

Ramping up CC-IN2P3 for the LSST challenge

fabio hernandez

CONTENTS

- IN2P3 computing center overview
- CC-IN2P3 current and foreseen contributions to LSST

IN2P3 COMPUTING CENTER

IN2P3

A DISTRIBUTED LABORATORY

*2500 researchers, engineers
and technicians*

*700 post-docs and PhD
students*

*25 laboratories and research
platforms in France, 16
international laboratories*

COMPUTING CENTER



IN2P3 COMPUTING CENTER

- **CC-IN2P3**

84 people, 80 FTE, 80% permanent positions

~15 M€ overall annual budget

scientific data center, high throughput computing


well connected to national and international networks

- **Shared computing facility** supporting the institute's research program

~70 projects in high energy physics, nuclear physics and astroparticle physics



- Operations: 24x7
unattended during nights and weekends
engineer on duty during off-hours



aerial view



building 2



machine room 1



machine room 2




HESS

Planck




Fermi





Pierre Auger

AMS



LHCb


LHC @ CERN




ATLAS




SuperNova Legacy Survey



CMS



Alice




KM3NeT



Virgo

LSST AT CC-IN2P3

- CC-IN2P3 preparing to be a **satellite data release production center**
under NCSA leadership
formal agreement signed by LSST Corp., NCSA and IN2P3 (2015)
- Principle of operations
*CC-IN2P3 to **process 50% of the raw data** and **store the full dataset**, both raw and reduced data (images and catalogs)*
both NCSA and CC-IN2P3 to exchange and validate the data produced by the other party
each site to host an entire copy of every annual data release
- Ongoing work to understand the scope and size of a LSST data access center



LSST Operations: Sites & Data Flows

HQ Site
Science Operations
Observatory Management
Education & Public Outreach

Base Site
Base Center
Long-term storage (copy 1)

Data Access Center
Data Access & User Services



French Site
Satellite Processing Center
Data Release Production
Long-term Storage (copy 3)

Archive Site
Archive Center
Alert Production
Data Release Production
Calibration Products Production
EPO Infrastructure
Long-term Storage (copy 2)

