



LSST Simulations on OSG

Overview

- The Open Science Grid (OSG)
- The Large Synoptic Survey Telescope (LSST)
- OSG User Support and LSST
- LSST Image Simulations on OSG
- Operational Experience

Feb 23, 2011

Gabriele Garzoglio for the OSG Task Force on LSST
Computing Division, Fermilab

Acknowledgments

- Bo Xin (Purdue)
 - LSST simulations on OSG and app integration
- Parag Mhashilkar (Fermilab)
 - LSST app integration in OSG
- John Peterson (Purdue)
 - Resp. for LSST production simulation
- Brian Bockelman, Derek Weitzel (UNL)
 - Job Management Infrastructure Administrators
- Ian Shipsey (Purdue)
 - LSST Management / LSST slides
- Ruth Pordes (Fermilab)
 - OSG Management / OSG slides

The Components of the Open Science Grid (OSG)

Consortium

Project

Collaboration



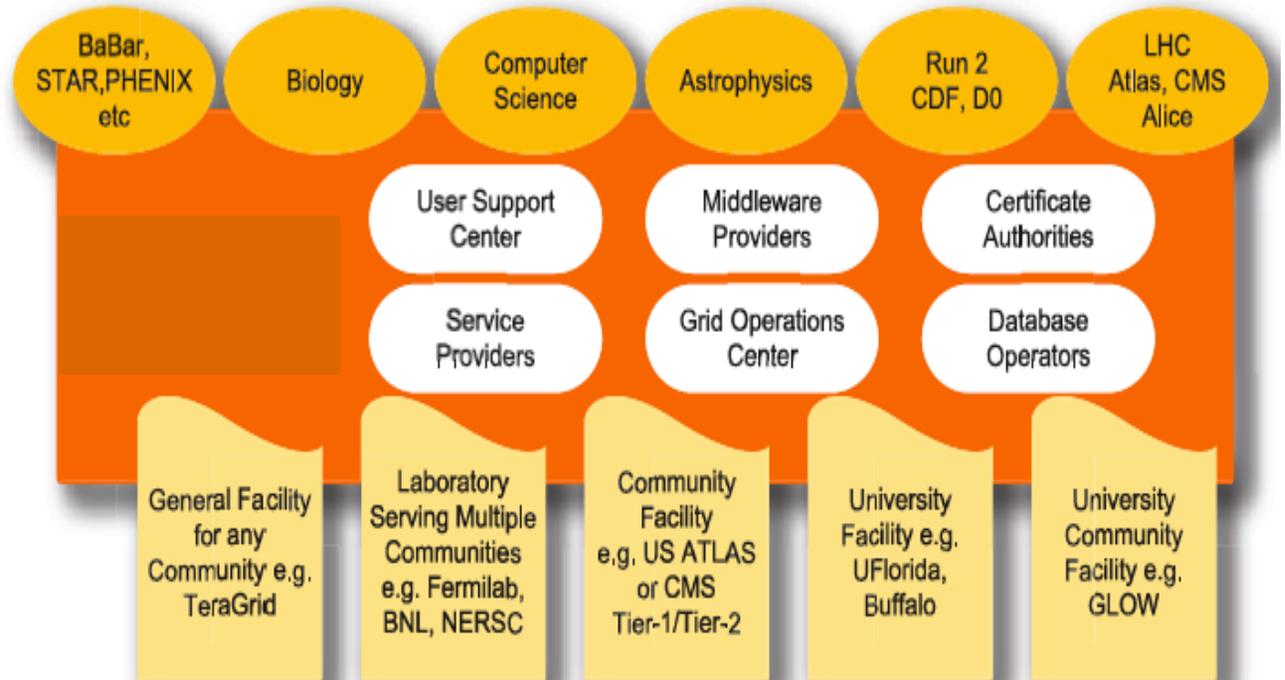
Mission: The Open Science Grid aims to promote discovery and collaboration in data-intensive research by providing a computing facility and services that integrate distributed, reliable and shared resources to support computation at all scales.

The Architecture of the Enterprise

Virtual Organizations

Common Services & Software

Resources, Sites, & Campuses



Resources, Sites & Campuses

Every site has OSG proponent and collaborators in place.

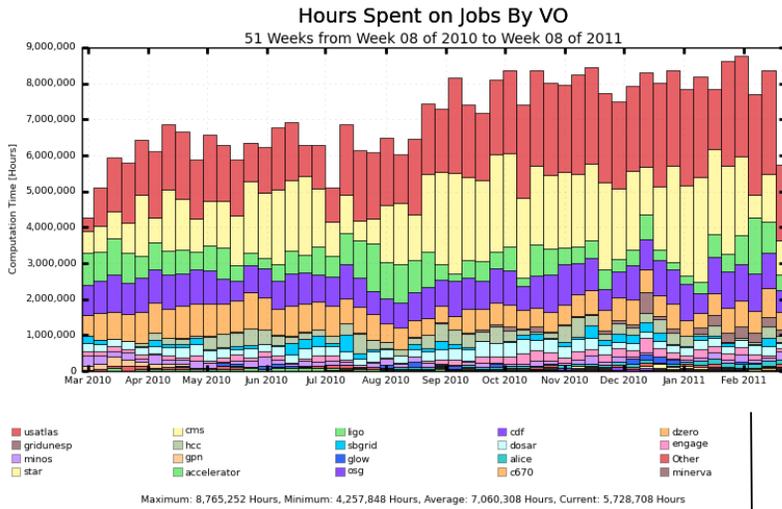
All sites support owner VO + OSG-mgmt VOs + at least one other VO.



