THE UNIVERSITY of EDINBURGH

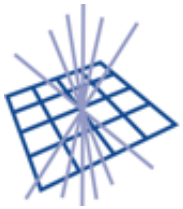

GridPP
UK Computing for Particle Physics

LSST and GridPP at Edinburgh

and

Local Group Computing Support

Marcus Ebert
University of Edinburgh
marcus.ebert@ed.ac.uk

- LSST:UK Science Centre
- GridPP in Edinburgh
- local group computing support

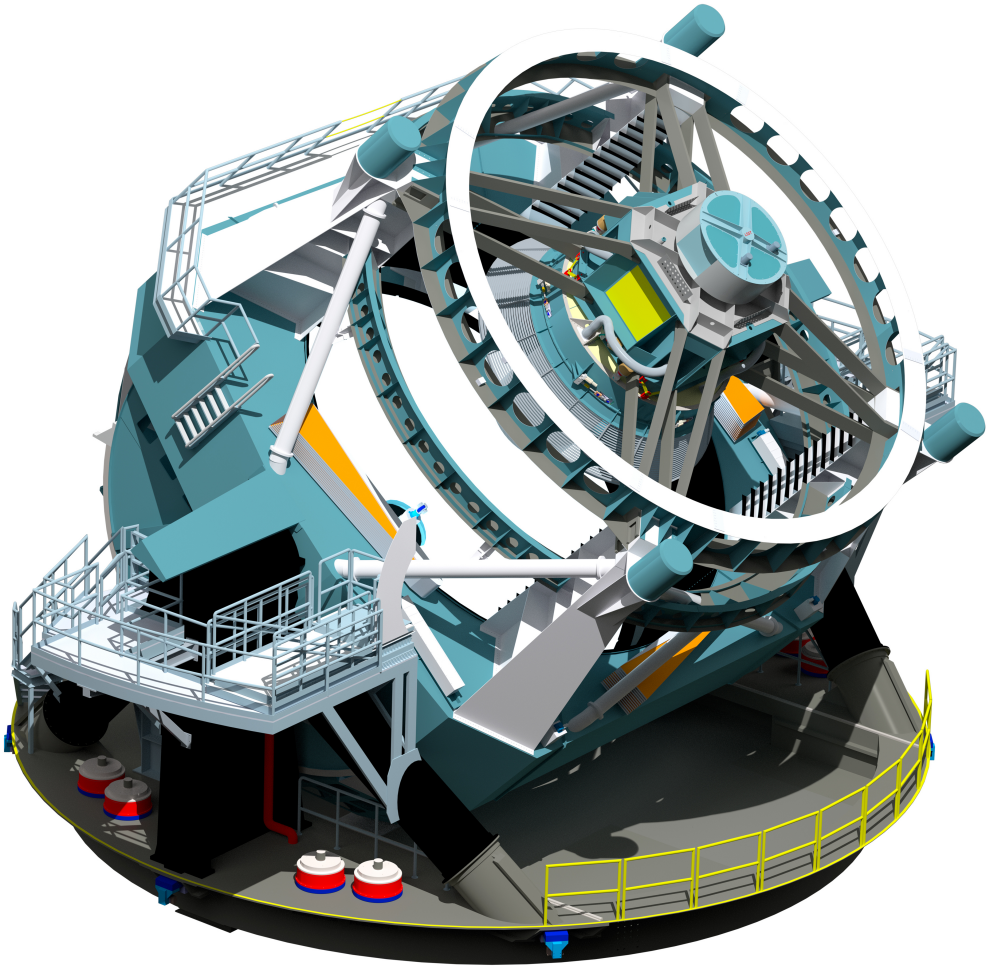

- LSST
- LSST Data Management
- LSST at Edinburgh

Large Synoptic Survey Telescope

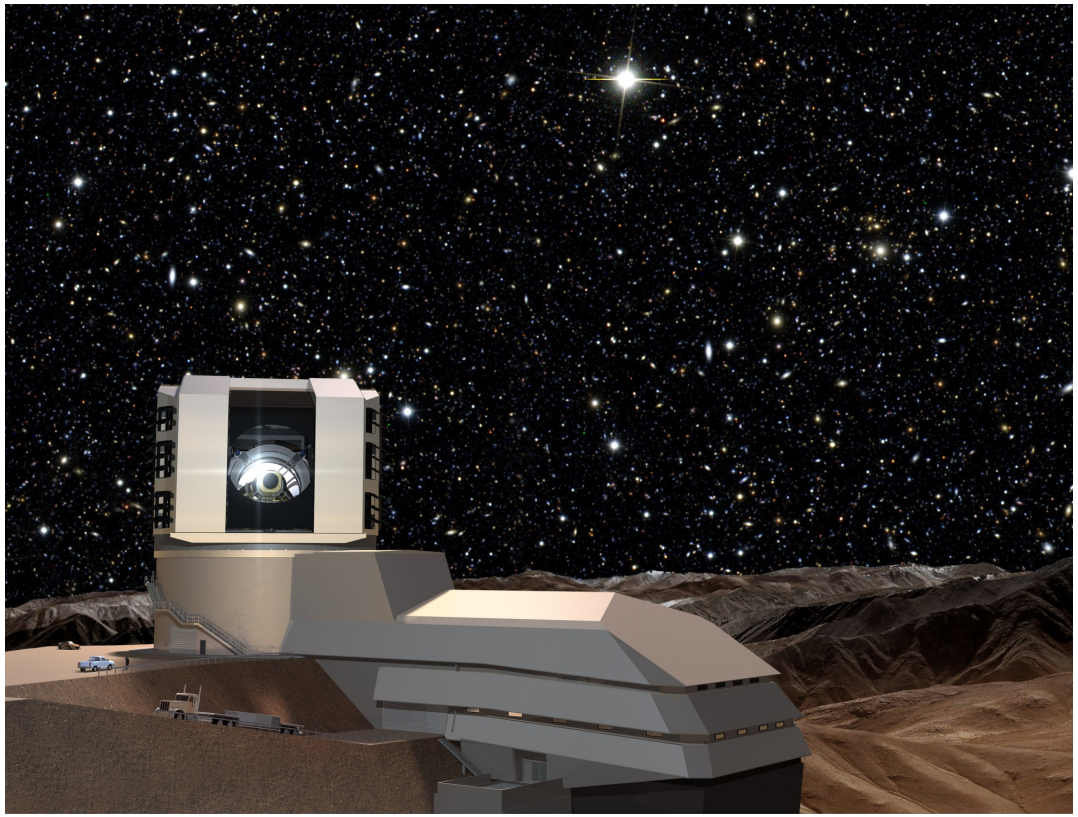

- decade long survey of optical sky
- earth-based telescope, located in Chile
- 3.2 giga pixel camera
- 3 nights to survey the entire sky
- gets built right now
  first images expected in 2021
- web cam on construction site

- cataloguing the Solar System (moving object)
- exploring the changing sky (time-domain, compare against reference in real time)
- Milky Way structure and formation
- Dark Matter and Dark Energy (weak/strong lensing effects, galaxy shape measurements)

LSST Data Management

- about 2 photos/visit, 1000 visits/night, 300 nights/year, 450 calibration images per day
- about 15TB each day
- annual database release with all observed objects
  - about 20×10^9 galaxies, 17×10^9 stars, 30×10^{12} measurements
- alert stream for transient events
  - 10 million transient events/night in alert stream
  - transmitted within 60s of observation
- after calibration and processing:
  - final image collection: 500PB
  - final database size: 15PB
  - uncompressed about 2x larger

QSERV, the LSST Database

- in MySQL format
- 15PB too large for single database server
- distributed, shared-nothing database model
- head node - worker nodes model
- headnode:
  - proxy interface for mysql user queries
  - rewrite query for the different worker nodes
  - send queries to worker nodes and receive query results
  - join all worker node results to final result to be send to user
- worker node: standalone mysql database
- transfer and communication mechanism between head node and workers: xrootd
- in heavy development (pre-beta state)

LSST UK

- 35 Universities participating
- interested in all science areas
- development: image analysis systematics (sensor characterisation)
- data management: Data Access Centre (Edinburgh)

LSST Data Access Centre

- will host database releases
  - need space for at least 2 releases
- receive and process nightly alert stream
- makes data available to UK astronomers
  - login machines
  - public database interface
  - maybe behind UK
- support of user analyses
  - compute resources
  - data storage

LSST Data Access Centre

- testbed with 4 machines available now
  - each with: 32 vcores, 128GB RAM, 12x4TB disks

- get experience with qserv
  - test installation in single and multinode installations using VMs
  - testing ingest of real data base (4TB UKIDSS release)
  - performance tests for qserv, mysql, mssql based on same UKIDSS db