Data Access Center
Data Access and User Services

Summit Site
Telescope & Camera
Data Acquisition
Crosstalk Correction


CURRENT STATUS

ENVISIONED ARCHITECTURE

22k → 122k CPU cores




Batch farm



Login farm



Data transfer nodes




Application servers




WAN


LAN



Catalog database



Disk storage



Mass storage

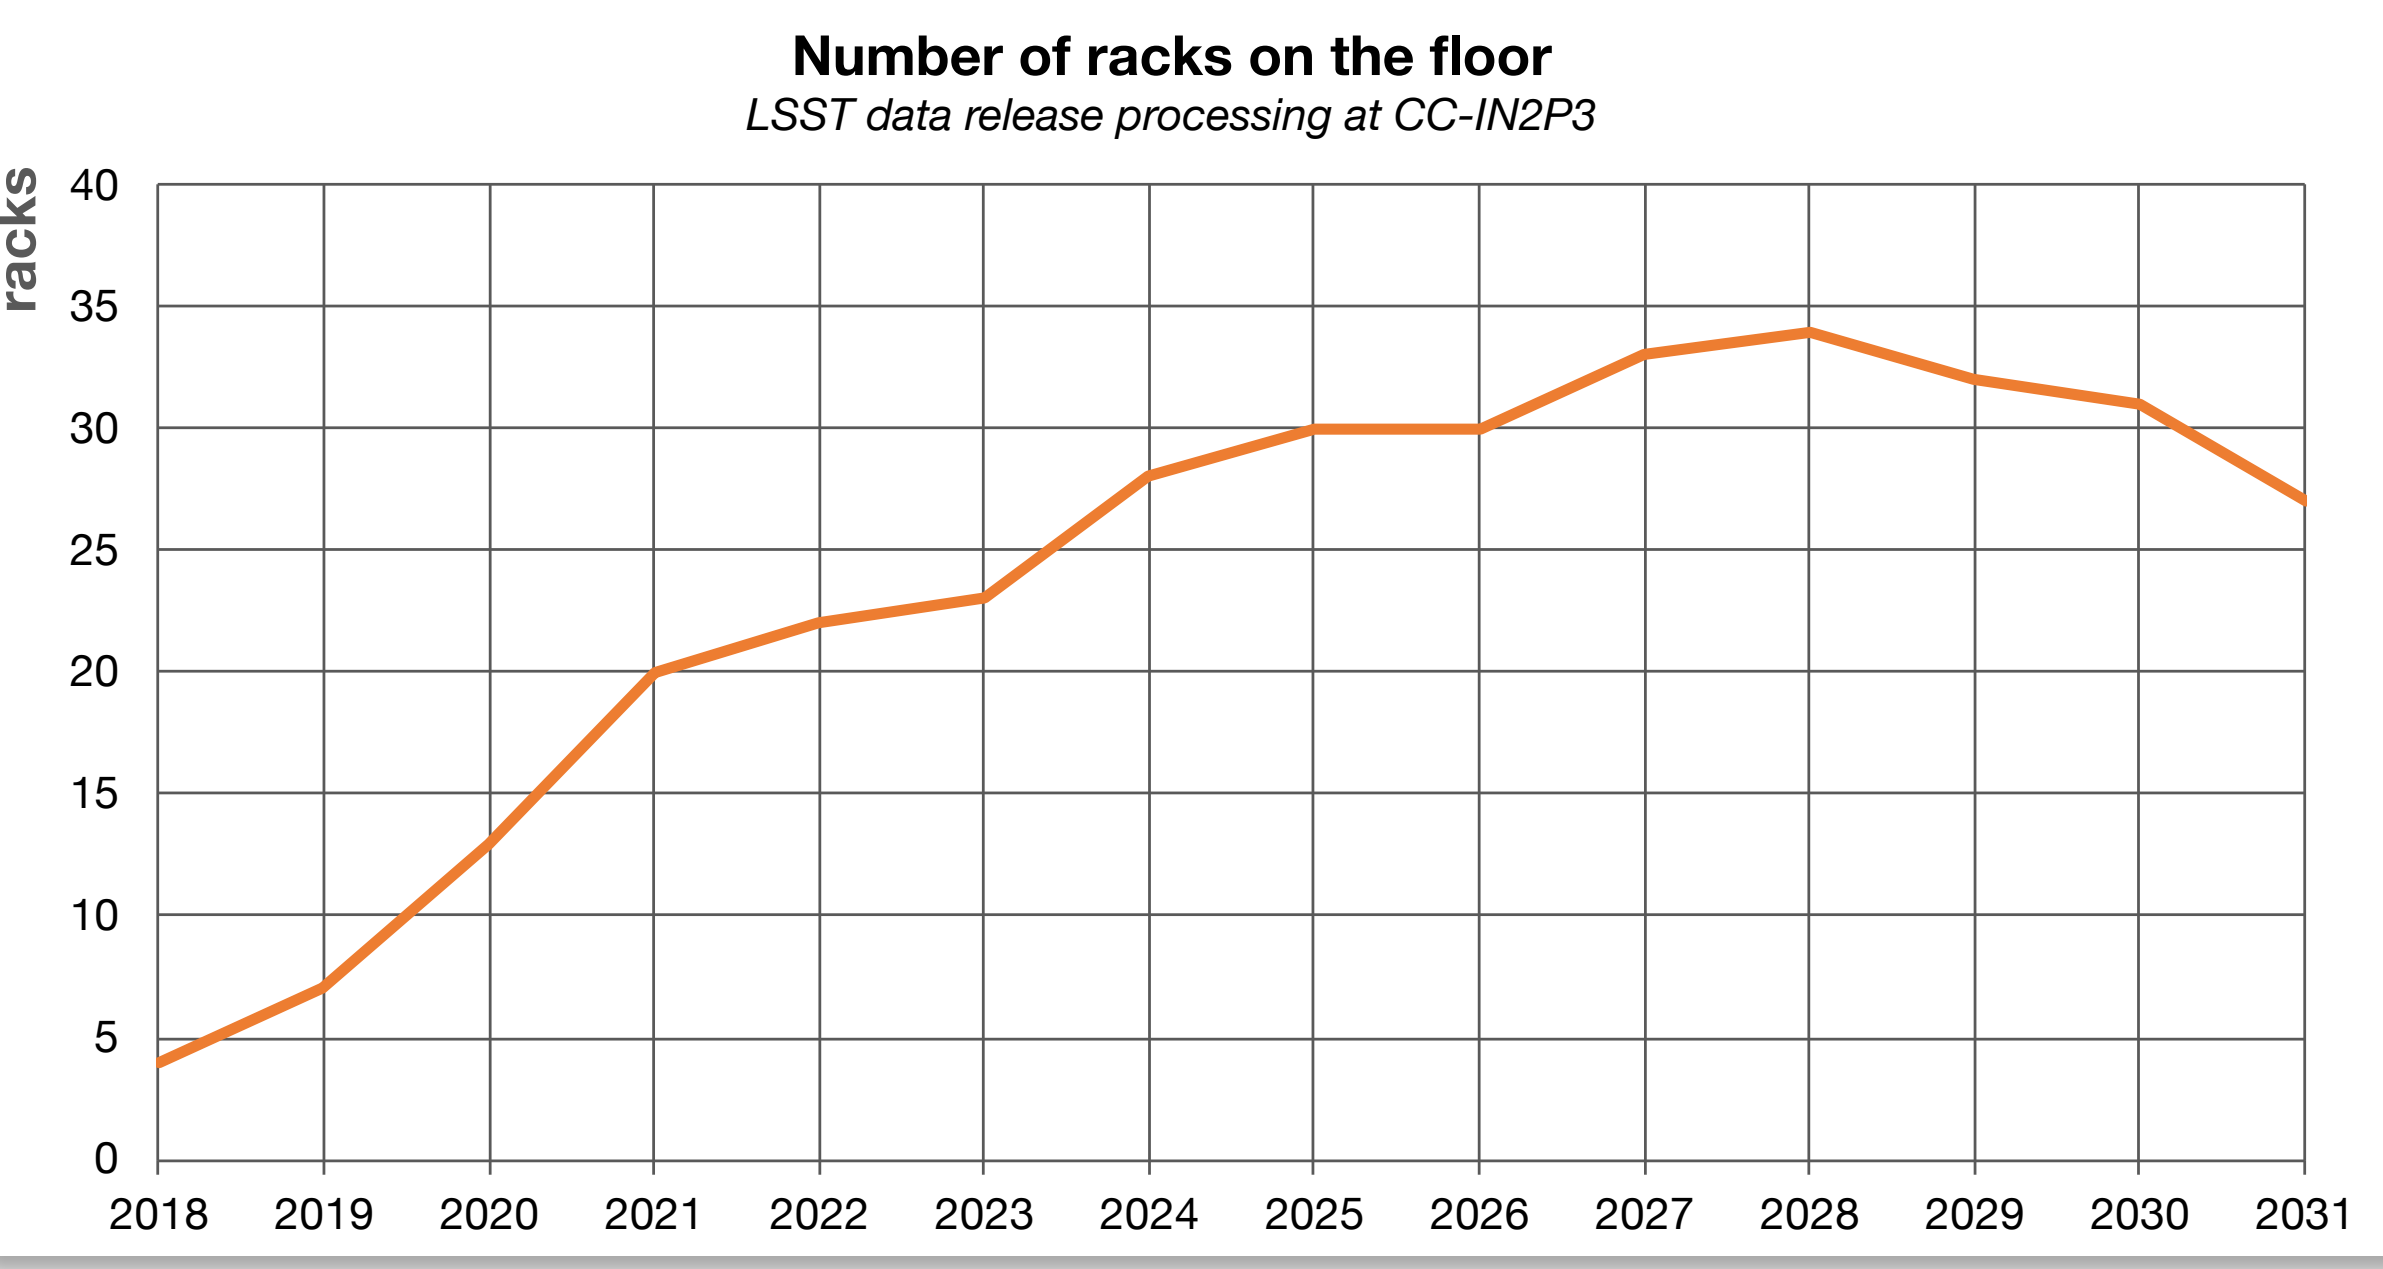
10 → 80 PB
120 → 550 GB/s

10 → 55 PB
50 → 300 GB/s


4 → 21 PB
8 → 30 GB/s

2022 → 2032

MACHINE ROOM INFRASTRUCTURE



Racks
peak 34 racks



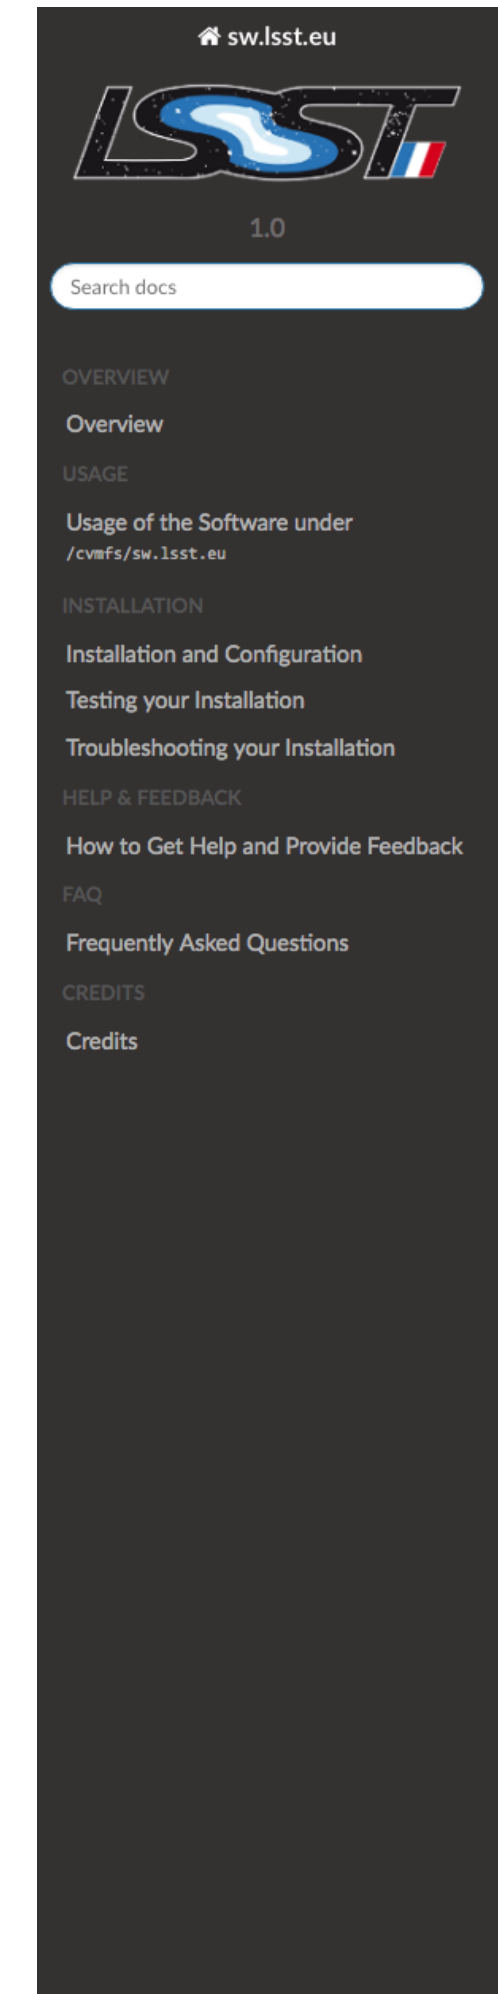
SIZING FOR DATA RELEASE PROCESSING

- Our sizing work is based on requirements as estimated by the project in 2013
at that time, the LSST software was significantly less developed
- It would be prudent we revisit the project's estimations in the light of the current state of the LSST software
in particular, we could measure the ability of the LSST software to exploit the computing capacity of modern CPUs
I expect we will find discrepancies which may have a significant impact on the computing capacity and machine room infrastructure required for DRP
*e.g. LHC experiments at CERN **measured** their simulation software uses **less than 10% of the FLOPs** theoretically delivered by recent CPUs (source: G. Stewart, slide 5 of [this presentation](#))*

SOFTWARE DISTRIBUTION

- LSST software **automatically delivered** to your personal computer *without your intervention*, both **stable** and **weekly** releases just appear as if they were installed on your laptop under `/cvmfs/sw.lsst.eu` same mechanism used for delivering the software to computers in both the login and batch farms at CC-IN2P3
- lower the barriers for end users to use the LSST software: the more people use it, the better it gets useful also for reproducibility*

