# LSST Data Management


## Jacek Becla
### (interim) DM Project Manager


**October 13, 2016**

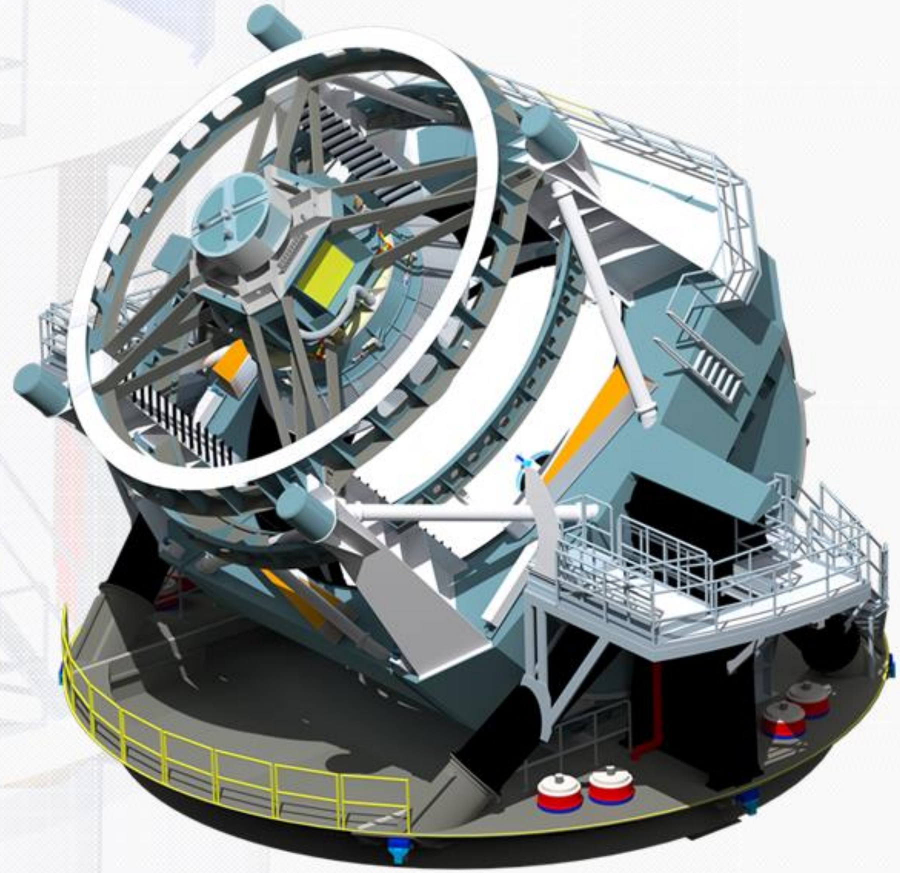

CHEP 2016
San Francisco, CA

# LSST in a Nutshell

- Decade-long, deep, wide, fast time-domain survey of optical sky
  - Highest ranked project in 2010 Decadal Survey
- Consists of
  - 8-meter class wide-field ground based telescope
  - 3.2 Gpix camera
  - Automated data processing systems
- Construction funding
  - Telescope & Site, DM, EPO funded by NSF, MREFC, $473M
  - Camera fabrication funded by DOE, MIE, $168M
- Science themes driving LSST requirements:
  - Taking a census of moving objects in the solar system
  - Mapping the structure and evolution of the Milky Way
  - Exploring the transient optical sky
  - Determining the nature of dark energy and dark matter

