# LSST and Astronomy Data in 2020

*A.S. Szalay (JHU), T. Axelrod (UA), J. Gray (Microsoft), R.H. Lupton (PU), R.Pike (Google)*

**Astronomy, like High Energy Physics, will bring GB/sec instruments on line within the next decade. As LHC shows, both technical and political forces imply a multi-tiered architecture in which all data is replicated for reliability and performance at several Tier-1 centers, and selected extracts will be stored at many Tier-2 sites that will subscribe to subsets of the data stream and derived products. LSST will be just one of several peta-scale astronomical observatories coming online with these data volumes. This suggests a world-wide federation of about 10 peta-scale sites, each associated with an LSST-class instrument, interconnected with high-speed networking, forming the core of the Virtual Observatory and each providing replication services for others in the federation. Unlike in HEP, there is great interest in cross-correlating astronomy data from multiple instruments and telescopes. Users need easy access to data from all sites and some will want to build their own Tier-2 archives. The real-time nature of LSST (and probably other instruments) forces substantial archive and processing facilities co-located with the telescopes.**

### Challenges—Size

We estimate the LSST will generate 5 petabytes of raw data every year (25% duty cycle). The on-line data rate from the detectors will be about 3 gigabytes/sec peak rate. This will be comparable to other experiments of the next decade, ranging from High-Energy Physics (LHC) to Oceanography and Remote Sensing. Efficient data management and access will be an essential part of these projects.

### Challenges—Data Transfer

One challenge will be moving these vast datasets around, largely due to the cost and speed of wide area networks. Today it would cost millions of dollars per year to support a reliable transcontinental data pipe with a throughput of 3 GB/sec over commercial networks. Although network speeds will increase and costs will decline, data transfer at these volumes will likely remain expensive. Therefore the architecture needs to minimize and optimize bulk data transfers in an end-to-end system.

### Challenges— Real-Time Detection

Many of the interesting LSST science problems involve transient phenomena. As a result, the data management must be an almost real-time system. Image quality will be adjusted, transients will detected, and an alert triggered within 30-60 seconds. This requires that the data stream coming off the telescope be processed almost instantaneously, presumably at close proximity to the telescope. The data stream needs to be compared to previous observations of the same area of the sky both to assess image quality and to detect transients. Thus the necessary archival data (the co-added sky) must be available in sync with the pointing of the telescope. The fast response also puts a heavy demand on the data archive, which must be able to ingest and index the incoming data and reprocess all the data at least once a year as the algorithms evolve.

### Challenges—User Access

Another challenge is to provide access to the wide spectrum of users who want to search and analyze the data. We expect that much of the analysis will be performed at science centers collocated with the peta-scale archives, in order to avoid the costs associated with network transfer. Much of the computing will be done on farms of compute bricks closely tied to (or overlapping with) a farm of data storage bricks.

### Challenges— Databases or Flat Files?

There is a lot of discussion whether the LSST data management should use flat files or standard databases or build a special purpose data management system. The commercial database systems and file systems are converging over the next decade: companies are adding file-system features to their databases and database-like features to their file systems. LSST expects to benefit from this convergence and to be able to use these systems to manage its multi-petabyte store and to define spatial, temporal, and associative indices over the data.

### Challenges— Fault Tolerance and Data Protection

One of the main issues the LSST (and many other peta-scale experiments) will have to address is fault tolerance. An archive of many petabytes will involve tens of thousands of disks, even if disks are terabyte scale. Since disk technologies will remain fundamentally similar, we can expect to see disk failures daily. The system needs to be able to recover from these without human intervention. Geoplexing datacenters should simplify this process – providing both disk fault-tolerance and also tolerance of site failures and catastrophes. The LHC experiments have already adopted a hierarchical, multi-tiered architecture for their data management (GriPhyN), where all the data coming off the instruments is stored next to the detectors in a Tier0 facility, then replicated to one or more Tier1 sites. The geoplexing of the data helps recovery from major disasters, also provides a relief on transatlantic networks. Various subsets of the data will then be filtered, processed and stored at various Tier2 sites, typically located at a research university. Such a hierarchical arrangement provides a nice impedance match between the different types of users and the data archives.