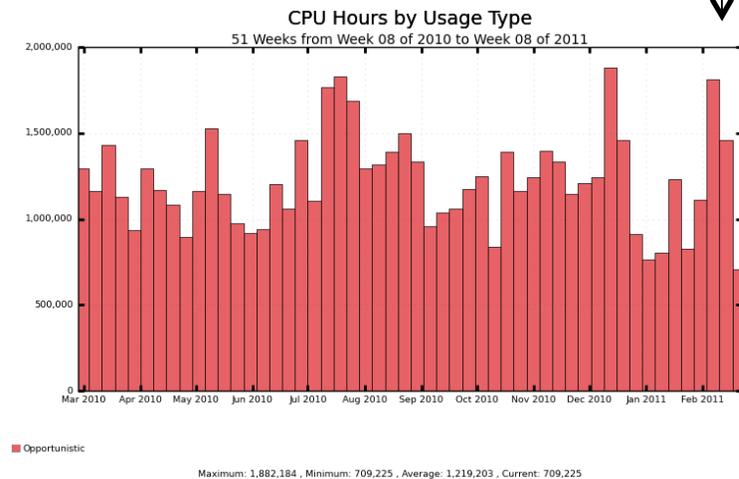
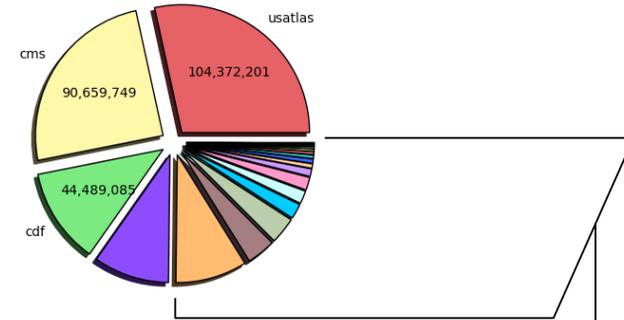
113 interfaces to clusters (small to large),
 55 interfaces to data caches,
 8 campus “grids”

<http://tinyurl.com/3x4gujs>

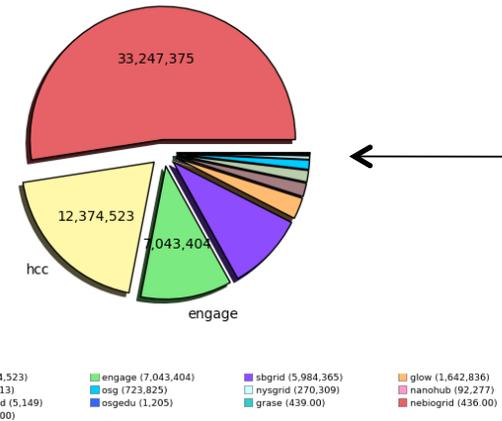
Mainly Particle Physics, Many non-Physics, Several multi-disciplinary Regional Communities



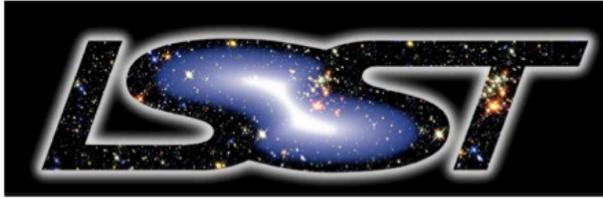
Wall Hours by VO (Sum: 367,136,004 Hours)
51 Weeks from Week 08 of 2010 to Week 08 of 2011



Wall Hours by VO (Sum: 63,397,869 Hours)
51 Weeks from Week 08 of 2010 to Week 08 of 2011



FEB 23, 2011



**8 meter wide-field
ground-based telescope
providing time-lapse
digital imaging of
faint astronomical
objects across the
entire visible sky
every few nights
for 10 years.**

