4 LSST REFERENCE DESIGN

4.1 THE LSST OBSERVATORY

The three subsystem teams have taken the technical specifications outlined in the previous chapter and generated the LSST baseline design. The baseline design is a flexible observatory system that we could build today. It will perform all of the science reference missions presented in Chapter 2.

The baseline design provides a basis for the Project Execution Plan presented in Chapter 7. We will continue to improve the design during the remainder of the design and development phase. At the same time, however, we are committed to holding mission scope constant and remaining within the current cost envelope while we work to enhance performance.

The LSST Observatory "complex" is distributed over four sites: the Summit Facility, the Base Facility, the Archive Center, and the Data Centers. Table 4-1 lists the location, features, and function of each site, the LSST team responsible for the subsystem components at each site, and the numbered sections of this Chapter 4 in which those subsystems are discussed.

Although the four facilities are distributed geographically, they are functionally connected via dedicated high-bandwidth fiber optic links. Real-time transient alerts will be broadcast over the Internet within one minute of data collection. Digitally processed images and feature catalogs will be available to the U.S. community via the Internet within 24 hours. All four observatory sites are integrally related and functionally part of this process.

The present chapter begins with a description of the observatory site (§ 4.1.1) and optical design of the telescope (§ 4.2). Subsequent sections deal with the three major subsystems: the telescope (§ 4.3), camera (§ 4.4), and data management (§ 4.5) systems. Chapter 5 describes system engineering, calibration and integration strategies, and commissioning plans. Chapters 4 and 5 together represent the technical description of all LSST equipment and facilities proposed for funding collectively by NSF, DOE, and private sources.

Site	Location	Features/Functions	Subsystem Team (Ch. §)
Summit Facility	El Peñón Peak on Cerro Pachón ridge, Chile	Telescope; service/support bldgs; calibration telescope Camera	Telescope & Site (§4.3) Camera (§4.4)
Base Facility	La Serena, Chile	Operations Center Pipeline computing facility; backup data system	Telescope & Site (§4.3) Data Management (§4.5)
Archive Center	NCSA, U of Illinois, Urbana-Champaign	Main computing center and data storage; data base archive	Data Management (§4.5)
Data Access Centers	Various locations in US and Chile	Data buffering and distribution centers	Data Management (§4.5)

TABLE 4-1 *The LSST Observatory Sites*

4.1.1 Observatory Site

The site chosen for the LSST observatory is El Peñón peak on the Cerro Pachón ridge in Northern Chile (Figure 4-1). El Peñón is located about 500 km north of Santiago and 57 km southeast of La Serena on property owned by AURA. The insert in Figure 4-1 shows the Cerro

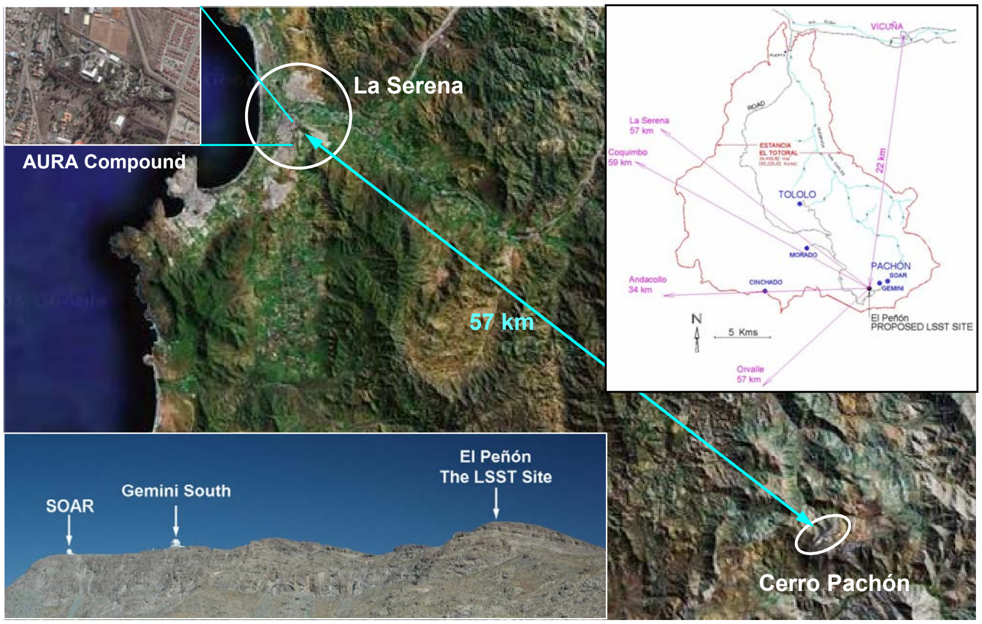

FIGURE 4-1 A satellite view showing the relation between Cerro Pachón and La Serena. The upper right panel shows distances to nearby towns and cities from the El Peñón site within the AURA property boundary (red). The inset (lower left) shows the Cerro Pachón ridgeline viewed looking southeast showing the LSST site at El Peñón in relation to SOAR and Gemini South.

Pachón ridge with El Peñón peak and the observatories already in operation there—the 8.2-m Gemini South and 4.3-m SOAR telescopes. The base facility for LSST operations in Chile will be located in La Serena within the AURA-operated compound.

The LSST project selected this site following a multi-year site selection campaign. The process was administered and the final recommendation made by an international committee of recognized experts. (The final report is available as LSST Document-1796.) Because we judged that several existing sites would meet LSST's requirements and wished to constrain costs, consideration was limited to developed sites. After reviewing extensive data sets from sites around the world and in both hemispheres, the site selection committee recommended that we narrow the search to semi-finalist sites in Spain, Mexico, and Chile. A campaign of data collection and further review of both scientific and logistical factors was conducted for these sites during a second year of analysis. Finally, the committee recommended the El Peñón site. The LSST Board of Directors accepted this recommendation in May 2006.

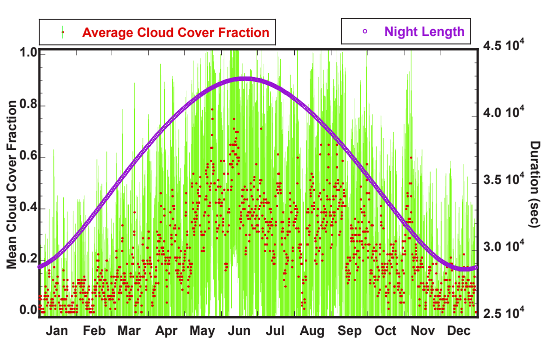

FIGURE 4-2 The daily 10-year mean (red) and standard deviation (green) of cloud cover fraction spanning 1995-2005 from the CTIO observatory logs. The amount of usable observing time between nautical twilight is also shown (purple).

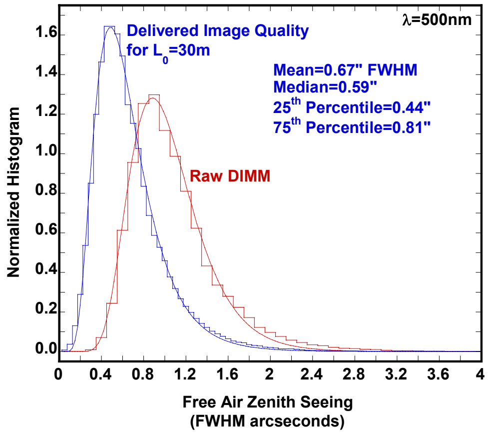

FIGURE 4-3 The distribution of measured Differential Image Motion Monitor (DIMM) seeing (red) and delivered image quality (blue) from Cerro Pachón

The excellent observing conditions at Cerro Pachón satisfy the LSST survey requirements. The weather patterns on Cerro Pachón are nearly identical to those on CTIO 10 km away. The average daily cloud cover fraction from the CTIO observatory logs for the past ten years is shown in Figure 4-2. These data show greater than 80% useable nights with excellent atmospheric conditions on Cerro Pachón. Differential image motion monitoring (DIMM) measurements made on Cerro Pachón show that the expected mean delivered image quality is 0.67 arcseconds (Figure 4-3). This image quality has been confirmed with observations from Gemini and SOAR. The data are also supported by measurements and logs from neighboring Cerro Tololo, where the records extend back many years.

Management of Cerro Pachón and neighboring Cerro Tololo is provided by AURA under contract to the NSF. AURA has the procedures in place to accommodate tenant telescope organizations and maintains the necessary agreements with the government of Chile. Through agreements already in place with AURA, LSST will enjoy official international organization status with tax and duty exemptions.

4.2 Optical System Design

The LSST optical design shown in Figure 4-4 is a modified Paul-Baker 3-mirror system. Conceptually, it is a generalization of the well-known Mersenne-Schmidt family of designs that produces large fields of view with excellent image quality (Wilstrop 1984; Angel et al. 2000; Seppala et al. 2002). The LSST étendue (including the effects of camera vignetting) is 319 m^2deg^2.

The effective focal length of the optical system is 10.3 m, making the final f/number 1.23. The plate scale is 50 microns per arcsecond at the focal surface. This choice of effective focal length represents an optimum balance of image sampling, overall system throughput, and manufacturing feasibility. The annular geometry of M1 produces 35 m^2 of on-axis collecting area, equivalent to a 6.7-m diameter unobscured clear aperture.

The primary mirror (M1) is 8.4 m in diameter with a 5.1-m inner clear aperture to accommodate the 5-m diameter tertiary mirror (M3). The relative positions of M1 and M3 were adjusted during the design process so that their surfaces meet with no axial discontinuity at a cusp, allowing M1 and M3 to be fabricated from a single substrate.

The 3.4 m convex secondary mirror (M2) has a 1.8-m inner opening. The LSST camera is inserted through this opening in order to access the focal surface.

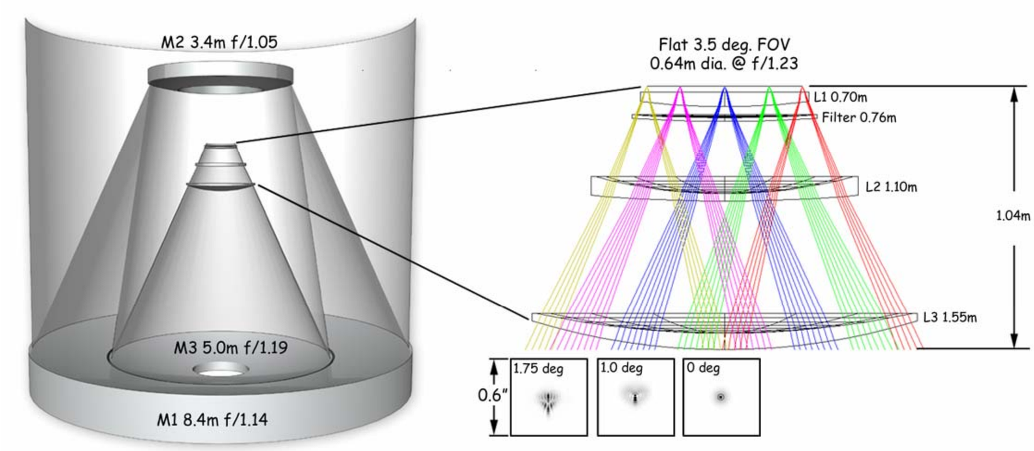

FIGURE 4-4 The optical design configuration showing the telescope (left) and camera (right) layouts. Diffraction images for three field radii, 0, 1.0 and 1.75 degrees, are shown in boxes 0.6 arcseconds square (3×3 pixels).

The three reflecting mirrors are followed by a 3-element refractive system that corrects field flatness and chromatic aberrations introduced by the filter and vacuum window. The 3.5-degree field of view (FOV) covers a 64-cm diameter flat focal surface. Spectral filters reside between the second and third refractive lens as shown in the right side of Figure 4-4.

The image brightness is constant to a field radius of 1.2 degrees and gradually decreases afterward by about 10% at the 1.75-degree field edge. The intrinsic image quality from this design is excellent. Figure 4-5 shows that the 80% encircled energy is <0.3” in all spectral bands and <0.2” in the *r* and *i* spectral bands. The design also has very low geometrical distortion, <0.1% over the full FOV, making the LSST an excellent system for positional astrometry.

There are five aspheric surfaces in the optical design: each of the three mirror surfaces and one surface each on two of the camera lenses. The asphericity on the two concave surfaces of M1 and M3 are well within standard fabrication methods used for astronomical mirrors. During

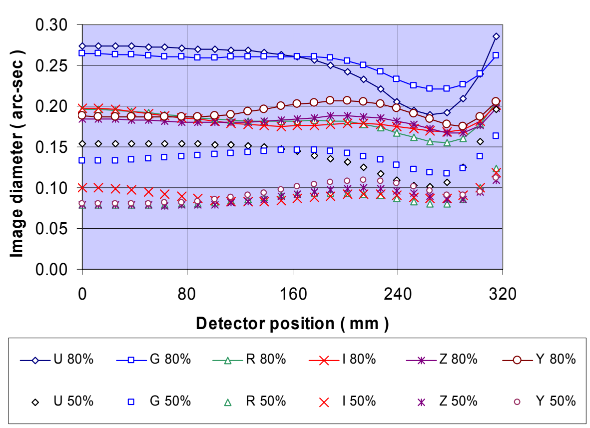

FIGURE 4-5 The image quality of the optical design versus field angle as measured by the 50% and the 80% diffraction encircled energy.

the optimization process, the asphericity of M2 was minimized to 18.9 microns of departure from the best-fit sphere in order to reduce technical challenge for this optic. The three fused-silica refractive elements, which have clear apertures of 1.55-m, 1.10-m, and 0.70-m, while large, do not present any particular challenge in their fabrication.

The 0.76-m diameter spectral filter is located just prior to L3. The filter thickness varies from 13.5–26.2 mm depending on the choice of spectral band, and is used to maintain the balance of lateral chromatic aberration. The zero-power meniscus shape of the filters keeps the filter surface perpendicular to the chief ray over the full field of view. This feature minimizes shifting of the spectral band wavelength with field angle. The last refractive element, L3, is used as the vacuum barrier to the detector cryostat. The central thickness of L3 is 60 mm to ensure a comfortable safety margin in supporting the vacuum stresses. The individual optical elements are discussed further in section 4.4.

4.2.1 Alignment and Surface Figure Control

A critical aspect of any large optical telescope design is the ability to sense and align the system to achieve the theoretical performance, given reasonable fabrication, assembly, and operational (i.e., gravity, thermal, and wind loading) tolerances. We have used Code-V and Monte Carlo techniques to evaluate the alignment and correctability of this particular design. These simulations also generate error budget allocations, system and component tolerances, and validation of the compensation strategy.

The LSST optical design is described by 77 parameters defining radii, conic constants, aspheric terms, thicknesses, and rigid-body displacements (LSST Document-1361). We have found that with the correct static and dynamic compensators, the LSST system is in fact quite tolerant. Of the 77 parameters, the system must be assembled with static control (i.e., one-time) over two axial element spacings and dynamic control (i.e., control during operation) over the secondary and camera assembly in 5 degrees of freedom each (X, Y, Z, θ_X, and θ_Y). In addition, the figures of all three mirror surfaces are actively controlled.

Figure 4-6 shows a typical result of the Monte Carlo alignment exercise. The optical system is initially uncorrected (i.e., misaligned) with an 80% encircled energy of 20 arcseconds. Using

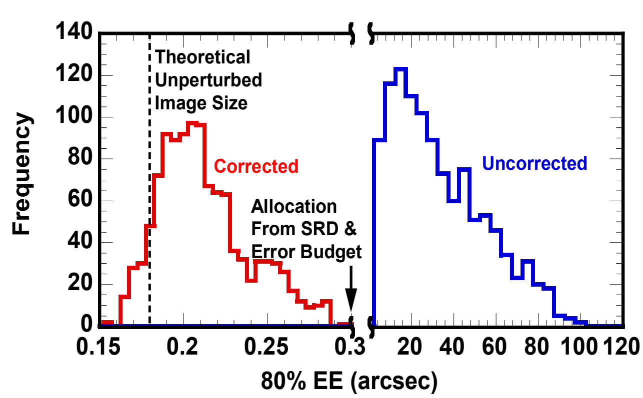

FIGURE 4-6 Distribution of the RMS 80% encircled energy (EE) in the i-band of the perturbed optical design before correction (right, blue) and after compensation (left, red). Note the change in scale by two orders of magnitude on the abscissa.

only the compensators discussed above, the system is corrected (i.e., aligned) to a median encircled energy of 0.2 arcseconds. This is comfortably below its allocation in the error budget (0.3 arcseconds for the system).

Control of PSF ellipticity is of vital importance in controlling systematic errors for the science missions. Ellipticity is tracked in the simulations referenced to 0.6″ arcsec atmospheric seeing by convolving the diffraction image with a 2-D Gaussian having a FWHM equal to the seeing. Figure 4-7 shows the cumulative ellipticity probability distribution in the i-band. For an optical system aligned with the compensators only, the residual ellipticity introduced by the optical system is well below the SRD specifications. Specifically, the SRD requires that 50% of the PSF signatures have a residual ellipticity of 0.04 or less due to all effects. The optics alone introduce an insignificant contribution to ellipticity.

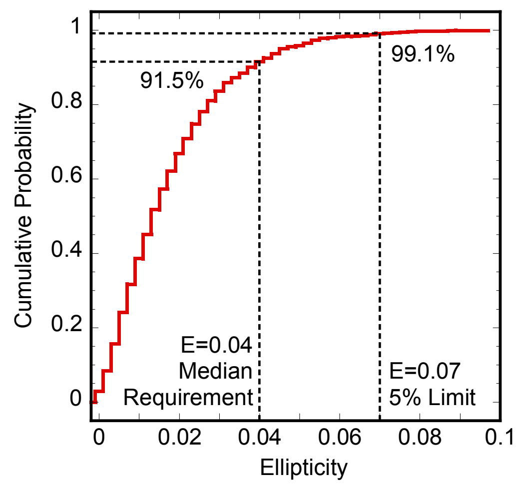

FIGURE 4-7 Cumulative *i*-band ellipticity probability referenced to 0.6″ seeing from Monte Carlo simulations of the corrected perturbed optical design.

We have also analyzed the optical design to determine how to sense the rigid body displacements and the mirror surface figure adjustments necessary to correct the imaging performance. We have shown that the errors can be determined unambiguously from four wavefront sensors positioned 90 degrees apart on the outer edge of the FOV. The corrective actions required for alignment and mirror figure control are described in Section 4.3.

4.2.2 Do We Need Atmospheric Dispersion Correction?

A critical question considered during development of the LSST optical design was whether to include an atmospheric dispersion corrector (ADC). An ADC entails additional optical elements and complexity in the camera subsystem, as well as the potential for reduced intrinsic image quality in the optical system. Without an ADC, image elongation in the direction of the parallactic angle results in a certain amount of PSF ellipticity. Our assessment of the need for an ADC involved the following considerations:

- Available survey area versus zenith angle and observatory latitude
- Dispersion-induced PSF ellipticity versus zenith angle and filter specifications
- Feasibility of optical design and image quality
- Impact on primary scientific methods for weak lensing and optical transient detection

From Cerro Pachón, the required 20000 square degrees of survey area are viewable above a zenith angle of 45 degrees. Using the Filippenko (1982) dispersion model, we find that the ellipticity error budget in the worst case for the filters we have adopted reaches its limiting requirement of e = 0.04 in 0.6 arcsec seeing at a zenith angle of 43 and 66 degrees in the *r* and *i* spectral bands, respectively (Figure 4-8). Just ~250 square degrees, i.e., ~1.25% of the survey area, would exceed the mean ellipticity requirement in the *r*-band, which is acceptable within the requirements definition on PSF ellipticity. The *g*-band image specifications have no constraint on PSF ellipticity.

During optimization of the optical design, one iteration included an ADC using available glass types, PSK3 and LLF6. This design showed that while an ADC was feasible, it would come at the expense of image quality. Analysis of the methods used in weak lensing analysis showed that the image degradation from an ADC more than offset the small amount of ellipticity caused by atmospheric dispersion. As a result of these analyses, we determined that the LSST reference design will *not* include an ADC.

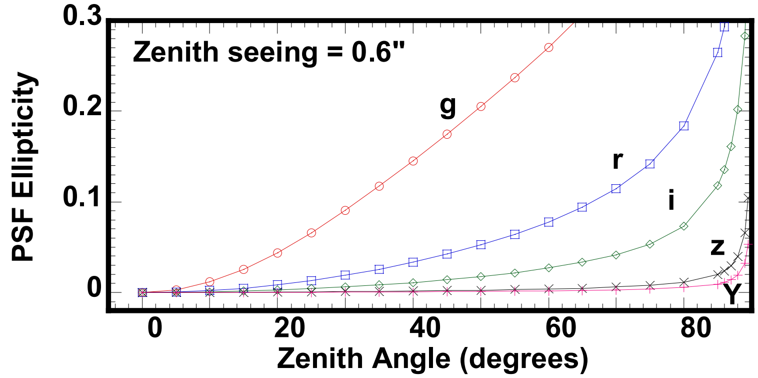

FIGURE 4-8 Expected PSF ellipticity in each of the LSST spectral bands as a function of zenith angle.

4.3 TELESCOPE AND SITE

The telescope capable of conducting the science missions described in Chapter 2 of this proposal must meet exacting requirements in terms of image quality, slew and settle time, and reliability. While these requirements have imposed significant engineering challenges, there is no doubt that the baseline design that has been developed can be realized with today's technology, expertise, and experience (Figure 4-9).

The aperture of the LSST and its optical field of view (FOV) combine to provide an étendue of 319 m^2deg^2. Each of the telescope's three mirrors—the 8.4-m primary, the 3.4-m secondary, and the 5-m tertiary—is well within current state of the art capability for optical fabrication. The proximity of the primary and tertiary surfaces enables fabrication of both surfaces into a single substrate. The result is a unique, monolithic mirror that will simplify alignment corrections during operation through a one-time factory alignment using existing metrology techniques. The secondary mirror, though traditional in design with a minimal aspheric departure of 19 microns, will be, to our knowledge, the largest convex mirror ever built.

Maintaining a continuous survey cadence over a ten-year period places a high premium on reliability and duty cycle. The LSST must be a very agile system, taking exposure after exposure, re-targeting roughly twice per minute all night long. This requirement puts an uncommon demand on observing automation, system coordination, and process efficiency. The relentless minute after minute, night after night, and year after year duty cycle of exposures for data collection puts a significant premium on the performance of the drive system (Neill and Sebag 2006c) and on minimizing the system slew and settle time. The LSST design addresses these demands by combining previous telescope experience with good industrial automation and life cycle engineering.

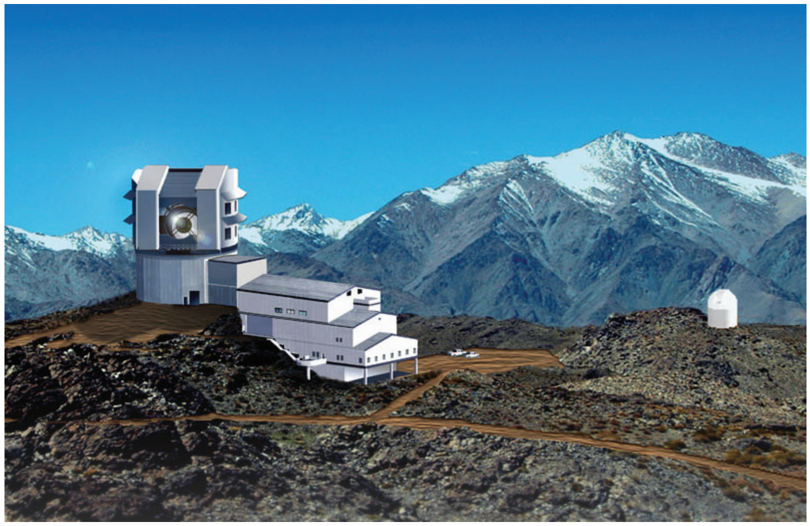

FIGURE 4-9 *Artist's rendering of the LSST and ome enclosure with the attached summit support building and the LSST calibration telescope shown on an adjacent peak.*

Image quality requirements call for an intrinsic 0.3 arcsec *r*-band performance across the full 3.5-degree diameter FOV with unprecedented control of PSF shapes and overall systematic errors. In order to achieve this performance, the LSST design incorporates robust wavefront quality assessment tools, efficient active optics control systems, an independent telescope alignment system, and facility support that provides daytime active cooling and significant natural ventilation during observing.

The current size of the telescope and site team—relatively small during the design and development phase—will increase to an estimated 25 telescope professionals for the construction phase. This team will focus on telescope system engineering and prime contracting, working with industrial vendors to produce the necessary subsystems. Systems will be procured in packages that align with industrial capability and experience, thereby limiting technical and cost risks. In-house design efforts will focus on the active optics system (AOS) and control software development; the project team will also be responsible for the integration of elements on site, commissioning the telescope, and handing over duties to an operations staff. Those team members hired late in the project for on-site support will be expected to remain with the LSST during operation. Staff continuity is the most efficient way to transition from project development to the operations phase of the project.

A top-level breakout of construction costs (Figure 4-10) shows that the largest fractions of construction costs are allocated to, respectively: the two mirrors (33%), the telescope mount (22%), and the facilities (13%).[1] The design, structure, and operational subsystems of the fully realized LSST described below incorporate and build on the lessons learned in the construction of major optical/infrared facilities during the last decade. Though today's telescopes do not have the étendue of the LSST, the experience, prototypes, and example of the predecessor projects—the Very Large Telescope (VLT), the Gemini telescopes, the Southern Astrophysical Research (SOAR) telescope, and the Large Binocular Telescope (LBT)—have proved invaluable to the design, development, and project planning of the LSST observatory.

1 Detailed cost estimates for the telescope and site are found in Chapter 7 of this proposal.

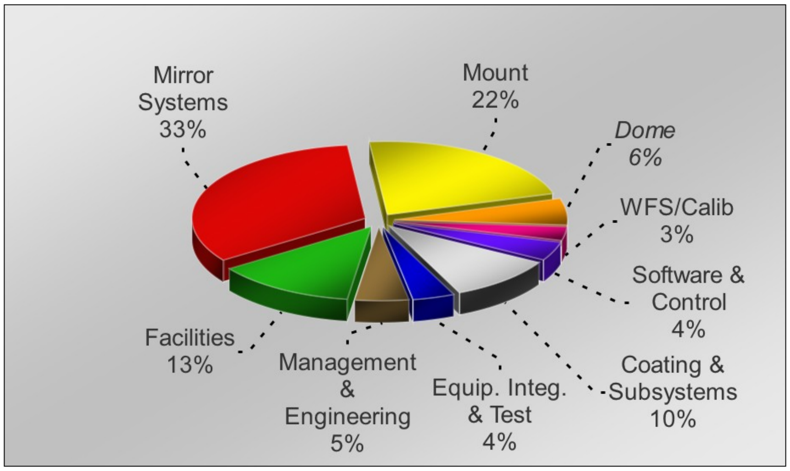

FIGURE 4-10 Percentage distribution of telescope subsytem and component costs

4.3.1 Telescope Structure and Subsystems

The proposed LSST telescope is a compact, stiff structure with a powerful set of drives, making it one of the most accurate and agile large telescopes ever built (CSA and VRSI 2006). The mount is an altitude over azimuth configuration (Figure 4-11). The telescope structure is a welded and bolted steel system designed to be a stiff metering structure for the optics and a stable platform for observing (Neill 2006a). The primary and tertiary are supported in a single cell below the elevation ring; the camera and secondary mirror are supported above it. The design accommodates some on-telescope servicing as well as efficient removal of the mirrors and camera, as complete assemblies, for periodic maintenance.