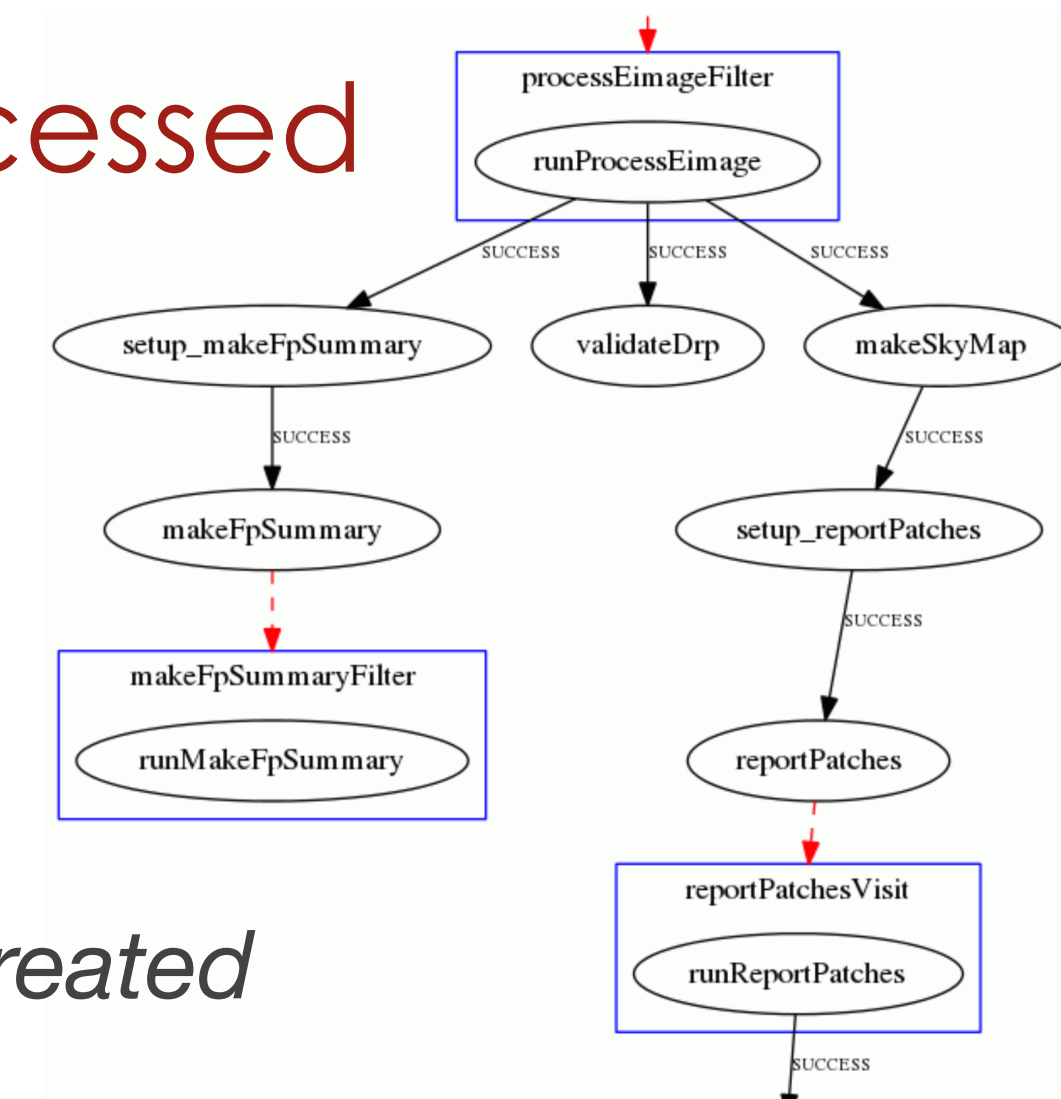
sw.lsst.eu



The screenshot shows the sw.lsst.eu website. At the top is the LSST logo with the version number 1.0. Below the logo is a search bar labeled 'Search docs'. A sidebar on the left contains a list of navigation links: OVERVIEW, Overview, USAGE, Usage of the Software under /cvmfs/sw.lsst.eu, INSTALLATION, Installation and Configuration, Testing your Installation, Troubleshooting your Installation, HELP & FEEDBACK, How to Get Help and Provide Feedback, FAQ, Frequently Asked Questions, CREDITS, and Credits. The main content area on the right is titled 'On how to get the LSST science pipelines delivered to your computer'. It contains text explaining that users will find information on how to get a binary distribution of the Large Synoptic Survey Telescope (LSST) science pipelines. It states that as a user of this software distribution service, both stable and weekly releases of the LSST software appear as if they were locally installed on your computer. Since new releases just appear under the local path /cvmfs/sw.lsst.eu without you to take any action, you can focus on using the software, instead of on the technicalities of installing and updating it regularly. It also mentions that users can visit the Overview to get more details on the benefits and intended audience of this service brought to the LSST community by LSST-France and CNRS / IN2P3 computing center (CC-IN2P3). Below this text are sections for OVERVIEW, USAGE, INSTALLATION, HELP & FEEDBACK, FAQ, and CREDITS, each with a list of links. At the bottom right of the main content area is a 'Next' button with a right arrow. At the bottom left of the page is the copyright notice: © Copyright 2018, LSST-France.

DESC DATA CHALLENGE 2

- Use the LSST software pipelines to process simulated LSST images