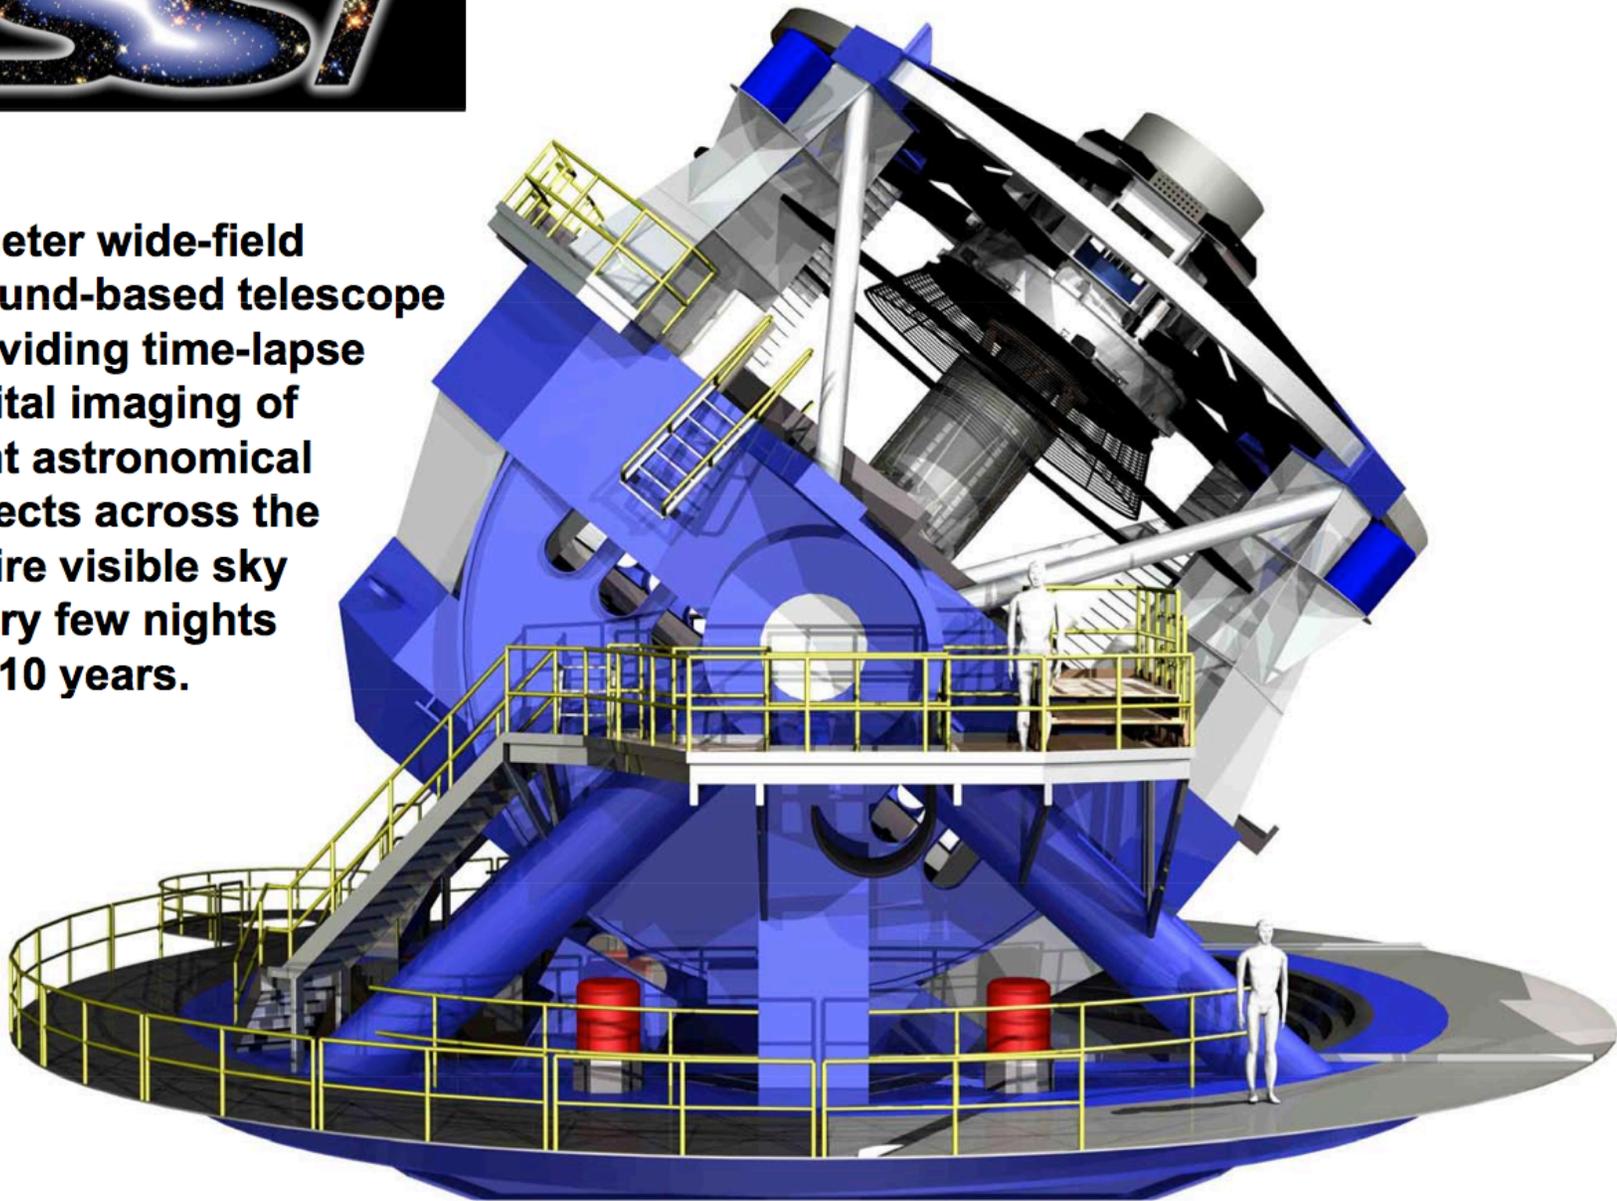


Image sizes LSST, Moon, HST

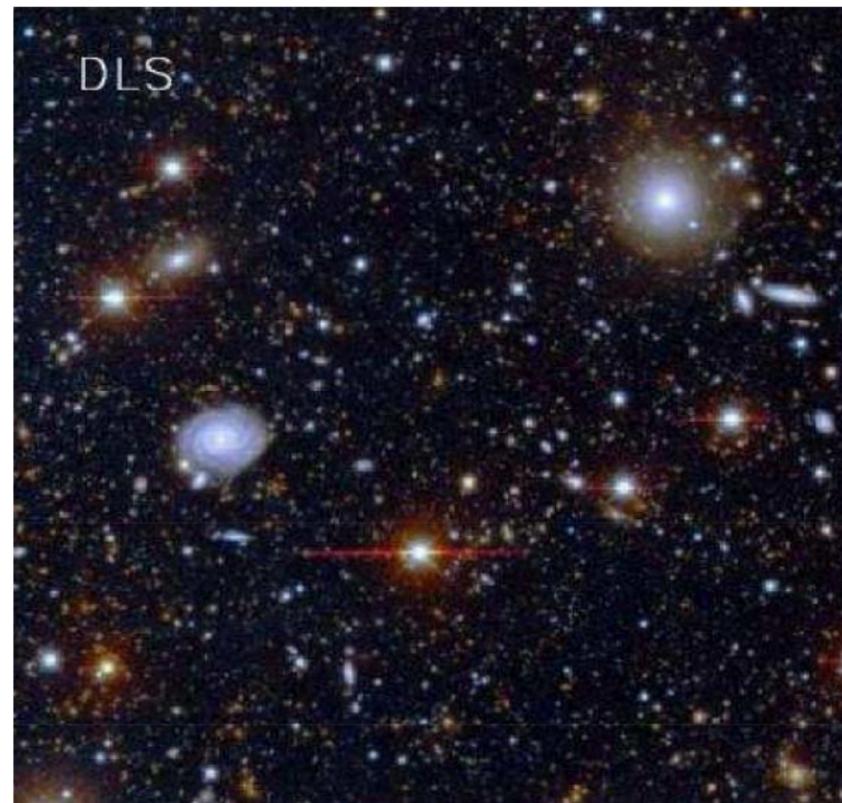


J. Shipsey OSG April 2010

Optical Quality at the LSST site

These two images are of the same patch of sky

LSST Chile , 0.67 arcsec seeing



SDSS Apache Point NM, 1.3 arc sec seeing

x2 better x5 fainter per image
(1,000 images at each sky location
will be obtained over 10 years)

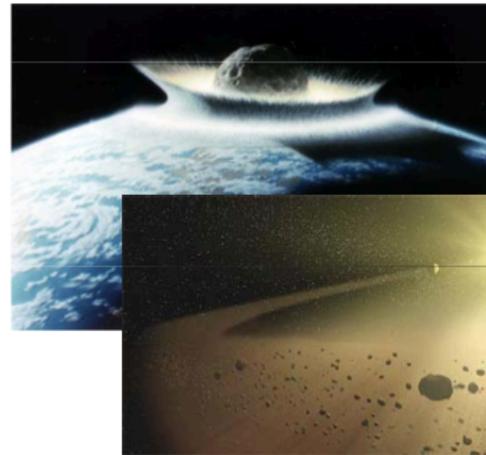
LSST 4 Science Missions

Dark Energy-Dark Matter



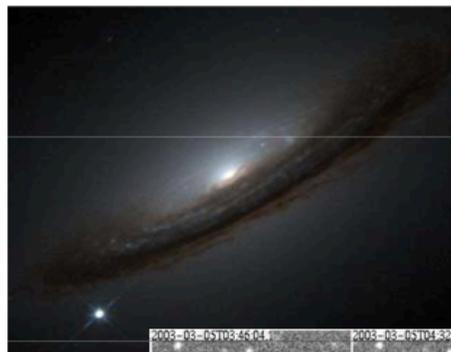
Multiple investigations into the nature of the dominant components of the universe