The stiffness of this innovative design is key to achieving a slew and settle time that is beyond the capability of today's large telescopes. The size and weight of the systems are a particular challenge, but the fast optical system allows the mount to be short and compact. Finite element analysis measures the first mode of the telescope system, including the concrete pier, to be 8.9 Hz and the critical support of the camera to have a 12.6 Hz mode. The top four modes are:

- — 8.9 Hz: Transverse telescope displacement
- — 9.0 Hz: Elevation axis rotation
- — 11.9 Hz: Top end assembly optical axis pumping
- — 12.6 Hz: Camera pivot

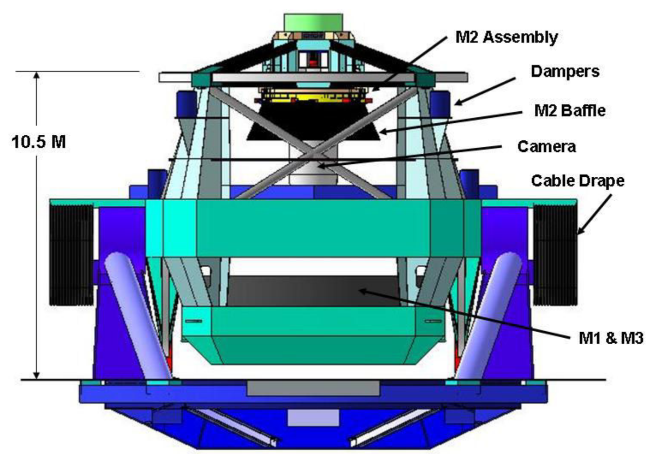

FIGURE 4-11 Configuration of the telescope mount showing several subsystems and key attributes

The motion time for a nominal 3.5-degree elevation move and a 7-degree azimuth move is five seconds. In two seconds, a shaped control profile will move the telescope, which will then settle down to less than 0.1 arcsec pointing error in three seconds. The mount uses 400 hp in the azimuth drive system and 50 hp in the elevation system. There are four motors per axis configured in two sets of opposing pairs to eliminate hysteresis in the system. Direct drive systems were judged overly complicated and too excessive, so the LSST design has each motor working through a multi-stage gear reduction, with power applied through helical gear sets. The 300-ton azimuth assembly is supported on hydrostatic bearings, while the 151-ton elevation assembly uses rolling element bearings. Each axis uses tape encoders with 0.001″ resolution. Encoder ripple from these tapes often dominates control system noise, so the LSST will include adaptive filtering of the signal in the control loop. All sky pointing performance will be better than 2 arcsec, an important efficiency parameter. Accurate pointing is key to tracking performance which, particularly for LSST's wide field of view, directly impacts trailing and imaging systematics. Traditional closed loop guiding will achieve the final level of tracking performance.

The top end assembly, shown in Figure 4-12 in cross section, provides the support system for the camera and the secondary mirror (Neill 2006b). These systems will be mounted on independent hexapods to adjust rigid body positions with 14 mm and 1.2-degree range, 10-micron and 15 micro-radian resolutions, providing high accuracy and repeatability. Design of the top end assembly has demonstrated that these systems will fit together in the compact space. The performance needed for the hexapods has been demonstrated in commercially supplied systems for other projects. The camera support will also include a rotator to provide +/- 90 degrees of motion for de-rotating during tracking.

Tuned mass dampers will be strategically placed on the telescope's top end assembly piers for increased damping performance. These relatively simple systems add a nominal 5840 lbs (2,651 Kg) of mass but will double the system level damping from 2.5% to 5%. The telescope has been analyzed for image motion caused by wind buffeting using a Gemini empirical data

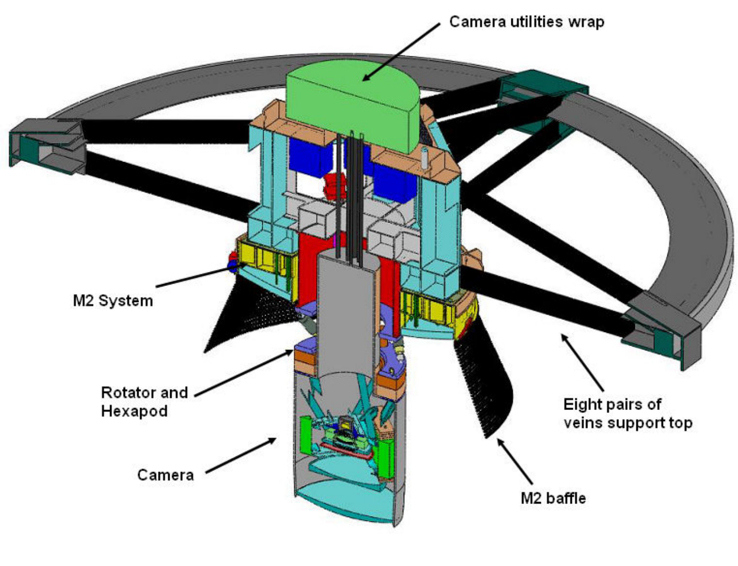

FIGURE 4-12 This 24-ton assembly, supported with 16 thin supports for minimal optical obstruction, carries the camera and secondary mirror.

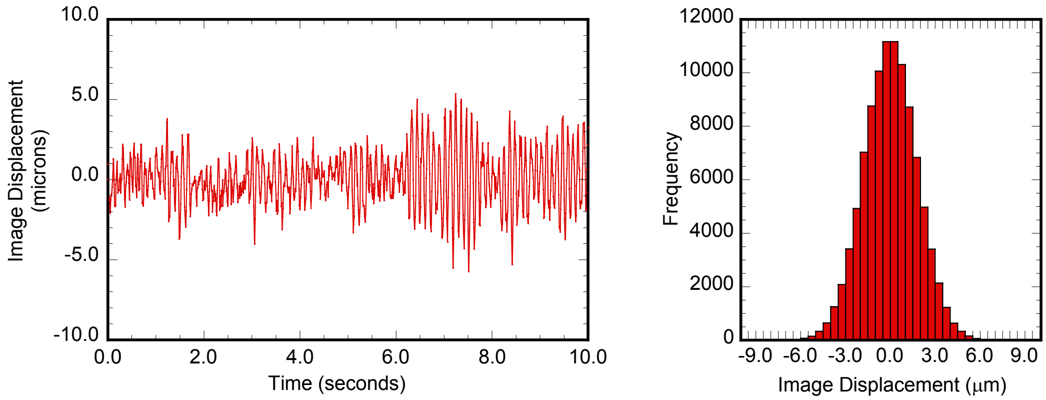

FIGURE 4-13 *Results of the wind induced image motion analysis*

set (Cho 2000) acting on the telescope. As a result of the added damping, the expected image motion should be reduced from 3 μm to 2 μm rms, which is below the image quality error budget of 2.2 μm rms for this component. Figure 4-13 shows the image motion time domain and histogram for this analysis.

The telescope is outfitted with light baffles and surface treatments necessary to control out-of-field and incident light from reaching the focal plane. Figure 4-14 shows the system stray light model, developed in FRED™, which was used to assess this critical wide-field issue. A point source transmittance calculation has evaluated the out-of-field stray light rejection and secondary ghost image performance. The design's critical surfaces have been identified, and the necessary baffles have been designed. Further results of this analysis are discussed in Chapter 5.

4.3.2 Active Optical System (AOS)

The telescope optical system is comprised of the three large, actively supported optical surfaces and a sophisticated control system. The primary and tertiary mirrors are being constructed as a single monolithic mirror by the University of Arizona Steward Observatory Mirror Lab (SOML); SOML will form the substrate and polish the surfaces to the final specification. The telescope project team will define and procure the mirror cell and support system, which will be provided

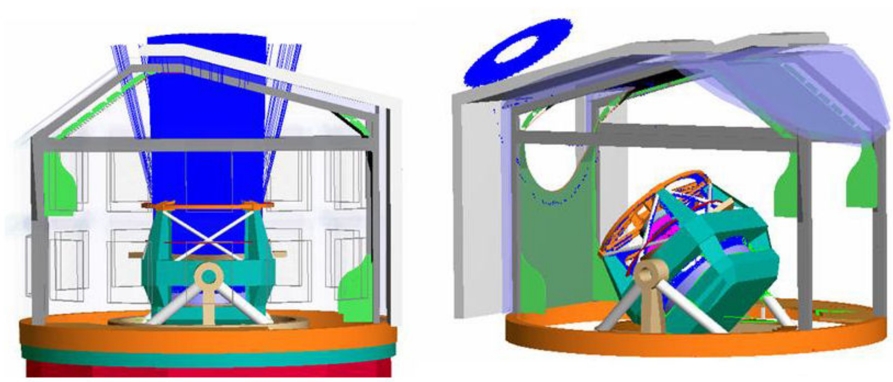

FIGURE 4-14 *Detailed LSST model analyzed for stray light and ghosting performance of the system. Dark blue in the figure represents projected light path as seen by the detector.*

for test integration at the conclusion of optical polishing. The secondary mirror will be procured commercially, with the expectation that different vendors will supply the substrate and optical polishing.

Primary "Monolithic" Mirror

The 8.4-m diameter primary mirror is a unique optic with both the primary reflective surface and the 5-m tertiary reflective surface fabricated into a single glass substrate (Figure 4-15). This "monolithic" mirror offers significant advantages in the reduction of degrees of freedom during operational alignment and better structural stiffness for the otherwise annular primary surface. It will be a spun-cast borosilicate substrate ground to form the final near net shape and then stressed-lap polished. This mirror fabrication and the support system have been demonstrated in several astronomical projects, including the Wisconsin-Indiana-Yale-NOAO (WIYN) telescope, MMT, Magellan, and the Large Binocular Telescope (LBT).

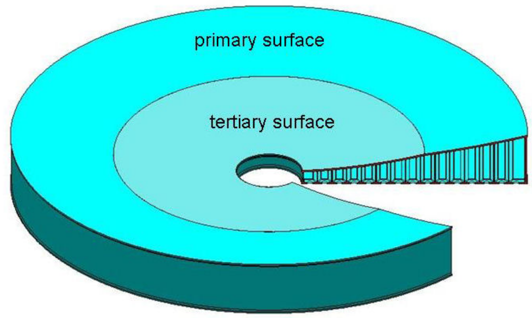

FIGURE 4-15 *The unique monolithic mirror with the primary and tertiary surfaces identified and a slice of the substrate removed for a view of the internal structure and surface*

Figure 4-16 shows the structural design of the mirror (Fern; LSST Document-2395 2006). The mirror will be fabricated by the University of Arizona SOML using their proven casting process, with some additional grinding to bring the cast parabolic shape to the multi-curve mirror. The thickness of the substrate at the center hole is thinner than previous cast mirrors, but the handling and support systems for the mirror have been designed to accommodate the lower stiffness in this region. The 2.1 cubic meters of glass that will "pool" over the tertiary mirror during spin casting have been analyzed for thermal and load impact, and each can be handled with minor adjustments to the casting process.

A key element in the performance of the mirror is the support system (Martin et al 2004). To limit the stresses and to provide appropriate figure control, the 17,600 kg mirror will be supported on 160 actively controlled actuators. These actuators will attach to the mirror through spreader bars, primarily three-point spreaders with some two-point spreaders near the outer

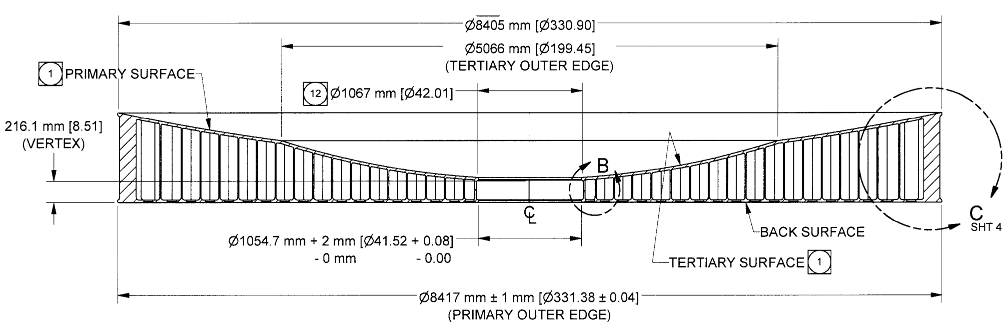

FIGURE 4-16 *Cross section showing the structural design of the M1/M3 monolithic mirror.*

edge and center hole for better support of these regions. Lateral stiffness in the support system is important in order to stabilize the mirror in the agile LSST. The mechanical connection to the glass, though not fully designed, will achieve this stiffness by reacting forces compressively at the back sheet, not through shear connections of RTV bond lines. Figure 4-17 shows the figure deformation of the mirror on its active supports for both the vertical and horizontal pointing positions. The mirror has two distinct surfaces but its eigenmodes are well behaved across the whole substrate, and controlled within specification by the actuators throughout the gravity variations on the telescope.

Thermal control is essential in order to maintain the optical performance of the primary/tertiary mirror (Cuerden and Sebag 2004). An active system will inject conditioned air into each cell of the mirror. Thermal control of this air is achieved with heat exchangers and circulating glycol positioned directly below the mirror. This approach is very similar to that taken on the Magellan telescopes (Perez 1994). There will be residual surface deformation from imperfect conditioning, so the active mechanical support system includes enough dynamic range to make corrections during operation. Thermal analysis indicates that both optical surfaces, after correction by the supports, will achieve their figure specifications in the presence of gradients developed from the mirror's radiative coupling to the sky. The LSST will also incorporate active daytime cooling of the telescope area within the dome to precondition the environment to nighttime observing conditions. Daytime cooling will provide a smooth thermal transition when the dome is opened and limit the mirror distortions caused by the thermal changes. There is additional discussion of the thermal control system later in this section.

Although the optical specifications of the monolithic mirror are a significant departure from previous large mirror designs, they are achievable with proven fabrication and metrology methods (Hill et al 1998). The mirror will be fabricated with loose abrasive slurries and an active figure lap. The M1 surface error will be better than that described by a structure function of 91.9 cm, the M3 to better than a 151 cm structure function. A single metrology tower will have relatively simple interferometric metrology stations to measure the primary and tertiary surfaces individually. The factor of two difference in the base radii of these surfaces puts these metrology stations at significantly different locations. During optical fabrication, the test procedure will

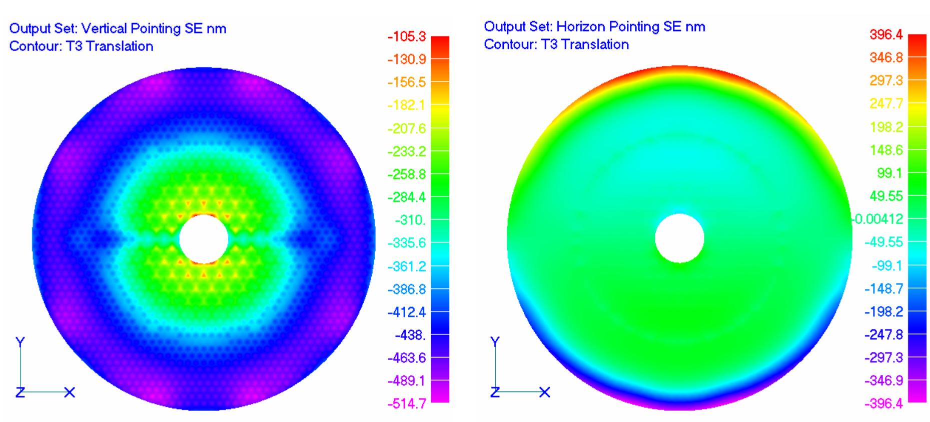

FIGURE 4-17 *Analytical assessment of mirror deformation (in nm) due to gravity for the mirror vertical pointing (left) and horizon pointing (right). The one fold symmetry (on the left) is due to the support layout and its reaction with the internal rib structure.*

include sufficient simultaneous tests of the two surfaces so that correctability calculations for appropriate errors can be handled together. To control the overall positions of the mirror surfaces relative to each other, Total Integrated Runout (TIR) measurements will be used to keep the two surfaces within 100 microns. A combination of laser tracker metrology to ensure the global position of the test sets and the coma measured in the tertiary will ensure that the aspheric displacement of the tertiary will be within 1 mm of the (TIR measured) best fit sphere.

Secondary Mirror

The LSST will require the largest secondary mirror ever fabricated. This low expansion glass mirror is 3.4 m in diameter, with a large 1.8-m diameter central hole (Figure 4-18). Analysis has shown that a solid meniscus substrate of 100 mm constant thickness will meet performance requirements (Neill 2005). This geometry is readily available in either ULE™ or Zerodur™ material and is very similar to primary mirrors of the 4.3-m SOAR and Discovery Channel telescopes.

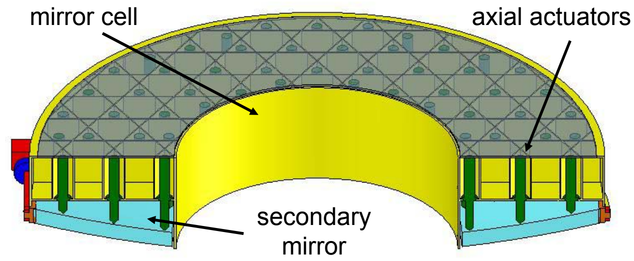

FIGURE 4-18 Cross section view of the secondary mirror system model showing the glass face sheet (blue), the axial actuators (green), and the support structure (yellow)

The secondary mirror will have its own active support system in order to deliver the required high-quality image performance. The mirror will be equipped with 102 axial and six lateral electromechanical, single axis actuators in order to achieve active figure control. In a meniscus mirror system, the axial stiffness is derived from the supports. In this case, the supports are 15,000 lb/in actuators using a stepper motor acting through a harmonic drive and lead screw to produce ±120 lbs of force in steps of 0.02 lbs. The six lateral supports have a 250,000 lbs/in stiffness and provide ±2000 lbs of force in steps of 3 lbs. The lateral supports are configured such that three define position and three respond with force for a well controlled kinematic system. Analysis has proven that this support configuration provides adequate figure correction of gravity sag in telescope use. Figure 4-19 shows the residual surface errors for the support system at vertical and horizontal pointing angles. The chosen actuator design has proven high performance with no risk of fluid leakage.

Optical testing of the convex annular shape is the challenging aspect of this mirror. Full aperture testing of the surface is cost prohibitive, so optical testing will be performed via sub-aperture stitching. Industrial design studies (L3-Brashear 2006; SAGEM 2006) have analyzed the configuration to validate our approach of taking twelve interferometric measurements with a 1.3-m aspheric test plate to form a synthesized, full-aperture, interferogram of the surface, as shown in Figure 4-20. A 14-inch prototype was also tested at the College of Optical Sciences at the University of Arizona to demonstrate the validity of the approach. The aspheric departure of the secondary mirror is just 19 microns, enabling the use of standard lap technology for optical finishing.

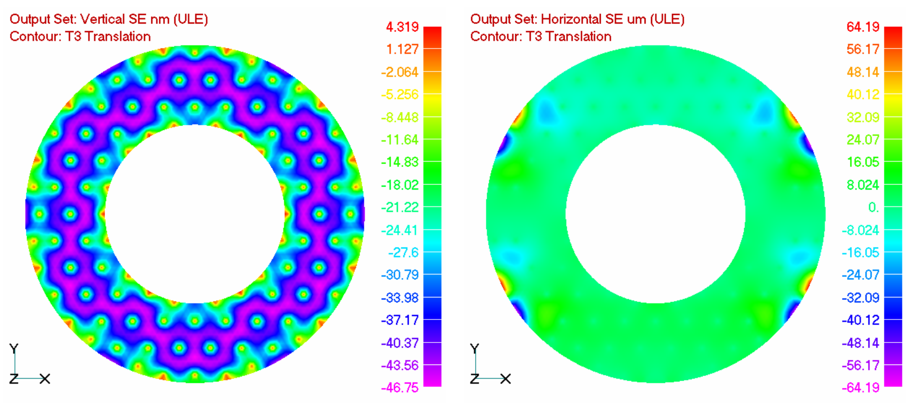

FIGURE 4-19 Actuator density is sufficient to provide the necessary support with limited gravity sag (in nm) in both the vertical (left) and horizontal (right) pointing positions of the mirror.

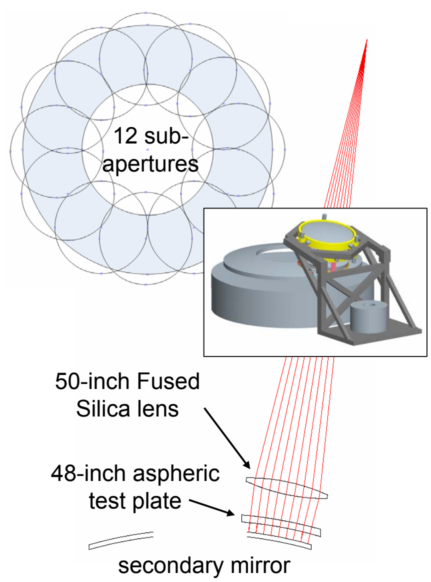

FIGURE 4-20 Composite diagram showing the metrology configuration for the secondary mirror, the footprint of the sub-apertures measured, and the concept model for the test station. A stitching technique was used at SAGEM for testing segments of a 10-m mirror, and QED Technologies uses it in their sub-aperture stitching interferometer commercial product.

Control of the Active Optics System (AOS)

Active control of mirror figure errors and telescope subsystem rigid body alignment will be accomplished via four wavefront sensors embedded in the focal plane periphery, look-up tables derived from a model of the mount, and an independent alignment system. The fast beam of the LSST optical design results in large single source beam footprints at each of the three mirror surfaces: 100%, 88%, and 74% of the clear aperture at M1 (the entrance pupil), M2, and M3, respectively. Thus low order surface errors on the mirror cause slowly varying aberrations across the field of view. This in turn allows wavefront sensors from the field edge to accurately determine the necessary corrective measures to maintain the interior image quality. Figure 4-21 shows the basic flow of information, feedback processes, and rates required to maintain high quality images. The Monte-Carlo analysis briefly discussed in Section 4.2 included wavefront sensing and correctability parameters to qualify the LSST hardware and control approach.

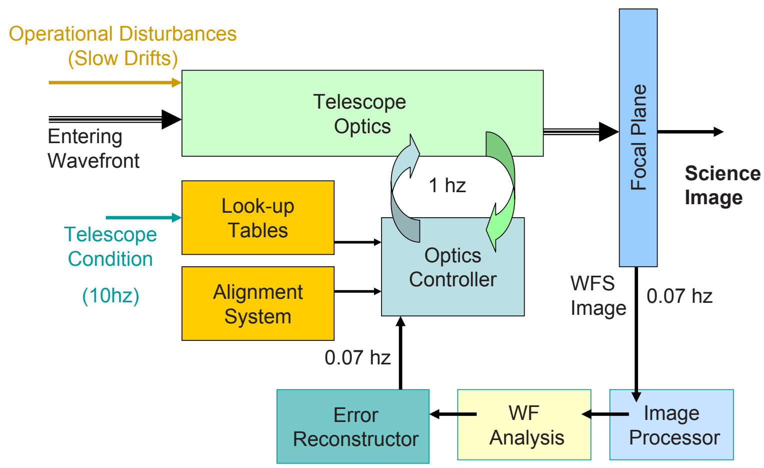

FIGURE 4-21 *Block diagram of the active optics system with major components and control channels*

Curvature sensing will be employed to measure wavefront aberrations via the irradiance transport equation (F. Roddier and C. Roddier 1991). The four wavefront sensor detectors will be positioned +/-1000 microns about best focus and will be sized to acquire 86 square arcminutes of field. This sky coverage will ensure 100% probability of capturing a star image of sufficient brightness in all filter bands (97% in the *u* band). An image processing algorithm will construct a master intra- and extra-focal image from the many sources that land on each half of the sensor sets. The optical aberrations from the telescope are nearly constant across the wavefront sensors, and the atmosphere is decorrelated from star to star. This allows the residual wavefront error caused by the atmosphere in the 15-second exposures to be effectively eliminated by combining multiple images. The performance of this approach has been demonstrated using data acquired on the sky with a similar focal ratio telescope, the Large Binocular Telescope (LBT) at prime focus. Figure 4-22 depicts the process and shows the very similar results compared to conventional single-source approaches.

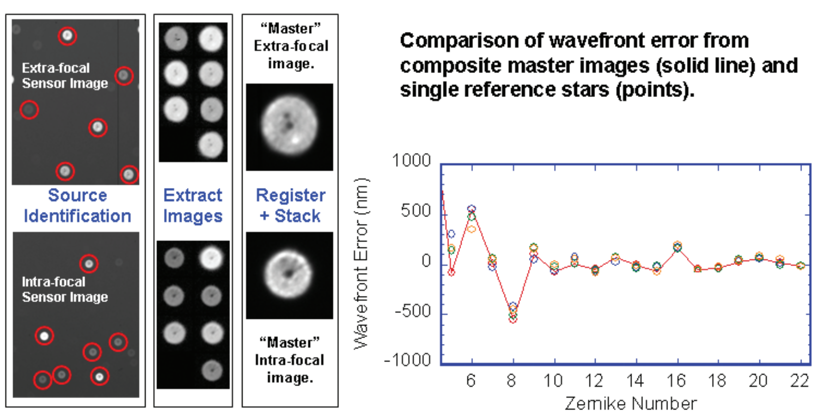

FIGURE 4-22 *Example of LSST-like wavefront sensing from 8.4-m LBT data*

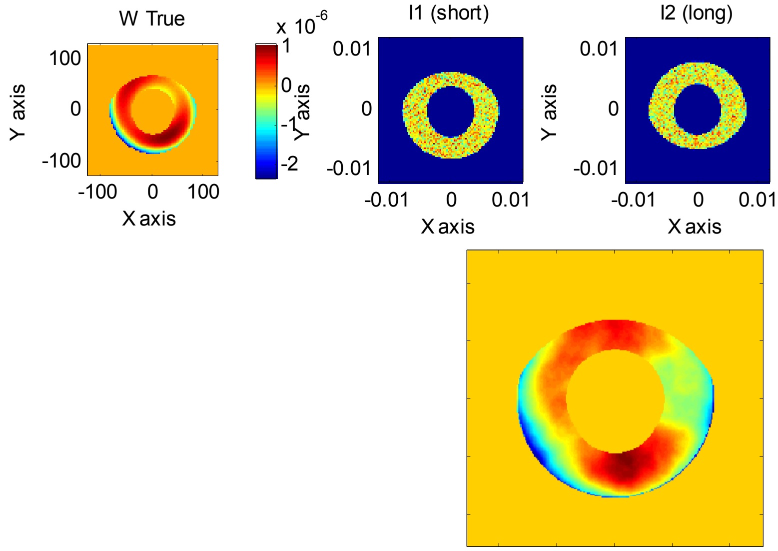

FIGURE 4-23 *Example of Matlab Curvature code for arbitrary pupil geometry*

Code for arbitrary pupil geometries, such as the LSST annular shape, has been developed (Lang 2006). The output of this Matlab library of routines, with the vignetted pupil that will be seen by the curvature sensors at the outer edge of the focal plane, is shown in Figure 4-23.

Computed wavefront aberrations at each sensor will feed an error reconstruction routine, which will determine both figure and rigid body corrections. This routine uses a tomographic approach to isolate the offending element in the optical system or to select the best option in the rare case that multiple options for correction are possible. Rigorous analysis has shown that four wavefront sensors on the focal plane periphery are sufficient to provide active alignment correction in a single iteration and to isolate figure corrections for each mirror. (Phillion et al 2006).

Initial telescope subsystem integration, subsequent realignment upon recoating or servicing, and beginning of night operations support will be provided by a commercial laser tracker. Within 60 seconds, the laser tracker embedded in the primary mirror center hole will query a set of three reference fiducials on each of the three principal optical elements (M1+M3, M2, camera) in order to enable alignment. This approximate alignment will produce images well within the capture range of the wavefront sensor system. In one more minute additional targets can be measured on each element to achieve ~ 1 arscec FWHM point spread function. Figure 4-24 shows the system and its line of sight to each target.

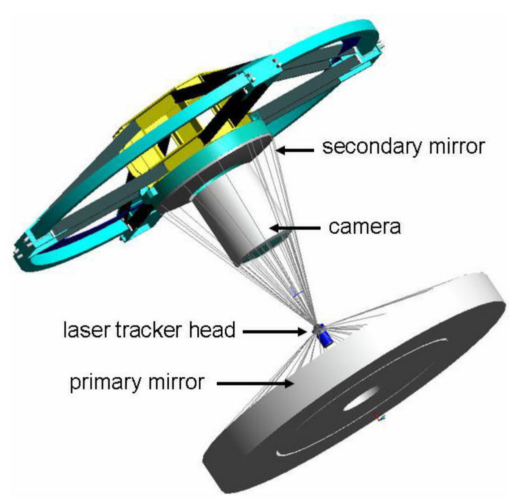

FIGURE 4-24 *Line of sight diagram for the laser metrology alignment system*

Optical Reflective Coating

The LSST requires maximum throughput for wavelengths from 350 nm to 1100 nm. Aluminum performs well from 350–400 nm but does not meet the requirements at longer wavelengths. Silver offers high reflectivity from 400 nm and higher but drops off very quickly below 400 nm. For LSST, we will use a DC magnetron-applied, multi-layer coating of aluminum, silver, dielectrics, and adhesion layers similar in design to that used at the Gemini Observatory (Boccas 2004; Sebag 2006).

The LSST observatory will include a mirror cleaning and coating system for application of the reflective coatings on site. A custom coating facility is required to support the 8.4-m LSST primary mirror size and the multi-layer coating requirements. This size chamber is not available commercially, and the Gemini chamber is not large enough to accommodate the LSST mirror. The design of the LSST coating chamber relies on experience with the Gemini and VISTA coating chambers. In order to minimize handling risks, the monolithic primary/tertiary mirror will remain in its cell for coating, and the magnetrons will rotate above its surface. The cell will be used as the bottom shell of the vacuum chamber and will be designed with suitable barriers to keep contamination in the support system away from the clean coating environment. The M2 mirror will be removed from its support cell for coating and installed in the regular bottom shell of the vacuum chamber. Cleaning/stripping of the mirror substrate will be carried out in an area adjacent to the coating chamber. Figures 4-25 and 4-26 show the coating design performance and the design of the chamber. Mirror washing and reflectivity monitoring will be standard maintenance procedures. Recoating is expected every two years.

