Data Management Pipeline: First Notes from LSST Data Management discussion with Phil Marshall

Michael Kagan

SLAC

SLAC MAGIS Meeting

September 2, 2021

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

LSST Data Management Model

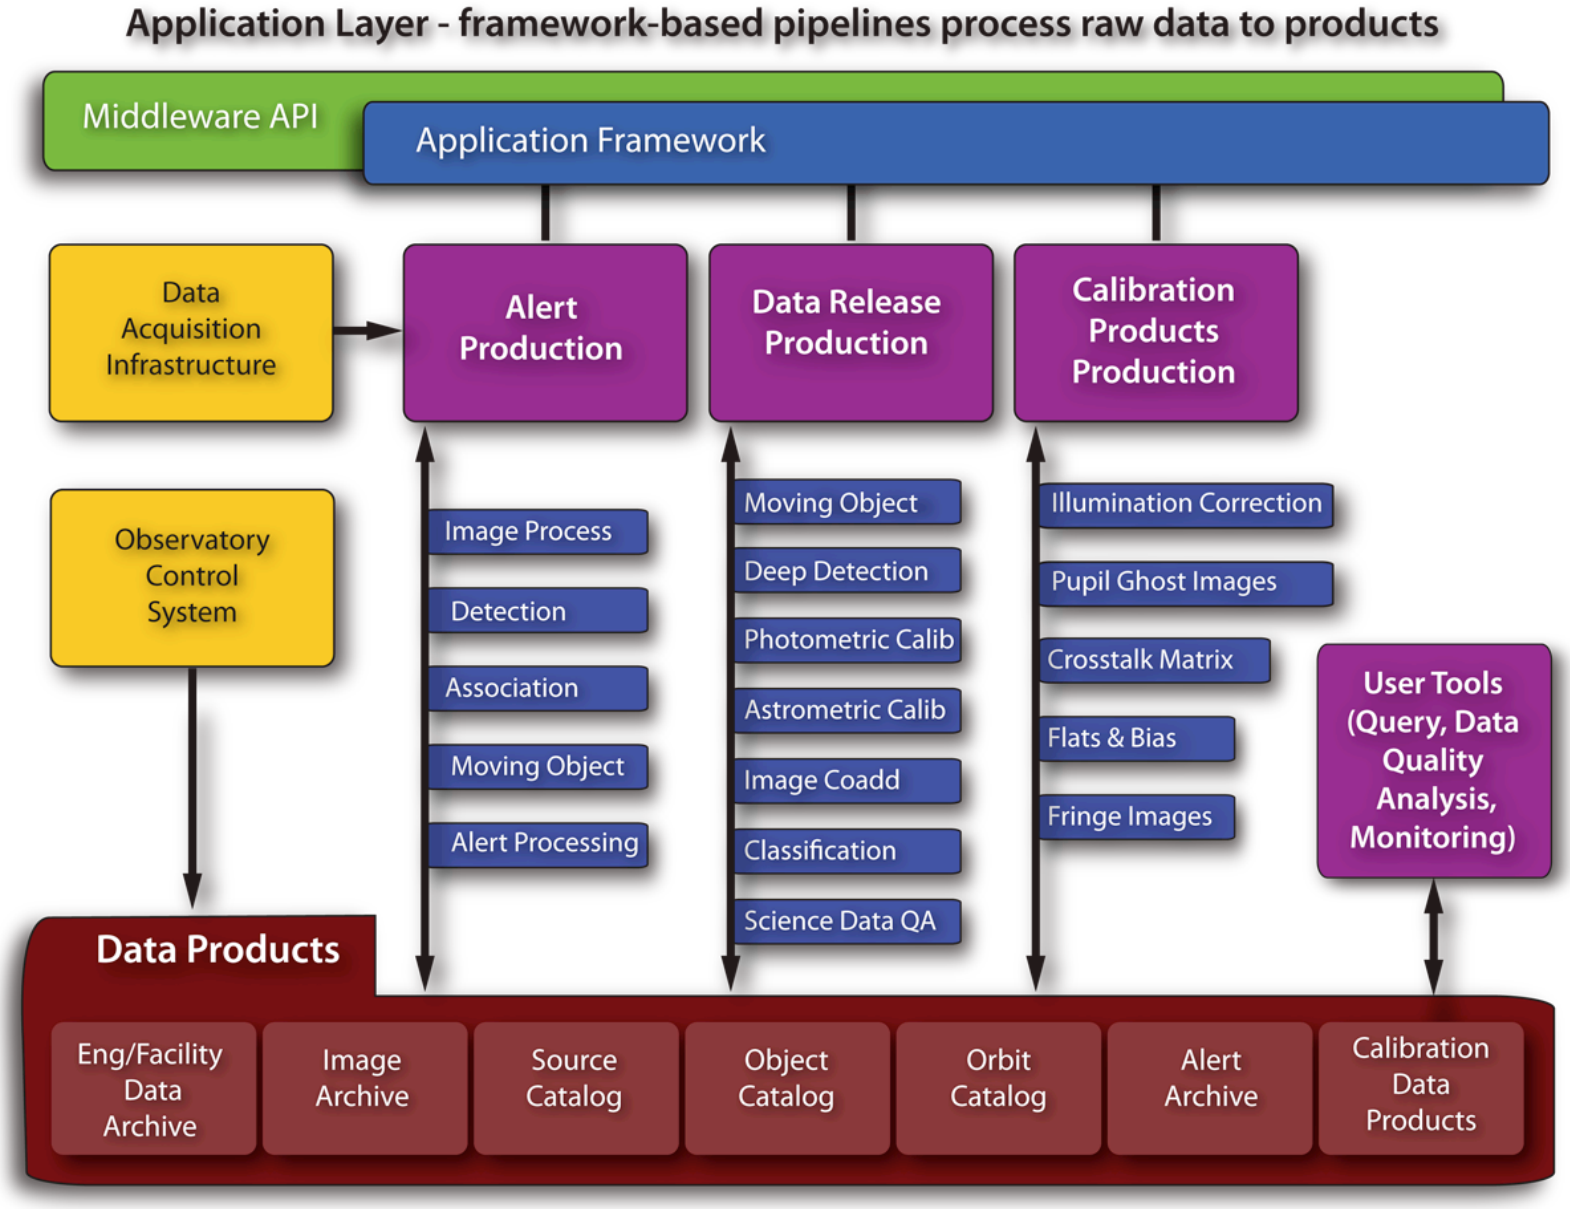

Image Credit: https://www.lsst.org/about/dm

LSST Data Management Model

https://www.lsst.org/about/dm

- Data Acquisition feeds into *3* pipelines
  - Alert Production
  - Data Release Production
  - Calibration Production

  - Raw data is replicated in several places in enormous databases
  - But analyzers don't typically use RAW data

- Analysis Pipeline: Data Release Production → Data Analysis
  - **Data Release Production:** does *Data Processing* which extracts key information from images and stores them in database
    - This is what Users typically use for analysis
    - If they need a new data processing, they write new algorithms and put them in the production pipeline for use in the next data processing

    - **Data analysis**: uses processed data to do analysis
      - Typically in Jupyter notebooks on a large Analysis cluster
      - They are figuring out if this will work, or if batch submissions will also be needed for analysis, as even the processed data is large.

- Pipeline code seems to be in Python, but many of the pieces are written in C++ (for speed) and then turned into Python objects with pybind (public: https://github.com/lsst)
  - This is NOT like ATLAS that uses Python for configuration
  - The main code seems to be in Python, but just with some objects written in C++ and converted

MAGIS Questions: Analysis Mode

- What is our real data rate? Analysis mode:
  - two 6-megapixel to 60-megapixel images, say 10MB-100MB each
    → 20MB-200MB per drop + metadata (time, configuration, etc… ignore for now)
  - Drop rate 1-3Hz → Lets use 3Hz for maximum
  - Running time per day: Assume 16 hrs ?
  - Running days per year: Assume 200 ?
  - Yearly data = $\frac{20MB-200\,MB}{drop} * \frac{180\,drop}{min} * \frac{60\,min}{hr} * \frac{16\,hr}{day} * \frac{200\,days}{yr} = 0.7 - 7\,\frac{PB}{yr}$
    - That is 30-100 million images → realistic?
    - More than 3-30TB per day! That is quite large → Better estimates?
  - If true, we need to consider if we do central data processing pipeline
    - I.e. a image processing to extract key physics and systematic information directly after acquisition
    - Can we have algorithms fast enough to keep up
      - Including non-running time for processing, need to process each image in ~2 seconds on one processor, or 1 image per minute on a 30 processor system

MAGIS Questions: Calibration Mode