valuable for exercising CC-IN2P3's infrastructure for LSST data release processing
we want to identify the limitations of our infrastructure and tools as early as possible
- Images produced at NERSC, transferred to CC-IN2P3, batch processed using LSST software
data products transferred back to NERSC
currently using SLAC's SRS Pipeline execution engine
- Preliminary lessons learned
work needed for making the data products easily available for analysis: DESC task force created
hoping that Gen3 Butler will enable easier creation of data-driven workflows and give us more flexibility with file-related management tasks
need to build workflow-aware tools for continuous data exchange with the partner site as soon as data is produced



CONNECTIVITY & DATA EXCHANGE

- **Allocated bandwidth between CC-IN2P3 and NCSA: 20 Gbps**
bottleneck link is currently 10 Gbps
to ask for an increase of this allocation we need to demonstrate that we are limited by the current available capacity
connectivity provided by Renater, the French academic and research network
- **We need to demonstrate capacity to import 15 TB of raw data per night from NCSA (RTT: 110 ms)**
in addition to capacity to exchange data products with NCSA
currently doing exercises with NERSC, in the framework of the DESC data challenge

CONNECTIVITY & DATA EXCHANGE (CONT.)

Data flow: **NERSC** (GPFS) → **CC-IN2P3** (GPFS) [3 servers, 4 clients]



Aggregated application-level network throughput: **1.5 GB/s (12 Gbps)**

pull model, disk-to-disk transfer, wide area network, 150ms RTT, secure HTTP

International connectivity provided by



LSST SCIENCE PLATFORM PROTOTYPE

Settings Help

display.ipynb x RGB.ipynb ● astrometry_cle x skyWcs.ipynb x test-butler.ipynb x Check_astrom x Example_note x multi_v