4.3.3 Telescope Control System (TCS)

The Telescope Control System (TCS) controls the telescope assembly and associated hardware. The TCS is responsible for the precise pointing and tracking calculations necessary to observe a certain field. The TCS provides a central process to orchestrate safe and efficient performance,

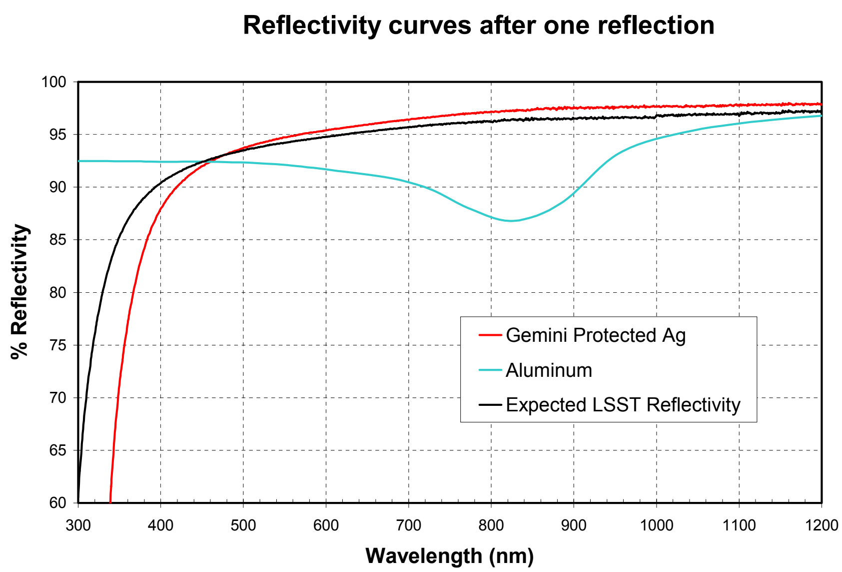

FIGURE 4-25 The reflectivity performance curves for silver, aluminum, and the LSST reference coating design

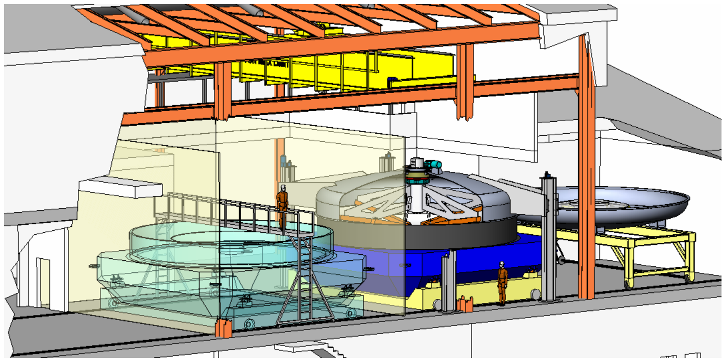

FIGURE 4-26 *The concept design for the coating chamber*

distributing commands and set points to the intelligent local subsystems where tight time critical loops are managed. Figure 4-27 shows a layout of the supervisory channels with the associated communication rates. The key components of the control system are 1) the software, 2) the pointing kernel, (3) the guider, (4) the AOS (described above), and (5) the telemetry system.

TCS Software

The Telescope Control System software supports the same Supervisory Control and Data Acquisition (SCADA) architecture and philosophy as the Observatory Control System (OCS) described in Chapter 5. The TCS provides the coordination, sequencing, and supervisory control of the many elements within the telescope subsystem along with the necessary graphical interfaces and data distribution. It must interact with diverse elements with a variety of data content and timing requirements. Two main principles incorporated in the design are fault tolerance and flexibility for evolving software and hardware.

The TCS follows a “publish/subscribe” approach to communication. It uses the same middleware as the OCS and shares several code components for maximum production and operation support efficiency. The TCS continues the coordination and sequencing initiated with

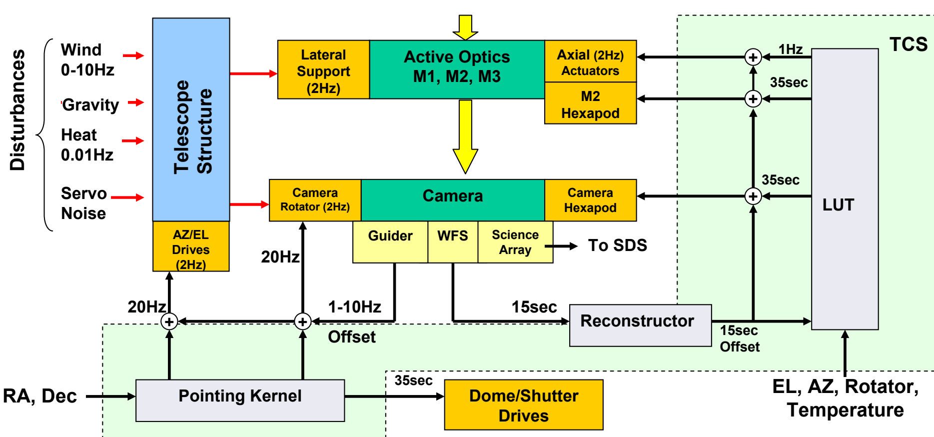

FIGURE 4-27 *Block diagram of the TCS architecture showing communication channels between internal and external components*

the OCS; it also provides the necessary supervisory control for operation of the telescope during engineering and maintenance operations when the OCS may not be engaged.

Implementation of the TCS will use industrial standards, commercially available, well-supported software products, and proven software components from existing telescopes, where applicable. The design will consist of a series of layers implementing successive levels of functionality and will provide a plug-in architecture to permit the incremental addition of capabilities as needed.

Pointing Kernel

At the heart of the TCS is the pointing kernel software. This code performs the astrometric transformations from the requested sky pointing position into the local telescope altitude, azimuth, and camera rotator positions. The kernel derives the pointing vectors from a pointing model and feeds the mount control system with 20 Hz updates for tracking during observations. The pointing model is based on a detailed mathematical and physical understanding of the effects that control apparent trajectories of astronomical objects. Taking advantage of LSST's large data set, the pointing model will be capable of continuous updates to maximize performance. The kernel for the LSST will be purchased from P. T. Wallace (CCLRC–Rutherford Appleton Labs). This package of TPOINT, TCSpk, and SLALIB software has a proven track record on many telescope projects, including Gemini and SOAR.

Guider

The guider component will minimize telescope tracking errors and tracking jitter by performing the closed loop correction of the errors. The guider will measure any telescope tracking error and send the correction signals to the camera rotator and telescope mount through the TCS.

The guider will process the information from guider sensors located in the focal plane and measure the mount x, y and Øz tracking errors. The high frequency component of the error signals will be sent directly to the telescope mount servo, where this information will be used to minimize tracking jitter. The TCS will receive the lower frequency components of the tracking errors and will convert these errors to the appropriate offset correction for the mount axes and the camera rotator.

The guide sensor and signal are provided by the camera control system (see Section 4.4); these will interface with the control software through the same communications middleware used throughout the system.

Observatory Telemetry System (OTS)

The observatory telemetry system (OTS) collects, stores, and serves engineering and system operation data from all observatory systems. It will monitor, analyze, and archive the conditions and state of the telescope, optics, enclosure, observatory environment, and camera systems. The OTS will also supply operational conditions for calibration and analysis of scientific data, will act as the database of conditions for real time and periodic analysis, and will record system actions to support troubleshooting and maintenance. The OTS is expected to accumulate 14.6 Tb of data per year; these data will be archived completely by the LSST data management systems and stored locally in ways suitable for engineering analysis.

The OTS will be an important resource during integration, commissioning, operations, maintenance, and data processing, and an essential tool for meeting LSST's goals of operational efficiency and data quality. This system will be fully operational during the early integration phase; it must operate in a manner consistent with a complex industrial machinery environment where minimizing downtime is critical.

4.3.4 Electrical Utilities

The LSST power system will use commercial power already available on Cerro Pachón; sufficient power exists to feed large peak demands. Figure 4-28 provides an overview of the power distribution, consumption estimates, and back-up assets. The LSST will bring 2400-volt service from the commercial power line over the 1.5 km to the Peñón site. For more efficient transfer and coupling and to mitigate the impact of LSST on other summit telescopes, a 500 KVA step-down transformer will be located at the telescope, rather than in the main switch yard. LSST will provide conditioning and back-up for its own systems. Power at the summit will be 380/220V 3Ø Wye with grounded neutral and 120V 50 hz.

Acceleration of the telescope and dome that is required for re-pointing every 30–35 seconds will impose high peak demands for power. To meet this demand, a flywheel UPS is included in the design. This system will store 1 megawatt of energy in a high-velocity spinning flywheel that will supply the 450 kW peak load of the mount and dome. This draws 8% of the available flywheel power in 2 seconds. It will regenerate in 12 seconds with commercial power. These flywheel systems currently cannot absorb power at the same rate they can provide it. When the telescope and dome decelerate, regenerated power cannot be stored back in the flywheel but rather must be dumped into a resistor bank. The average power going into the resistors is 16.3 kW, a manageable heat source.

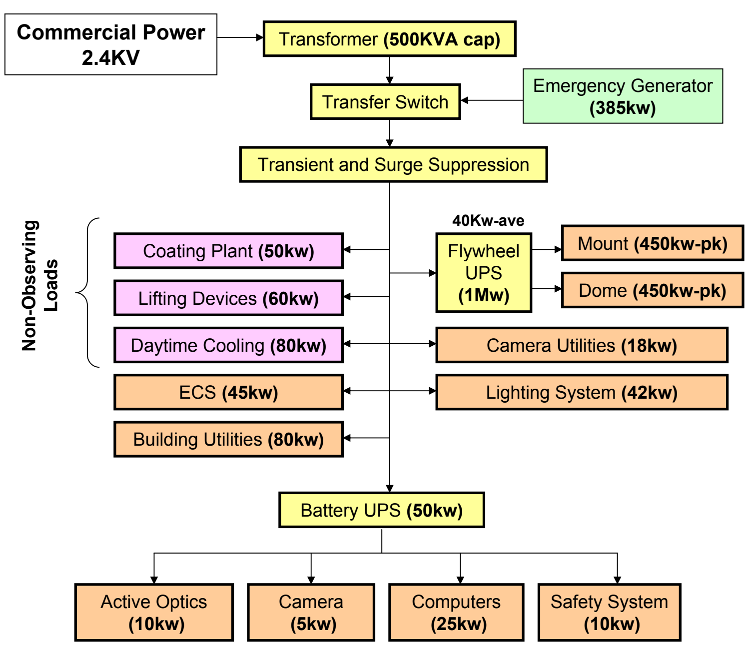

FIGURE 4-28 Diagram of the power distribution and major electrical loads on the summit

A grounding grid embedded in conduction-enhancing salts will be located under the LSST building foundation. All subsystems will be separately grounded directly to the grid. All interface signals will be differential or optical, and signal shields will be left open at the receiving end. Lightning protection for the entire facility will use a conventional Franklin Lightning system that will guide the static electric flow down a path away from people, structures, and electrical equipment. The lightning protection system is composed of air terminals, overhead main conductors, downcomers, a grounding system, and a system of transient voltage surge suppressors (TVSS), all fully compliant with U.S. standards.

Utility Systems

The telescope system will include all the auxiliary hardware, utility subsystems, and service features necessary for safe and proper operation and maintenance. Key items include light baffles, in situ mirror washing, mirror cover, elevation and azimuth cable wraps, utility and power distribution, sensor and telemetry for system monitoring, personnel access and safety interlocks, vibration suppression and a balancing system.

Calibration Equipment

The telescope system will provide much of the hardware necessary to perform the calibration plans outlined in chapter 5.3. Significant in this area is an IR camera bore-sighted on the telescope, a calibration screen mounted in the dome, and the 1.5-m class robotic telescope sited on the hill adjacent to the primary summit facility.

4.3.5 Summit Facility and Site Infrastructure

The LSST facility will be sited on AURA property atop the El Peñón summit of the Cerro Pachón ridge approximately 1.5 km SSW of the Gemini and SOAR observatories (Barr 2005). Proximity to these two recently completed observatories affords access to developed site infrastructure, proven bearing conditions, and predictable site construction logistics. Access to the site is via the

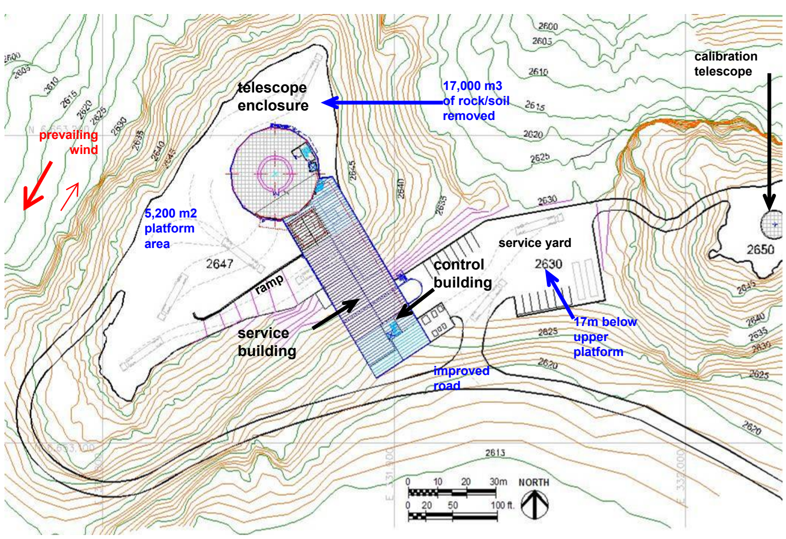

FIGURE 4-29 *Preliminary summit facility site plan*

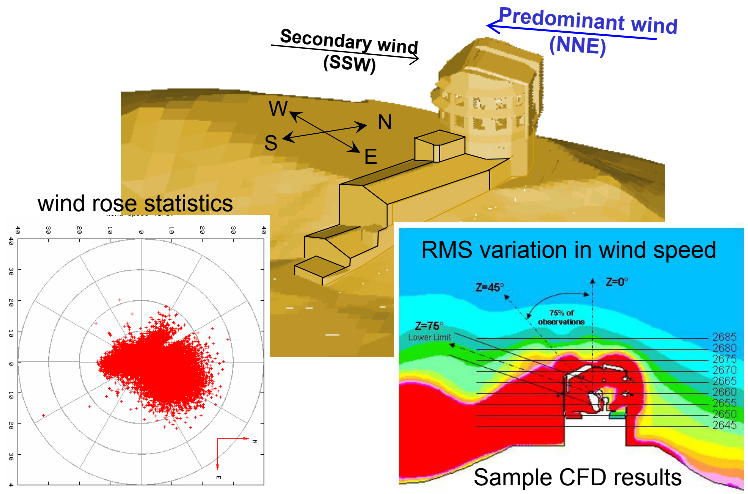

FIGURE 4-30 *Site CFD study: model, wind plot and air turbulence profile.*

existing main road from below and a 1-km provisional access road; the latter will be improved as required. Utilities (electric power, communications, and water) will be extended approximately 1.3 km from the common Pachón utility yard near Gemini. A new shared dormitory and dining facility on Cerro Pachón will also accommodate LSST requirements (AURA 2006).

Creation of an adequate platform for the telescope enclosure requires removal of approximately 17,000 cubic meters of rock from the El Peñón peak, down to approximately the 2647 m elevation level. A similar, smaller pad will be created on the adjacent hill for the calibration telescope. The rock and soil that are removed will be used as fill for the roadbed and to create a level service yard in the saddle between the two hills (M3 2006a; M3 2006b). To reduce the extent of the required cut and to minimize thermal turbulence around the telescope enclosure, the support facility structures will be stepped down the hill toward the saddle and road below (Figure 4-29).

To understand the site-specific airflow conditions, alternate building layouts were analyzed using computational fluid dynamics (CFD) modeling, which accurately depicted topography, prevailing wind conditions, and mass of the buildings. A composite illustration of the site model (Figure 4-30) shows the wind direction plot and profile of the air turbulence (variation in wind speed) caused by wind passing around and over the buildings. To minimize potential image degradation caused by air temperature variation, the orientation and shape of the support building was modified to reduce the height of the induced turbulence (Barr 2006). Calculated image degradation based on turbulence, convection, and radiation, is quantified in Table 4-2; these data assume that the telescope is pointed perpendicular to the prevailing wind, which the operations simulator shows to be the most common observing direction.

TABLE 4-2 Impact of thermal turbulence on imaging.

Telescope Zenith Angle (deg)	Total Contribution (milli-arcsec)	Outside the Enclosure (milli-arcsec)	Inside the Enclosure (milli-arcsec)
20	101	73	60
40	107	93	40
75	127	58	105

4.3.6 Dome

The telescope dome will be a rotating cylindrical structure, 30 m in diameter, supported by a fixed lower enclosure (Vertex 2006). In addition to its usual functions of protecting the telescope against weather and providing housing for support, maintenance, and other equipment, the LSST dome satisfies the critical need for extreme light baffling necessitated by the telescope's wide observing angle and fast tracking speed. The major components of the dome are:

- — Carousel: The primary protective shell of the dome; a faceted, cylindrical housing made of insulated metal panels supported on a structural steel framework
- — Observing aperture shutters: Two L-shaped sliding doors that open to provide a 10-m observing slit
- — Wind screen/light baffle: A robust set of folding or bi-passing panels that move along the slit to form the horizontal edges of the observing aperture. The wind screen serves as a light baffle to form a 10-m × 10-m aperture and provides the first level of protection against wind and extraneous light sources.
- — Vent openings: Windows in the vertical faces of the carousel with roll-up gates to allow controlled natural ventilation of the telescope area during observing and with baffle structures to prevent light passage
- — Rotation system: The drives, wheels, and control system that rotate the dome slit to any required azimuth position. The observing cadence, which is dictated by the science requirements, demands faster than normal dome drive speeds and special strategies, as described below.
- — Dome crane: A 20-ton bridge crane mounted to the arch girders of the dome, which is used to service the camera and other telescope-mounted equipment

The dome design provides adequate space for internal handling operations and interfaces with an exterior platform lift, while minimizing interior air volume and keeping the structure lightweight and agile. The cross section views of the dome depicted in Figure 4-31 show the telescope in the horizontal pointing position used for removal of the camera (left) and the vertical position used for removal of the primary mirror (right); these are the two maintenance activities requiring the most radial space and shutter opening, respectively.

Weighing approximately 435 tons, the dome is supported on 16 steel wheel carriages with suspension and damping support for smooth operation. The drives for the system consist of eight 50-HP, liquid-cooled motors mounted to the fixed enclosure. These drives provide sufficient

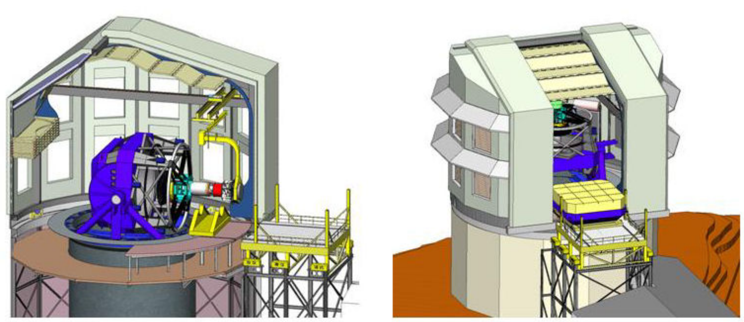

FIGURE 4-31 Cross-section of the LSST dome showing the telescope in the horizontal and vertical position for camera and primary mirror removal respectively.

power to move the dome in pace with the telescope for adjacent field observations. (Vertex 2006). A key strategy is to oversize the dome aperture by ~1 m and slew the dome only as far as needed for the oversized opening to satisfy a clear path for the telescope. During the exposure, the dome will creep to a new position where the extra dimension of the aperture is toward the next field position. This allows a nominal dome azimuth move of 7° (3.5° on sky when the telescope is at 60° elevation) to be accomplished with a negligible slew, followed by another 30-second slow move during an observation.

Effective wind flushing and minimal thermal turbulence are vital factors in the design of the dome. CFD analysis was used to optimize beneficial air flow characteristics (Fluent 2006); the wind velocity plot for the selected dome design is shown in Figure 4-32. Twenty-two vent openings are distributed around the vertical faces of the dome, providing a maximum open area of 390 m^2 (4200 ft^2). Each opening has a variable roll-down gate to throttle back the air flow in high wind conditions and a fixed louver panel to block stray light from directly entering the dome. The optimal shape of these louvers (for minimal pressure drop) was evaluated as part of the CFD analysis. Even with the restriction of the louvers and in a relatively low (1 m/sec) wind condition, the extensive vent area was shown to be effective in achieving the target minimum of 30 air changes per hour in the telescope space.

4.3.7 Support Facilities

The support facility consists of three primary elements: the telescope lower enclosure, a platform lift, and an attached service and operations building. The overall layout is designed to keep heat-producing equipment and conditioned spaces, such as the control room and computer rooms, away from the telescope area; the unconditioned service space between them acts as a buffer. The linear arrangement of the support facility allows carts for large equipment (e.g., mirrors and camera) to be moved within the building on rails. The buildings are connected for efficient equipment and personnel circulation.

The lower enclosure has minimal structure and few functional spaces in order to minimize heat capacity near the telescope. An enclosed stair and personnel elevator connect the levels within the enclosure: the telescope area at the top and an open high-bay ground floor with some space for equipment.

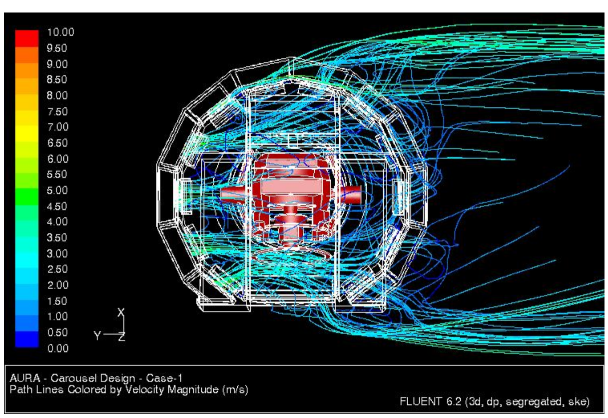

FIGURE 4-32 The carousel dome in plan view with CFD air flow path lines. Result is for wind approaching orthogonal to viewing direction.

The telescope area has two levels. The main floor is 3 m below the top of the telescope pier to allow access to the dome drive and support systems. There is also a partial section of raised floor, even with the top of the pier and with the highest level of an equipment lift adjacent to the lower enclosure. Carts transition directly from the telescope to the lift, which provides the primary means of moving large equipment between the telescope and the service building.

The equipment lift is a commercially available, vertical reciprocating conveyor with 80-ton capacity. It provides access to three levels: the telescope level, the base of the lower enclosure, and the main floor of the service building. It has a deployable upper weather cover that rises up with the lift to the telescope level and when not in use, stows below the dome aperture to minimize the wind profile of this structure.

The layout of the service building is shown in Figure 4-33 (M3 2006c). Almost all maintenance of the optics and camera for LSST will be performed here. Mirror recoating and major camera maintenance can be carried out concurrently in order to minimize downtime. A 60-ton overhead bridge crane serves a receiving area and the spaces where major handling operations are performed. The coating chamber and its ancillary spaces occupy the left side of the plan in Figure 4-33. In addition to the 10 m × 10 m area for the chamber itself, there is space to store the lower coating vessel when it is not in use and an adjacent area for mirror stripping and cleaning just prior to coating. The camera work area occupies the upper right area of the plan, approximately 280 m^2 (3010 ft^2) of the service building. This includes progressively cleaner rooms specified ISO class 8 and class 7 required

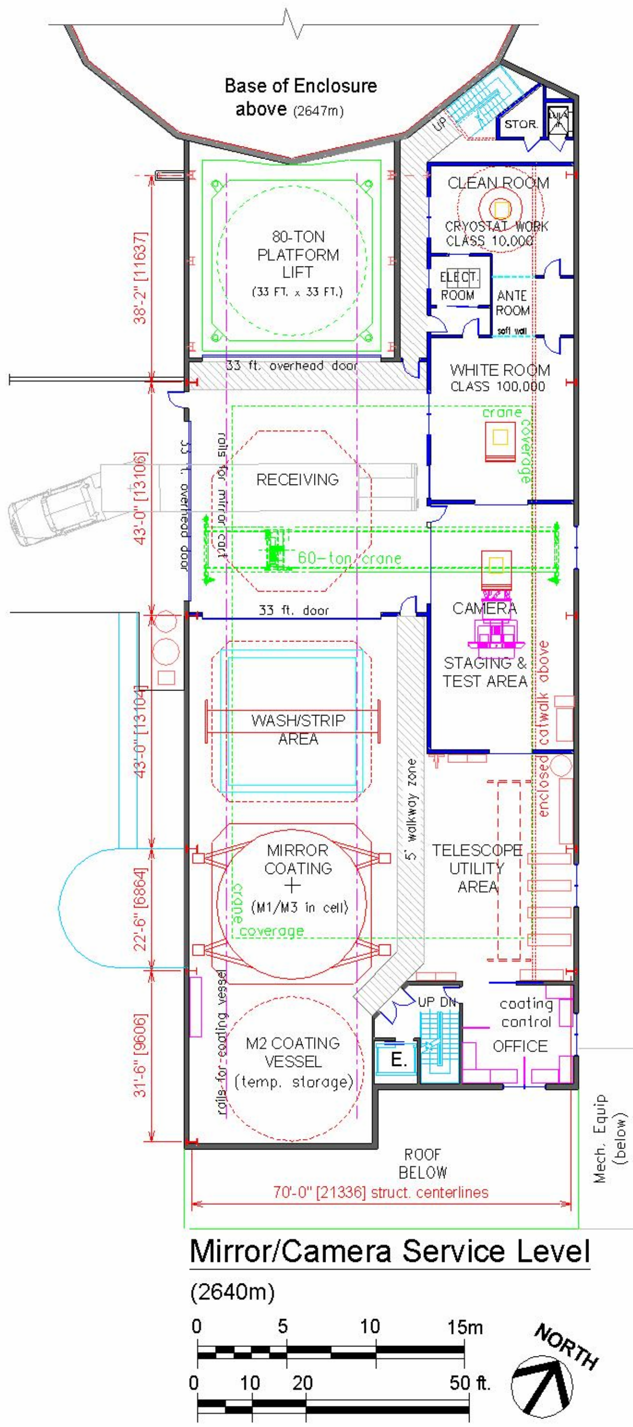

FIGURE 4-33 Floor plan of service building

for work on the camera and the focal plane, respectively (Nordby 2006).[2] Access ports will be installed in the ceilings of some of the camera work areas to allow the main overhead crane to be used there as well. Where possible, the major equipment to support these rooms and maintenance activities will be located in the main utility yard of the facility. Some equipment space is provided within the service building. Although equipment located here will generally not be in use during telescope operations, fans and ducting will be installed to prevent any build-up of warm air and to exhaust air from the service building away from the telescope.

The summit facility includes the control room, data processing room, and offices. These conditioned spaces are located at the farthest end of the building, 70 m away and 25 m below the telescope base. A total area of 470 m^2 (5050 ft^2) is provided for this part of the building. A personnel elevator provides access to the service building above, both at the main floor level and at a mezzanine level where an enclosed corridor leads directly to the telescope enclosure.

The entire summit facility will be constructed using standard steel frame structure with insulated metal siding and roofing. The exterior surface of the dome will have a finish (reflective aluminum or high-titanium white) that minimizes both solar absorption and nighttime super cooling to reduce thermal turbulence in the light path of the telescope. The fixed facility has a manufacturer's standard white finish for low solar heat absorption. The entire dome and building are designed for low thermal inertia and clean unobstructed air flow over the surfaces for fast equilibration to ambient air temperature and effective flushing of any heated air above the building at night.

The LSST is designed with a comprehensive thermal management system operating both day and night. A central controller for the thermal management system utilizes a network of sensors installed throughout the facility and on the telescope to monitor the temperature of critical surfaces and areas. The key environmentally controlled spaces are shown in Figure 4-34.