- What is our real data rate? Calibration mode:
  - O(30) 6-megapixel images, say 10MB each.. But only readout 1 station: 2 images → 20MB per drop + metadata (time, configuration, etc… ignore for now)
  - Drop rate 1 per minute → assume slower for calibration process
  - Running time per day: Assume it takes a 16hr day?
  - Calib data $= \frac{20\ MB}{drop} * \frac{1\ drop}{min} * \frac{60\ min}{hr} * \frac{16\ hr}{day} = 20\ \frac{GB}{day}$
    - Manageable size, but will we keep this data? Will we want it later to check calibrations?
    - Maybe drop rate is much lower? Or Faster?
    - Maybe it doesn't take a day?
  - Even considering a single shot, we need to extract information from cameras for calibration feedback. how fast can we do it?
    - Can we process 20MB data in 1 minute?
      - Probably don't need too many processors for this
      - Fits + feedback to lasers & drop system must all be done

Important questions, if the data rate is close to this scale

- Do we need a data processing pipeline to extract key information before data analysis?
  - Maybe we don't need it in year 1, but we will by year 2 at these rates
  - Can algorithms keep up with data rate?
  - Central reprocessing needed
  - How big a system do we need? CPU or GPU?
  - How can users contribute?
  - Community SW development standards needed
  - Helps standardize physics parameter extraction for images
- RAW data may be to large for every user to store, do we need central processing facilities? How big does it need to be?
- If we have a central processing, may still need a central analysis facility
  - If no central processing, seems like analysis facility is critical
- Do we need a data analysis framework?
  - To work on processed information?
  - To work on raw images?
- Do we need to develop a calibration pipeline, for use in real-time?
  - What size system do we need? Can we use the same system as data processing pipeline?
  - What real-time-feedback framework do we need to develop?

Next Steps

- This came from a short conversation with Phil Marshall and from reading: https://www.lsst.org/about/dm
  - Phil is not the expert on this, but it was clear I did not have many of the answers needed to help define a system

- The "Architect" of LSST data management system is KT Lim, a SLAC employee
  - I would like a little help on the questions before speaking to KT Lim
  - Then I would like to start the conversation with him to start to get a more detailed understanding together.

- LSST / Ruben Observatory seems like a good model to start from
  - Need better understanding of our real data rates
  - Need to understand what pipelines, hardware, and frameworks needed