### Hierarchical Data Usage

We expect that most science queries to the data will follow relatively simple patterns, which can be answered with a small subset of the data (aggregated time-series, co-added object catalog, etc). If 90% of the users only touch 1% of the data, it may be sensible to extract these "popular" data products and provide a very fast access by replicating the data many times in server farms at Tier2 centers. A smaller number of more sophisticated users will want to perform more elaborate analyses of the data which still can be answered from no more than 10% of the data. Such data sets can be replicated to Tier-2 centers, distributed across the collaboration, according to local expertise. User analyses would be performed right on top of the data.

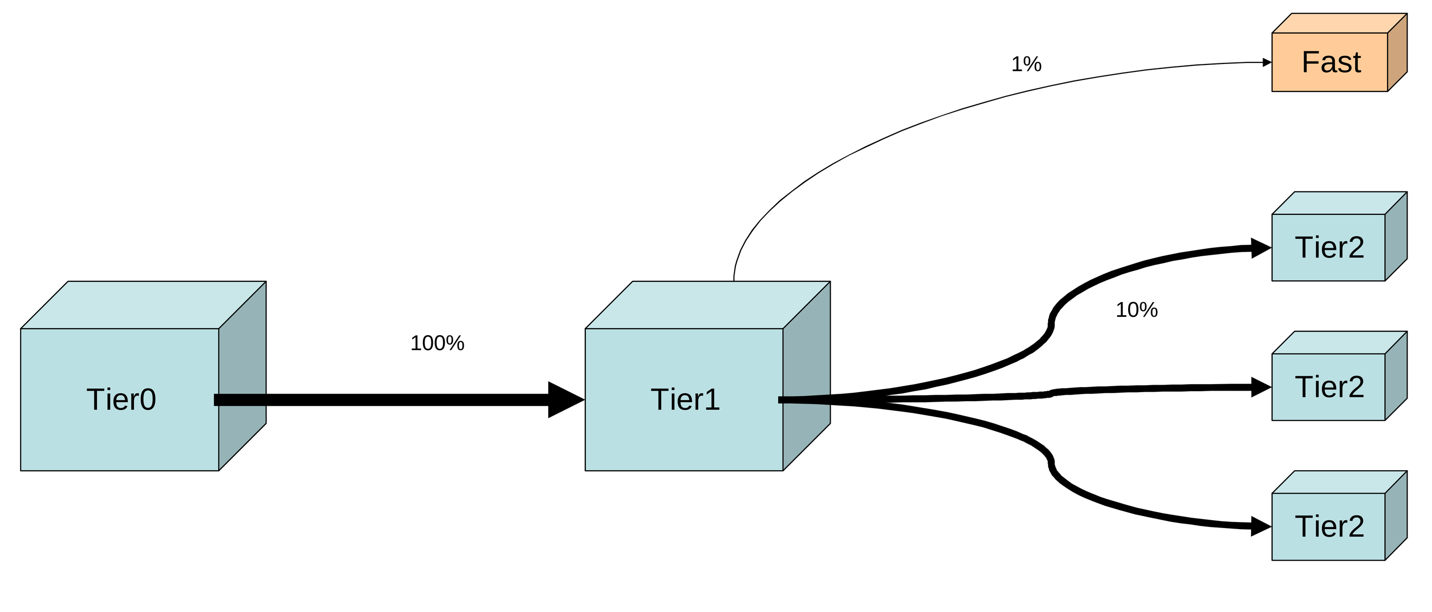


**Figure 1.** The LSST conceptual organization of the LSST data federation. Data are replicated to Tier1 sites, and then "popular" data product subscriptions are published to Tier2 sites. Both Tier1 and Tier2 sites have enough computational facilities to support user queries and data analysis.