# Located in Central Chile

# In Construction

# Commissioning in < 3 years

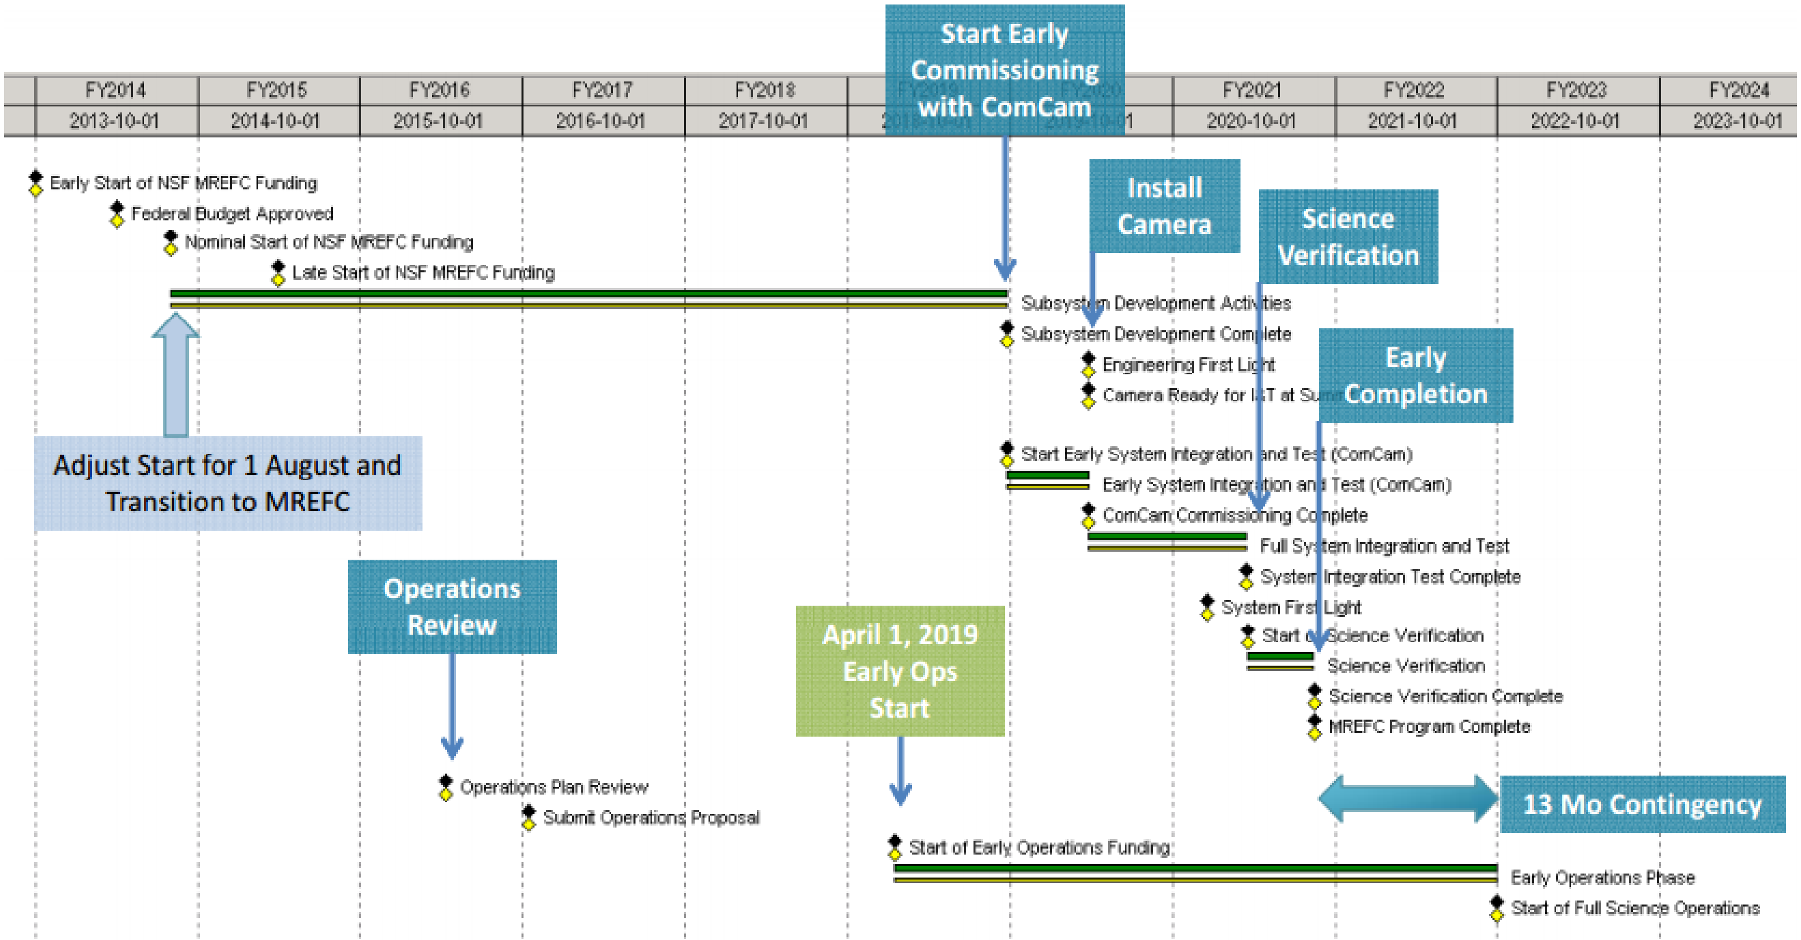

# LSST DM From a Scientist's Perspective

- A stream of ~10 million time-domain events per night, detected and transmitted to event distribution networks within 60 seconds of observation.
- A catalog of orbits for ~6 million bodies in the Solar System.

**Level 1**

- A catalog of ~37 billion objects (20B galaxies, 17B stars), ~7 trillion observations ("sources"), and ~30 trillion measurements ("forced sources"), produced annually, accessible through online databases.
- Deep co-added images.

**Level 2**

- Services and computing resources at the Data Access Centers to enable user-specified custom processing and analysis.
- Software and APIs enabling development of analysis codes.

**Level 3**

# Infrastructure: Petascale Computing

*Running processing pipelines*

- *Dedicated data center*
- *Process images in real-time to detect changes in the sky*
- *Produce annual data releases*

*Data Access Centers will provide end-user analysis capabilities and serve the data products to LSST users*

# Setting up a Prototype DAC at NCSA

1. Compute and storage resources in OpenStack cluster.
2. 1000-core processing cluster and 2.5-PB GPFS storage for QA and verification testing, stack development.
3. 30-node qserv database and SUI service integration environment
4. Infrastructure for testing object storage technologies and inter-site data transfer.

# Computing Infrastructure and Networks

**Long Haul Networks** to transport data from Chile to the U.S.

- *200 Gbps from Summit to La Serena (new fiber)*
- *2x100 Gbit (minimum) for La Serena to Miami, FL (protected, existing fiber)*
- *Add'l 100 Gbit link Santiago. Chile to Boca Raton, FL*

## *Satelling Processing Center*

***(CC-IN2P3, Lyon, France)***

*Data Release Production (50%)*
*French DAC*

## Archive Site

### Archive Center

Alert Production
Data Release Production (50%)
Long-term Storage (copy 2)

### Data Access Center

Data Access and User Services

## Summit and Base Sites

Telescope and Camera
Data Acquisition
Crosstalk Correction
Long-term storage (copy 1)
Chilean Data Access Center