Inventory of the Solar System

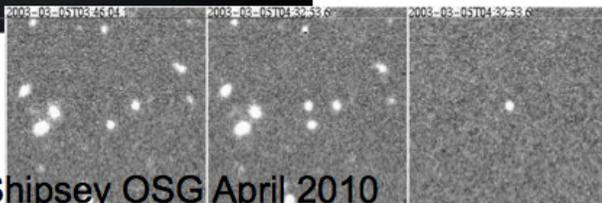


Find 90% of hazardous NEOs down to 140 m over 10 yrs & test theories of solar system formation

“Movie” of the Universe: time domain

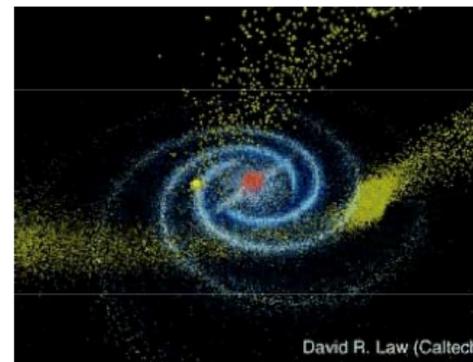


Discovering the transient & unknown on time scales days to years



I. Shipsey OSG April 2010

Mapping the Milky Way

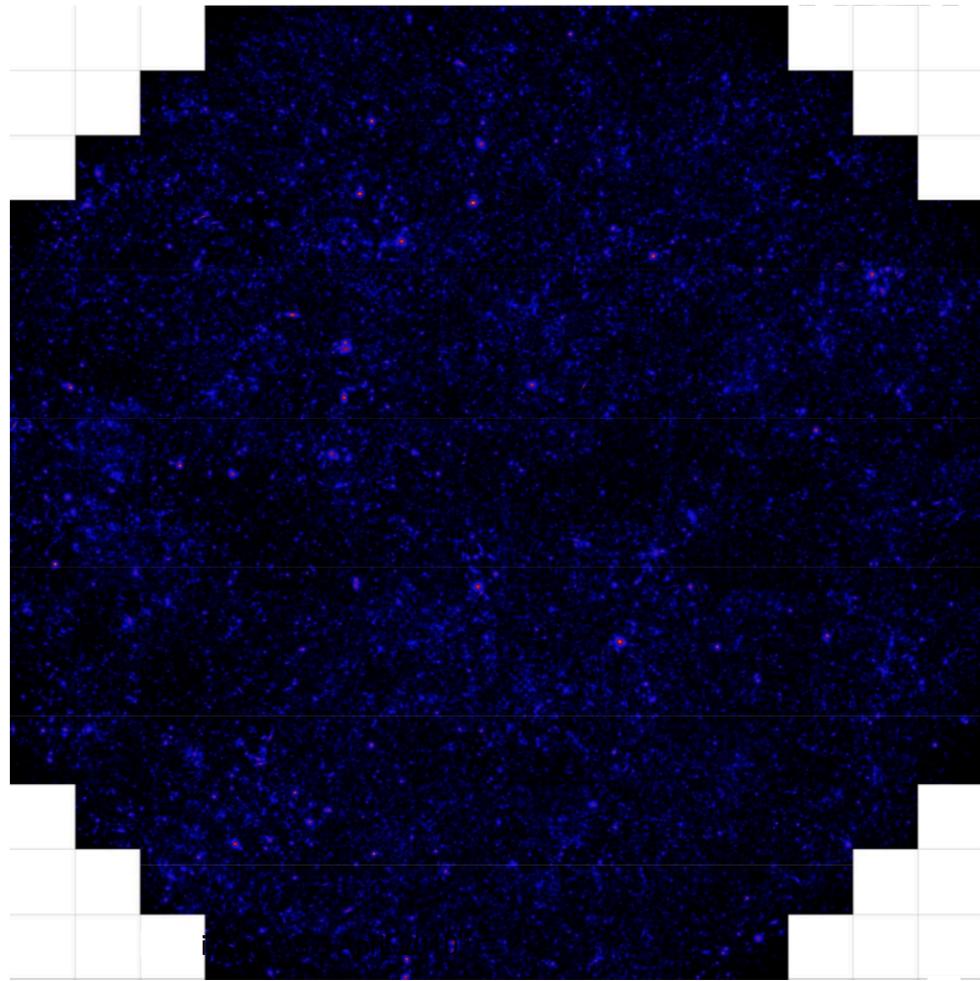


Map the rich and complex structure of the galaxy in unprecedented detail and extent