### Scaling to 2012—Data Size

The size of the LSST data set appears overwhelming today, but in perspective it should be manageable. In 2001 the aggregate volume of all astronomical data was estimated to be at around 100 Terabytes, doubling every 18 months, roughly paralleling Moore's law. At this rate the projected data volume in 2012 would be about 30 petabytes. The current doubling is due to the fact that the astronomy community is spending an approximately flat budget related to data acquisition. Every new instrument that comes on line has an instrument growing by Moore's law. Once an instrument is built, it acquires data at a linear rate. The LSST will generate about 5 petabytes per year, about 1/6th of the total data available in 2012.

### Scaling to 2012—Number of Pixels

In 2000 the total number of pixels available for astronomy was at about 800 megapixels, doubling every 2.25 years. If this trend continued, in 2012 the total number would be 20 gigapixels. The LSST camera with its 3 gigapixels fits this trend very well. It is about 1/6th of the total, not an overwhelmingly large fraction, indicating little technological risk. This fraction of the total is similar to what the SDSS camera was in 2000, and similar to the initial fraction of the data size.

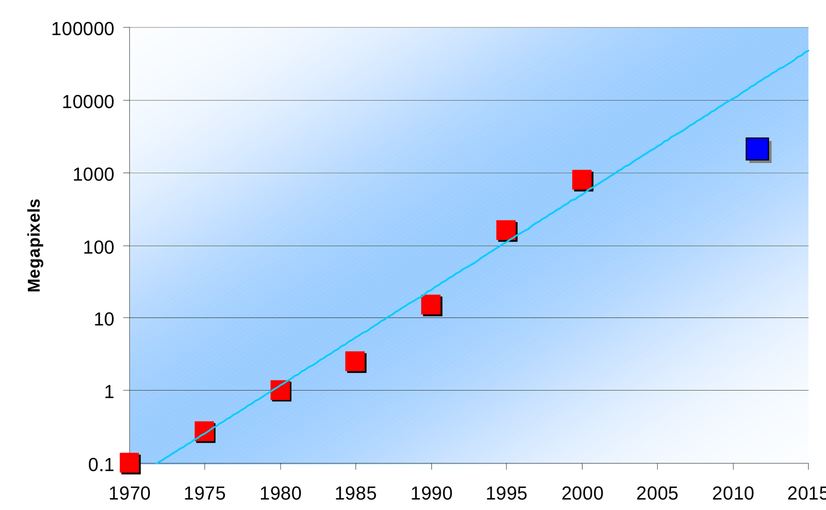


**Figure 2.** The projected number of pixels between 2010 and 2015. The red marks indicate existing data, the blue mark represents LSST. The light blue line corresponds to a doubling in 2.25 years.

### Data Inflation

The LSST will produce approximately 5PB/y of raw image data (uncompressed), or 2 PB/y compressed. In Physics (e.g. BaBar, LHC,…) the instrument simulation data generates 2x this raw data. In Astronomy, based on our experience with the SDSS, the derived data products (images, catalogs, derived datasets) are about the same size as the base data once the software is fully developed – we call this "data inflation". So we expect the actual LSST demand to be about 4PB/y per Tier1 site when all data are included.

### LSST and the Virtual Observatory

Today the Virtual Observatory is emerging as a way to cope with the explosion of these new data sets. Unlike in High Energy Physics, astronomers are interested in cross-correlating data acquired in different wavelengths and at different times, since the telescopes point to the same regions on the sky. In today's world there are a few large catalogs (2MASS, USNO, SDSS) but most of the other collections are rather small. At the same time, the data doubling is coming from many new instruments and telescopes coming on line. In the next decade the Virtual Observatory will be increasingly dominated by a handful of "peta-centers", of which the LSST will be one. We believe that these peta-centers will replicate all the small (less than a petabyte) datasets locally and only need to go "outside" for new data and for data from other peta-centers. Most likely these centers will be connected to each other via high speed links, and will form a geoplexing alliance replicating each other's data for redundancy.

**Figure 3.** LSST Data Processing Infrastructure – circa 2012