## HQ Site

Science Operations
Observatory Management
Education and Public Outreach

# State-of-the-art Scientific Algorithms

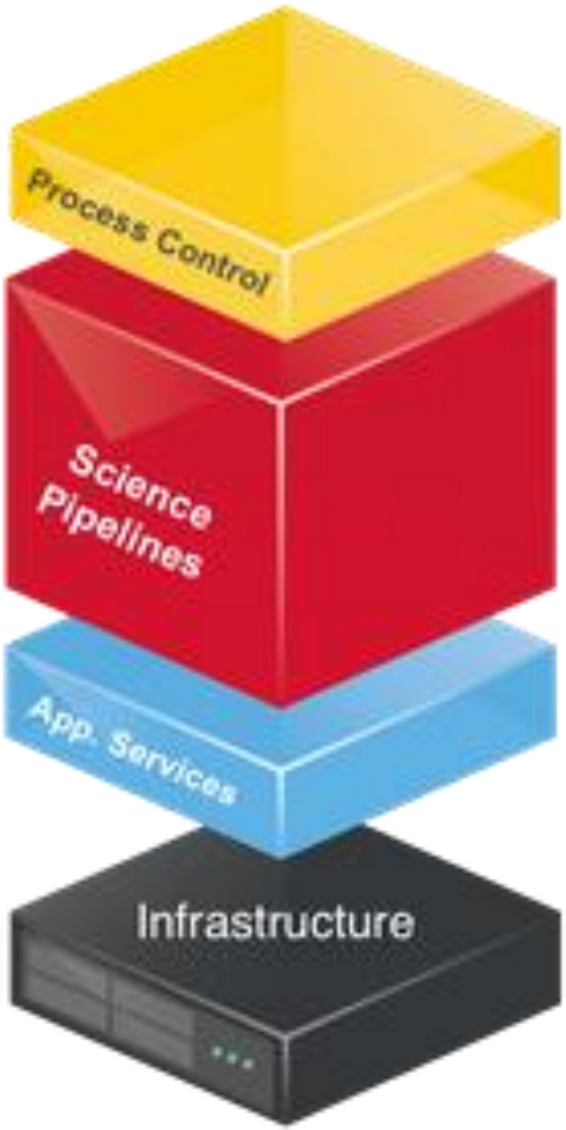


- Science Pipelines carry core scientific algorithms that process or analyze raw LSST data to generate output Data Products

- Variety of processing
  - Image processing
  - Measurement of source properties
  - Associating sources across space and time, e.g. for tracking solar system objects

# Prototype LSST Science Pipelines Are Running on HSC Survey …

HSC "ultra deep" gri imaging in COSMOS, with a total of 1.5 hours in g and r and 3 hours in i; (280/550 LSST visits).

The visits were processed, calibrated, registered, added, and the resulting coadds processed using the LSST stack.

These catalogues are being used to carry out first-year HSC science.

***Credit: HSC collaboration, Robert Lupton and LSST DM @ Princeton.***

Processing data at scale:

**HSC Survey S16A Data Release**

**Exposures:** 8192 images
(note: deeper than LSST, exposure time is ~minutes)
**Areal coverage:** 174 $deg^2$
**Total data volume:** ~200TB

Notes:

- Using a fork of LSST science pipelines adapted to HSC camera (will not be necessary as of Dec'16)
- Using HSC-Survey's orchestration middleware

# … and enabling science.

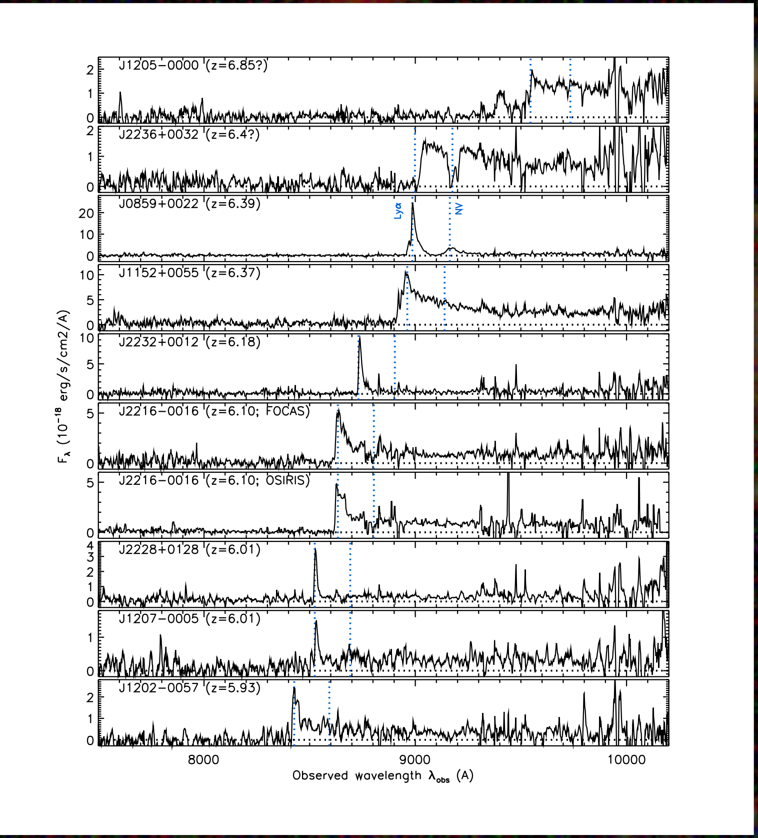


*High-redshift quasars selected from the HSC `wide' data (10 minutes per band in gr and 20 minutes in /izy/; equivalent to 30 and 60 LSST visits, respectively).*