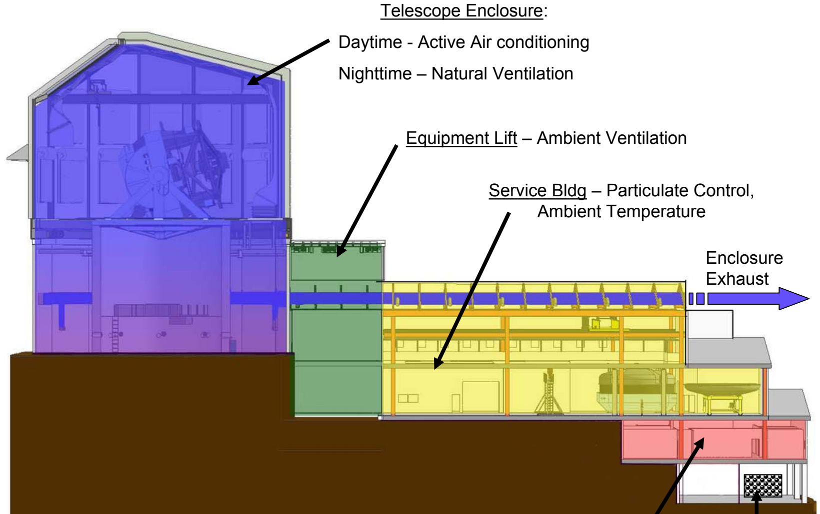

FIGURE 4-34 LSST environmental management. The service building buffers the telescope from heated control rooms.

2 International Standards Organization classes 8 and 7 correspond to classes 100,000 and 10,000 of the former U.S. Federal standard for clean rooms.

The telescope dome is naturally ventilated. During observing, the roll-down gates on the dome vent openings are controlled by the thermal management system to effectively equilibrate interior and exterior temperature without causing unacceptable wind pressure on the telescope and optics. During the day, the telescope area is actively cooled by fan-coil units supplied with chilled glycol to precondition the dome air and the telescope to just below the anticipated ambient temperature at the onset of observing.

Fans and ductwork extract building heat away from the telescope. Air is drawn from the dome area, the lower enclosure, and the service building and is exhausted 50 m away from the telescope. The ducts for this system run through the roof structure of the service building to fans at the south end of the building. The calculated fan power for this system is 5 hp, providing 10,800 CFM flow with a duct size approximately 1.3 meters. Chilled glycol is provided throughout the facility to extract heat from point sources. Any heat-producing equipment that must be located near the telescope is housed in an insulated enclosure with conductive cold plates or glycol-to-air heat exchangers to extract the heat.

4.3.8 Telescope Integration and Testing

The strategy for efficient on-site assembly of the LSST has had a strong influence on the design. To the extent possible, the telescope will be procured in subsystems that are pre-assembled and fully tested at vendor factories. Significant de-bugging and subsystem commissioning can thus be performed long before the parts arrive at the remote summit location. Systems will be constructed with alignment and interface guides in place as required elements for integration support. At the summit, integration will be staged with clearly defined tests and commissioning tools to verify that each step meets its performance goals, thereby limiting the complexity of the final on-sky testing and commissioning phase.

The telescope will be designed and developed as a single, integrated system and procured as large subsystems. Each subsystem will have normal acceptance criteria, and for certain systems such as the telescope mount and the large mirrors, the contractors will pre-assemble the system in their factories for additional testing and performance verification.

Factory assembly and testing have been successfully employed in many telescope development efforts. Complete assembly at the fabricator and testing with a complete complement of surrogate masses allows extensive testing and efficient resolution of numerous integration and commissioning issues. In the case of the SOAR telescope, for example, factory commissioning allowed for rapid assembly and very short summit integration efforts. A similar approach will be taken with LSST.

The major telescope subsystems, including the camera, will include measurement fiducials that reference the critical dimensions of the part or subsystem. The fiducial targets will be laser tracker retro-reflectors. Measurements will be made with a commercial laser tracker and the data processed with Spatial Analyzer software. This analysis and comparison with mechanical CAD file geometries will enable efficient integration.

Summit Integration

The basic integration plan is shown in Figure 4-35. This chart shows the general flow and the potential critical paths for integrating the system and reaching first light on schedule. To realize this plan, several key tasks will be completed in the design and development phase. Both mirror fabrication efforts will be well underway, a benefit of LSST non-federal funds, and the

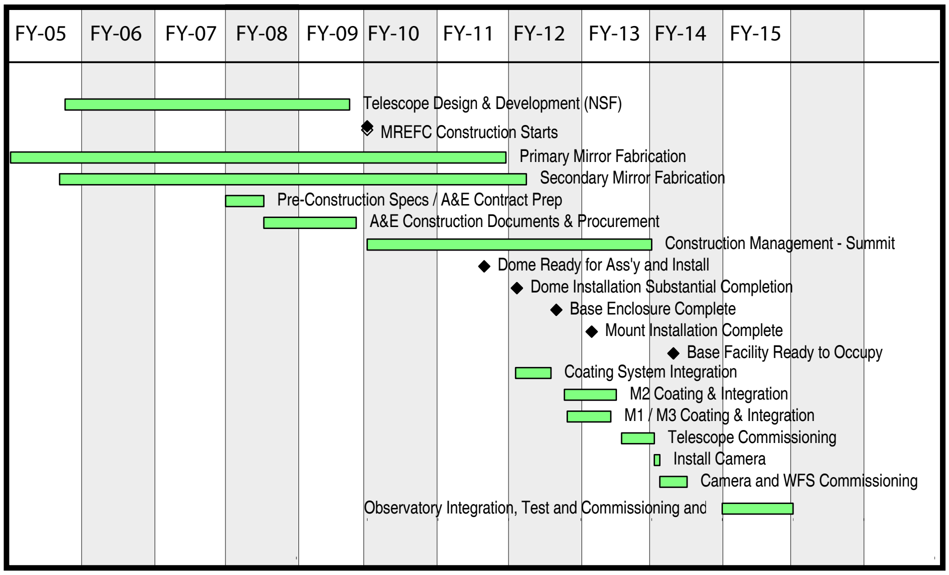

FIGURE 4-35 *LSST summary integration chart for critical tasks*

design and procurement packages (Summit Facility, Dome, Mount) for several key elements will be complete to enable rapid starts. Facility construction must commence early and support staged completion, so that the coating support facility is available for occupancy earlier than the remaining facilities. In addition, the dome must be started during the facility construction, so that telescope mount integration can begin immediately following completion of the facility and substantial progress on the dome construction.

Commissioning Cameras

The integration plan relies on a set of commissioning camera systems to aid in initial integration, operation, and debugging of the telescope mirrors after they are mounted within the telescope system. Substantial final optical testing of the mirror surfaces will be performed at the optical fabrication vendor's facility during final acceptance testing under an interferometer in their respective telescope mirror cells. These tests will validate single elevation (zenith) control of the optical surface figures. Some small variations of these control matrices are, however, to be expected due to elevation/gravity effects. To address this and to support our short (6-month) integration schedule, we will measure and validate the ability for each mirror to maintain its optimal shape in use on the telescope individually, prior to installation of the LSST Science Camera. Performing these mirror tests separately allows identification of error sources more rapidly than if all three mirrors were installed first, thus requiring decoupling of the various error influences.

To measure the M1/M3 monolithic mirror after its integration on the telescope, a prime focus camera system will view images reflected off the M1 surface, while a center of curvature null lens simultaneously views the optical surface figure of M3 (the M3 null lens assembly sits in the shadow of the M1 obscuration). A prime focus camera for M3 is impractical due to the large spherical aberration caustic and prohibitive cost for such a system. This design choice also affords the possibility of packaging both systems together on the telescope mount. A surrogate

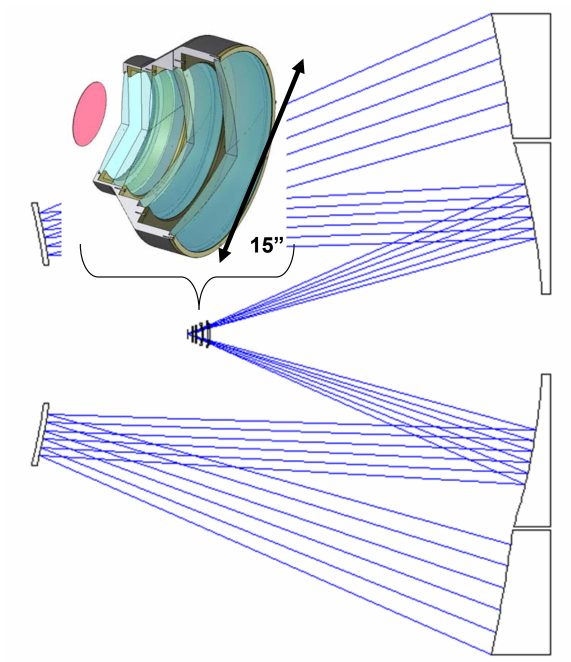

FIGURE 4-36 Commissioning camera lens for use with the 3-mirror telescope. Red disk is the focal plane, which is easily captured with a commercial CCD sensor.

Top End Assembly (TEA) with appropriate dummy masses to emulate the real telescope conditions will be assembled to support the M1 and M3 metrology stations.

After performance of M1 and M3 is validated over varying elevation ranges on the sky, the surrogate TEA will be removed and M2 will be integrated on the telescope mount. A prime focus camera system supported off the LSST camera hexapod will view images reflected off the three-mirror telescope system. Similar dummy mass will be included to replicate the LSST science camera geometry. As designed, the three-mirror telescope provides a good on-axis image and enables a relatively simple camera design (size of lenses driven by the FOV desired). Figure 4-36 shows the four-element commissioning camera lens designed for use with the three telescope mirrors.

A core integration team will be on the summit for the duration of assembly, testing, and commissioning efforts. Extra teams will augment this group to focus on specific subsystems. The specialized teams will have been part of the subsystem assembly and testing at the factory so they will arrive at the summit with direct experience and will be well prepared to carry out the reassembly. The LSST summit team will be staffed with support personnel that are positioned to remain with the project from initial installation, through integration, commissioning, and into operations.

Base Facility

Remote support functions for the telescope will take place in a base facility to be located in the AURA compound (*recinto*) in La Serena (M3 2006c). The base facility will house the local data processing center, a remote observing room, administrative offices, conference room, work spaces, and storage. Plans and cost estimates have been developed for this 1460 m^2 facility.

4.4 CAMERA

The LSST camera contains a 3.2-gigapixel focal plane array comprised of 189 4K × 4K CCD sensors with 10 µm pixels (Figure 4-37). The sensors are deep depletion, back-illuminated devices with a highly segmented architecture that enables the entire array to be read out in 2s or less. The detectors are grouped into 3 × 3 arrays called "rafts." All the rafts are identical; each contains its own dedicated front-end and back-end electronics boards, which fit within the footprint of its sensors, thus serving as a 144-megapixel camera on its own. The rafts and associated electronics are mounted on a silicon carbide grid inside a vacuum cryostat, with an intricate thermal control system that maintains the CCDs at an operating temperature of −100 °C. The grid also contains four sets of guide sensors and wavefront sensors at the edge of the field.

The entrance window to the cryostat is the third of three refractive lenses. The other two lenses are mounted in an "optics housing" at the front of the camera body. The camera body also contains a mechanical shutter and a filter exchange system holding five large optical filters, any of which can be inserted into the camera field of view for a given exposure. A sixth optical filter will also be fabricated; this filter can replace any of the five via an automated procedure accomplished during daylight hours.

Specific features of the LSST camera leading to the reference design concept are discussed in more detail in the following sections. A breakout of estimated costs among the various camera subsystems is shown in Figure 4-38.

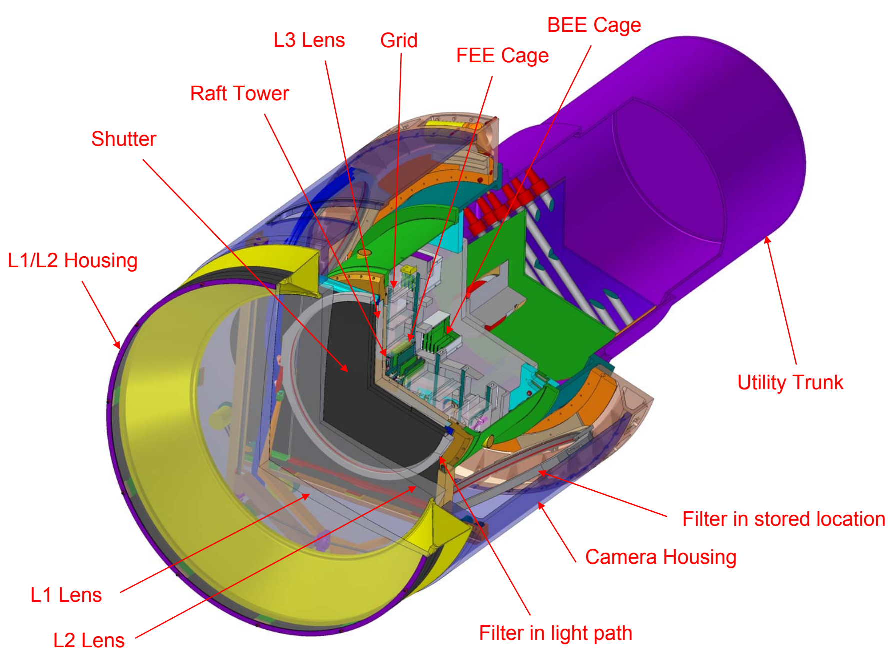

FIGURE 4-37 Cutawsay drawing of the LSST camera. The camera body is approximately 1.7 m in diameter and 3 m in length. The optic, L1, is 1.57 m in diameter.

4.4.1 Focal Plane Components

The 3200 cm^2 unvignetted field of view of the camera is tiled with a mosaic of 21 identical square rafts configured with science sensors and an additional four triangular rafts carrying wavefront sensors and guiders (Figure 4-39).

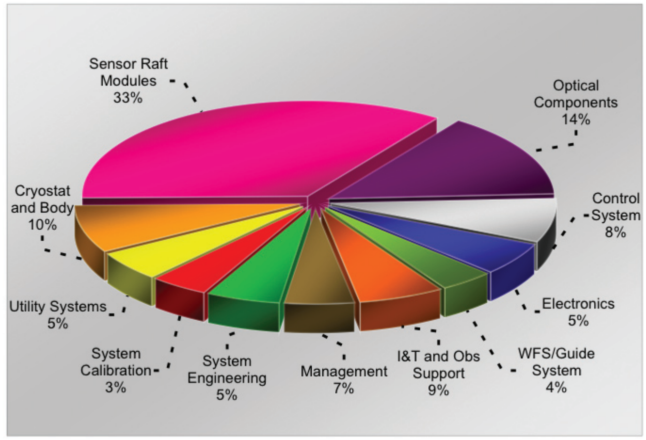

FIGURE 4-38 Percentage distribution of camera subsystem costs

Sensor Design

The "heart" of the camera is the science sensor. Its design is driven by the following key science and engineering requirements (O'Connor et al. 2006):

- **High quantum efficiency** from the blue to 1000 nm. For silicon sensors, this implies a depletion depth >75 µm and implementation of the sensor in a back-illuminated configuration with a thin entrance window.

- **Minimal detector contribution to the point spread function.** To reduce charge diffusion, the sensor must be fully depleted, and a high internal field must be maintained within the depletion region. That, in turn, requires the use of high resistivity substrates, high applied voltages, and back-side contacts. To reduce ~1 µm light spreading prior to photoconversion, the depletion depth cannot exceed 150 µm.

- **Fast focal ratio (f/1.2).** The implied depth of field is short, requiring <10 µm peak-to-valley focal plane flatness with piston, tip, and tilt adjustable to ~1 µm.

- **Wide field of view.** 189 4K × 4K sensors are required to cover the 3200 cm^2 focal plane. To maintain high throughput, the sensors must reside in 4-side buttable packages and be positioned next to each other with gaps of less than a few hundred µm.

- **Fast readout.** The camera must be read out in less than 2s. To reduce the read noise associated with higher readout speeds, the sensors must be highly segmented. The large number of I/O connections then requires that the

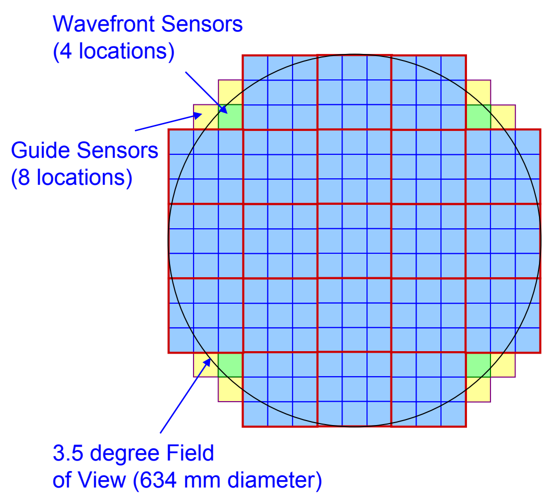

FIGURE 4-39 With its 189 sensors, the focal plane of the camera images 10 square degrees of the sky per exposure.

detector electronics be implemented within the cryostat to maintain a manageable number of vacuum penetrations.

Our reference sensor design is a charge-coupled device (CCD) with a high degree of segmentation, as illustrated in Figure 4-40. A 4K × 4K format was chosen because it is the largest footprint consistent with good yield. Each amplifier will read out 1,000,000 pixels (2000 × 500 sub-array), allowing a pixel readout rate of 500 kHz per amplifier. Key parameters for the sensor design are listed in Table 4-3.

While some CCD fabricators have achieved one or more of these parameters in individual cases, the LSST sensors will be the first to incorporate the full set of performance parameters in a single sensor optimized for surveys.

A 3½ year development program is underway with several industrial vendors to demonstrate a robust manufacturing process and test capability for the unique LSST sensors. This will be followed by a two-year production phase. Figure 4-41 is an image of an LSST wafer from a vendor that shows the highly segmented partitioning.

The sensors are mounted in aluminum nitride (AIN) packages as shown in Figure 4-42. Traces are plated directly to the AIN insulator to route signals from the CCD to the connectors on the back of the package. The AIN strongback provides a stiff, stable structure that supports the sensor, keeps it flat, and extracts heat via a cooling strap.

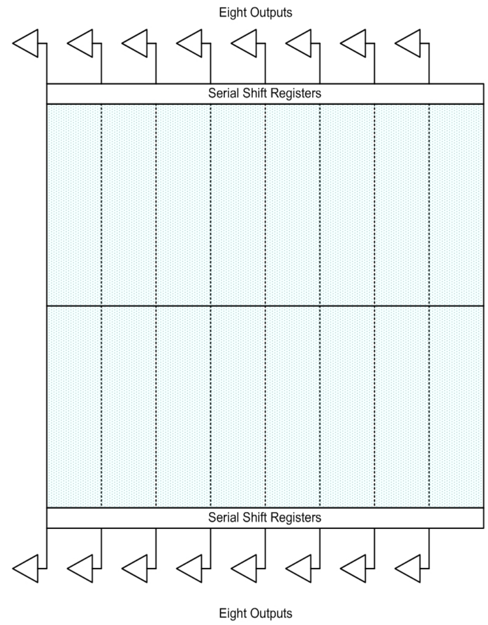

FIGURE 4-40 The camera CCD layout, showing the 16 readout ports per sensor. This high degree of segmentation enables fast readout with moderate clocking speed.

TABLE 4-3 LSST Sensor Design Key Specifications

Parameter	Value
Pixel size	10 μm
Depletion Depth	75-150 μm
Format	4000 × 4000 pixels
Segmentation	Sixteen 2000 × 500 pixel sub-arrays
Total no. of output amplifiers	16
Anticipated gain	3–5 μV/e$^-$
Parallel clocking	4-phase (4 poly layers)
Serial clocking	3-phase
Contiguous column length	500 pixels (100 arcsec)
Guard ring	100 μm
Read Noise	5 e$^-$

Raft Towers

The sensor packages are assembled into raft "towers" (Figure 4-43). Each raft consists of a 3 × 3 array of CCDs and an electronics/cooling package contained in a supporting cage located directly behind the sensor array. A set of high-density flex cables and cooling straps connects the sensors to the

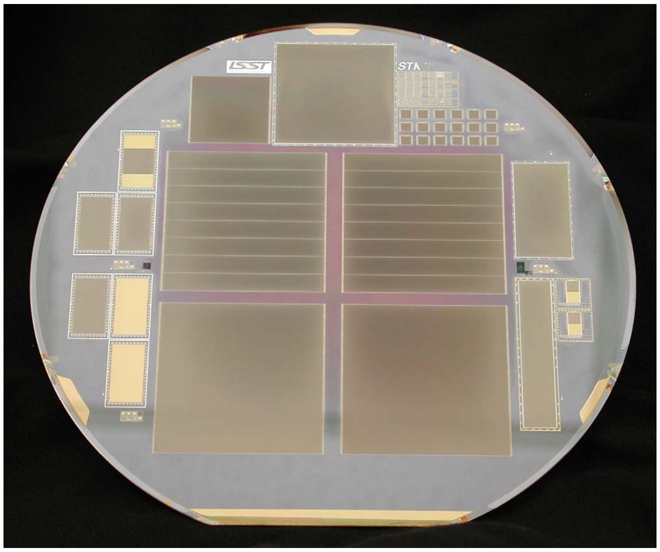

FIGURE 4-41 *An image of an LSST wafer fabricated during the sensor study phase of development. The top two die show the segmented architecture. Additional die on the wafer are used for testing purposes.*

front-end boards and cooling sinks. A candidate sensor concept has the packages attached to the raft structure with a 3-point kinematic mount, which provides fine tip, tilt, and piston adjustment to align the nine sensors parallel to a fiducial reference plane.

Wavefront and Guide Sensors

Four special purpose rafts, mounted at the corners of the science array, contain wavefront sensors and guide sensors. Wavefront measurements are accomplished using curvature sensing, in which the spatial intensity distribution of stars is measured at equal distances on either side of focus. Each curvature sensor is composed of two CCD detectors, with one positioned slightly above the focal plane, the other positioned slightly below the focal plane. The CCD technology for the curvature sensors is identical to that used for the science detectors in the focal plane, except that the curvature sensor detectors are half size so they can be mounted as an in-out defocus pair. Detailed analyses have verified that this configuration can recover the wavefront to the required accuracy.

These four corner rafts also hold two guide sensors each. The guide sensors monitor the locations of bright stars at a frequency of ~ 30 Hz to provide feedback for a loop that controls the guiding of the telescope. We are investigating the use of hybrid-CMOS detectors for this guide sensor application.

4.4.2 Electronics

The main subsystem of the camera electronics is the science array readout. Other main functions are the readout of the guide sensors, wavefront sensors, and the camera control electronics.

Science Array Readout

The architecture of the science array readout is dictated by the four science requirements: 1) low read noise, 2) high dynamic range, 3) fast readout, and 4) low and stable cross-talk.

A salient feature of the readout is its raft-based structure. Each raft has an associated electronics package with no interconnection to any other raft.

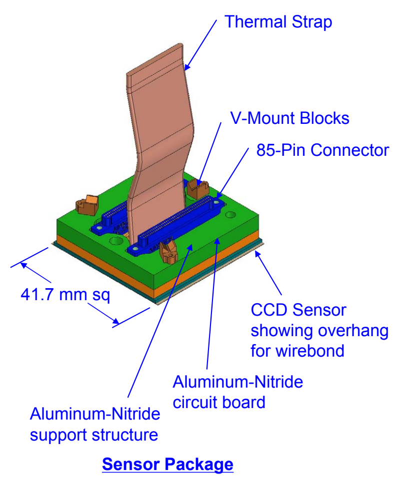

FIGURE 4-42 *The sensor package with details of the connections and mounting fixtures*

Each of the rafts is connected to a single timing and control module. In this way, strict synchronicity is ensured across the entire focal plane array.

The raft-based architecture implies that the readout is located within the cryostat. This choice was motivated by the high degree of segmentation in the sensors, which leads to approximately 150 bond pads per sensor. The array of 189 sensors implies that the focal plane has some 30,000 connections to the readout. Since the cryostat is a high-vacuum vessel with strict contamination requirements, the self-contained architecture we have adopted eliminates the need for, and risk of, high-vacuum feed-through on all those connections. To control contamination due to outgassing from cables and electronics boards, we will carry out extensive materials testing of all component parts.

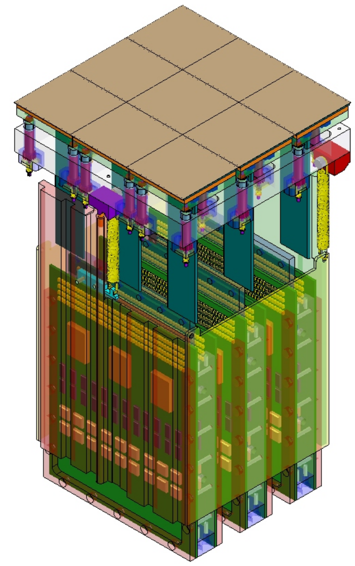

FIGURE 4-43 The raft tower is the main building block of the camera. The front end electronics boards are mounted in a cage directly under the 3 x 3 array of sensors.

Each raft electronics package is subdivided into two distinct units referred to as front-end electronic (FEE) units and back-end electronic (BEE) units. The division of functionality between the FEE and BEE is listed in Table 4-4.

The FEE unit is very close to the focal plane array and operates at a temperature that is near, and perhaps below, the focal plane temperature of –100 C. The BEE is located at the back flange of the cryostat and operates at a much warmer temperature between –40 °C and –20 °C. The FEE/BEE flex cable is sufficiently long, approximately 25 cm, to minimize the heat load to the FEE by the large temperature differential across it. Existing large mosaic cameras have a much smaller number of readout channels and do not use custom-designed application-specific integrated circuits (ASICs).
The sensor segmentation proposed by the LSST requires the use of these ASICs in order to accommodate the large electronics channel count.

TABLE 4-4 Division of Tasks between Front-End Electronic (FEE) and Back-End Electronic (BEE) Units

FEE Tasks	BEE Tasks
Analog signal processing	Analog to digital conversion
CCD clock line drivers	Data collection and formatting
CCD bias generation	Data transmission via optical fiber to Data Acquisition (DAQ)
Temperature sensors	Copper cable interface to timing & control module
Sensor heaters	Miscellaneous “housekeeping” or “slow-control” functions, such as generation of programming voltages for FEE bias and clock levels and ADCs for temperature monitoring

FEE and BEE Packaging

Each raft of nine sensors has a total of 9 × 16 = 144 parallel readout channels in an area of approximately 12 cm × 12 cm. This is an extraordinarily high channel density by standards of astronomical cameras constructed to date. The maturity of custom CMOS

ASIC design allows for the implementation of state-of-the-art signal processing in packages small enough to meet the requirements imposed by these channel densities. The BEE utilizes commercial 16 bit, 1 MSPS (mega samples per second) analog to digital converters (ADCs) to digitize the signal-processed signals from the FEE. Since all channels operate in parallel, 144 ADCs are required per raft FEE package. Commercial-off-the-shelf ADCs are available in chip-scale packages that simultaneously fulfill all of these requirements. The BEE package consists of six printed circuit boards, each of which contains twenty-four ADCs. Early FEE prototyping will concentrate on achieving low temperature operation with low crosstalk using a discrete front end board; high density prototypes using ASICs will be tested in a second phase. Prototypes of the BEE will demonstrate state-machine generation of timing signals and high density using advanced commercial ADCs, field-programmable gate arrays (FPGAs), and memory.

4.4.3 Cryostat Design and Assembly

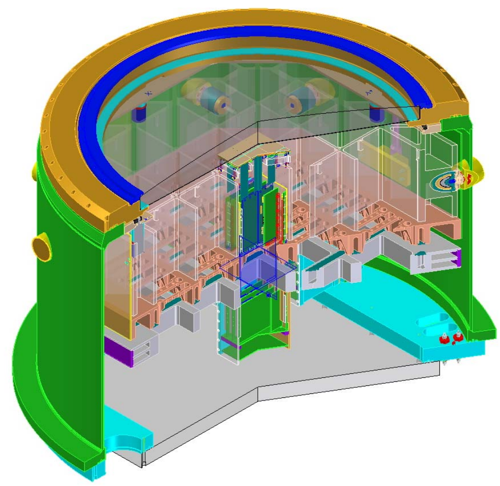

FIGURE 4-44 The cryostat housing holds the grid assembly. The raft towers are mounted into the square slots of the grid.

The cryostat body (Figure 4-44) forms a vacuum vessel operating at a base pressure of <10^{-6} torr, containing all the optical readout elements of the LSST camera with the L3 optic forming the front window. The cryostat houses the cold focal plane array, the raft tower, and the support, alignment, and thermal control hardware. The rafts are supported on a silicon carbide (SiC) grid structure. Finite element analysis of the grid structure, illustrated in Figure 4-45, has shown that the grid design maintains gravity flexure within tolerances at all angles of elevation. The cryostat also contains a copper cryogenic plate (–120 °C) for heat extraction from the focal plane and the front end electronics, and a second copper cold plate (–40 °C) for heat extraction from the back end electronics modules. There is additional instrumentation for sensor alignment and miscellaneous thermal, vacuum, and contamination monitoring.

