LSST Data Management Infrastructure

*D. Dossa (LLNL), C. Matarazzo (LLNL), S. Marshall (SLAC),
C. Smith (NOAO), R. Lambert (NOAO), M. Butler (NCSA),
C. Cribbs (NCSA), R. Plante (NCSA), D. Sweeney, J. Kantor (LSST)*

Large Synoptic Survey Telescope

A possible computing infrastructure that can support the extremely high LSST data rates, aggressive alert times, and data archive needs are described for the telescope site, base camp, and data archive center. A common computing architecture was chosen for all 3 locations with the systems differing principally in size. Each computing system has redundant computing nodes and sufficient disk space to buffer data in case of network outages. The computing infrastructure design will evolve with advancing technology.

LSST data rates are much higher than any other existing or planned optical telescopes. The camera will deliver data to the Data Management system at a rate of 6 Gbytes every 20 seconds. The computing systems must process and archive raw images, including calibrations, at a rate of 7 Pbyte/year. The image catalog will grow by 0.3 Pbyte/year. To support these data rates, the computing infrastructure will be composed of computing systems at the telescope, the base camp, and a data archive site. If both the base camp and data archive computing systems were built today, both would be among the 10 fastest computers systems in the world. When the systems are built in 2012, they are still likely to be among the 1,000 fastest computers in the world. The final design for the computing infrastructure will be defined about the time telescope construction begins in order to allow sufficient time for software development and debug activities to be completed by first light.

Telescope Site Computing

The raw image data is sent from the camera to data acquisition nodes. Metadata from the Observatory and Telescope Control Systems are associated with the appropriate raw image by the computing system at the telescope site and transmitted to the base camp during the following exposure. This cluster will be approximately 40 nodes on an Infiniband switch, with 10 1 Gigabit Ethernet lines to the base camp, which depending on final site selection, will be located between 80 km to 200 km away. The telescope system will have sufficient disk space to completely buffer 2 nights worth of data in case of a network failure.

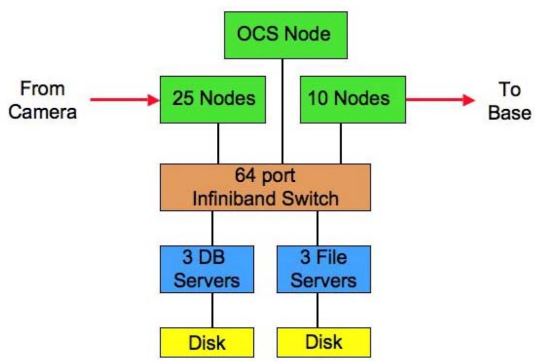

Base Camp Computing

The base camp computing system must perform all of the time critical calculations from image calibration, quality assessment, object detection/differencing, and moving object detection. Alerts will be automatically generated by the base camp computers. All of this work must be done during an image exposure time in order to not fall behind the incoming data. During the following day, the entire night's data will be transmitted to the data archive site in the United States for further processing, including non-time critical pipelines.

This system will be a large Infiniband-based cluster of 1800 CPUs used for running the pipelines, and 128 CPUs each to support the data base server and image file servers. The system storage contains the catalogs and template images for differencing, moving object ephemerides, and sufficient disk space to store 4 nights of raw images in case of a failure of the network connection to the archive site.

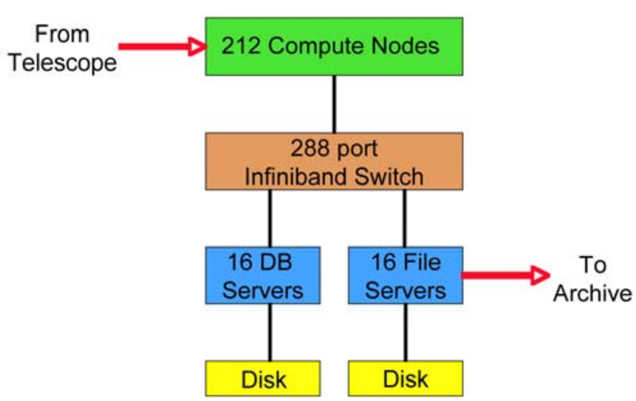

Archive Center Computing

The archive center will have a system built from the same components as the base but approximately 3 times the size. This is done to allow for additional processing pipelines, such as weak lensing and re-processing of the entire data set each year. The archive computing system is capable of re-processing multiple years worth of data in a reasonable time while keeping up with the incoming data rate. Since the entire 10 years worth of image data will be maintained here, the archive site will have a 3 tiered storage structure of fast disk for the most recent and most used data, a secondary store of less expensive disks for frequently requested data, and a large tape library for all remaining data. Over the 10 year life of LSST, this data will grow to at least 100 Pbytes.

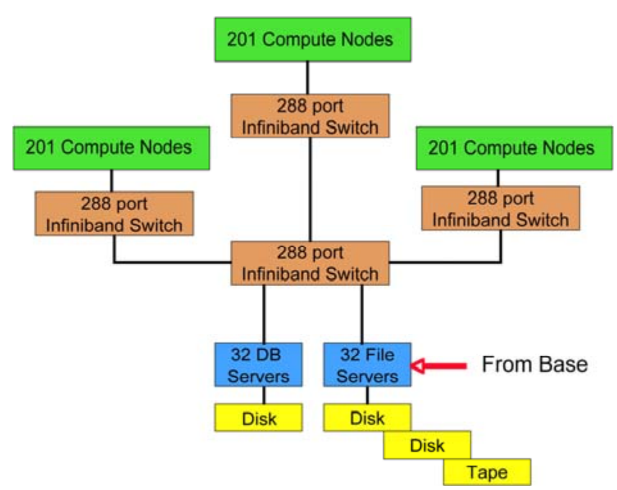

The LSST research and development effort is funded in part by the National Science Foundation under Scientific Program Order No. 9 (AST-0551161) through Cooperative Agreement AST-0132798.
Additional funding comes from private donations, in-kind support at Department of Energy laboratories and other LSSTC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