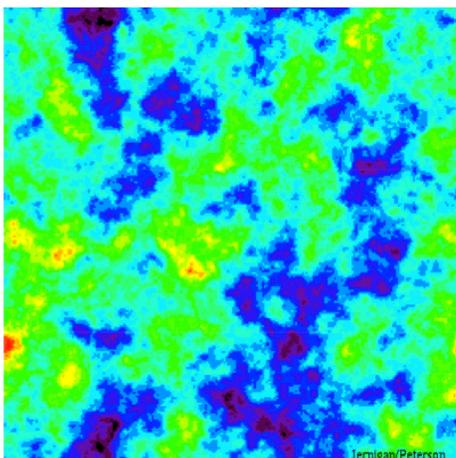
All missions conducted in parallel₉

Full LSST end-to-end photon Simulation

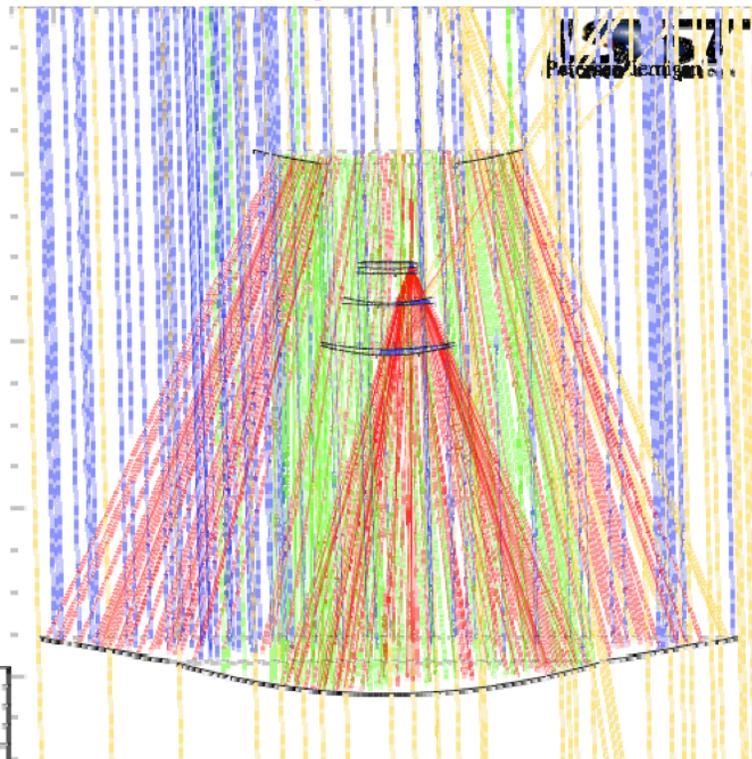
- Collaborators: Berkeley, Purdue, U. Washington, SLAC
- Simulation:
 - Cosmological Models
 - Galaxy Spatial Models & Spectra
 - Atmosphere
 - Optics
 - Detector
- In 1 simulated image...
 - 189 chip images
 - 3 billion pixels
 - 12 million objects
 - Billions of ray-traced objects



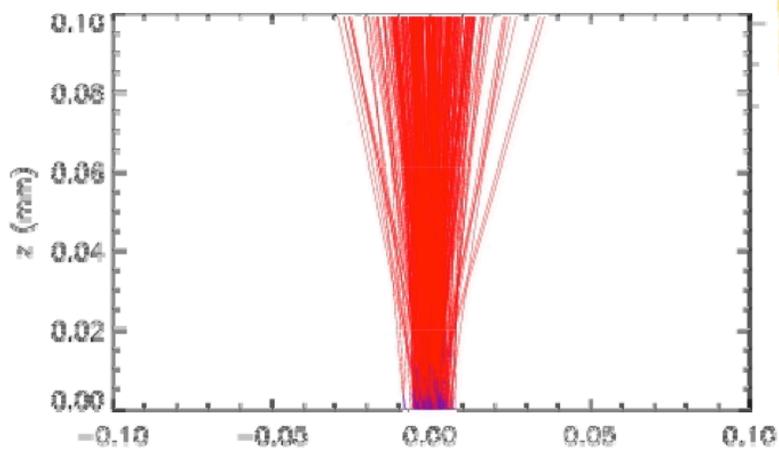
Simulation Stages



Atmospheric Turbulence



Optics Raytrace



Detector Simulation

I. Shipsey OSG April 2010

OSG User Support for LSST

- OSG and LSST are working together to enable LSST applications to run on OSG. Started focus on simulations.
- OSG bootstraps communities with a “phased approach” based on well-scoped focused task forces

Phase	Preparation	Jobs	Data	Job Management Effort	Docs	Support	VO Services Responsib.
Proof-of-Principle	1 mon.	O(100)	O(GB)	OSG staff for 1 day	Basic	Dedicated	Lead by OSG w/ VO contrib. (job sub., user reg., monitoring, ...)
Production Demo	A few mon.	O(100,000)	O(TB)	VO staff for 1 month	Production-like	Dedicated	Lead by OSG w/ more VO contrib.
Production	Several mon.	Production goals	Prod. goals	VO staff for production goals	Production quality	Towards Std Methods	Lead by VO w/ support from OSG

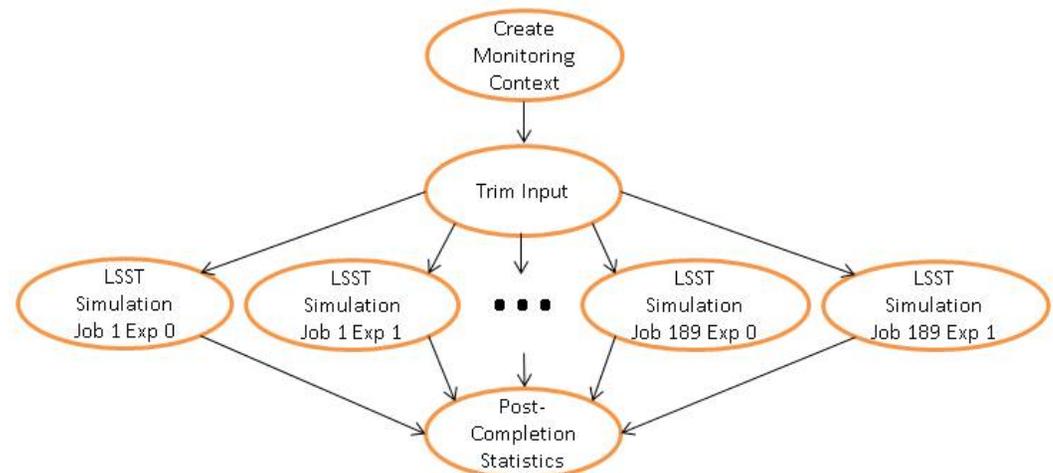
- Final goal of the activity is to empower the community to manage their computing activity using OSG services and platforms.

History

- Feb 2010 – Proof of principle: 1 img on OSG
- Jun 2010 – OSG EB forms task force for production scale: 500 img pairs i.e. 1 night of observation
- Jul 2010 – Commissioning of the LSST submission system on OSG
 - 1 person produced 183 times the same image in 1 day
- Sep 2010 – LSST operator (Bo Xin) ran operations to produce 529 pairs
- Oct 2010 – LSST validated results
- Nov 2010 – Feb 2011 Preparation for production phase

Workflow Requirements

- LSST simulation of 1 image: 189 trivially parallel jobs for the 189 chips
- Input to the workflow:
 - SED catalog and focal plane conf. files: 15 GB uncompr., pre-installed at all sites
 - Instance Catalog (SED files + wind speed, etc.): 500 MB compr. per image pair
- Workflow:
 - Trim catalog file into 189 chip-specific files
 - Submit 2 x 189 jobs: 1 image pair (same image w/ 2 exposures)
- Output: 2 x 189 FITS files, 25 MB each compr.