The body of the cryostat is a cylindrical stainless steel vessel (~94 cm outer diameter, ~74 cm long) with front and rear stiffening flanges. Contamination due to outgassing by the various components contained in the cryostat is a concern because of the possibility that outgassed materials will migrate to the cold focal plane surface and modify the quantum efficiency of the detectors. To minimize this effect, the vessel is electro-polished on all inner surfaces, and all flanged penetrations are either welded or contain double Viton O-ring seals. In addition, traditional gettering techniques (e.g., molecular sieve, charcoal, ion pumping) are incorporated, and the focal plane is cryogenically baffled to avoid a direct path for contaminants.

The front flange of the vessel supports the L3 optic (acting as a vacuum window), the instrumentation ring, and the Z-support flexures for the grid. The rear flange is comprised of an outer annulus containing glass-metal seal instrumentation and power feed-throughs, fiber optic feed-throughs, and inlets and outlets for vacuum-jacketed coolant lines. The central-octagonal part of the rear flange supports a turbo-molecular pump and vacuum valves and

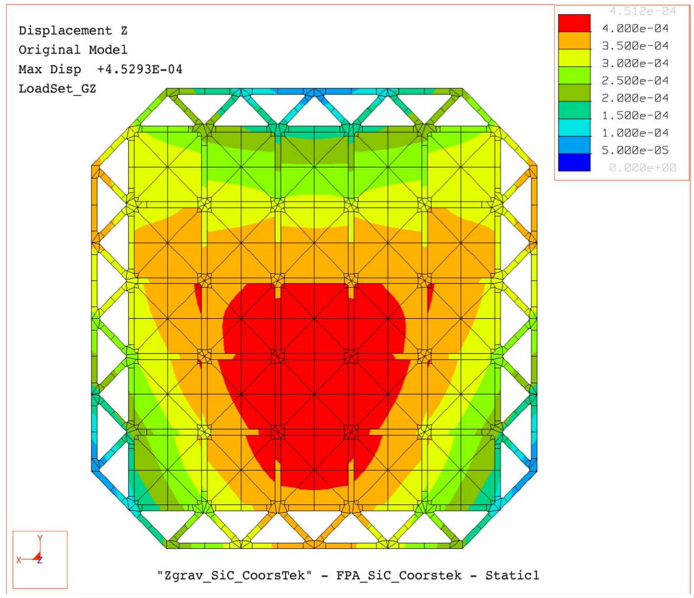

FIGURE 4-45 The results of the FEA of the SiC grid structure which shows the z-displacement under 1 G gravity. The color scheme indicates the z-displacement in microns within the active part of the grid, assuming three support flexures.

instrumentation. It is removable for raft loading and access. The sidewalls of the cryostat support the cryogenic and cold plates.

A full thermal model of the cryostat has been performed to verify that the thermal gradients across the sensors are within tolerances required to maintain flatness and uniformity of the quantum efficiency at the red end of the spectrum. A sample output from this calculation is shown in Figure 4-46, where we plot the temperature distribution that is realized across the SiC grid. In Figure 4-47, we list the predicted thermal loading of the sensor due to radiation, conduction, and dissipated power for a model in which each individual detector is cooled via a single thermal strap attached to the cryogenic cold plate held at -120°C. This figure also shows

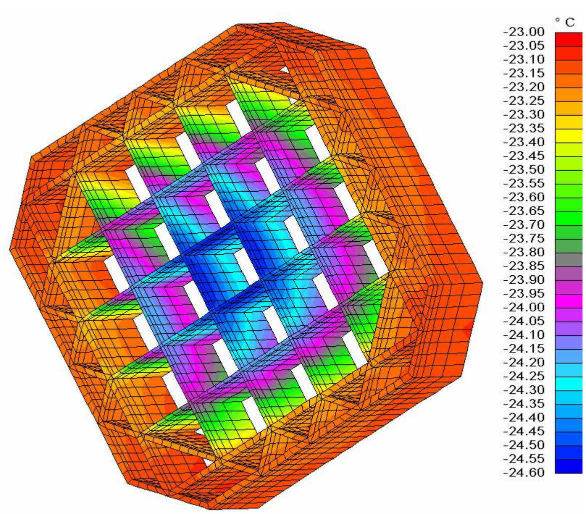

FIGURE 4-46 The SiC grid shown with the temperature distribution realized across the grid with radiative cooling from the shroud and raft electronics. The SiC grid is slowly cooled from ambient temperature to -100 deg. C by this method and when thermal stabilization is reached, a maximum temperature deviation across the grid is ~1.5 deg. C which is within the tolerance specified.

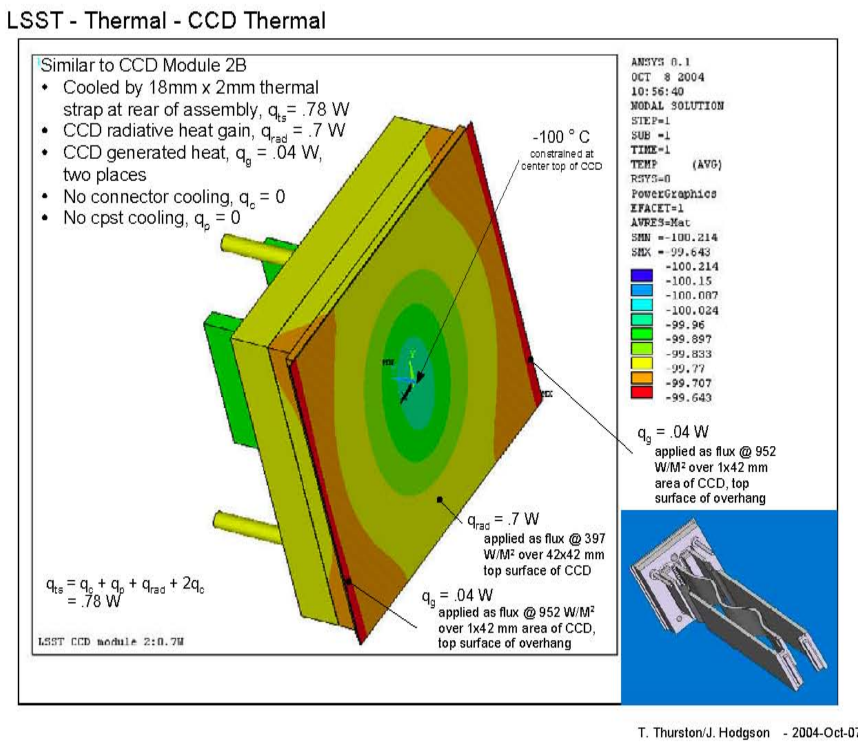

FIGURE 4-47 FEA model of thermal gradients with one cooling strap. Heat loading by radiation, conduction, and generated heat is indicated.

a plot of the thermal gradients that are generated across the sensor. The predicted temperature variation is less than 0.3° K, which leads to thermal distortions less than 0.1μm, well within our flatness tolerance.

Focal Plane Assembly and Metrology

The cryostat design allows for rapid and safe assembly of the imaging focal plane, while meeting metrological-flatness requirements (overall 10 μm P-V) at –100 °C. It also permits close proximity of the FEE and BEE cards to the focal plane sensors. Its modularity allows easy removal of raft towers and/or BEE boxes for repair/replacement.

After each raft insertion, non-contact metrology (described below) is performed across the partially filled focal plane to verify flatness. Raft and back end electronics insertion proceeds until complete; this is followed by insertion of ancillary sensors and other instrumentation and the remaining thermal insulation. When complete, the central back flange plate is closed. The front is closed with a temporary flat optical window. The cryostat can then be purged, evacuated, and processed to operating temperature, while monitoring contaminants. Testing allows the verification of the focal plane flatness and alignment and the performance of all sensor electronics under normal camera operating conditions. In the final assembly step, the temporary window is replaced with L3.

The stringent focal plane flatness requirement coupled with the fast assembly and testing time of the camera is met by engineering precision and modularity in advance of the assembly into each of the sub-components (sensor, raft, and grid), and by selecting materials and mounting techniques that introduce negligible distortions under gravity and in the transition from 20 °C to −100 °C. In addition, techniques have been developed to non-invasively verify flatness at all stages of the assembly and testing and to monitor any changes over the lifetime of the instrument. To achieve these requirements, sensors are kinematically mounted to the raft plates, with individual adjustment, to make each raft identical ("standard raft"). Rafts are then mounted kinematically to the grid without further adjustment, having been pre-adjusted with the grid mounts at each location to accept a standard raft.

In advance of cryostat assembly, the silicon carbide (SiC) grid structure is ground and polished to a modest surface flatness (~10 μm). Each raft bay contains three balls brazed in place to form the mating points for the raft's vee mounts. Pre-measurement of the SiC surface makes it possible to pre-select balls that are well within the adjustment range of the final step of processing: with a raft-like fixture and a coordinate measuring machine, the brazed balls are measured and adjusted (shimming, lapping, etching, or plating) to accept a standard raft and place it in the reference plane. During assembly of the focal plane, rafts are inserted onto the grid and lie within the reference plane.

This process relies only on the reproducibility of the kinematic mounting scheme itself. The choice of SiC for the grid and AIN for the rafts ensures that the impact of gravitational deflection of the grid or rafts on the sensor surface is small (<0.25μm). While the kinematic mounts are nearly immune to small differences in thermal expansion, the excellent thermal properties of AIN and SiC (close CTE and excellent conductivity) make it possible to cool down and warm up the complex structure in a modest amount of time while introducing mechanical distortions across the grid (<0.3 μm) that are well below the expected reproducibility of the kinematic mounts.

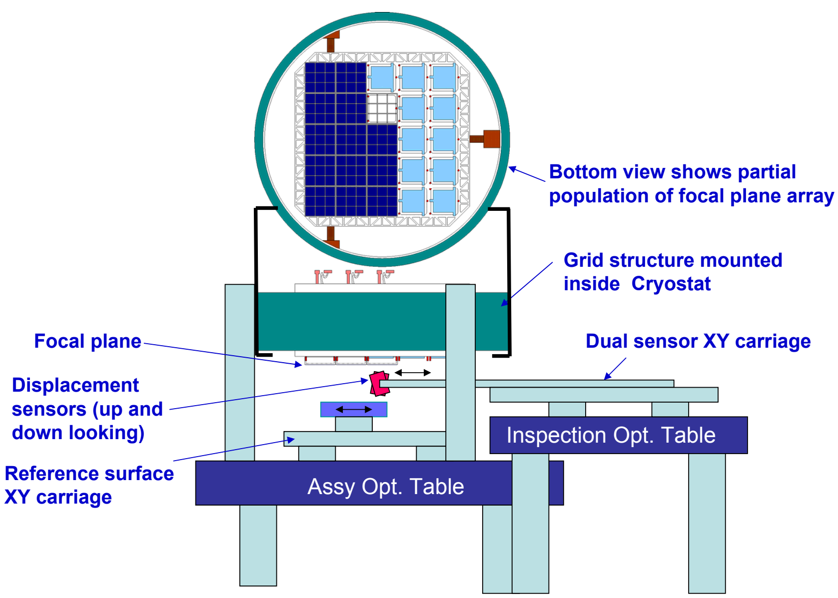

FIGURE 4-48 *The non-contact metrology station, used to check the alignment of the focal plane during raft tower installation. This facility also has the capability to verify alignment after L3 is mounted on the cryostat, and the system is pumped down and cooled.*

External metrological techniques are used to rapidly verify focal plane flatness at all stages of the cryostat assembly, including subsequent closure and cool down. In addition, an in-situ system is present in the cryostat to continually measure changes in the sensor flatness and/or the raft alignment.

External Metrology

During integration of the rafts on to the grid, a non-contact metrology station, shown schematically in Figure 4-48, is utilized in the clean room. Two commercial non-contact laser triangulation heads are mounted together on an XY stage, itself mounted to an optical bench to the side of the cryostat. These sensors measure displacements to <100 nm, over a range of several cm. One sensor points up to the focal plane surface while the other looks down at an optical flat of 20-cm diameter, which is itself translatable on a larger reference surface on a separate optical bench. The XY stage is programmed to move over the region where the flat is located, taking 300 data points in 10 min. Summing the sensor outputs removes most mechanical and thermal effects of the stage. The flat is moved to a new position, overlapping the previous one, and the process is repeated. In this manner, the surface of the focal plane can be fit and "stitched" together using the multiple overlapping references. Reconstruction of the surface figure to 0.4 μm P-V with 0.1 μm repeatability has been obtained in the lab.

After closure and pumpdown of the cryostat, measurements are taken through either the flat optical window or through L3, substituting one head for a longer range laser-head with 1-μm precision and applying fixed corrections for the optical path.

In-Situ Metrology

An in-situ metrology capability is also built into the cryostat for independently tracking changes to the focal plane and raft alignment during camera assembly and during its operational lifetime. This system utilizes a direction pattern generator, which provides a direct measurement

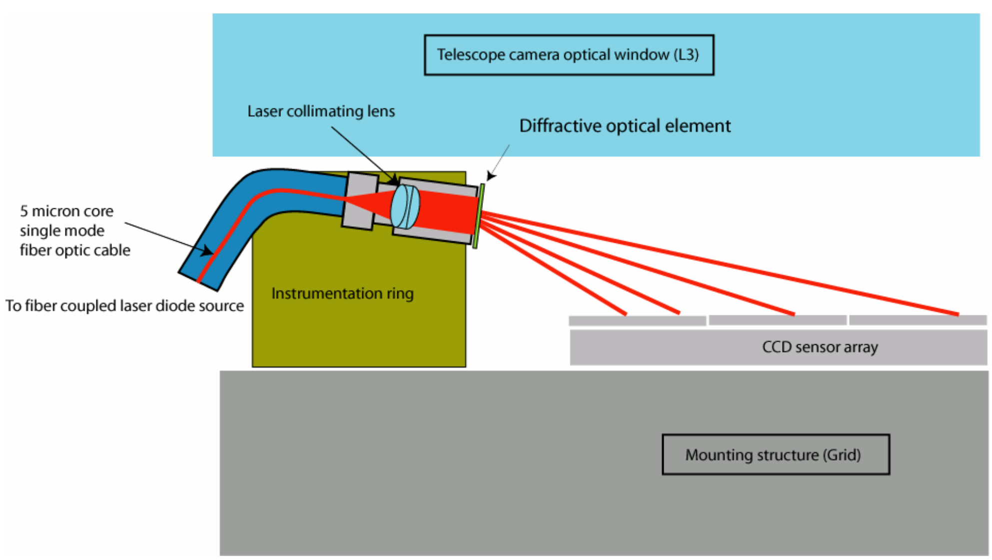

FIGURE 4-49 The diffraction pattern generator, which projects an array of laser spots at grazing incidence to the focal plane. The positions of these spots on the array provide a sensitive measure of the focal positions of the CCD's.

of the sensor flatness "on demand" by projection of a known and fixed diffraction pattern of spots onto the sensor surfaces at a steep grazing angle. The pattern is generated by a set of four pairs of projector heads mounted in the instrumentation ring over the focal plane (Figure 4-49). Each head has a single mode 5-µm core fiber optic input that is coupled monolithically to a focusing element and a diffraction grating.

4.4.4 Camera Body and Mechanisms

The camera body is shown in Figure 4-50. The camera and back flange, which support the cryostat, L1/L2 subassembly, shutter, and filter exchange system, are the main structural elements of the camera. Near the front end of the camera, a stiffening ring in the housing provides the interface to the L1/L2 subassembly and support for the shutter and changer. This also reinforces the housing to reduce gravity-induced deformations, because it must support the camera optics to tight tolerances.

At the back end, the stiff, back flange of the camera provides the mount for the cryostat stand-off structure. The cryostat is inserted into the camera housing from the back end, allowing for servicing without disturbing the rest of the camera. The filter carousel is also mounted to the back flange, and camera body cabling and purge lines are routed through the flange on the way to the utility trunk. The back flange also serves as the mechanical interface to the telescope through double bolt circles to the camera rotator.

Shutter

The shutter mechanism controls the length of time that the sensors on the focal plane are exposed to an image. For the LSST, the camera shutter must precisely control the exposure time for all pixels on the focal plane, while operating continuously in a tightly packaged space between the filters and L3 lens. Because the LSST camera is within the optical path of the

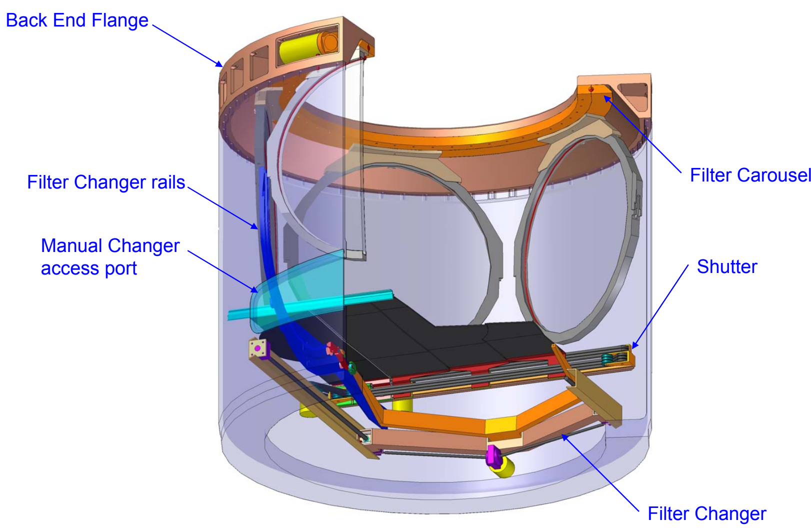

FIGURE 4-50 A schematic of the camera body showing the filter exchange mechanism, the large optical filters, and the shutter assembly

telescope mirrors, there is not adequate room for a large, one-piece flat shutter. The LSST design simply segments the one-piece shutter into three flat blades, which are stacked up when the shutter is open, then slid over the field of view in a coordinated motion to close. To maintain exposure time uniformity across the FOV, two stacks of three blades each are used, as shown in Figure 4-51.

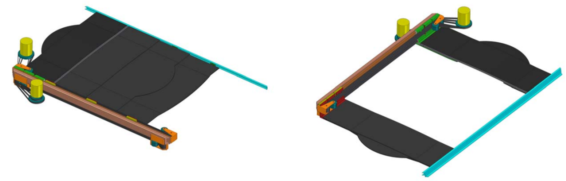

FIGURE 4-51 The shutter assembly includes three flat blades which stack when the shutter is open.

This design is a modification of that of existing shutters, which have been developed by a group in Bonn, Germany. They have extensive experience building large focal plane shutter systems. Our concept is similar to the Bonn approach in that it uses conventional, off-the-shelf parts and materials commonly found in industries requiring clean, fast, high-cycle-life automated systems.

Filter Exchange

The filter exchange system provides the capability to change filters remotely during the course of nighttime telescope operations. Given the survey nature of LSST and the multiple science programs that the observatory is supporting, filter changes will be a regular occurrence during a typical night of operation. Thus, the filter exchange system must reliably and safely change filters quickly and with minimal impact on the nightly observing schedule.

The LSST filter exchange system (Figure 4-50) is designed to store five filters within the clean environment of the camera housing and to exchange any of them into the camera field of view with a two-minute turn-around time. In addition, the system can cleanly and safely swap out any of the five on-board filters with a sixth during part of a daily servicing access. The exchange system includes three major sub-assemblies to perform these functions. First, the filter carousel stores the filters within the camera housing where they are available for use in the field of view. Second, the changer mechanism remotely actuates to pick a filter out of the storage carousel and move it into the field of view for a particular observational task. Third, a manual changer is used for storing unused filters and swapping them into the carousel during a daytime access. While each of these sub-assemblies is a physically separate component, they all work together to perform the necessary exchange functions.

4.4.6 Camera Optics

Refractive Optics

The camera refractive optics consist of three fused silica lenses (L1, L2, L3) with clear apertures of 1.55 m, 1.10 m, and 0.69 m, respectively. L1 has an edge thickness of 3.4 cm and a center thickness of 8.2 cm. L2 has a central thickness of 3.0 cm. The space between L2 and the filters is 36 cm, which is adequate for accommodating the filter interchange. L3 is the vacuum barrier for the cryostat containing the detector array. The central thickness is specified in order to provide a significant safety margin for potential fracture of L3 due to the pressure differential. Empirical data show that a thickness ratio of ~12 is adequate to provide this safety margin, which yields a thickness of 6.0 cm for the 69-cm diameter lens. An finite element calculation of L3 as a vacuum barrier is shown in Figure 4-52.

The camera optics design can meet all design requirements using fused silica substrates. The Corning manufacturing process for fused silica can produce glass of the required size and quality. (Corning's estimates of cost and schedule to produce the required fused silica glass have been used as input for the LSST camera optics schedule and budget.) Although the required lenses are large, we have confirmed with potential vendors that a substantial industrial base exists for fabricating them; we have also verified that our requirements are consistent with optics that have been produced for other programs. A key aspect of the fabrication involves the use of relatively simple null tests that have been factored into the optical design. All null tests use a retro-reflecting mirror to test the lenses in double pass transmission. A perfect point source is re-imaged onto itself with a wavefront error less than 1/20 wave at 633 nm.

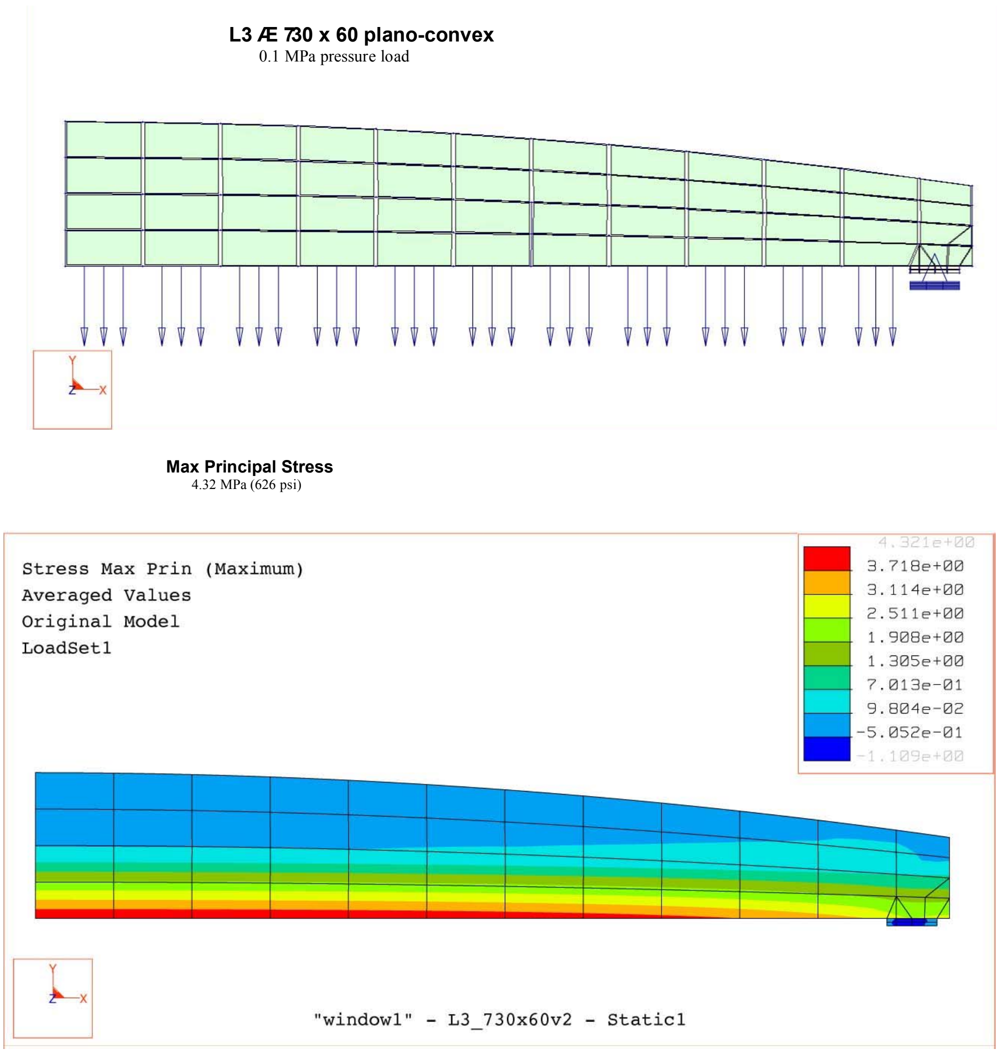

FIGURE 4-52 Finite element model (upper) and plot of maximum principal stress (lower) for L3 acting as a vacuum seal. Positive stress is tensile in the normal sign convention used here.

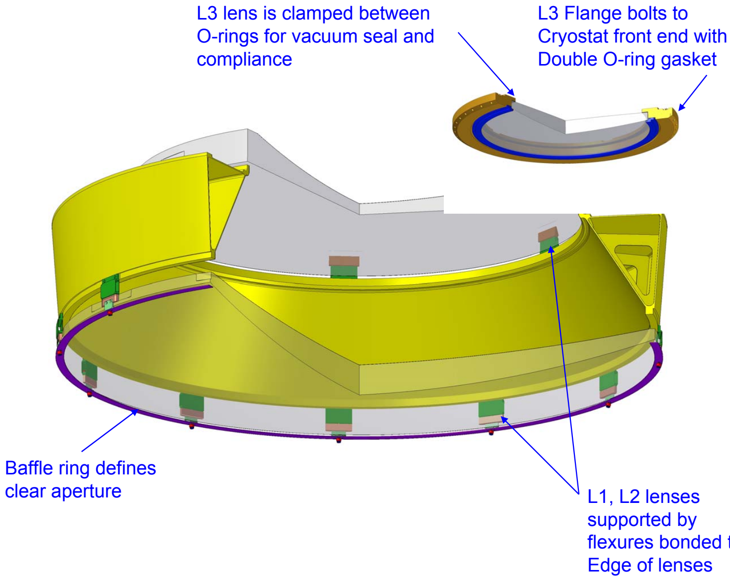

FIGURE 4-53 The optics L1 and L2 are contained in a housing equipped with centering and alignment fixturing. L3 is the entrance window of the cryostat.

The optic mounts provide the interface between the lenses and the structures that support them. These interfaces may include adjustment capability for alignment purposes. If it proves cost effective, potential vendors may be asked to supply the optics already installed and qualified in the mounts. In terms of optical correction, L1 and L2 are the most significant optics in the camera. Their alignment to one another is therefore more critical, and the pair forms logical datums to use in aligning the focal plane array, L3, and the filter. In addition, the size and location of L1 and L2 on the front end of the camera (away from the shutter and filter changer) make an aligned and tested subassembly a reasonable package for an optics vendor to produce (Figure 4-53). A repeatable kinematic interface allows the front-end assembly (L1 and L2) to be easily removed from and reinstalled on the camera housing to provide ample access to components.

Finite element modeling confirmed the feasibility of creating an independently stiff housing on which to attach both L1 and L2 optics. Between the housing and L1 and L2 are, respectively, 18 and 6 radial-motion flexures, which allow for differential thermal expansion. This number of supports limits gravity-induced edge ripple to less than 1 µm P-V. It is expected that each optic will be measured in the mount during fabrication and that repeated polishing-measurement steps will require simple and repeatable remounting.

Filters

The filters consist of multi-layer dielectric interference coatings deposited on fused silica substrates. The baseline design has the first surface of the filters concentric about the chief ray in order to keep the angles of the light rays passing through the filters as uniform as possible over the entire range of field positions. The central thickness and the curvature of the second surface are optimized for image quality.

TABLE 4-5 Filter Bandpass Definitions

	Half-Maximum Transmission Wavelength		
	Blue Side	Red Side	Comments
U	350	400	Blue side cut-off depends on AR coating
G	400	552	Balmer break at 400 nm
R	552	691	Matches SDSS
I	691	818	Red side short of sky emission at 826 nm
Z	818	922	Red side stop before H20 bands
Y	948	1060	Red cut-off before detector cut-off

The current LSST filter complement (*u,g,r,i,z,Y*) is modeled on the system used for the Sloan Digital Sky Survey (SDSS), which has been successful in a variety of applications (e.g., photometric redshifts of galaxies, separation of stellar populations, photometric selection of quasars, among others). Extension of the SDSS system to longer wavelengths (*Y*-band) is mandated by the increased effective redshift range that will be achievable with the LSST due to deeper imaging. The addition of a *u*-band improves the robustness of photometric redshifts of galaxies, stellar population separation, and quasar color selection, and will provide significant additional sensitivity to star formation histories of detected galaxies.

The current LSST baseline design has a goal of 1% relative photometry and as such defines the general features of the filter set. The filter set wavelength design parameters and the approximate FWHM transmission points for each filter are shown in Table 4-5. The selected science filter vendor will be required to:

1. Establish procedures to distribute a uniform coating over the entire filter surface.
2. Determine optical quality of the coatings necessary for rejecting out-of-band transmissions.
3. Develop techniques to ensure wavelengths of pass band are met.
4. Monitor and record techniques to reduce variations in coatings.

The process and equipment used to fabricate the prototype filters and deposit the optical coatings will be developed so that the full size LSST Filters (750-mm diameter) can be coated. Colored glass filters or laminates must be scalable to the full size filters. Several vendors have produced interference filters that approximate the LSST filter size. Initial testing of these filters indicates that vendors can produce filters with the coating uniformity required by the LSST scientific programs.