Code ▾

```
exposures = {}
for bandpass in bandpass_color_map.values():
    dataId['filter'] = bandpass
    # refs[bandpass] = butler.get('deepCoadd_ref', dataId=dataId)
    exposures[bandpass] = butler.get('deepCoadd', dataId=dataId)

rgb_im = rgb.makeRGB(*(exposures[bandpass_color_map[color]].getMaskedImage().getImage()
                        for color in ('red', 'green', 'blue')), Q=8, minimum = -0.1, dataRange=1.5,
                        saturatedPixelValue=100,
                        xSize=None, ySize=None)

fig = plt.figure(figsize=(12,12))
#rgb.displayRGB(rgb_im)

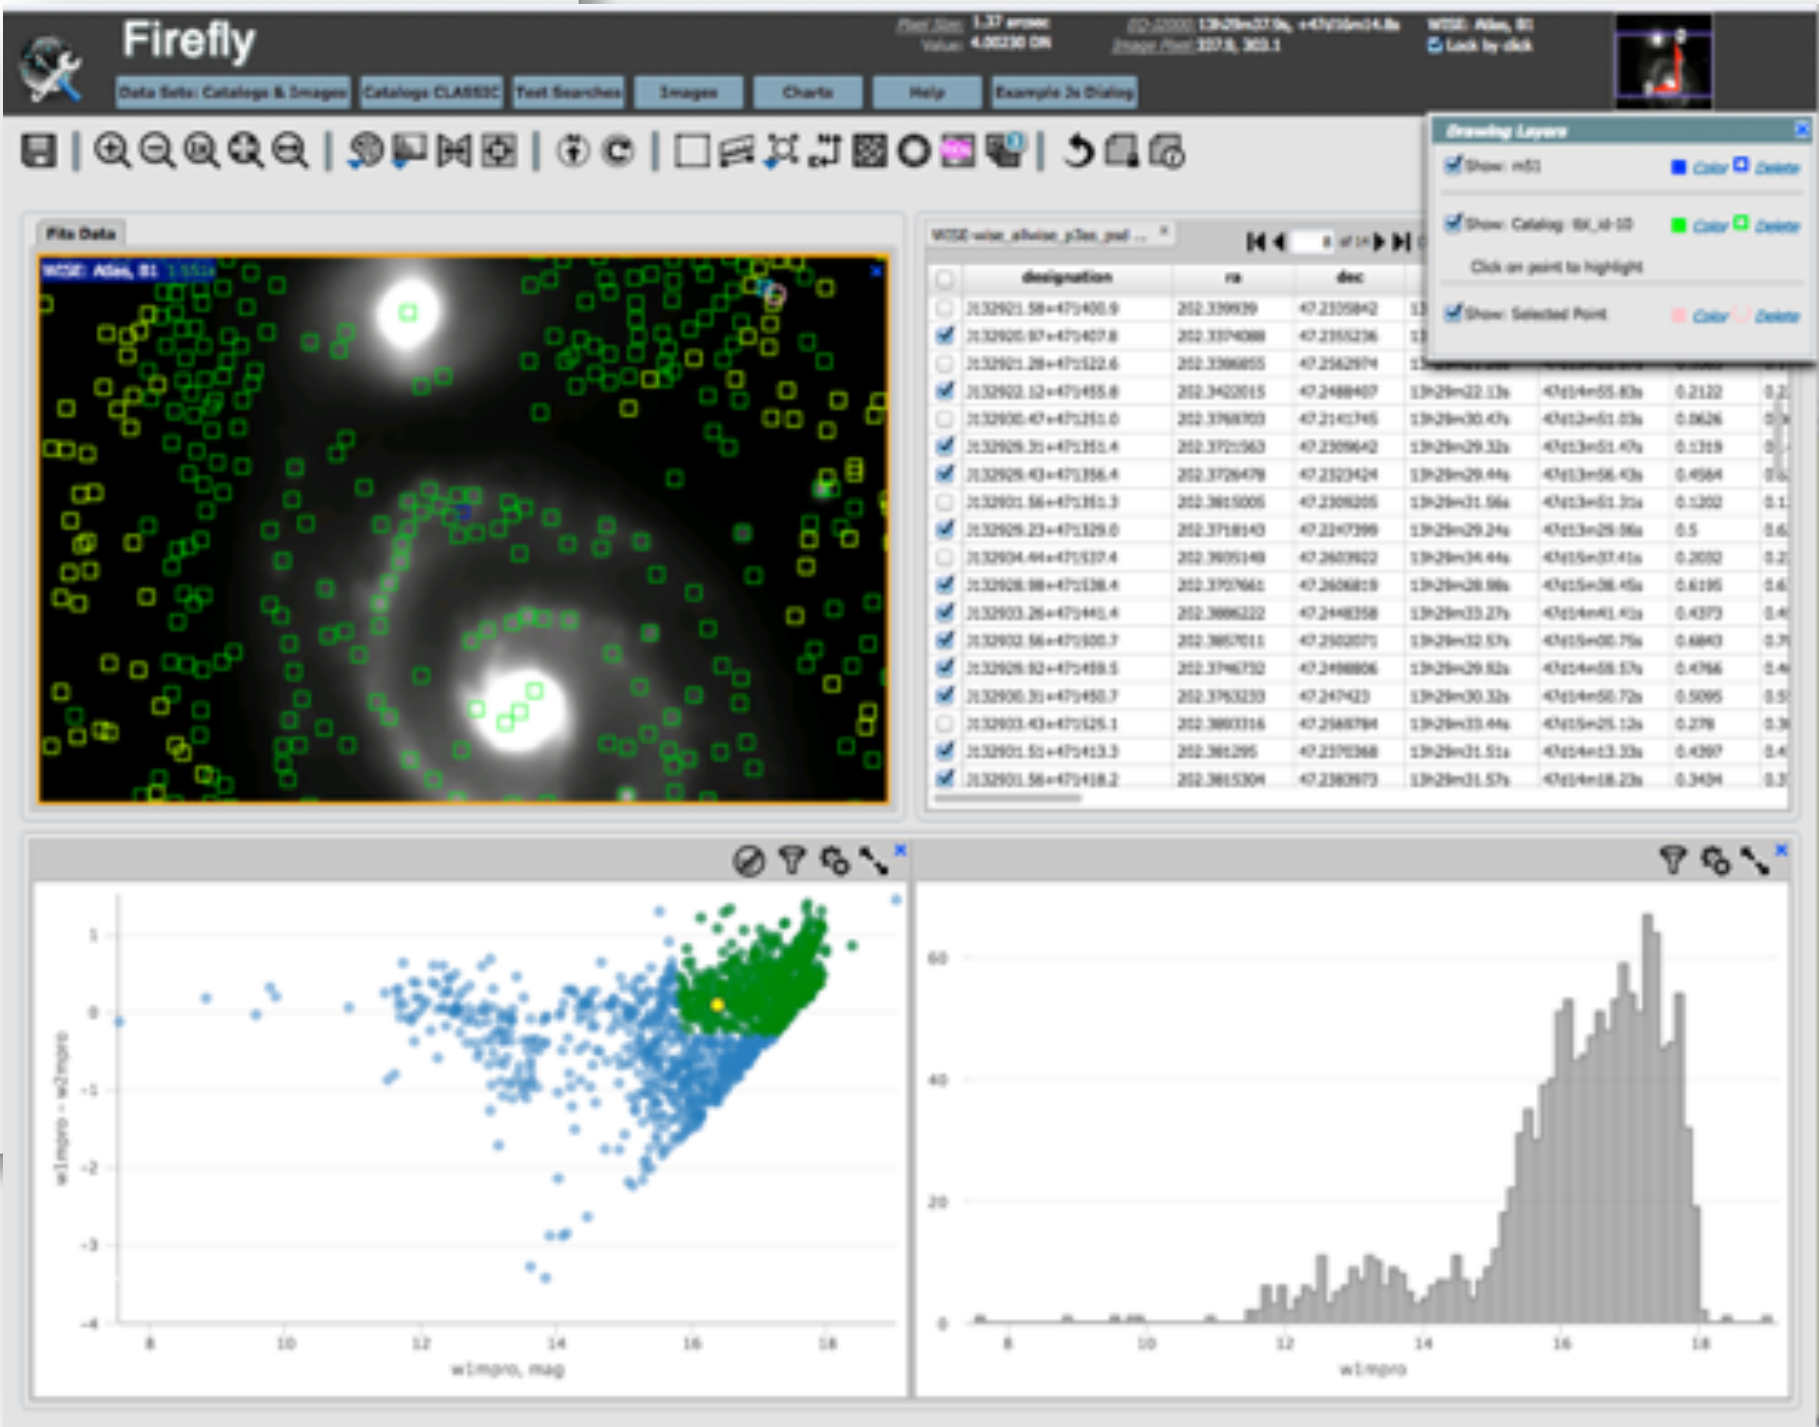
In [3]: fig = plt.figure(figsize=(10,10))
rgb_im = rgb.makeRGB(*(exposures[bandpass_color_map[color]].getMaskedImage().getImage()
                        for color in ('red', 'green', 'blue')), Q=8, dataRange=1.0,
                        xSize=None, ySize=None)
#saturatedPixelValue=130,
rgb.displayRGB(rgb_im)
```



Sources: D. Boutigny, X. Wu et al.

LSST-ready Python notebooks, routinely used by scientists in CC-IN2P3's login farm

Currently exploring how to integrate a Firefly server without requiring users to go through a VPN



CATALOG DATABASE

- We continue supporting the Qserv development team by operating a dedicated development cluster
*50 hosts, 400 CPU cores, 800 GB RAM, 500 TB raw storage capacity
hardware lent by Dell in the framework of an institutional partnership*
- A fraction of that cluster was recently prepared for ingestion exercises of data products of DESC data challenge
*tools for multi-threaded ingestion were developed
work lead by S. Elles (IN2P3 LAPP)*

Visit to CC-IN2P3 scheduled on
Wednesday afternoon

[click here for a virtual visit](#)

QUESTIONS & COMMENTS