*Followup spectra were taken at Subaru, **resulting in the discovery of second-highest known QSO redshift**.*

*Credit: Yoshii Matsuoka et al. 2016.*

HSC has been a good test bed for early LSST technology. Both HSC and LSST benefited.

# Complex Middleware

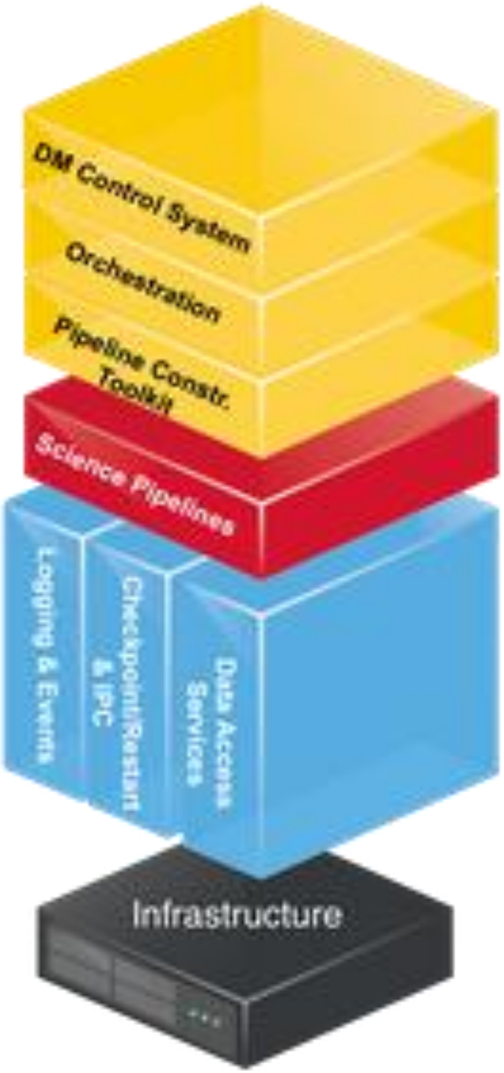


*Orchestrating execution of science pipelines on hundreds of thousands of cores*

- Frameworks to construct pipelines out of basic algorithmic components
- Orchestration of execution on thousands of cores
- Control and monitoring of the whole DM System

*Isolating science pipelines from underlying hardware*

- Services used by pipelines to access/produce data and communicate
- "Common denominator" interfaces handle changing underlying technologies