4.4.7 Camera Utilities

The camera utilities systems provide the infrastructure required to maintain camera components in suitable environments to meet all performance requirements. The environments needing control are the thermal environment, vacuum environment of the cryostat, and ambient gas environment of the camera housing. The thermal, vacuum, and purge systems also provide capabilities for operations and servicing needs of the camera.

The three utilities systems are comprised of two segments: components mounted to the camera or installed on the utility trunk on the back of the camera and hardware installed on the floor of the dome in the observatory mechanical room. The ground hardware includes chillers, pumps, and a reservoir for the cryogen circuits that feed the camera with temperature-controlled fluid. The cryogen fluid is transported to the camera by vacuum-insulated transfer lines. These are a combination of rigid lines running the telescope structure and flexible transfer lines that run through the three wraps on the telescope. A custom transfer line is used to route the fluid lines through the Top End Assembly spiders without cooling the spiders or the surrounding air. Figure 4-54 shows the different thermal network zones and the interactions between them.

Thermal System

The camera thermal system consists of seven discrete thermal zones, each characterized by a different operating temperature. All zones use one source of cold fluid for heat rejection, 3M Novec HFE7200, in combination with resistive heaters. The cryogen fluid circuits are flow-balanced on the back of the camera to ensure that all thermal zones receive adequate cooling capacity.

- ***Ground thermal zone.*** The ground thermal zone includes two Polycold chiller units, coupled to a plate-and-fin heat exchanger used to cool the Novec working fluid. The end product of the entire ground thermal zone is a supply of temperature-stabilized cryogen to the camera.

- ***Cryostat thermal zones.*** There are three thermal zones within the vacuum of the cryostat. The warmest zone is the cold plate and BEE modules, which operate at approximately –30 °C. The second and coldest zone is that of the cryoplate, grid, and FEE modules. These nominally operate at –120 °C and are thermally stabilized with trim heaters to minimize temperature-induced distortions in the grid. The third zone is that of the sensor packages, which operate at approximately –100 °C.

- ***Camera thermal zones.*** Outside of the cryostat, there are two thermal zones in the camera housing. First, the temperature of the outer cylinder of the cryostat is held stable with resistive heaters. This protects the cryostat and L3 lens from developing cold spots for either heat leaks or larger-than-predicted radiative heat transfer and may not be implemented if testing during early production shows it is not required. The second thermal zone is that of the camera purge gas itself. Purge gas is actively circulated throughout the camera volume to eliminate hot spots from motors and other heat sources and to reduce the possibility of natural convection cells, which could affect the seeing conditions within the camera. The gas is cooled with a simple fan heat exchanger, and the gas temperature is adjustable to keep the camera temperature as close to the dome air temperature as possible.

- ***Utility trunk thermal zone.*** The final thermal zone is the gas-cooling system in the utility trunk on the back of the camera. This removes process heat from the electronics that are resident in the trunk and maintains the trunk temperature near the dome ambient temperature, as well.

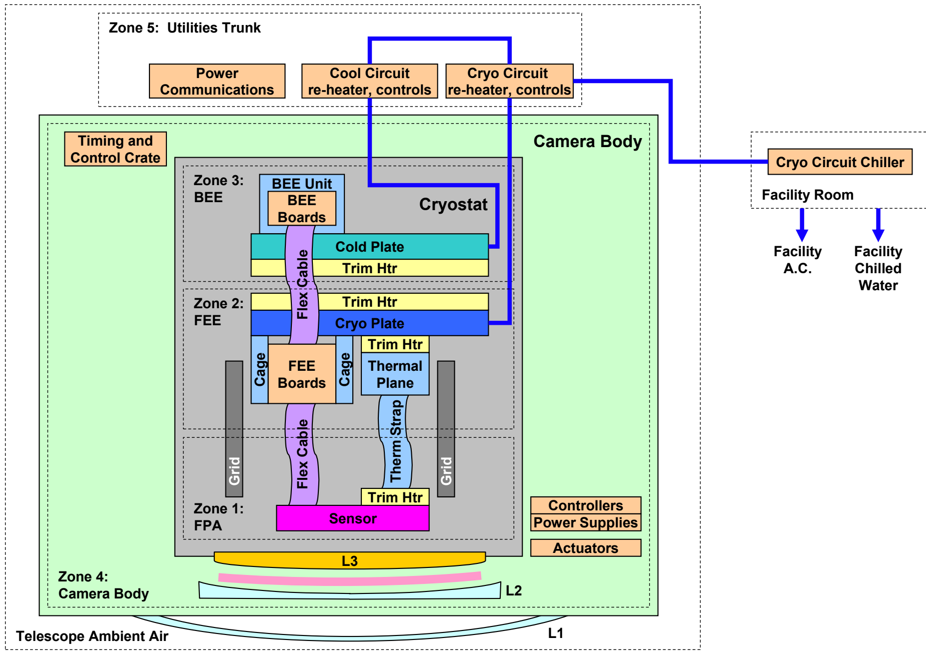

FIGURE 4-54 *Thermal network zones and associated connections.*

Vacuum System

The camera includes three vacuum systems. The first system maintains the cryostat at or near a high-vacuum state during operation. This system provides continuous pumping on the back of the camera and is supplemented by the significant cryo-pumping capacity of the cold surfaces in the cryostat during operation. The camera is pumped down and much of the water vapor removed by way of additional pumps that mount to the back end of the camera. These are then valved off and removed during operation, and holding pumps maintain the clean, high vacuum.

The second vacuum system provides the insulating vacuum for services on the telescope that are not part of the cryostat vacuum. This includes the insulating vacuum for the cryogenic valve box on the back of the camera and the pump-out grooves between the double O-ring seals on the cryostat. None of these volumes is connected to the cryostat vacuum, but all require active pumping and monitoring. The vacuum pump for this system is installed on the ground, with a vacuum line running up the telescope. This will provide low- to medium-vacuum at the camera.

The third vacuum system provides the insulating vacuum for cold fluid systems on the ground. It is also used for evacuating portions of the cryogen liquid transfer lines in preparation for disconnecting them prior to removing the camera from the telescope for servicing.

Purge System

There are three purge systems on the camera. First, the camera housing is purged with clean, dry nitrogen. This is temperature-controlled to provide a constant-temperature environment for the camera and telescope optics. The temperature is slaved to the outside ambient temperature to minimize thermal effects in the optical beam. Purge gas is plumbed up to the camera from gas cylinders on the ground, then cooled as needed on the camera.

The second purge system maintains a constant pressure head on the cryogenic fluid reservoir on the ground. This helps stabilize the fluid flow rates and pressures despite the varying hydrostatic head of the fluid due to telescope re-pointing.

The third purge system is used for purging the cryogenic fluid from the transfer lines and for warming them up prior to removing the camera. This system ties into the lines, is manually controlled completely from the ground, and is only used sparingly.

4.4.8 Camera Control System (CCS)

The camera control system (CCS) manages the activities of the various camera subsystems and coordinates those activities with the LSST observatory control system (OCS). The CCS consists of a set of modules (nominally implemented in software), each of which is responsible for managing one camera subsystem. Generally, a control module is a long-lived “server” process running on an embedded computer in the subsystem. Multiple control modules may run on a single computer or a module may be implemented in “firmware” on a subsystem. In any case, control modules must exchange messages and status data with a master control module (MCM). The main features of this approach are: 1) control is distributed to the local subsystem level; 2) the systems follow a “master/slave” strategy; and 3) coordination is achieved by the exchange of messages through the interfaces between the CCS and its subsystems.

The LSST camera is a complex, heterogeneous device consisting of many subsystems that have significantly different behaviors. Nevertheless, the subsystems must act in concert if the camera is to perform adequately. While the focal plane is controlled electronically, the shutter, filter changer, cryogenics, and vacuum subsystems require motors, encoders, pressure sensors, temperature sensors, limit switches, voltage monitors, relays, and other electro/mechanical devices. The camera control system is designed to play several roles in the operation of the camera.

- The CCS maintains the state of the camera as a whole, and thus is able to orchestrate the sequence of operations that enables data collection.
- It receives status information from the subsystems and is able to detect and report faults. (This does not include the detection of safety-related error detection, which is the responsibility of individual subsystems.)
- It provides the camera interface to the observatory control system (OCS), responding to commands from and reporting status information to the OCS;
- It is the interface between the human operators and the instrument. It can respond to reconfiguration commands, ensuring orderly transitions between different modes of camera operation;
- It provides the hardware and software necessary to receive the data streams generated by the camera and transmit them to downstream clients, such as the Data Management system (DM).

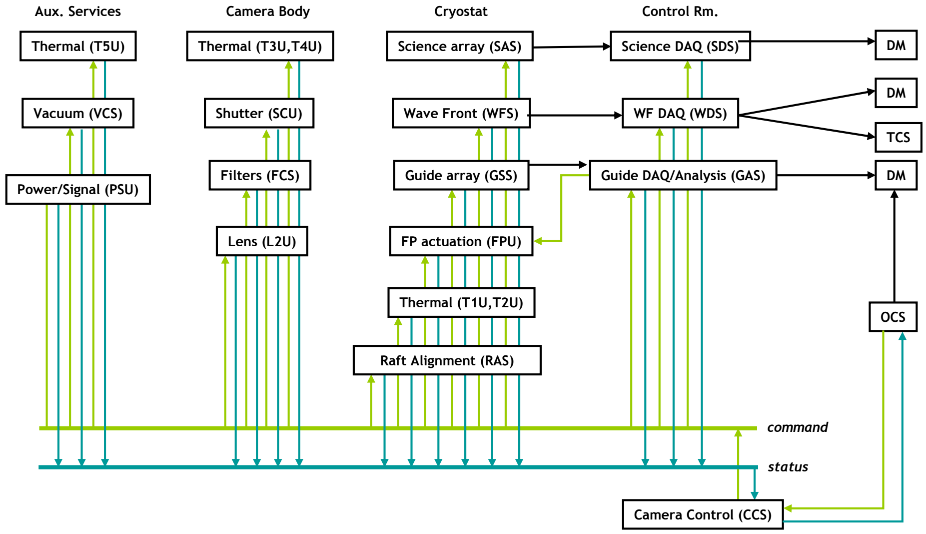

FIGURE 4-55 *Flow chart of the camera sub-systems and their interactions.*

The logical relationship between the CCS and other subsystems is illustrated in Figure 4-55. (Subsystems that produce data to be analyzed are distinguished from those that simply produce status information.) Data do not flow through the CCS, but directly from producers to consumers.

The most demanding component of the data acquisition and control system for the LSST camera is the subsystem that receives the pixel data streams transmitted from the CCD raft readouts of the Science Array Subsystem (SAS).

The Science Array Data Acquisition Subsystem (SDS) has the principal responsibility for receiving science data from the camera and presenting these data to a set of downstream clients. The SDS is both a hardware and software system. The SDS is partitioned into two major deliverables. First, the client interfaces are "data sinks" and are logically part of the SDS; second, the Camera Data Source (CDS) provides the camera data to the clients. The CDS consists of three types of components: four Raft Network Adapters, which do the speed and format interfacing for the science data coming from the rafts, form the interface to the camera; one RAID storage; and one network switch to interface the above components to the DAQ subnet. (The Raft Network Adapter is the only component of the CDS not available as a commodity product.)

4.4.9 Integration and Test

Camera integration and testing will be carried out in a clean room and high-bay facility required by the deep access needed to install and replace raft towers. The facility has been identified at SLAC and includes approximately 5000 square feet of clean room with anteroom, 1000 square feet of white room lab and office space, and 3000 square feet of high-bay shop floor with crane access outside the clean room. The clean room includes a dedicated air-handling system, HEPA filtration, and top-down air flow. The air-conditioning system also includes a re-heater for humidity control, and the room is held at positive pressure. The clean room also includes a full-coverage anti-static floor and ESD personnel and equipment grounding throughout. Other features include nitrogen and air purge systems, a dry-pipe sprinkler system, Web-accessible temperature and humidity monitoring with automated alarm system, and key-

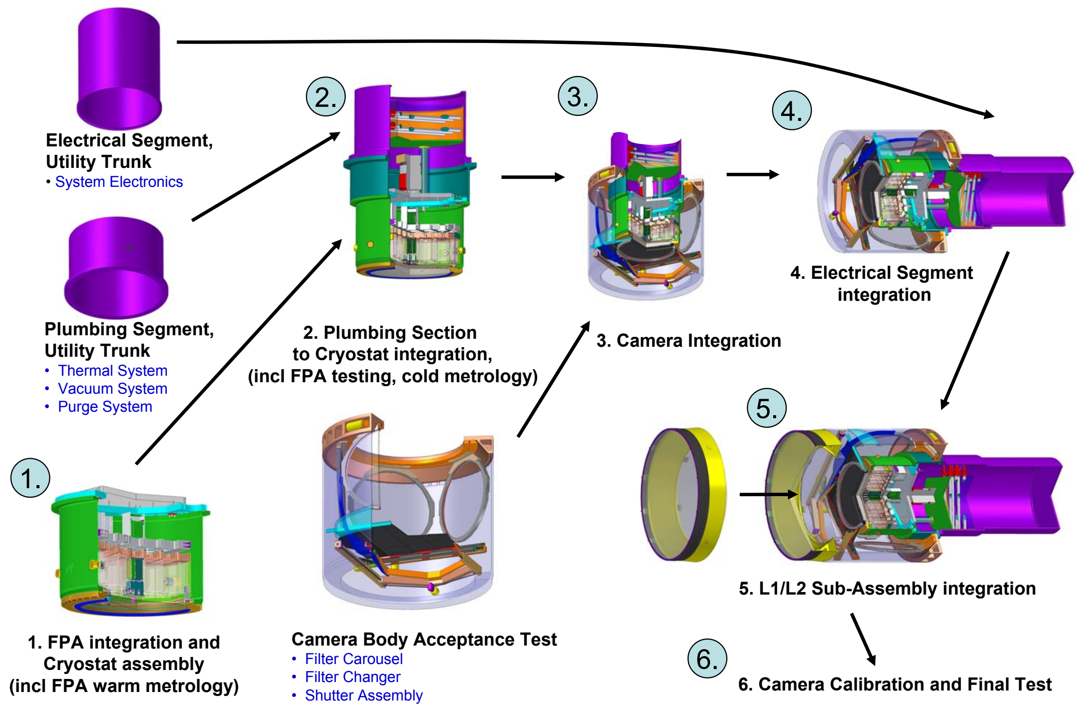

FIGURE 4-56 *LSST camera main sub-components and assembly sequence*

code entry to control access. The building also contains its own subnet with firewall to the main SLAC computer network, and Ethernet drops throughout the clean rooms. Integration-and-test advance planning has been used to identify requirements imposed on facilities by the large physical size of the camera.

As discussed in previous sections, the camera will be assembled from pre-tested and verified sub-assemblies and components. Figure 4-56 reviews the main sub-assemblies and the integration sequence into the camera assembly. The flow of camera integration is generally a serial flow of work, divided into three large time groupings: The first phase involves loading the raft towers in the cryostat and completion and testing of the cryostat and focal plane assembly. The second phase is the completion of the camera assembly with its mechanisms, environmental control and monitoring systems, and camera optical lenses and filters. The final phase is the full system-level testing and calibration activities.

After tests and calibration are completed at SLAC, the camera will be shipped to Cerro Pachón. The timeline for installation and testing at the telescope is shown in Figure 4-57.

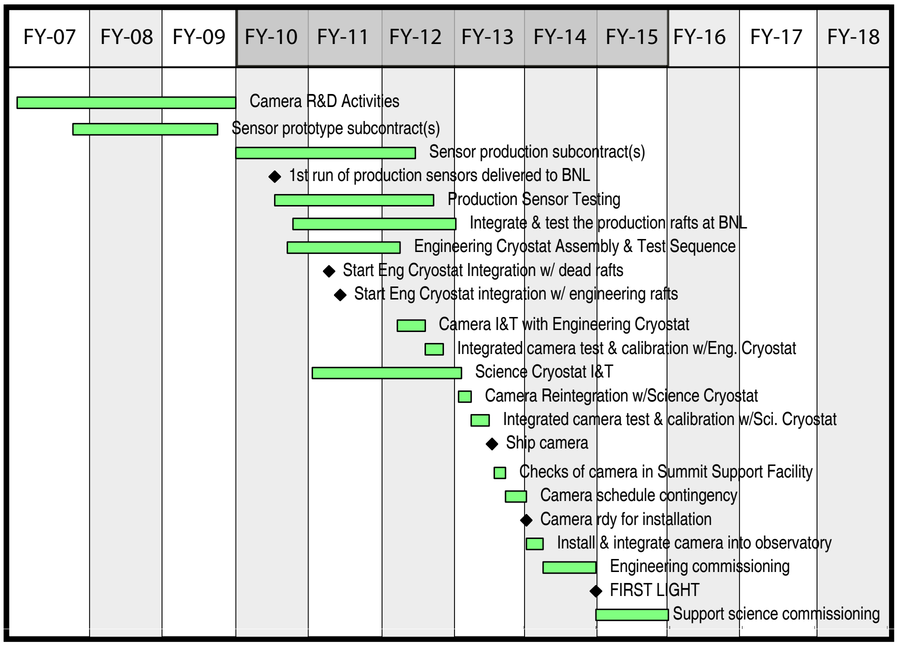

FIGURE 4-57 Timeline for onsite commissioning of the LSST camera

4.5 DATA MANAGEMENT SYSTEM (DMS)

The unique but enabling challenge for the LSST Observatory is to provide fully calibrated public databases to the user community in a way that facilitates the frontier science described in Chapter 2—but without precluding new lines of research that cannot be anticipated today. The nature, quality, and volume of LSST data will be unprecedented. New algorithms will have to be developed and existing approaches refined in order to take full advantage of this resource. In this section, we describe the reference design for the LSST data management system (DMS).

The LSST data management system will be a world-class computational facility that will process and deliver LSST data to meet these challenges. The computational facilities and data archives of the LSST DMS will rapidly make it one of the largest and most important facilities of its kind in the world. To tackle these challenges and to minimize risk, LSSTC and the data management subsystem team have taken four specific steps:

1. The LSSTC has put together a highly qualified data management team. The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Champaign-Urbana will serve as the primary LSST archive center; the Infrared Processing and Analysis Center (IPAC) at the California Institute of Technology will develop custom software and review software baseline costing models; Lawrence Livermore National Laboratory (LLNL) and Stanford Linear Accelerator Center (SLAC) led the effort to size and design the computational infrastructure and database system. Google Corporation is providing advice and collaborating on file system structures and public data access.

2. The DM team has closely examined and continues to monitor relevant developments developments within, and recruited key members from, precursor astronomical surveys

including 2MASS, SDSS, the ESSENCE and SuperMACHO projects, the Deep Lens Survey, and the NVO.

3. The DM team is monitoring relevant developments within the physics community and has recruited members with experience in comparable high energy physics data management systems. Specifically, members were recruited from the BaBar collaboration at SLAC and the Atlas collaboration at BNL.

4. The DM team defined four basic “Data Challenges” (DC) to evaluate and certify the LSST DM architecture. Success with these challenges prior to the beginning of construction will provide confidence that DM budget, schedule, and technical performance targets will be acheved.

The DM subsystem risk assessment, which was based on the project and risk management processes described in Chapter 7, resulted in 39% contingency at this early stage in the D&D phase. While contingency is controlled at the project level, project management is fully aware of the uncertainties involved in generating large custom software packages.

The DM construction budget distribution by area is shown in Figure 4-58. Forty-five percent of the budget is associated with infrastructure, which primarily involves purchasing computational, storage, and communications hardware. The 27% applications component is dominated by labor to code data reduction and database related algorithms. Detailed budget and schedule for data management are presented in Chapter 7 (Project Execution Plan).

4.5.1 DMS Data Challenges (DC)

A substantial portion of the data management tasks involves custom software and algorithms scaled to process and store petabytes of data quickly. Traditionally, such data management projects are not as crisply defined as construction projects and are well known for being difficult to scope. This is particularly true for astronomical algorithm development, which is dependent on characteristics of the image data that cannot all be known in advance. To meet this challenge, we have assembled a team that includes both astronomers and software professionals. During the D&D phase, this team is tasked with developing prototype algorithms

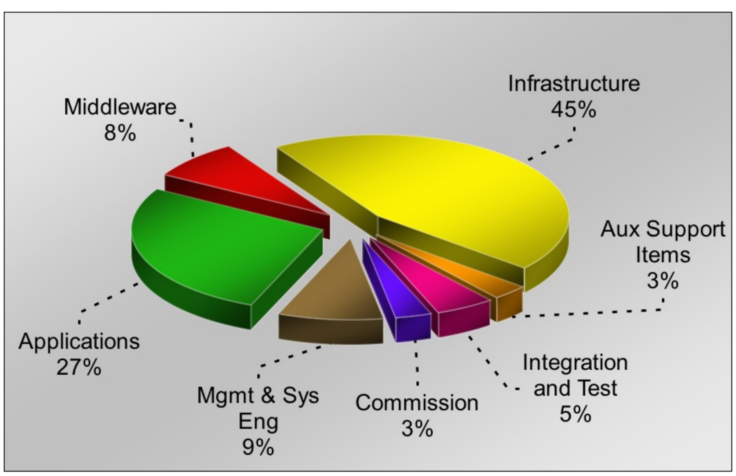

FIGURE 4-58 *Data management budget estimates by area*

TABLE 4-6 Goals, Configuration, and Completion Dates for Four LSST Data Challenges (DC)

DC No.	Completion Date	Data Challenge Goal	Configuration
DC 1	July–Oct 2006	Validate infrastructure and middleware, scalability to 5%	Simulated data and applications running on 16 nodes on TeraGrid. Simulated real-time data flows from Mountain to Base through Nightly Pipelines. Ingest into database, transfer to Archive, re-run Nightly Pipelines. Purdue cluster represents Mountain, NCSA represents Base Facility, SDSC represents Archive Center
DC 2	Sept–Nov 2007	Validate nightly pipeline algorithms, scalability to 10%	Same as above plus: 64 nodes, prototype pipeline environment, Image Processing, Detection, Association pipelines. Pre-release Moving Object Pipeline (from Pan-STARRS collaboration)
DC 3	Sept–Nov 2008	Validate science pipelines, data quality, reliability, scalability to 15%	Same as above plus: 128 nodes, prototype Deep Detection and Calibration Pipelines
DC 4	May–July 2009	Validate open interfaces and data access, scalability to 20%	Same as above plus: 256 nodes, VO interfaces, query services

for all stages of the DM pipeline, which will be tested on precursor and simulated data. The team will be expanded and continue with algoithm development during construction.

To exercise and validate the prototype algorithms and to evaluate and verify the LSST plans for data management, the team has adopted an early test and validation strategy that involves scaled functional and performance tests called Data Challenges (DC). The four DCs are described in Table 4-6. While the scaling goals of 20% or less shown in the table are modest and constrained by available D&D resources, the promising results of DC 1 (see Section 4.5.6) demonstrate the efficacy of this approach.

Each Data Challenge is a scaled version of the LSST DMS construction phase and represents a project in its own right. Each DC includes the following:

— Scoping document establishing the objectives, metrics, and technical requirements for the DC
— Project plan with a schedule and resources
— Unified Modeling Language-based design of the system
— Standard development environment with software standards and coding conventions
— Formal review, version, and change control process
— Formal, automated build and test environment
— Report analyzing the process, technical design, and technical results[1]

4.5.2 Requirements and Deliverables of the Data Management System

The LSST Data Management System (DMS) will deliver:

1 These documents are available in the LSST project archive (Collection-472).

- **Image archives:** An image data archive of over two million raw and calibrated 12 GB scientific images, each image available within 24 hours of capture
- **Alerts**: Verified alerts of transient and moving objects detected within 60 seconds of the image capture (30 s goal)
- **Catalogs**: Astronomical catalogs containing billions of stars and galaxies, richly attributed in both time and space dimensions and setting a new standard in uniformity of astrometric and photometric precision. This includes a query toolkit with associated computational, storage, and communications infrastructure
- **Quality Assurance**: A process including documentation to ensure data integrity with full tracking of data provenance

All data products will be accessible via direct query or for fusion with other astronomical surveys. The user community will vary widely. From students to researchers, users will be processing up to terabyte-sized sections of the entire catalog on a dedicated supercomputer cluster or across a scientific grid. The workload placed on the system by these users will be actively managed to ensure equitable access to all segments of the user community. LSST key science deliverables will be enabled by computing resources co-located with the raw data storage. Table 4-7 shows the content of the data products and the cadence on which they are generated.

The system that produces and manages this archive must be robust enough to keep up with the LSST's prodigious data rates and will be designed to minimize the possibility of data loss. This system will be initially constructed and subsequently refreshed using commodity hardware rather than "cutting edge" technology to ensure affordability, even as technology evolves.

Functions of the Data Management System

The principal functions of the DMS are to:

- Process the incoming stream of images generated by the camera system during observing to generate and archive the nightly data products.

TABLE 4-7 Key Deliverables and Processing Cadence of the LSST Data Management System

Processing Cadence	Image Category	Catalog Category	Alert Category
Nightly	Raw science image Calibrated science image Subtracted science image Noise image Sky image	Source catalog (from difference images) Object catalog (from difference images) Orbit catalog	Transient alert Moving object alert
Data Release	Stacked science image Template image Calibration image RGB JPEG Images	Source catalog (from calibrated science images) Object catalog (optimally measured properties)	Alert statistics and summaries

- Periodically process the accumulated nightly data products to measure the properties of fainter objects and to classify objects based on their time-dependent behavior. The results of such a processing run form a data release (DR), which is a static, self-consistent data set for use in performing scientific analysis of LSST data and publication of the results.

- Make all LSST data available through an interface that utilizes, to the maximum possible extent, community-based standards such as those being developed by the Virtual Observatory (VO).

LSST images, catalogs, and alerts are produced at a range of cadences in order to meet the science requirements. Alerts are issued within 60 seconds of completion of the second exposure in a visit. Image data will be released on a daily basis, while the catalog data will be released twice during the first year of operation and once each year thereafter. Section 4.5.5 provides a more detailed description of the data products and the processing that produces them. All data products will be accompanied by a record of the full processing history (data provenance) and a full range of metadata services. A unified approach to data provenance enables a key feature of the DMS design: data storage space can be traded for processing time by recreating derived data products when they are needed instead of storing them permanently. This trade can be shifted over time as the survey proceeds to take advantage of technology trends, minimizing overall costs. All data products will also have associated data quality metrics that are produced on the same cadence as the data product. These metrics will be fed back to the Observatory Control System for use in schedule optimization and are available for scientific analysis as well.

The latency requirements for alerts determine several aspects of the DMS design and overall cost. Two general classes of events can trigger an alert. The first is an unexpected excursion in brightness of a known object or the appearance of a previously undetected object such as a supernova or a GRB. The astrophysical time scale of some of these events may warrant follow-up by other telescopes on short time scales. These events must be recognized by the pipeline, and the resulting alert data product sent on its way within 60 seconds (30 s goal). The second event class is the detection of a previously uncatalogued moving solar system object. The LSST field scheduler may not generate another observation of this object for some days; again, prompt follow-up may be appropriate, but in this case, a one-hour latency is acceptable.

Finally—and perhaps most importantly—automated data quality assessment and the constant focus on data quality leads us to focus on *key science deliverables*. No large data project in astronomy or physics has been successful without key science actively driving data quality assessment and DM execution. For this activity, metadata visualization tools will be developed or adapted from other fields and initiatives such as the Virtual Observatory (VO) to aid the DM operations and the science collaborations in their pursuit of systematic errors system-wide.

4.5.3 Sizing the Data Management System

A fundamental question is how large the LSST data management system must be. To this end, a comprehensive analytical model has been developed driven by input from the requirements specifications. Specifications in the science and reference designs translate directly into numbers of detected sources and astronomical objects, and ultimately into network bandwidths and the size of storage systems. Specific science requirements of the survey determine the data quality that must be maintained in the DMS products, which in turn determine the algorithm requirements and the computer power necessary to execute them (LSST Documents 1193, 1779, 1989, 1990, 1991, and 2116).

Key input parameters include the processing operations per data element, the data transfer rates between and within processing locations, the data ingest and query rates, the alert generation rates, and latency requirements.

- **Processing requirements** were extrapolated from the functional model of operations/element and from existing pre-cursor pipelines (SDSS, DLS, SuperMACHO, ESSENCE, and Raptor) adjusted to LSST scale.
- **Storage and input/output requirements** were extrapolated from the data model of LSST data products, pre-cursor schemas (SDSS, 2MASS), and existing database management system (DBMS) overhead factors (SDSS, 2MASS, BaBar) in pre-cursor surveys adjusted to LSST scale.
- **Communications requirements** were developed and modeled for the data transfers and user query/response load, extrapolated from existing surveys and adjusted to LSST scale.

