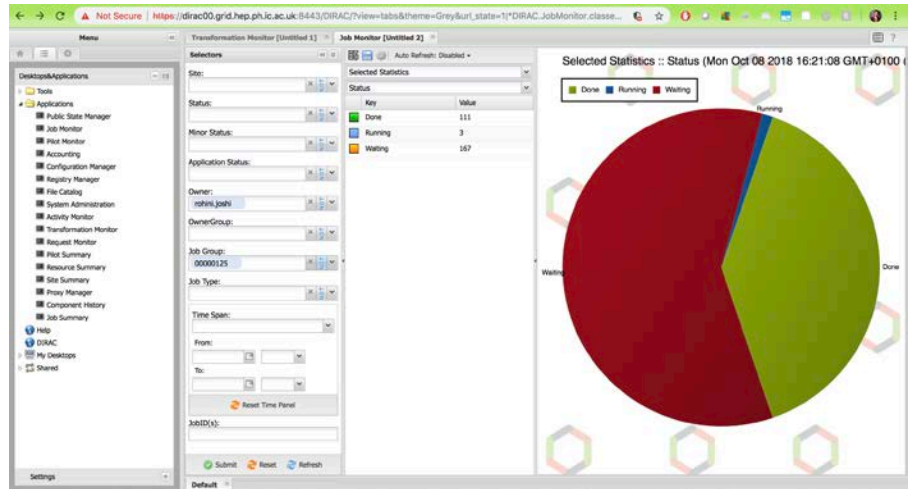


**Prerequisites**

- Learn the middleware (DIRAC) API
- Upload data

# DIRAC

- Workload and Data management System
- Developed at the LHCb project
- Communities using it include CTA, CMS, PNNL, JNR
- multi-VO instance on GridPP
- Python, REST API, CLI and Web portal for monitoring and accounting

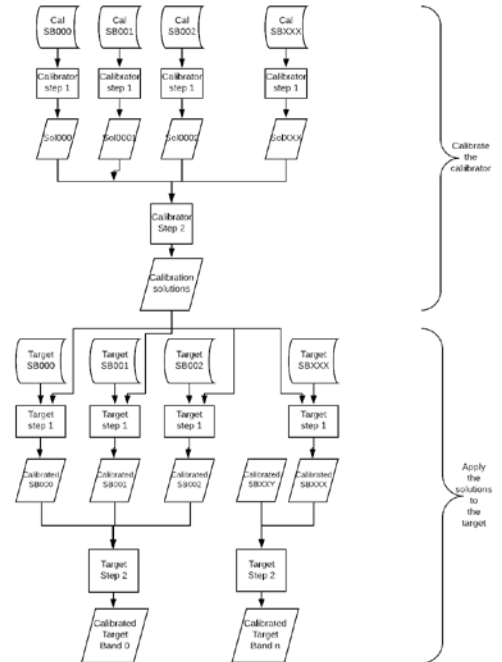


# Direction Independent Calibration of Visibility Data

- LOFAR GOODS-North data
- Prefactor pipeline
- Distribute the workload as described here

[http://liacs.leidenuniv.nl/~plaata1/ISGC2\\_017\\_apmechev\\_final.pdf](http://liacs.leidenuniv.nl/~plaata1/ISGC2_017_apmechev_final.pdf)

- Work on direction dependent calibration ongoing...



# Other use cases run on the Grid

## Pulsar timing

- Scrunch pulsar data in frequency and time
- Calculate ToA (times of arrival) for pulse

## Automated Object Classification with Supervised Learning

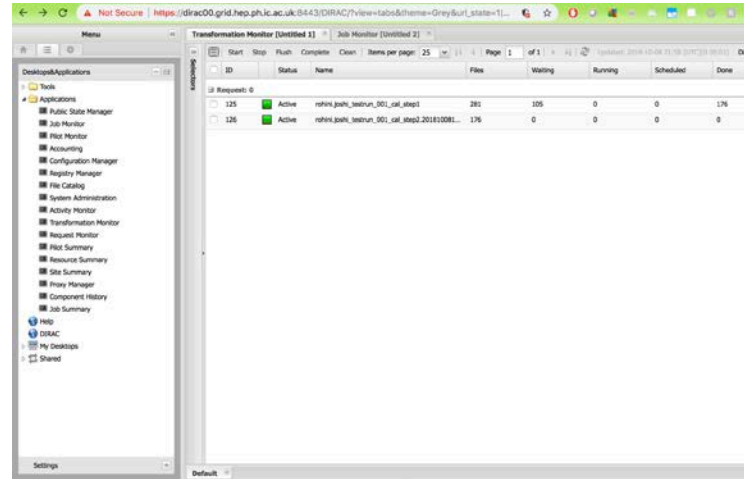
- Crossmatch sources with other surveys – obtain multiwavelength photometry
- Train a Random Forest and implement it to classify new sources



# DIRAC Transformation system

For production workflows of the type do 100 stage 1 jobs, use their output to do 50 stage 2 jobs, then aggregate outputs using one job, then 100 stage 4 jobs etc.