# Massively Parallel, Scalable Database for Spherical Data

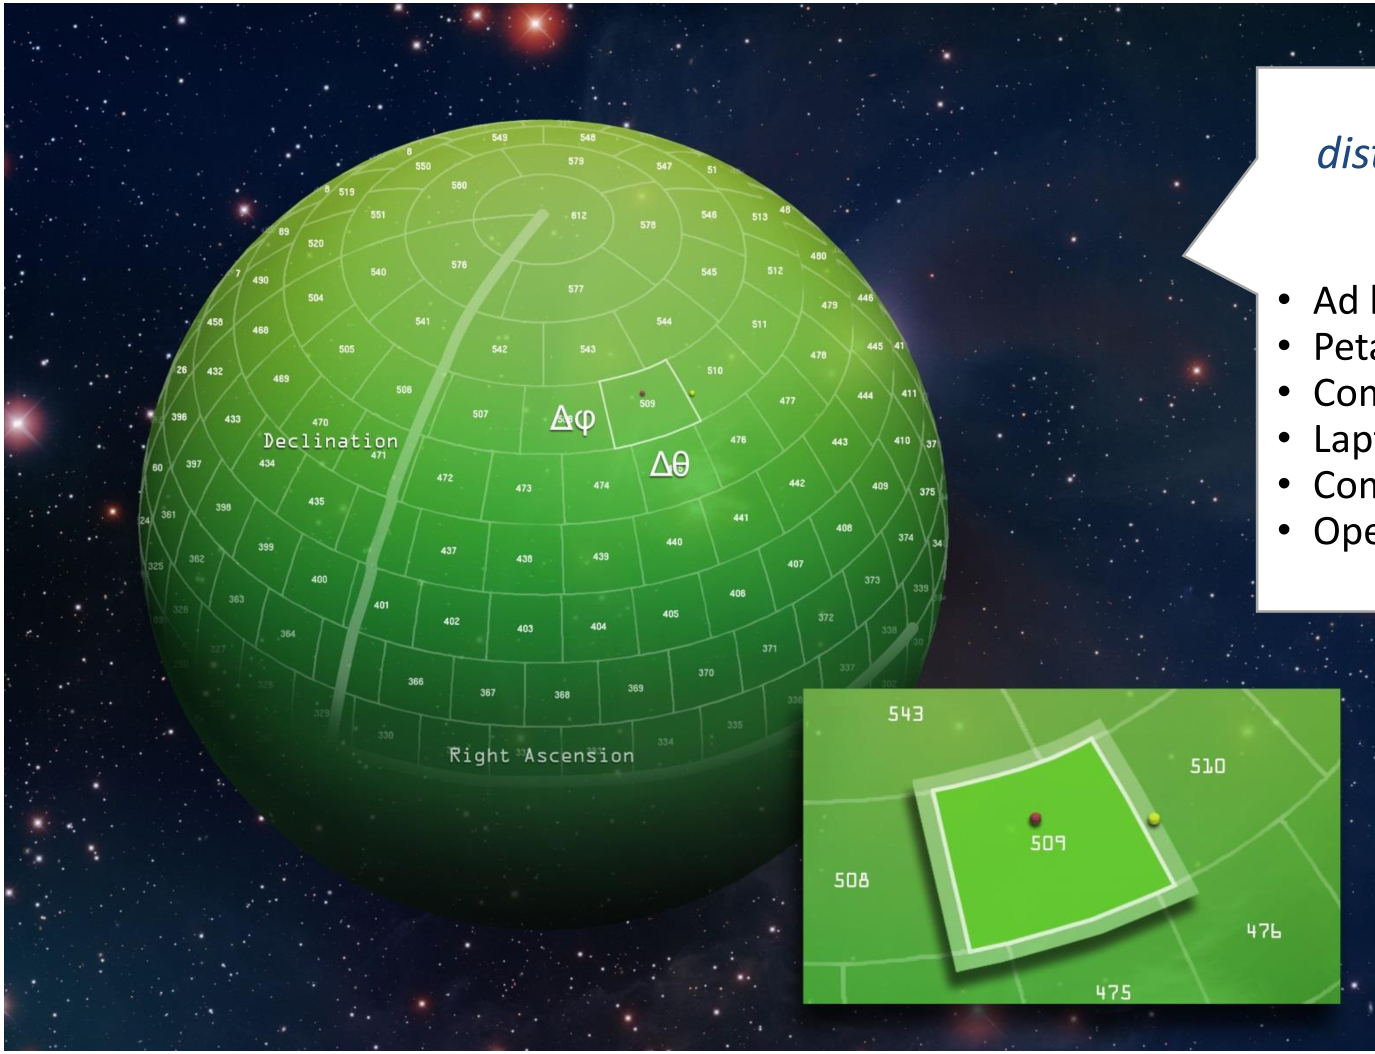


*Massively parallel, distributed, fault-tolerant* ***database***

- Ad hoc user query analysis
- Petabyte level scans
- Complex correlations
- Laptop scale to petascale
- Commodity hardware
- Open source

# Qserv

- Home-grown, specialized, shared-nothing query service
- Relies on MySQL and XRootD
  - XRootD used for scheduling queries and communication
- Unique state-of-the-art features:
  - Special optimizations for concurrent scans
  - Spherical geometry
- Reaching beta state