- work on specification and handling of alert streams
  - simulate different amount of data per broadcast and latency
  - endpoints at different global locations

- user analysis support
  - possibilities for database ingests of user data
  - testing possibilities to use GridPP

LSST and GridPP

- galaxy shape measurement
  - first usage of Grid
  - based on DES data which is similar to what LSST will see
  - analysis by Joe Zuntz (Manchester, now Edinburgh)
  - good candidate since galaxies are independent of each other
  - job submission/handling based on Dirac and Ganga
  - successfully finished porting and running earlier this year

- preparing for new project
  - development of light curve measurement algorithms (flux over time, by Mark Sullivan and Rob Firth at Soton)
  - useful for supernovae characterisation
  - simulation of supernovae, populate sky images, reconstruct (MC)
  - will start early next year to submit jobs to GridPP

- people from the DES Collaboration in US want to use GridPP infrastructure based on our work

- local installation
- VO support
- local group support

GridPP installation at Edinburgh

- 11 storage server, 396 disks (2TB, 4TB, 8TB)
  - 1.1PB usable space using ZFS
  - on-disk compression adds extra space (4% for Atlas, 3% LHCb data, 18% for LSST)
  - 214TB additional RDF space to be added soon
- 2 new servers in high availability configuration to host VMs for middleware services
- on Eddie3 352 CPUs reserved for GridPP compute jobs, more through opportunistic usage of free slots on whole cluster
- 100 additional CPUs in 50VMs available on OpenStack installation
- development of distributed storage network (=GridPP storage) which makes access easier for non-LHC groups
  - similar like CMS' AAA or Atlas' FAX
  - but single network for all GridPP supported VOs
- local support contact: Andy, Rob, and me

group support

- supported since long time: Atlas, LHCb, CMS, GridPP
- newly supported: LSST, LZ, HyperK
- possible more VOs supported in near future (e.g.Dune)
- HyperK/Titus:
  - data files transferred to (local) Grid storage from datastore (about 10TB)
  - usage on worker nodes through gfal-fs (POSIX mount similar to NFS but through fuse, also possible on desktop machines)
  - can explore future support by job submission through a grid interface

Local Group Computing Support

Thanks to Andy for the slides!

Prominent Issues

Procurement

- Source of frustration for many
- Long lead time to buy even standard computing equipment
  - Many reasons, most notably supplier issues
- Recent discussions with computing support on how to streamline the purchasing process
- New school purchasing system under being considered - includes a self service portal
  - Aim to minimise number of manual / paper steps

PPE Datastore Area

- Inconsistent file permissions model has led to recurring problems for users
- New simplified storage space being created
- Aim for approximately 20TB to match our pooled free allocation
  - **Note**: large PPE datasets can (and should) be stored on our 1 PB of Grid storage (e.g Titus)
- Clean up and migration of share scheduled for early next year
- Are there any significant non-Grid data storage requirements in 2017?

Development Areas

School Computing Cluster

- Collaboration with computing support to extend the school cluster computing capacity
- We can also buy priority access to the Eddie 3 cluster for short term large capacity requirements

JCMB Public Displays

- Opportunity to showcase PPE material on public displays
  - Can be anything from live event displays, slides, seminar notifications
  - Will schedule material onto display presentation systems on 4th floor

Video Conferencing

- Computing support may have the capacity to upgrade video conferencing facilities in 5202
  - What is top of our wishlist? New projector? Improvements in call quality?

Cloud Printing Service

- New service trial to send files to any printer in the JCMB, Appleton Tower, Main Library etc.
  - Not limited to the 5301 workhorses

School Gitlab Area

- Used for GridPP and PPE computing project management for over a year now
  - Any other use cases within the group? (CERN offering may suffice)

Support Helpdesk

For help with any general computing issues:

- Computing Support Helpdesk Room **4210**
- Email: **sopa-helpdesk@ed.ac.uk**
- Web: https://www.wiki.ed.ac.uk/display/PandAIntranet/Computing+Services
- **NEW**: Trial of school computing support rota in a Slack/Mattermost channel

For PPE-specific issues and help with procurement:

- Room **5412**
- Email: **awashbro@ph.ed.ac.uk**
- **NEW**: Will have **ppe-computing@ed.ac.uk** email address to increase coverage

Thank you!

Any Questions?