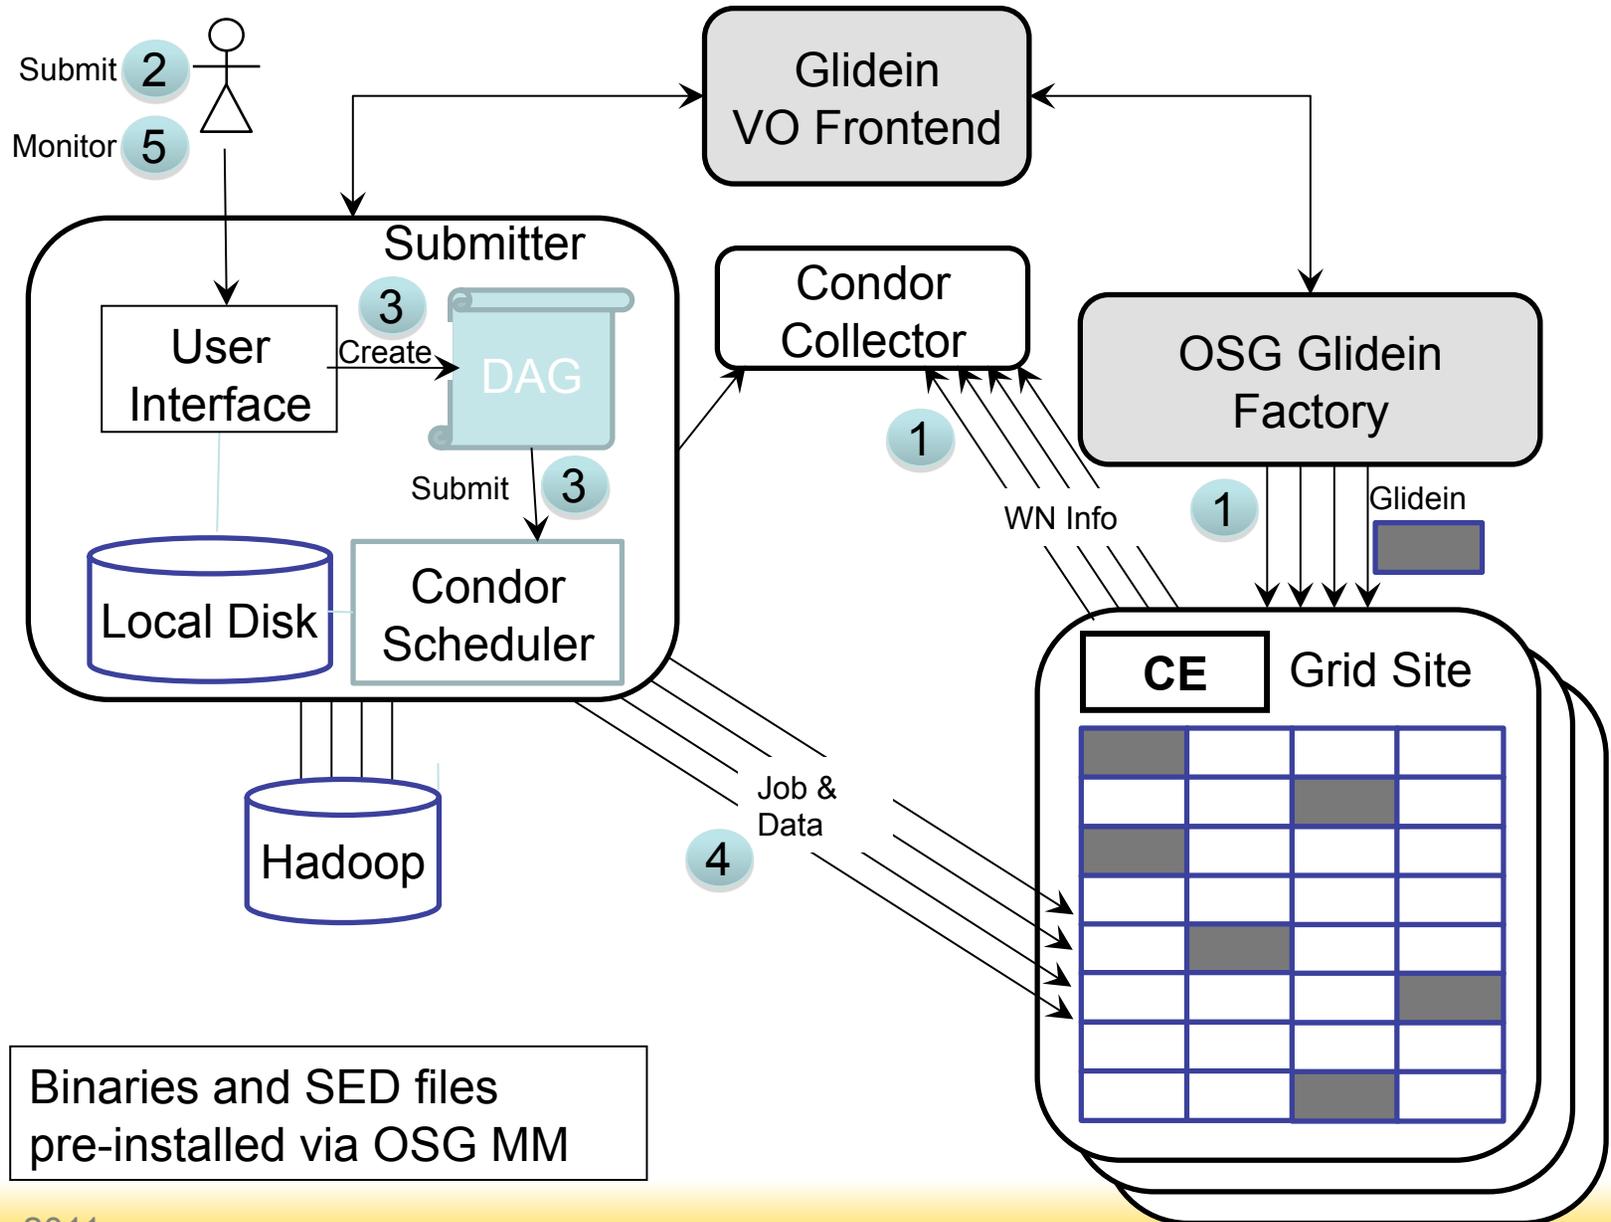


Production by Numbers

- Goal: simulate 1 night of LSST data collection: 500 pairs
- 200k simulation jobs (1 chip at a time) + 500 trim jobs
- Assume 4 hours / job for trim and simulation (over-est.)
→ 800,000 CPU hours
- Assume 2000 jobs DC → ~50,000 CPU hours / day