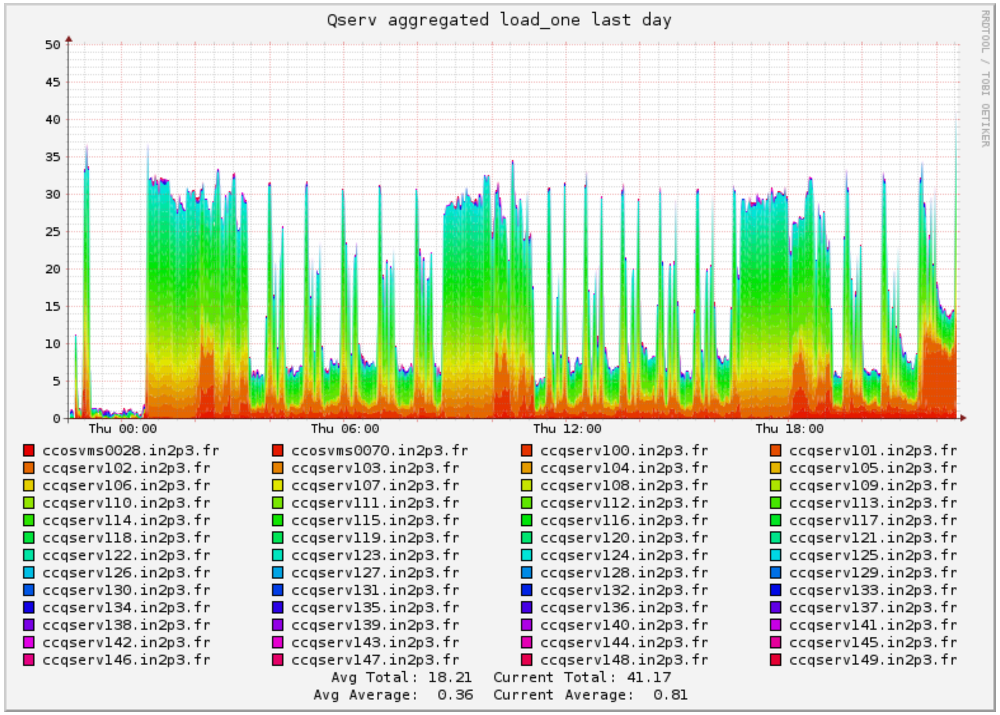

# Cutting-edge Visualization

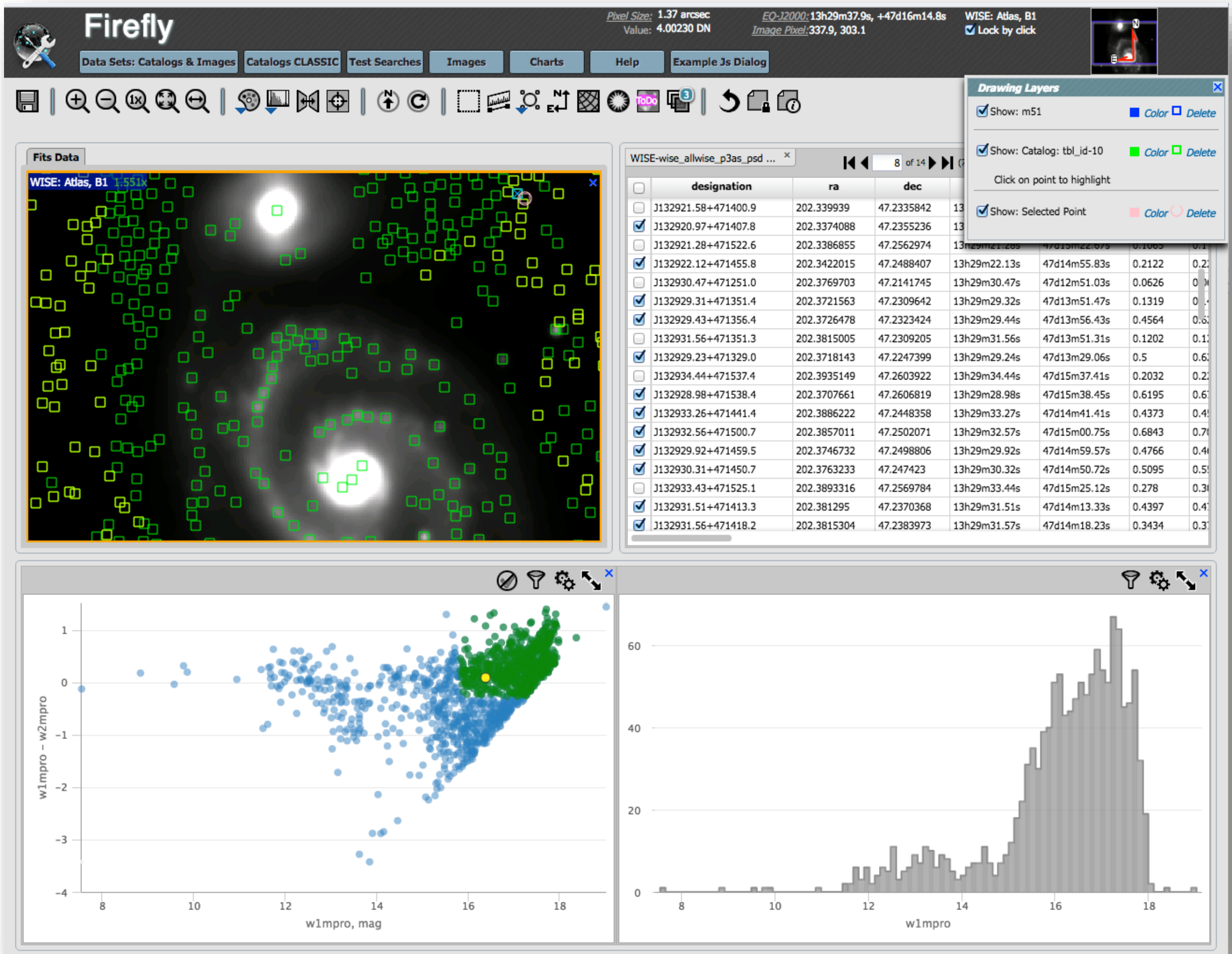


**Firefly** – Open source WEB UI Framework, for building web-based interactive, exploratory analytics for large science archives

Exposing native JS API, can easily embed Firefly widgets into HTML pages (e.g. IPython notebooks)

Industry standard build/deployment tools (npm, webpack)

**Note: To be deployed in production at IRSA in November 2016**

# It Has to Run. Reliably.

UI

Database

Core Algorithms

Middleware

Infrastructure

Science QA & Reliability Engineering (SQuaRE)

# DM Team

- Distributed team
- ~60 FTEs / 80+ people
- Mix of domain-experts and software engineers
- Hiring and on-boarding almost complete
- ~5x expansion comparing to start of MREFC

# Computing in Astronomy and HEP – what is similar and what is not?

# Data

- Size
  - HEP produces more data, Astro is catching up (LSST, SKA)
- Model
  - Astro data structures simpler, fit RDBMS
  - But some queries don't!
- Cross-survey
  - Astro data from other surveys correlate well
- Standardization
  - Significant efforts in Astro (VO, ADQL, TAP)

- LSST:
  - O(100) PB all databases
  - Largest data release o(10) PB
  - Images O(400) PB
    - Virtual and persisted

# Software

- Many similarities
  - Open source & tools
  - Data reduction
  - Pipelines
  - Data releases
  - User-facing frameworks, well defined APIs

# Infrastructure

- Tiered data centers
  - HEP ahead, Astro starts to follow
- Dedicated long-haul networks
- Cloud anyone?

# Data Analysis in Astro

- Wide range of users in Astro
  - Professional astronomers, amateur astronomers, classrooms, public
- Wide range of questions
  - Complex full-sky correlations, scans, simple selections
- Wide range of tools
  - SQL, iPython, Jupyterhub

# Summary

- Computing in astronomy
  - Large scale / distributed
  - Some unique challenges
  - Many commonalities with HEP computing

Let's collaborate more!