


The LSST Data Management System



J. Kantor, T. Axelrod (LSST Corp) and the LSST Data Management Team

The LSST scientific requirements for a very deep, wide field survey with controlled systematics and rapid transient alerting drive the LSST Data Management system capabilities and performance. This system is a highly distributed, scalable, and reliable system that acquires the scientific data, transports it from mountain to base to archive center, reduces the data and assesses its quality, and publishes the data for broad access. Designing and implementing such a system requires advanced scientific algorithms, a layered architecture providing for extensibility and technology evolution, high-performance computing, storage, and networking technologies, and the rigor of a formal development process.

Scientific and Technological Evolution

Scientific Evolution

Depicted below are the Sloan Digital Sky Survey (SDSS) instrument, and same image processed with different algorithm/parameters by the SDSS image processing pipeline. SDSS went through many iterations of software/algorithms after first light before achieving the "optimal" PSF variation estimation.



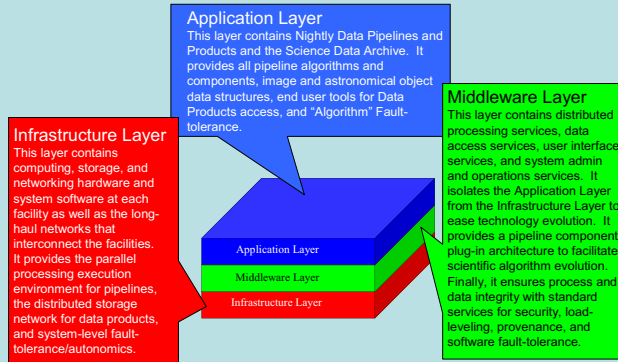
Images courtesy of Sloan Digital Sky Survey and International Business Machines Corporation. Unauthorized use not permitted.

Technological Evolution


If the useful life of the LSST Data Products is at least two decades, the raw data will be re-processed at least 20 times and the computers and software this is done by will be changed at least 4 times!

LSST will be updating software and data very frequently. This must be a largely automated process and highly documented as it occurs. We cannot afford to rewrite the entire software at one time or port to entirely new hardware en masse, so evolutionary change must be supported.

Layered Architecture



Estimated Nightly Data Volume



LSST Estimated Nightly Data Volumes

The nightly data volumes generated by LSST will be an order of magnitude larger than those estimated for Pan-STARRS 4 and 2 orders of magnitude larger than SDSS. The data archive will grow at a rate of roughly 7 PB/yr. This requires scalability and reliability in LSST data management systems.

Infrastructure Layer

Long-Haul Communications

Mountain/Base to Archive and Archive to Data Centers Networks are 2 - 10 Gbps protected clear channel fiber optics with protocols optimized for bulk data transfer

Archive/Data Access Centers

In the United States. Nightly Data Pipelines and Data Products and the Science Data Archive are hosted here. Supercomputers capable of 60 TFLOPS provide analytical processing, re-processing, and community data access via Virtual Observatory interfaces to a 7 PB/yr archive.



Base Facility

In Chile, Mexico, or the United States. Nightly Data Pipelines and Products are hosted here on 25 TFLOPS class supercomputers to provide primary data reduction and transient alert generation in under 60 seconds.


Mountain Site

In Chile or Mexico. Data acquisition from the Camera Subsystem and the Observatory Control System, with read-out in 2 seconds and data transfer to the Base at 10 Gbps

Application Layer

Nightly Pipelines and Data Products

Nightly Pipelines are executed and Data Products are produced within 60 seconds of the second exposure of each visit.



Science Data Archive

These pipelines are executed on a slower cadence and the corresponding data products are those that require extensive computation and many observations for their production.

LSST Data Management Design & Development Process

Data Management follows a rigorous, proven software/system engineering methodology based on the Iconix Process®, combined with extensive prototyping. This ensures that the design is validated and end-to-end requirements are clearly traceable all the way to code.





Image courtesy of Iconix Software Engineering, Inc. Unauthorized use not permitted.

Middleware Layer



Distributed Processing Services

The Pipeline Manager is a process that handles overall control of executing pipelines. The Event System is for publishing & distributing asynchronous messages between processes and to system logs. The Pipeline Configurator is a service that creates an instance of a pipeline from a template. A Pipeline is an instance of a Pipeline class that carries out a logical set of data processing (e.g. Processing Nightly Observations).



Data Access Services

The Data Stager copies/organizes data & software in preparation for and clean up from a pipeline execution. The Data Replicator is an inter-site data mirroring system. The Database API is an interface for inserting/extracting data to/from the database. The Ingest Service is a service that formally incorporates data into the archive, including replicating data to long-term storage and inserting associated metadata in the database



The LSST research and development effort is funded in part by the National Science Foundation under Scientific Program Order No. 9 (AST-0551161) through Cooperative Agreement AST-0132798. Additional funding comes from private donations, in-kind support at Department of Energy laboratories and other LSSTC Institutional Members.

National Optical Astronomy Observatory
Research Corporation
The University of Arizona
University of Washington

Brookhaven National Laboratory
Harvard-Smithsonian Center for Astrophysics
Johns Hopkins University
Las Cumbres Observatory, Inc.

Lawrence Livermore National Laboratory
Stanford Linear Accelerator Center
Stanford University
The Pennsylvania State University

University of California, Davis
University of Illinois at Urbana-Champaign