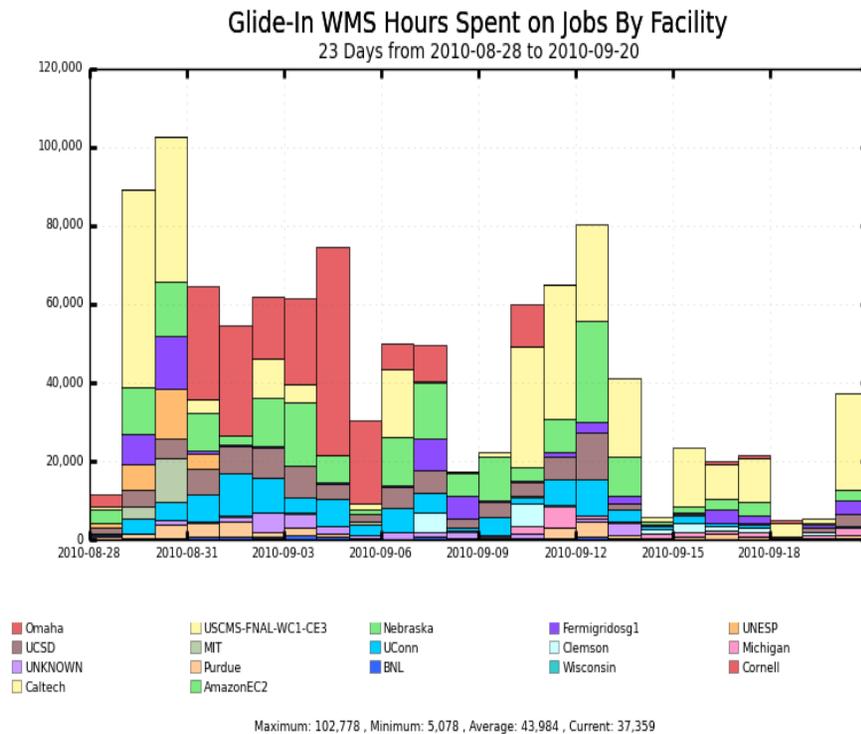
- 17 days to complete (w/o counting failures)
- 12,000 jobs / day i.e. 31 image pairs / day
- 50 GB / day of input files *moved* (different for every job)
- 300 GB / day of output
- Total number of files = 400,000 (50% input - 50% output)
- Total output compressed = 5.0 TB (25 MB per job)

Architecture

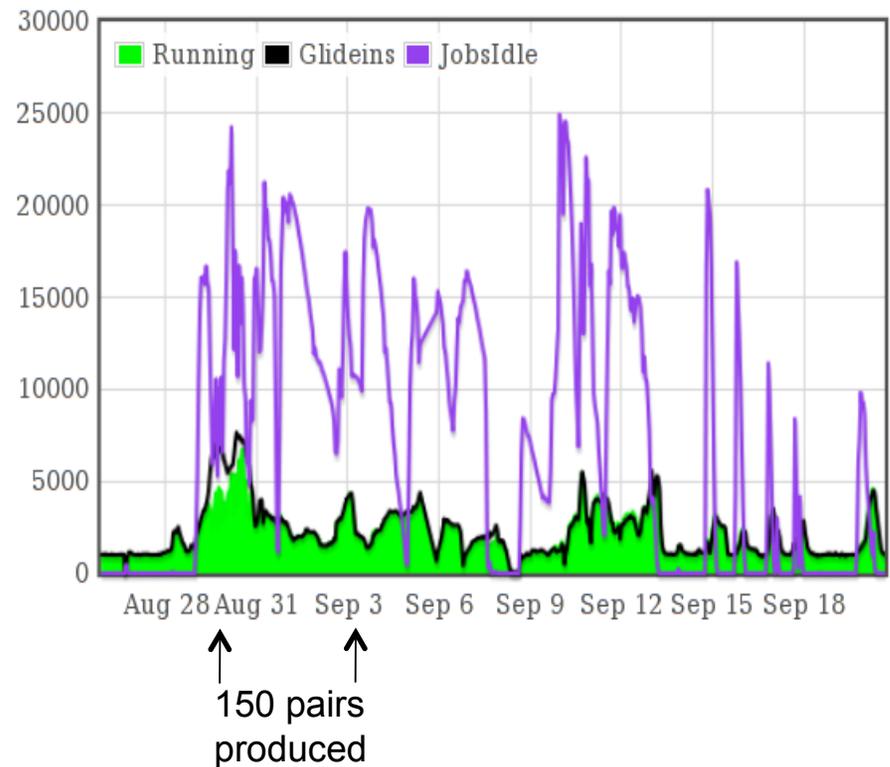


Resource Utilization

- By Sep 3, produced 150 pairs in 5 days using 13 sites.