The results of the functional and performance flow down exercise are presented in Table 4-8 and discussed below for processing, storage, and communications. The resulting performance and sizing requirements show the DMS to be a supercomputing-class system, with over 100 TFLOPS of aggregate processing power and correspondingly large data input/output and network bandwidth rates.

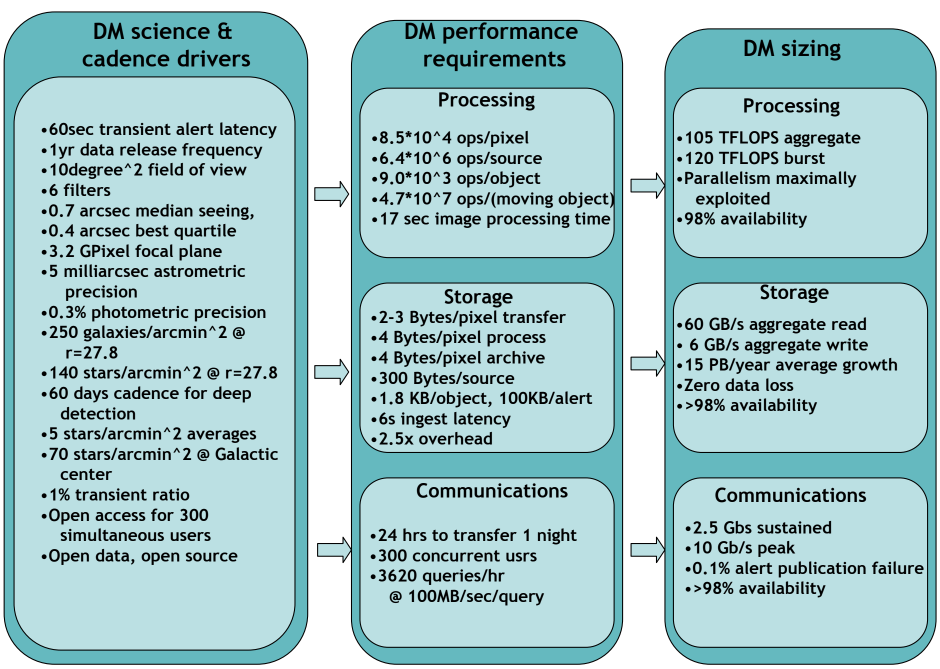

TABLE 4-8 The science and cadence drivers flow-down to produce the final requirements for sizing the DM system.

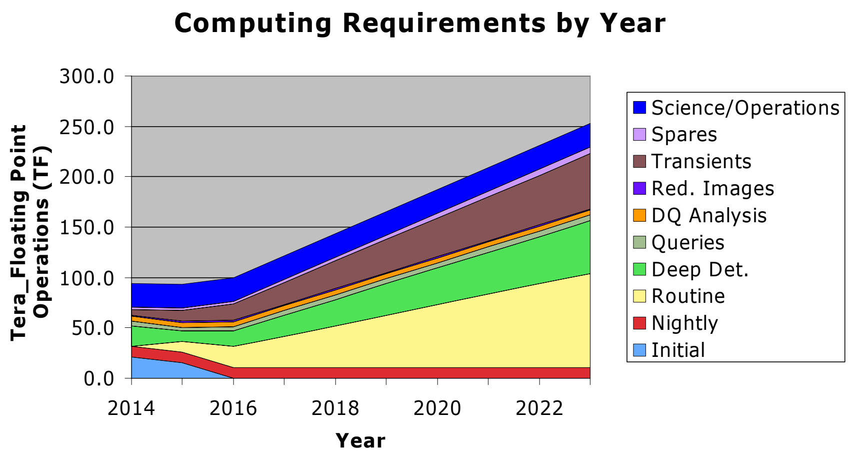

FIGURE 4-59 *The computational performance requirement for the DMS grows during the life of the survey*

Computational Sizing

In order to balance between near-real time and non-real time requirements, we have divided computing across four specialized facilities: Mountain Summit, Base Facility, Archive Center, and Data Access Centers. (Each facility is described in detail later in this section.) As shown in Table 4-8 and Figure 4-59, the total processing requirements for LSST start at approximately 100 TFLOPS in the commissioning period and first survey year and then grow linearly. These requirements are technically feasible now. (The growth trend is less steep than industry trendsin computing system throughput.) Though LSST processing requirements are considerable by the standards of current optical/infrared astronomical surveys, they are fairly modest when compared to the requirements of high energy physics, radio astronomy, and other scientific programs driving the top 500 supercomputing sites (see Figure 4-60).

The sizing of the computational requirements involved a conscious trade-off between the cost of saving all computed data products at all stages of processing vs. re-computing data products on demand. This resulted in a CPU-intensive model because image data expands from 15TB to over 100TB per night upon processing. Transferring/storing 100TB per night is cost-prohibitive; computing system throughput trends support the required re-processing cost effectively.

The strategy adopted requires that data be reduced at the Base Facility in order to meet alert latency requirements; raw image data and metadata are transferred to the Archive Center, where they are processed again. Only certain data products and the provenance needed to re-create the other data products on demand are stored permanently. Before every data release, all data previously processed are re-processed with the latest algorithms, calibration products, and parameters. Finally, the system is sized to support on-the-fly, “smart” re-processing of as yet unreleased data for other needs, such as end-user science. This includes regenerating individual reduced frames when necessary.

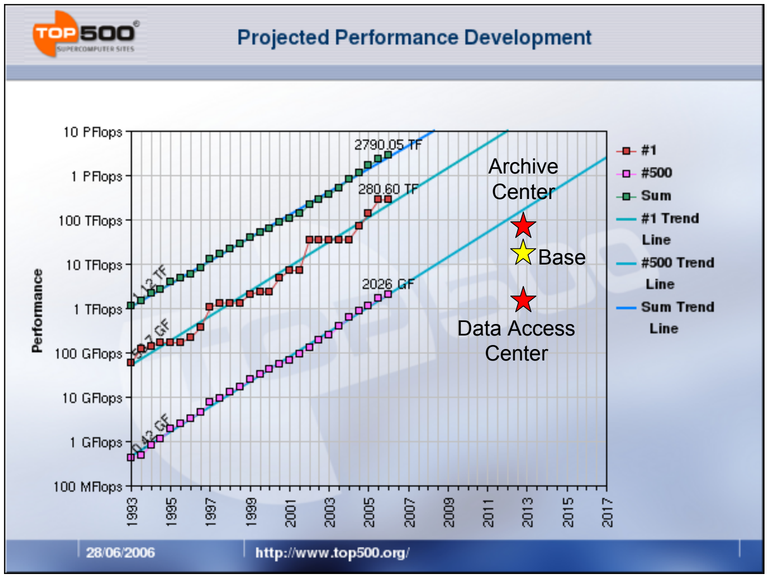

FIGURE 4-60 While the LSST computational requirements are large, growth trends for the top 500 computers show that by first-light, the LSST DMS will not be in the top 500. The central facility at NCSA will likely not be their largest computational cluster.

Input/Output Sizing

LSST data storage requirements are substantial: LSST will reach multi-petabyte scale during the first year of operations (Figure 4-61). While the image data will grow at a steady 6 petabytes/ year based on observing cadence and image size, a significant portion of the catalog data growth is a direct result of observing near or in the Galactic plane. The largest of the catalogs, the Source Catalog, will hold about 250 billion source measurements in the first data release (DR1) and 6.5 trillion in the last data release (DR10). Image data will be released on a daily

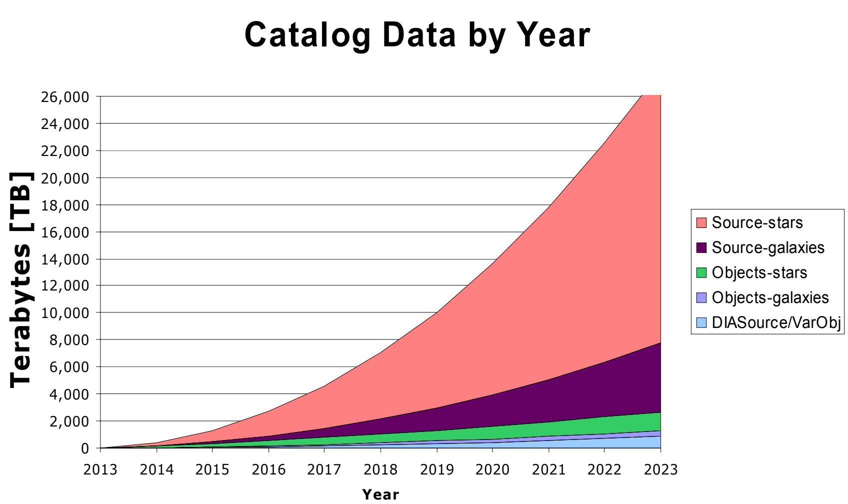

FIGURE 4-61 The DMS storage requirements are substantial, especially for astronomy, though well within technology limits.

basis and catalog data will be released twice in the first year and once each year thereafter, with DR1 projected for release in 2014. Once released, all data products will be permanently stored and immutable. Unreleased data collected and processed between releases will be updated nightly and will be immediately available; these data will carry the caveat that analytical results based on that data may be non-reproducible once the data are re-processed for inclusion in the new DR.

Row sizing for database storage estimates uses an interim schema based on experience with SDSS, 2MASS, and SuperMACHO. For purposes of expansion, 50% additional space has been reserved for as yet unknown attributes. A 150% factor has been applied to accommodate "persistent overheads," i.e. database indices, database overheads, and data replication. An additional 50% disk space is reserved for database swap space. Estimates of overheads and swap space are based on the experience of SDSS, 2MASS, and BaBar, and input from external database experts.

The two most recent data releases and the as yet unreleased data, along with corresponding persistent overheads and database swap space, will be served from fast disks, while older releases will be served from more economical slower disks and tape. Data indices will be kept for the most recent releases; older indices are re-created as needed from stored releases.

The LSST disk input/output requirements are estimated based on pipeline data access/ingest rates (including re-processing), indexing rates, and hypothetical end-user query load. The query load (interactive or automated) modeled is:

- — 50 simultaneous low-volume queries against 10 million objects in the catalog; response time 10 s; resulting data set: ~0.1 GB
- — 20 simultaneous high-volume queries against all objects in the catalog; response time 1 hr; resulting data set: ~6 GB

The system will support many times more users who may log in, but typically most users are in "think time mode" rather than constantly executing queries.

Long-haul Communications Sizing

Massive volumes of LSST data and data products must be moved from the Mountain Summit to the Base Facility in Chile, from the Base Facility to the Archive and Data Access Centers, and finally from those centers to the users.

The key driving requirements for the Mountain Summit to the Base Facility are the bandwidth required to transfer the raw and crosstalk-corrected image data for alert processing (10 Gbps) and reliability. Because this link closely follows the path of existing CTIO networks, LSST will manage and maintain this mission-critical link, providing the reliability and availability necessary to run the combined Mountain-Base infrastructure as a single component of the LSST system.

The required bandwidth from the Base Facility in La Serena to the Archive Center in Illinois is only 2.5 Gbps because the data only need to arrive at the destination before the next night of observing starts. However, this link must also be reliable to avoid data transfer bottlenecks, and there must be spare capacity to "catch up" in the event that a failure or slow-down occurs. Commercial carriers and research networks are the only realistic source of this capacity and reliability over this distance.

Within the U.S., LSST must provide open access to LSST data—not only to the scientific community, but also to diverse segments of non-specialist users whose curiosity may be triggered by LSST educational and public outreach programs. LSST will have portals to the public Internet and Internet2 for such educational and outreach efforts. For internal high-speed data transfers from the Archive Center to Data Access Centers, 10 Gbps will also be required; commercial and research-funded networks are both capable of this capacity and reliability.

4.5.4 System Reliability Requirements

Based on the survey reference design, we established an overall availability percentage for the DMS as a constraint for the design effort, because this drives the amount of spare capacity and redundancy that must be included in the design. We defined the availability in terms of data preservation and the ability to keep up with the data volume.

The requirement for overall system availability is 98% and the justification is as follows. First, based on statistics from CTIO, CFHT, and other observatories, we anticipate losing approximately 20% of total available observing time to weather and overall operational considerations. (This lost survey efficiency is included in the operations simulations.) We have set the goal that no more than 10% of this loss will be attributable specifically to the Data Management System, or 2% of total available observing time, thus leading to 98% availability. Our reliability architecture is based on two fundamental principles: *data preservation* and *keeping up with the data*, including producing transient alerts within 60 seconds.

Data preservation. At any given point in time, at least two copies of the data are maintained in separate geographic locations. Caches are large enough to handle any projected system downtime, and multiple recovery options exist for any failure mode. That is, once the data leave the Data Acquisition System on the Mountain Summit, we immediately store it to disk (two copies) on the Mountain Summit and stream it to the Base Facility. We cache each night's data for four nights; once we have passed the fourth night, the oldest night is deleted from the Mountain Summit.

As the data are received at the Base Facility, they are buffered and processed for alert generation. They are simultaneously streamed to the Archive Center over the course of the night and the next day; the transfer is completed before observing is started for the next night. Each night of processing at the Base Facility is also cached for four nights. Once it is confirmed that the data have reached the Archive Center, the data from the oldest of the four nights are deleted from the Base Facility.

Data are processed again at the Archive Center nightly (to avoid transferring all the processed data), and new additions to the catalogs are retained at the Archive Center and sent back to the Base Facility for the next night's processing. At the same time, data are replicated to the Data Access Centers on a daily basis. Thus, the Mountain Summit, Base, Archive, and Data Access Centers together always contain both the full data set and a backup copy.

Keeping up with the data. The established design goal is a small fraction (2%) of total nightly data not processed within 24 hours. In order to achieve this, we employ a mix of redundant systems, hot spares, and cold spares to ensure a minimum number of single points of failure, while staying within our cost constraints. We will employ "autonomic" computing systems that "heal" themselves quickly without human intervention, and in the case of failure, we resume processing and deliver the alerts as soon as the system has recovered.

In case of storage failure in any processing center, the backup copies of the data are used to perform the restore. At the Mountain Summit and Base Facility, the copies are local. In the case of the Archive Center, the backup copies are replicated from the Data Access Centers, and vice versa. In all cases, there is sufficient spare computing, storage, and network capacity to continue nightly processing while restoration is performed.

4.5.5 The DMS Reference Design

The collective data management experience from a wide range of institutions in astronomy and high energy physics was used to translate the DMS capabilities and performance specifications into a reference design. Experience with such systems demonstrates that these levels of performance, scalability, and reliability can be achieved through a planned, documented, systems-level approach to design and development, with process and specification reviews at each major milestone. A formal design process captures all design elements into specifications and models. The DMS requirements have been captured in the Data Management Requirements Specification (Axelrod and Kantor 2006a) and a Unified Modeling Language (UML) model (Axelrod and Kantor 2006b).

DMS requirements and design specifications use the Iconix process (Rosenberg and Scott 1996). This widely used process is based on Unified Modeling Language (OMG 2004). Costing estimates were obtained from the following:

- — An industry-standard software cost estimating process based on the Constructive Cost Model (Boehm 1981) and Function Point Analysis (Jones 1991)
- — Vendor-provided quotes and price/performance trends for hardware and off-the-shelf software
- — An independent cost estimate generated by the California Institute of Technology Infrared Processing and Analysis Center (IPAC)

Three-Layered Architecture of the DMS

The Data Management System architecture partitions the DMS into individual centers optimized for specific responsibilities, as already described. This approach has been adopted in order to maximize the throughput of the DMS pipelines, to meet the latency requirements, and to distribute resources to serve the data close to the users. This "vertical" partitioning is, however, only half the story. The LSST requirements and challenges for scalability, reliability, modularity, and evolutionary design are best addressed by what is commonly known as a "layered" architecture, a concept that is widely employed in information systems with these requirements. (The Open Systems Interconnect (OSI) Mode and MVC and Layered Architectures in Java give examples. See reference list.) This horizontal layering separates the primarily custom application software from the physical infrastructure (i.e., hardware and system software) with "middleware" software technology. Figure 4-62 shows the layering structure with associated components. This layering is ubiquitous in the DMS and is incorporated into the design of each of the individual DMS centers.

- — The **Application layer** embodies the fundamental scientific roles and responsibilities of the LSST DMS, i.e., the scientific algorithms, pipelines, and data product implementations. This layer is implemented as custom developed, open source software.
- — The **Middleware layer** enables "portability," by providing a "thin" abstract and uniform interface to the hardware and system software; it also provides standard services to the

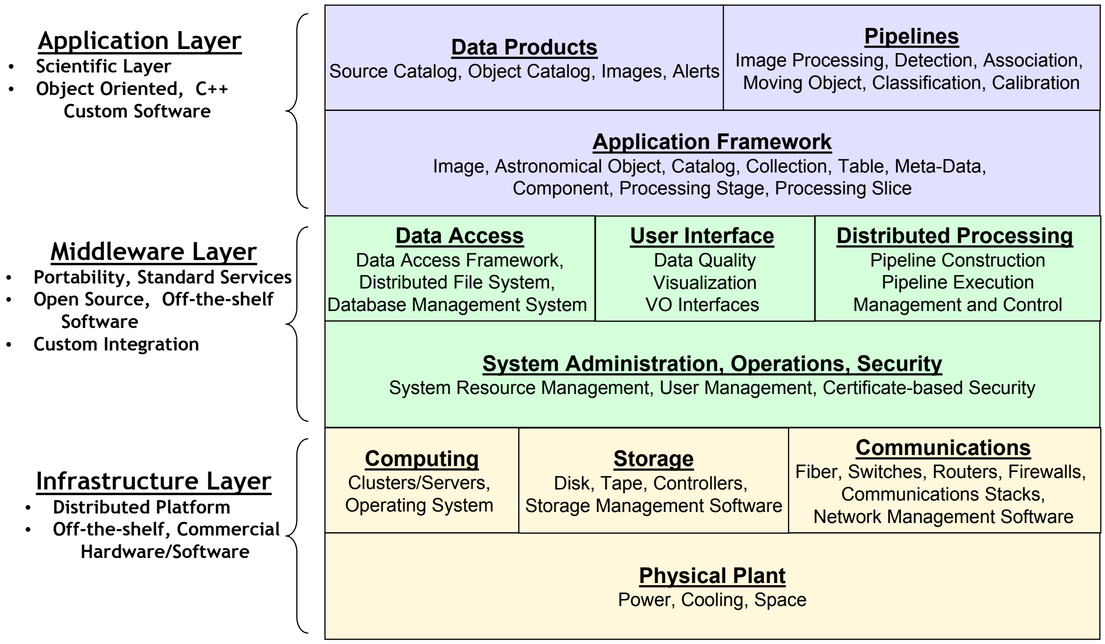

FIGURE 4-62 *The three-layered architecture of the data management system enables scalability, reliability, and evolutionary capability.*

applications for security, reliability, plug-in pipeline algorithms and stages, and extendable data types. This layer is implemented by custom integration of off-the-shelf, open source software.

- — The **Infrastructure layer** provides all the hardware and system software in the DMS computing, storage, and network resources that are the environment in which the Application and Middleware Layers are hosted and execute. This layer is implemented as off-the-shelf commercial hardware and system software.

The UML specifications, combined with industry-standard estimating methods (CoCoMo and Function Point Analysis) provide the basis for estimating the effort and schedule for the software elements of the DMS.

In the remainder of this section, we describe the reference design for each of these layers. The Applications layer is treated first because the scientific value of the LSST resides in this layer. Then we discuss the Infrastructure layer, which hosts the applications and the middleware. Next we discuss the middleware, which provides for evolution and scalability. Finally, we conclude with a section on validating the design and integrating and testing the DMS.

1) Application Layer

The Application layer functions have been specified via UML use Cases. A total of approximately 100 use cases have been defined that cover the required capabilities. (See the Supplementary Information section of this proposal.) In general, the Application layer will be developed in an object-oriented fashion as custom C++ class modules. The classes fall into three major categories: 1) those that provide pipeline processing components, 2) those that manage the stored data, and 3) those that provide re-formatting for transmission or presentation and visualization of the data. In formal Iconix process terms, these are known as *control*, *entity*, and *boundary* classes.

The Application Framework forms the base classes for the entire layer. Specific algorithms and data products are derived from the base classes of the framework. The framework also encapsulates all the interfaces to the Middleware layer. For example, there are data classes called “Catalog” and “Astronomical Object” that encapsulate interfaces to the database management system (DBMS) that ultimately stores the catalogs. Similarly, there is a class called “Image” that encapsulates the distributed file system (DFS) that ultimately stores the image data. This encapsulation ensures that all classes provide the necessary indexing, provenance, security, and fault tolerance operations and also facilitates evolution of the underlying Middleware and Infrastructure.

All of the Application layer classes have been modeled using UML; approximately 200 classes have been defined.[2]

The Application layer is described in the remainder of this section in terms of algorithms and pipelines and the database schema.

— Application Layer: Algorithm and Pipelines

The requirements for the algorithms employed in the DMS pipelines are set by several factors. The first are the data quality metrics that must be achieved for the LSST science missions. An example is the requirement that photometry must be consistent across the sky and between filters to 0.005 mag. An equally essential requirement derives from the fact that LSST will observe the entire visible sky under a variety of diverse observing conditions—i.e., seeing, sky brightness, atmospheric extinction—and field characteristics (particularly crowding). An algorithm able to deliver the required photometric precision in excellent conditions may be unable to do so when conditions deteriorate even slightly. The enormous data rates of the LSST rule out routine human intervention to tune up an algorithm that is functioning badly. Such tuning must be performed autonomously by the DMS itself.

Algorithm development for LSST builds on the experience base and expertise already gained in large astronomical surveys (e.g., SDSS, Super/MACHO, ESSENCE, the DLS, CFHTLS, UKIDDS) and will benefit from the advanced surveys planned for the near future, such as Pan-STARRS, the Dark Energy Survey (DES), and SkyMapper. However, while existing algorithms have been shown to meet the LSST data quality requirements, they are generally unable to meet them over the full range of sky and field characteristics, and do not have self-tuning or self-repair capabilities. In addition to these general areas of development, specific algorithms may require particular performance improvements. To carry out the needed algorithm development, we have undertaken early prototyping of our pipeline software and will test the prototypes on both existing survey data and simulated LSST data.

The LSST algorithms are executed by a set of pipelines that are structured to partition the overall functionality of the DMS cleanly. Each pipeline has sole responsibility for producing one or more data products. Several pipelines are run on a nightly basis, first at the Base and later at the Archive Center, producing the Nightly Data Products (see light blue elements in Figure 4-63).

- The **Image Processing Pipeline** transforms the raw science image from the sensor into a calibrated science image, which has had the instrumental signature removed and is both astrometrically and photometrically calibrated. This calibration will later be improved at the

2 For a complete description of all use cases and classes, see the LSST DM UML Model (Axelrod and Kantor 2006b).

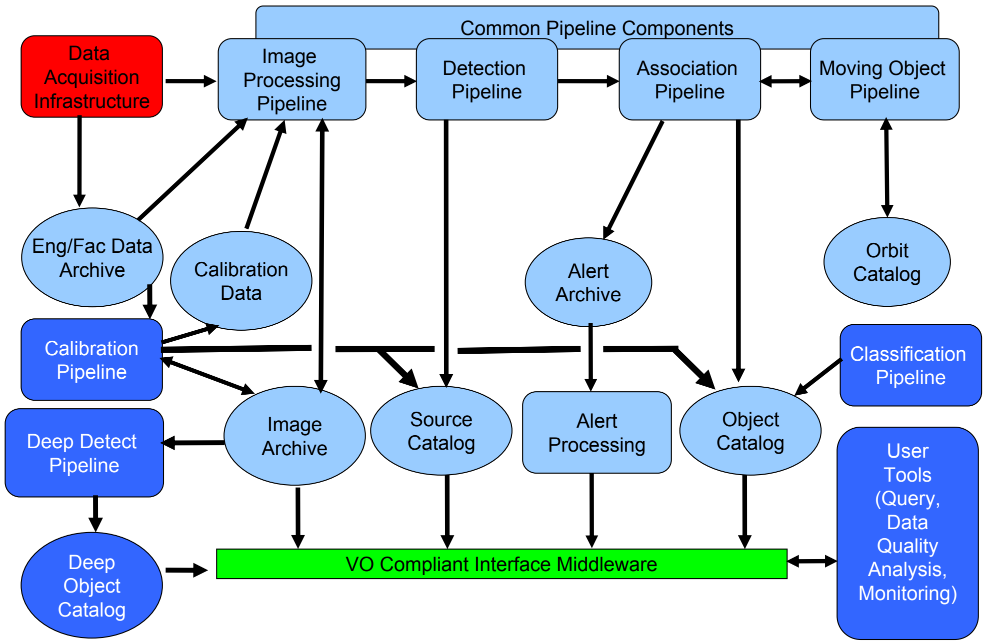

FIGURE 4-63 *Pipeline layout to produce the data products at nightly or longer cadences*

Archive center. From the calibrated science image and a template image obtained from the Archive center, this pipeline produces the subtracted science image.

- The **Detection Pipeline** is responsible for detecting and measuring sources in the subtracted science image, producing the Source Catalog.
- The **Association Pipeline** is responsible for associating new entries in the Source Catalog with the information on previously known astrophysical objects contained in the Object Catalog. Sources that do not correspond to a previously known object cause a new entry to be made in the Object Catalog. Alerts are generated for those objects whose behavior meets the alert criteria, and these are forwarded to the Archive Center for distribution to the community.
- The **Moving Object Pipeline** is responsible for the Orbit Catalog. It takes as input the entries in the Source Catalog that were not matched with known static objects by the Association pipeline. From these, and the existing Orbit Catalog of solar system objects, it determines which of these unmatched sources correspond to already cataloged orbits and whether new orbits need to be formed. In the latter case, it can send an alert to the Archive Center.

The nightly processing of a visit concludes with Alert Processing, which is performed on the full object entry. Each object's characteristics, including its prior history, are checked against a set of rules that determines whether or not an alert will be issued, and if so, the type of alert. If an alert is to be generated, the object information, together with the sub-images containing the source, is packaged into an alert structure, which is sent to the Archive center for dissemination to the community using the VOEvent mechanism (IVOA 2006a).

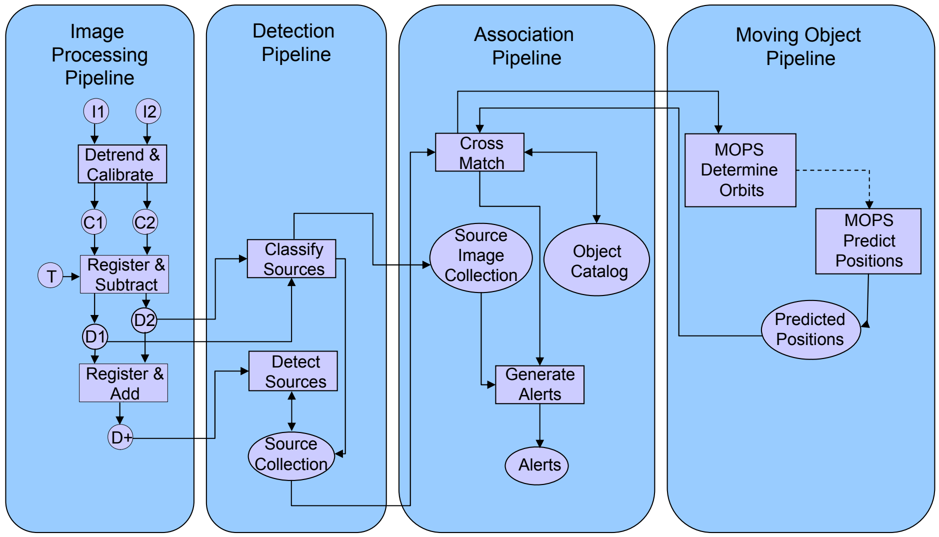

FIGURE 4-64 *First level of detail inside each of the nightly pipelines*

For brevity, Figure 4-64 omits much of the complexity of the nightly processing. Figure 4-64 shows the next layer of complexity specified in the DMS UML model. These UML models, carried to their full complexity, generate code that provides the structure of the programs. This code is expanded via conventional programming and then reverse-engineered back into the UML model to maintain the consistency of the code with the design.

The remaining pipelines are run on a less frequent cadence (as needed in the case of the Calibration pipeline, roughly every six months for Classification and Deep Detection) only at the Archive Center and produce the Data Release Products (see dark blue elements in Figure 4-63).

- The **Calibration Pipeline** is responsible for producing calibration products such as flats and darks that are needed for the nightly data processing.
- The **Deep Detection Pipeline** is responsible for optimally combining information from multiple images to produce the best possible measures of object properties, including shapes, photometry, proper motions, and parallaxes. This information is stored in the Object Catalog. Additionally, the Image Processing Pipeline produces image stacks for use as template images. Some LSST key science teams will utilize these data, and their own derived research data products, to pursue a science product and thereby continuously probe the system for systematics.
- The **Classification Pipeline** utilizes the multi-epoch, multi-filter information produced by the Deep Detection Pipeline to classify objects into astrophysically meaningful categories and summarize properties of their time history, such as the waveform and period of periodic variable stars.