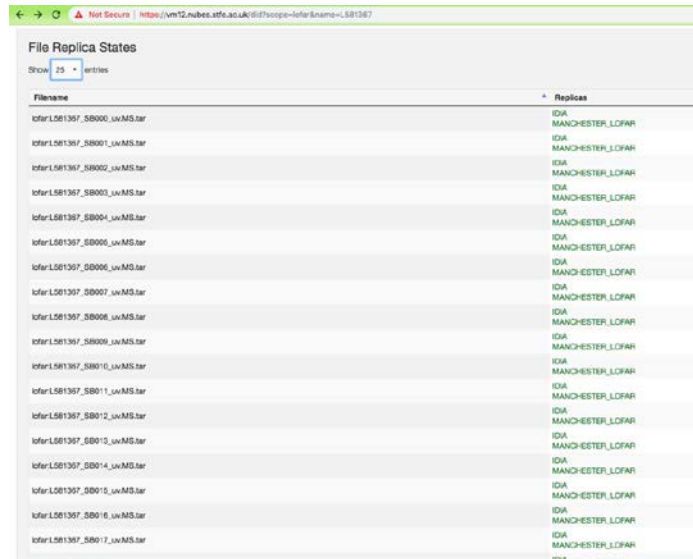
Data driven model that relies on defining metadata to identify files as an input for a set of jobs.



Thanks to Daniela Bauer and Simon Fayer

# Rucio: Scientific (Distributed) data management

- ATLAS experiment
- skatelescope.eu VO instance set up at RAL
- Rucio Storage Elements (RSE) set up at Manchester, IDIA, RAL, Queen Mary, Cambridge
- Shown to handle 400 PB, 800 million files (average replica factor  $\sim 1.3$ )
- Supports S3-style object stores directly (Ceph, OpenStack Swift, Amazon S3, Google Cloud Platform)



The screenshot shows the 'File Replica States' interface in a web browser. The browser address bar shows 'https://vm12.nubex.sfs.ac.uk:8443/voops-efsr&name=L581367'. The interface has a search bar with '25' and 'articles' entered. Below the search bar is a table with two columns: 'Filename' and 'Replica'. The table lists 18 files, each with a corresponding replica location. The replica locations are either 'IDIA' or 'MANCHESTER\_LCFAR'.


Filename	Replica
lcifar.L581367_S8000_1wMS.tar	IDIA
lcifar.L581367_S8000_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8001_1wMS.tar	IDIA
lcifar.L581367_S8001_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8002_1wMS.tar	IDIA
lcifar.L581367_S8002_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8003_1wMS.tar	IDIA
lcifar.L581367_S8003_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8004_1wMS.tar	IDIA
lcifar.L581367_S8004_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8005_1wMS.tar	IDIA
lcifar.L581367_S8005_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8006_1wMS.tar	IDIA
lcifar.L581367_S8006_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8007_1wMS.tar	IDIA
lcifar.L581367_S8007_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8008_1wMS.tar	IDIA
lcifar.L581367_S8008_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8009_1wMS.tar	IDIA
lcifar.L581367_S8009_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8010_1wMS.tar	IDIA
lcifar.L581367_S8010_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8011_1wMS.tar	IDIA
lcifar.L581367_S8011_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8012_1wMS.tar	IDIA
lcifar.L581367_S8012_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8013_1wMS.tar	IDIA
lcifar.L581367_S8013_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8014_1wMS.tar	IDIA
lcifar.L581367_S8014_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8015_1wMS.tar	IDIA
lcifar.L581367_S8015_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8016_1wMS.tar	IDIA
lcifar.L581367_S8016_1wMS.tar	MANCHESTER_LCFAR
lcifar.L581367_S8017_1wMS.tar	IDIA
lcifar.L581367_S8017_1wMS.tar	MANCHESTER_LCFAR

# PanDA: Production ANd Distributed Analysis

- A pilot-based WMS (similar to DIRAC)
- PanDA throughput was at about a million jobs per day, with about 150,000 jobs running concurrently
- Generalization of PanDA to BigPanDA that allows users to use supercomputing resources with PanDA (in addition to traditional grid and cloud resources)
- Initial PanDA tests have been done in the field of human brain studies, biology, and molecular dynamics, with varying kinds of payloads
- Initial astronomy tests successful - more investigation ongoing



WallClock HEPSPROC6 (Sum: 26,472,437,322)



[http://dash-atlas-job.cern.ch/dashboard/request\\_py/resourceutilization\\_individual?sites=All%20Sites&sitesCat=All%20Countries&resourcetype=All&sitesSort=0&sitesCatSort=0&start=2018-01-01&end=today&timeRange=daily&granularity=Weekly&generic=0&sortBy=20&diag1=0&diag2=0&diag3=0&diag4=0&diag5=0&diag6=0&diag7=0&diag8=0&diagT=0&diagPl=0&series=All&type=wchspie](http://dash-atlas-job.cern.ch/dashboard/request_py/resourceutilization_individual?sites=All%20Sites&sitesCat=All%20Countries&resourcetype=All&sitesSort=0&sitesCatSort=0&start=2018-01-01&end=today&timeRange=daily&granularity=Weekly&generic=0&sortBy=20&diag1=0&diag2=0&diag3=0&diag4=0&diag5=0&diag6=0&diag7=0&diag8=0&diagT=0&diagPl=0&series=All&type=wchspie)

# To note

- Grid like and cloud like models will require production managers to liaise with sites, maintain middleware server instances and do the job of translating science talk to compute talk and vice versa (as mentioned in D3.1)
- Middleware i.e. software that will sit between the sites and users will require some amount of customization. As long as we know what we want, this shouldn't be an issue

**Thank you!**