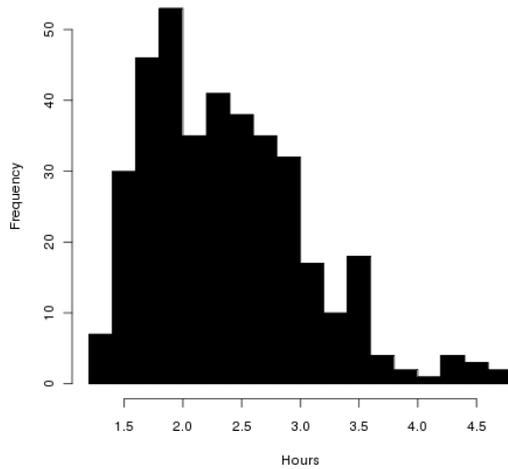
Gratia Resource Utilization plots



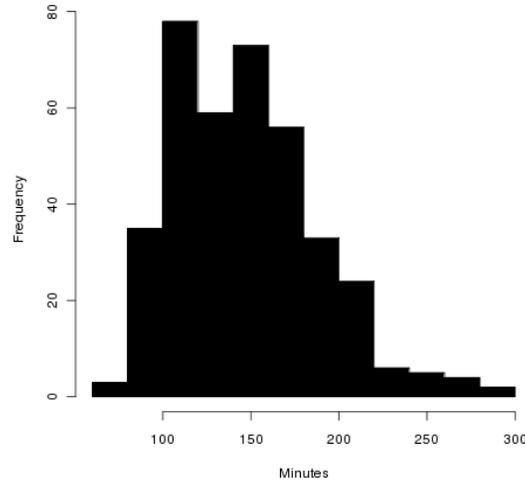
Frontend Status: Jobs & Glideins

Typical Workflow Statistics

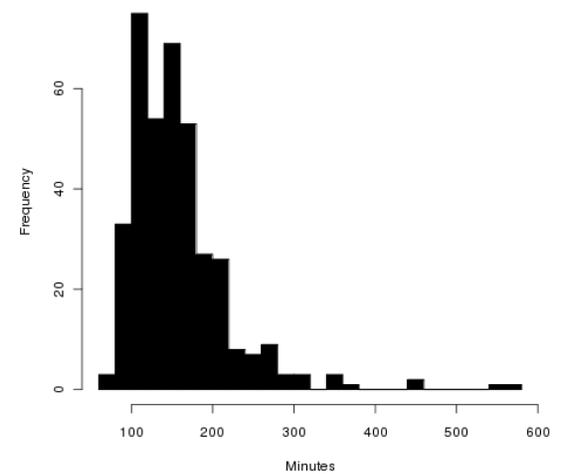
CPU Usage - Id LSSTsim_20100824_182721



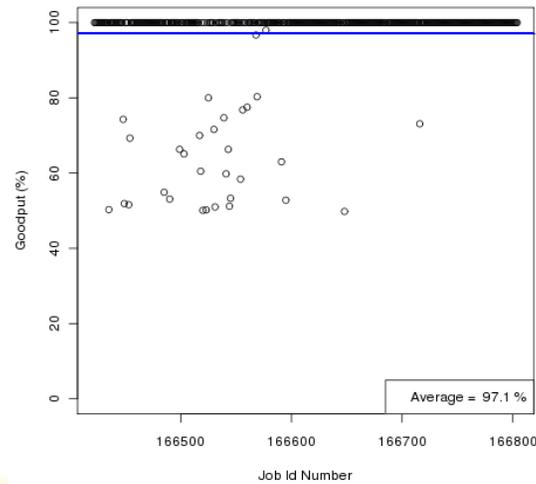
Successful Job Duration - Id LSSTsim_20100824_182721



Overall Job Duration - Id LSSTsim_20100824_182721



Job Goodput - Id LSSTsim_20100824_182721

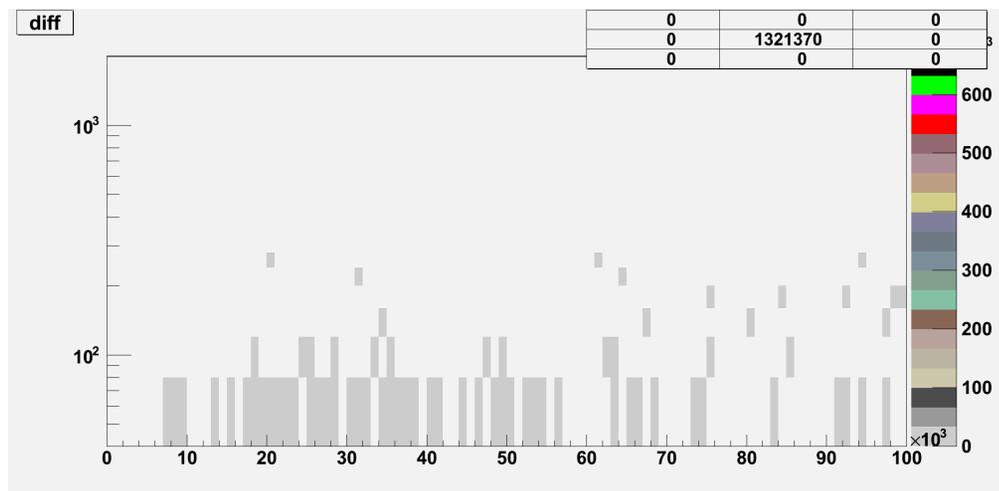
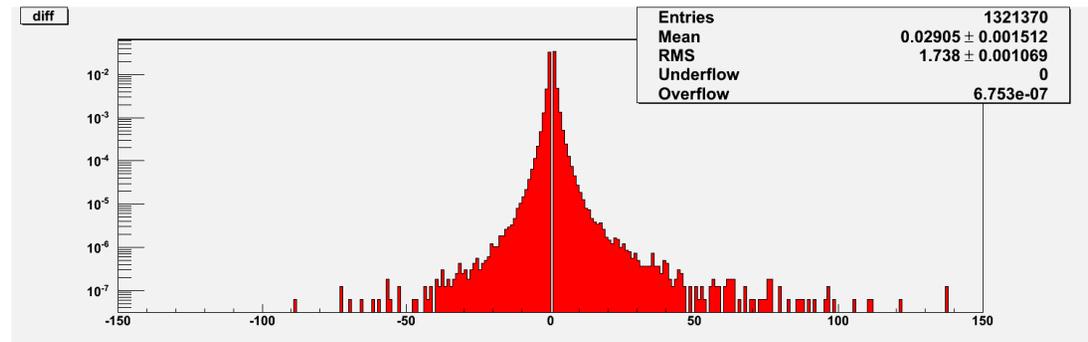


Operational Challenges

- LSST binaries were not Grid-ready
 - Application assumed writable software distribution
 - Application assumed path-lengths too short for the Grid
 - Orchestration script did not exit with failure upon error (required manual recovery until fixed)
- Typical failures at sites:
 - Job required more memory than the batch system allotted
 - Storage unavailable due to maintenance at some of the most productive sites
- Limited disk quota on the submission machine
- After the production of 150 pairs, the operator was mostly traveling and had limited time to dedicate to the operations

Validation: comparison with PT1

- We compared 6 image pair (2268 chips) with “official” references (PT1).
- 99% of the chip imgs are identical (pixel-subtraction is consistently 0)
- 14 chips are identical except for a few pixels: negligible
- Images produced on OSG are reproducible



Conclusions

- The Open Science Grid promotes data-intensive research providing a computing facility and services that integrate distributed, reliable and shared resources
- The Large Synoptic Survey Telescope (LSST) will capture 1000 panoramic sky images each night with a rapid fire 3.2 Gigapixel camera, covering the sky twice per week.
- OSG User Support and LSST are working together to run Image Simulation and Data Management applications on OSG.
- 1 LSST person, Bo Xin, with expert support has simulated 1 night of observations (529 image pairs) in 4 weeks, using 50,000 CPU h / day.
- This project has demonstrated that OSG is a valuable platform to simulate LSST images