— Application Layer: Database Schema

Working in a petabyte regime frequently requires non-conventional techniques and innovative solutions. Optimizing queries is extremely important in LSST's world of multi-billion-row tables. As an example, the Source table is expected to reach over 200 billion rows in the

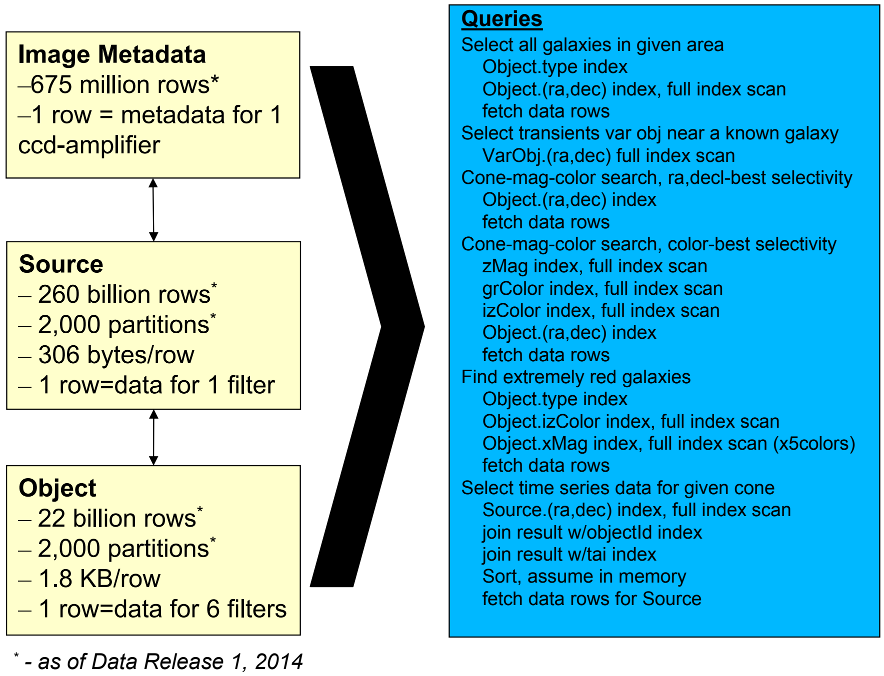

FIGURE 4-65 *The data query schema is optimized for spatial and temporal queries. A subset of the reference design query set is shown on the right.*

first year of data taking (Figure 4-65). Triggering a single scan through this table (which can be done through a simple *"SELECT * FROM Source"* query) would result in fetching 100 terabytes of data from disks. Further, having to support simultaneous multi-dimensional searches (spatial and temporal) significantly complicates the database design.

The set of queries used for sizing and costing the system is representative enough to drive the schema design and cover both the public access as well as the professional astronomer's access. They were generated based on experience from other astronomical experiments and archives (SDSS, MACHO, IRSA), carefully selected such that both spatial and temporal aspects are well covered. They are aligned with the precursor schema that was used for sizing and testing and were optimized using rewriting, clustered indices preferred over non-clustered ones, and other well-known techniques. The key schema optimizations include:

- Pre-calculating most frequently used fields to minimize full table and/or full index scans. Example: storing colors in addition to bandwidths.
- Extracting most frequently used information into summary (*tag*) tables to de-randomize disk I/O. Example: extracting information about variable objects and replicating it in a much smaller table than the original table containing all objects.
- Preparing summaries to reduce the size of searched data set. Example: extracting time-dependent information like wavelets from Source table and storing it in smaller Object table.

- Maintaining and using optimal indexes. This includes preparing *covering indices* and building highly specialized *spatial indices* (based on R-trees and the Hierarchical Triangular Mesh developed by the SDSS team)[3].
- Employing query plans and estimates and row/byte/time limits to queries to ensure that a few users do not dominate the resources to the exclusion of others.

The key layout optimizations include appropriate data clustering, table partitioning, and choosing optimal data and index block size (based on extensive database size and input/output models (Becla 2006a, b, c, d; see also section entitled Middleware Layer: Database System Design).

2) Infrastructure Layer

The infrastructure layer provides the total hardware and system software for DMS computing, storage, and networking. Processing and storage activities are distributed across four DMS facilities: the Mountain Summit, Base Facility, Archive Center, and Data Access Centers (Figure 4-66). The Archive and Data Access centers also provide data access for End User sites that are external to the LSST. Connecting all the facilities are long-haul communications networks.

Each DMS facility operates 24/7 and is staffed for both nighttime (observing) and day/swing (non-observing) shifts. Distribution of processing and storage activities among separate facilities promotes efficient performance of the system via separation of real time and non-real time resources.

In the Supplementary Information section of this proposal, we document the reference design for these facilities as they would be implemented at a 25% scale using 2006 era technology. The actual implementation will use 2010 and later technology, which at constant dollars will provide the performance needed to meet 100% of the LSST requirements. The system performance projection is based on industry-provided trends for price/performance and capacity improvements between now and 2010 (Beldica, Dossa, and Plante 2006a, 2006b). The trends used are conservative.

The remainder of this section describes each facility reference design and the long-haul networks design.

— Infrastructure Layer: Mountain Summit Site

The Mountain Summit site is on Cerro Pachón. Its sole DMS responsibility is to capture data and move it down to the Base Facility for nightly pipeline processing. The Data Acquisition System (part of the Camera Subsystem) and the Observatory Control System (part of the telescope subsystem) interface to the DMS on the summit, with image read-out in two seconds and immediate data transfer to the Base at 10 Gigabits/s on fiber optic lines dedicated to this traffic. Metadata from the Engineering and Facility database in the Observatory Control System (OCS) move to the Base on a dedicated 1 Gigabit/s line on a nightly basis. The Mountain Summit also has a data buffer sufficient for four nights of data in case of communications failure to the Base Facility.

Using standard Linux systems at the Summit permits a relatively simple data acquisition (DAQ) output interface, lowering software development costs and long-term hardware support

3 Maintaining an index is expensive, so each index should yield good read-write ratio.

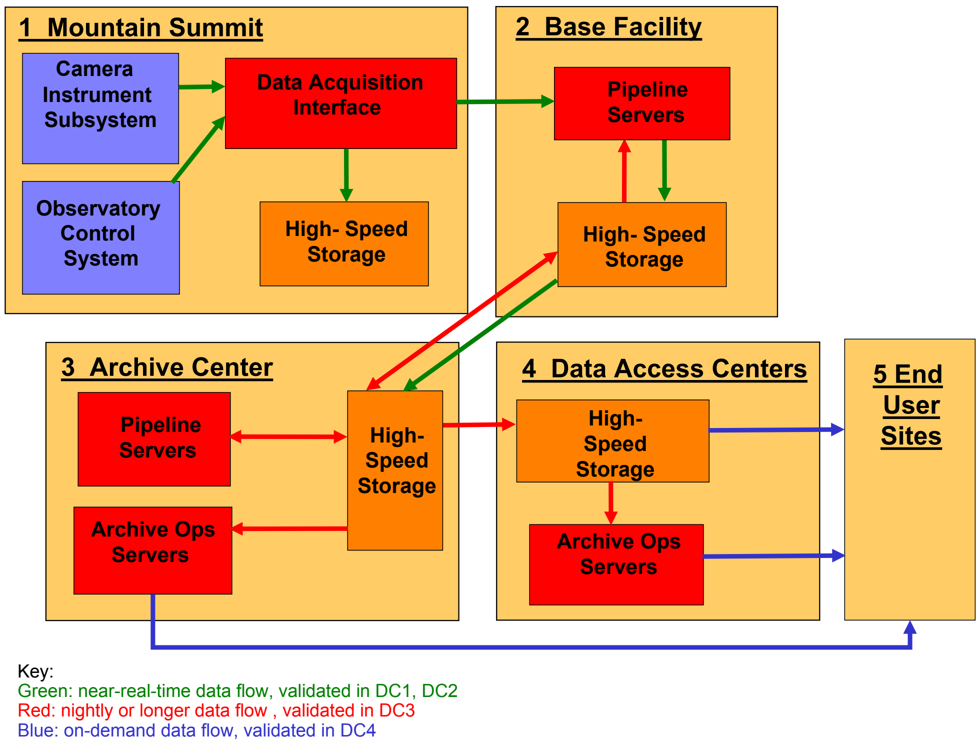

FIGURE 4-66 *Facility distribution for the DMS showing data flows*

costs. These commercial Linux systems are known to be highly reliable. The chosen Summit architecture results in an easily upgradeable Summit computing system that does not impact other telescope systems.

Moving cross-talk corrected and raw images to the base consumes 10 Gigabit/s network bandwidth. This is implemented as one long-haul 10 Gigabit/s link from the summit Ethernet switch to the base Ethernet switch. A second 10 Gigabit/s link is used as a backup. This network choice is lower risk compared to other less mature high-speed communication methods and permits easy technology upgrades without impacting other systems at the summit and base. The firewall and optical switch connecting the Base to the Archive Center site are easily accommodated with current era technology. The Base file system design uses RAID-8+2 disks to store multiple copies of image data to further enhance data preservation. There is a redundant link to the Base Facility for fail-over and in the extremely unlikely event of simultaneous failure in the primary and backup links it is possible to transport one set of the backup data drives to the base.

— Infrastructure Layer: Base Facility

The Base Facility is at the Cerro Tololo Inter-American Observatory compound in La Serena, Chile. The Base Facility’s primary role is processing the image data to generate transient alerts within the required latency. The Nightly Data Pipelines and the most recent previously processed co-added sky templates and catalogs are hosted here on a 25 teraflops-class computing cluster to provide primary data reduction and transient alert generation within the required 60 seconds.

Similar to the Mountain Summit, the Base Facility has capacity to store four nights worth of data in the case of communications failure to the Archive Center.

We considered locating this equipment at the Summit, but the operating reliability of the system would be decreased at an altitude of 9,000 feet. Therefore, the smallest possible computing system with sufficient disk space has been specified for the Summit, and nightly processing for the production of timely alerts will be performed at a Base at low altitude, where servicing is easier and the expertise to maintain this system is already located.

The Base to Archive Center Network is 2.5 Gigabits/s full duplex, protected, clear channel fiber optic circuit, with protocols optimized for bulk data transfer. The term protected means that the link is a ring with traffic moving in both directions; only a dual outage on both sides of the ring between source and destination will prevent data delivery. Clear channel means that LSST will have a dedicated portion of the network bandwidth guaranteed by an Indefeasible Right to Use (IRU). This link is used both for high priority command and alert traffic, as well as to “trickle” the raw image data over a 24-hour period and to send processed data back to the Base Facility. In the event of an outage in this link, once service is restored the link supports a 400% on-demand “burst” increase in capacity to ten Gigabits/s for transmission of buffered data.

— Infrastructure Layer: Archive Center

The Archive Center will be located at the National Center for Supercomputing Applications (NCSA) on the campus of the University of Illinois at Urbana-Champaign (BOX 4-1). This center will publish the alerts via public alerting protocols based on the Virtual Observatory VOEvent standard. It will also repeat the Base Facility processing (recall that only the raw image and metadata are transferred from the Base in order to keep the data rates feasible) and will merge the processed data into the Science Data Archive. The Archive Center computing resources will handle the processing needed to generate Nightly and Data Release data products and will support the activities of the Key Science teams and the planned re-processing. The Archive Center will be the source of replication operations that deliver the data to the Data

Box 4-1 Leveraging NSF Facilities at NCSA

In its selection of the Archive Center site, LSST has made it possible to build its construction and operations plans on the extensive infrastructure which has been laid by NSF investments in U.S. supercomputing centers and research networks.

The LSSTC will not build or operationally staff the LSST Archive Center. Theis proposal calls for our partner, the National Center for Supercomputer Applications at the University of Illinois-Champaign Urbana (the NCSA) , to be our primary archive center in the United StatesU.S.. NCSA is one of the five original centers in NSF’s

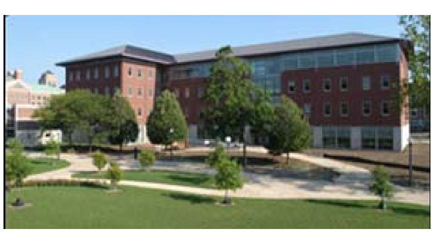

The NCSA Building *(left)* is located on the campus of the University of Illinois, Urbana-Champaign; the Advanced Computations Building (center) houses NCSA’s large computational and storage platforms *(right)*.

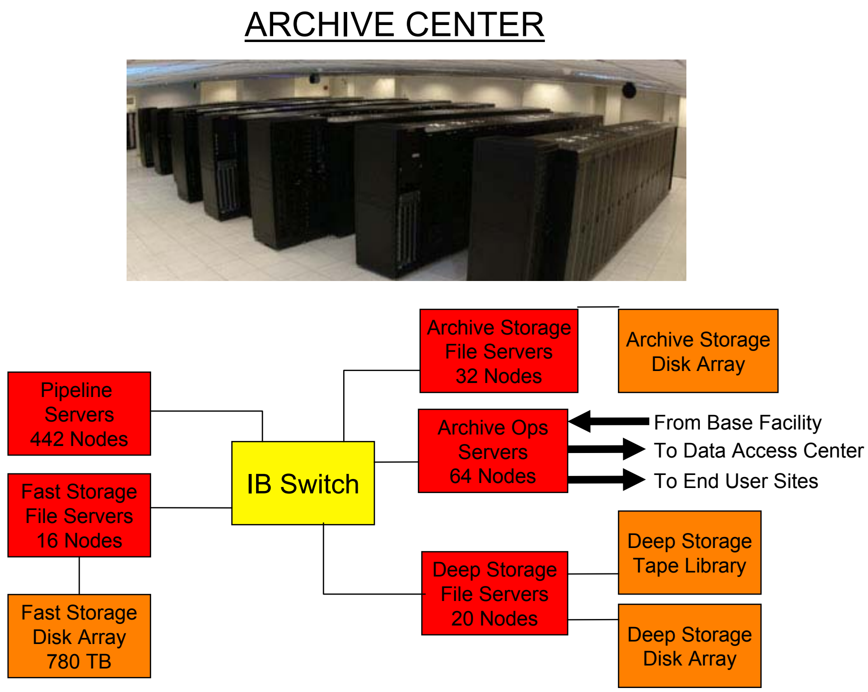

FIGURE 4-67 *The reference design for the NCSA archive center*

Access Centers, and along with the latter, will support end user data queries and data access. Finally, the data releases are published from the Archive Center. These responsibilities require supercomputers capable of 100 TFLOPS (in Year 1 of the survey) interfaced to a 15 petabytes/year data archive.

As depicted in Figure 4-67, the design includes two types of computing systems (pipeline and archive operations servers) and three types of storage systems (fast, deep, and archive storage). Disk space will be used for processing, serving data, deep storage cache, and other archive operations. The pipeline servers provide cluster nodes for running nightly processing, analysis pipelines, and reprocessing jobs. These servers are attached to a fast storage system that supports only this processing. The archive servers provide computing support for all services that run the archive, including orchestrating pipelines, managing data transfers, ingesting new data products, and running data query and access services for outside users. These servers are attached to a large archive storage system, sized to hold the most recent versions of the catalogs, the latest co-added sky images, and a cache for intermediate products. Finally, a deep storage system is configured to provide archival storage and data access for older catalog versions and image files. All of these systems grow in capacity and power over the life of the LSST survey.

— Infrastructure Layer: Data Access Centers

The Data Access Centers provide a backup copy of the data archive and additional access capacity for external users. The Data Access Centers will be located, at a minimum, in the U.S. and Chile. No data reduction processing beyond servicing data access requests is performed at the Data Access Centers. As such, the Data Access Center design is a subset of the Archive

Center design, and the details are not repeated here. The Centers are connected to Archive Center via a high-speed 10 Gigabits/s network (TeraGrid in the U.S., REUNA in Chile).

The first Data Access Centers will be deployed in 2014 and 2015, coinciding with the commissioning phase of the project. At present, the baseline calls for Data Access Centers at the San Diego Supercomputing Center and in Chile at either La Serena or Santiago, depending on further discussions with the Chilean astronomy community. We anticipate that the scientific collaborations may wish to establish additional Data Access Centers to host selected data products based on relevance to the communities served. Funding for these additional data centers will be provided by the host institutions.

— Infrastructure Layer: End User Sites

With its huge repository of data, LSST will greatly accelerate the advance of "e-Science" for professional astronomers through the use of the Virtual Observatory tools and interfaces, combined with the advanced computational resources available both at large computing sites and more generally, through the Grid. In addition to the data, the data processing and analysis software will be openly available for deployment not only on proprietary resources (for groups who have funds to build and manage significant computational centers), but also for deployment onto NSF-funded supercomputing and Grid resources. While the primary data processing will be executed on dedicated computing resources, and additional computing resources are budgeted for query support, the DMS is being designed to allow expansion to Grid and other available resources as needed to provide scalability as the user community grows.

Managing access to core LSST resources will be important for the project to achieve its mission-critical goals, both in day-to-day operations and in support of the key science programs. The concept of "service levels" will be used to manage resources effectively, and to distribute access loads over a variety of shared or collaborative sites. This approach has been quite successful in the high-energy physics context where "tiers" of access and shared responsibilities are common. Level 1 sites are relatively rare and typically have their own supercomputers dedicated to limited groups, while Level 2-5 sites have more limited resources but will be more common. Combining these provides flexible access to cover the gamut of scientific and EPO users.

— Infrastructure Layer: Long-haul Communications

We will install a wholly owned private fiber link from the Mountain Summit to the Base Facility (80 km). LSST will manage and maintain this mission-critical link, providing the reliability and availability necessary to run the combined Mountain-Base infrastructure as a single component of the LSST system. This is the only new fiber that must be installed.

Existing and planned upgrades to existing fiber optic networks provide the remainder of the solution. We have received cost estimates for configuration and maintenance from both commercial and Research and Education Network (REN) consortia for the remaining links. The commercial Chilean fiber backbone provides the needed bandwidth from the Base Facility in La Serena to Santiago. One of two independently owned fiber rings surrounding South America provides the necessary network connectivity from Santiago to the U.S. mainland. Both of these rings have current capacities in excess of 40 Gigabits/s (compared to the LSST need of 2.5 to 10 Gigabits/s) and are expandable to two terabits per second with long-distance Dense Wavelength Division Multiplexing (DWDM) technologies and offer protected circuits with IRUs.

Within the U.S., LSST will use NSF-funded research networks and/or their successors to provide open access to LSST data. For internal high-speed data transfers from the Archive Center to U.S.-based Data Access Centers, the TeraGrid (or its successor in 2013) provides the needed connectivity, and the National LambdaRail (NLR) provides external high-speed access for other computational centers. NLR and TeraGrid already link the major U.S.-based LSST research centers.

Two additional considerations are relevant to future DMS communications. Because they are not fully defined as of this time, we have not yet incorporated them into the reference design.

First, given the significant role being played by DOE national laboratories in the LSST, we anticipate that DOE high-speed research networks such as ESNET will also play a role in providing access to LSST resources for the HEP community; integration of DOE and NSF supercomputing grids via the Open Science Grid consortium will facilitate this.

Second, multiple carriers operate on this communications infrastructure, and competition tends to encourage innovation on these links, increasing aggregate bandwidth and driving prices down. The LSSTC and its Chilean-based partners (AURA/NOAO) are actively engaging in technology road maps with NSF-funded programs such as WHREN-LILA and the South American RENs, REUNA, and other non-profit groups. These initiatives have the goal of forging partnerships to enhance high-speed networking for research and education-based connectivity and achieving economies of scale in Chile and from South America to the U.S. Thus, our communications cost estimates should be considered conservative, and the net cost to LSST will be reduced to the extent that these initiatives provide synergies with other projects and communities.

3) Middleware Layer

The Middleware layer is key in terms of support for scientific evolution by providing services that enable plug-in pipeline algorithms and stages. It also supports technology evolution by providing an abstract and uniform interface to the hardware and system software. Beyond that, it provides standard services to the Application layer that enable extendable data types and data scalability. Lastly, it provides uniform services for system administration, operations, and security. In this section, we describe the design in terms of support for each of the areas mentioned above.

Middleware is also the key to controlling the cost of maintaining a robust and flexible DMS. We have adopted a strategy of building our middleware system from existing technologies, drawing heavily on proven technologies from the grid and information technology communities. Middleware development during the construction phase thus becomes primarily a custom integration activity, greatly reducing risk.

In all cases in the reference design, the middleware solutions are open source (refer to the Supplementary Information section), but LSST DM will use proprietary software as well (e.g., in the infrastructure layer), where the licensing is not overly restrictive and the cost is reasonable. The primary barrier to using commercial middleware is typically cost and the need to rely on the supplier if customizations are needed. LSST has factored in the effort of maintaining the open source software into the total project cost.

Between now and the start of construction, we anticipate advances in middleware technology from such initiatives as the Common Component Architecture (CCA), and we will remain open to incorporating those advances throughout the D&D phase, where they show particular value to the LSST design.

— Middleware Layer: Support for Scientific and Technology Evolution

We have planned the evolution of the underlying hardware infrastructure to take advantage of technology advances during the 10-year duration of the LSST survey. Accordingly, we require that the middleware provide flexibility that allows changes to hardware while protecting the applications from those changes.

First, the middleware provides flexibility in where we run our applications. Our pipeline execution middleware leverages grid, messaging, and distributed logging technology to provide this flexibility. Not only are we free to deploy the programs to any configuration of Linux-based cluster or parallel processing system, we can move or replicate processing between facilities, and we have made it possible for users to deploy our system on platforms beyond those supported by the project. This allows our users to tap into large community computing systems, such as those at the national supercomputing facilities, to assist with extraordinary processing that may be part of advanced science analysis.

Second, the middleware provides the ability to “plug-in” modules as algorithms improve. This plug-ability is essential for allowing new versions of algorithms to be flexibly swapped in as they are improved over the life of the project.

Figure 4-68 illustrates how several of the components of the Middleware layer work together to create and execute a pipeline. The Pipeline Construction System is used to form a sequence of processing stages into a pipeline, assembled from pluggable components that apply astronomical algorithms to the data. The components plug into a framework that applies an algorithm in a data parallel fashion—where each processor applies the algorithm on a different portion of the data. This framework, which organizes work into “processing slices,” can adapt not only to the number of independent data chunks but also to the number of processors available. We have demonstrated the adaptability of the processing stage/processing slice design in DC 1, with an implementation based on Python and the Message Passing Interface (MPI).

The Pipeline Control and Management system then configures and deploys that pipeline on a particular platform and monitors its progress. The adaptation to a physical platform is configurable, so not only can it potentially be deployed on a variety of systems, the deployment policies can be updated simply to adapt to the evolving infrastructure.

— Middleware Layer: Services for Provenance and Re-processing

It is critical that the DMS be able to track the full processing history, or provenance, of any data product. Provenance is used in two ways. First, many of our intermediate data products, such as individual calibrated frames and difference images, will not be preserved long-term due to their large volume. Once such products expire from the data cache, they will be regenerated on the fly as needed, either because of a user request or because they are needed for subsequent processing. Second, provenance will be important for the planning of complete re-processing of the archive required for each Data Release. In both cases, the regeneration of data products will be achieved via the use of “smart” re-processing” techniques, in which an analysis of both the provenance of the requested data and an assessment of what data currently exist, will result in a workflow that recreates the product in the most efficient manner.

In the DMS, provenance is implemented as a set of custom classes that provide an application programming interface to a set of machine-interpretable recipes. Provenance is created and updated in several places in the processing flow:

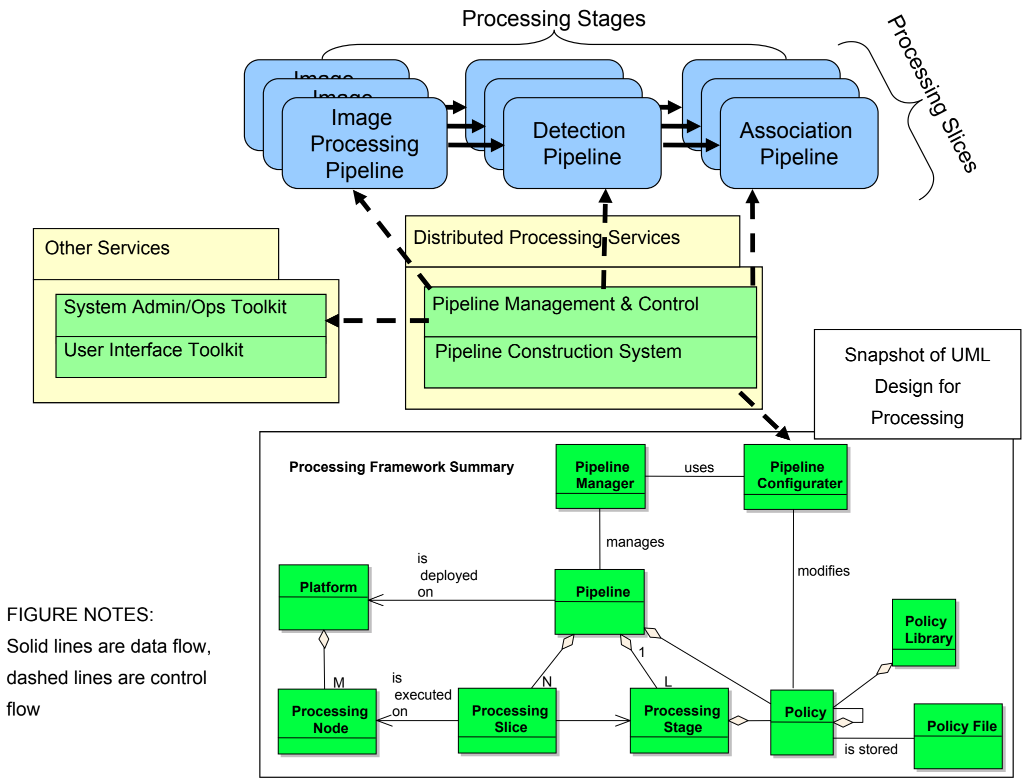

FIGURE 4-68 *Middleware design allows pipelines to be built and executed in processing slices. (Solid lines are data flow, dashed lines are control flow.)*

- Provenance is pre-configured during the component and pipeline construction processes to establish the execution parameters, input and output data streams, and error and logging parameters.
- Provenance is captured at pipeline run time to document the results of execution.
- Provenance is captured at data ingest and data release time to document the data state.

In Data Challenge 1, a simple, flat, file-based provenance implementation was created; the final design includes an implementation based on eXtensible Markup Language (XLM) and a database management system (DBMS).

— Middleware Layer: Services for Open Interfaces and Extendable Data

The Data Access Framework (DAF) is the middleware that provides uniform, open interfaces to all LSST data, copies and organizes data in preparation for pipeline execution, and incorporates features for data and metadata extensions.

The Archive Service incorporates data into the Science Data Archive from transfers or pipelines. Thus, this service is used primarily by internal applications, although external use is also possible. Primary drivers for this service are high performance and scalability to keep up with LSST data volumes. Underlying this service are a distributed file system (DFS) for image and other files, and a DBMS for record-based catalog and metadata. Both are selected from off-the-shelf solutions and integrated with the rest of the middleware layer. Grid-enabled technologies are employed for data transfer as well. This design was prototyped in Data Challenge 1 using the IBRIX DFS, the MySQL DBMS, and various portions of the Globus Toolkit and the Storage Resource Broker.

Virtual Observatory (VO)-compliant interfaces are provided to all data products to support responding to VO protocols requests for data. The VO Architecture Overview (IVOA 2006b) provides detailed information on VO protocols and architecture. LSST members from NCSA, Johns Hopkins University, Lawrence Livermore National Laboratory, and NOAO lead and/or participate in VO workshops and summer schools; they also contribute to and review VO standards. By ensuring VO-compliance, LSST enables access via the many VO-community based tools that are being created.

The DAF employs advanced software engineering techniques to support extendable data and metadata. These techniques have been utilized by LSST designers in previous scientific applications to achieve data extensibility (Kantor 2006) and include:

— Providing services that expose logical paths to data and map to physical paths transparently
— Employing context-specific, externally visible labels distinct from internal identifiers
— Defining pre-allocated spare attributes, variant records, and pre-defined extension tables in the database
— Utilizing data type handlers that map data types into tools and visualizations dynamically
— Providing deep and shallow cloning operations
— Employing abstract data types and object-oriented data access objects that encapsulate the physical implementation of the data type

Incorporation of these techniques into the DAF design allows the DMS to rapidly evolve to accommodate new data requirements when the LSST needs to enable astronomical analyses that have not been previously conducted.

— Middleware Layer: Database System Design

The LSST data ingest rate at the Archive Center will be 30MB/s from pipelines, or 1 TB in each 10-hour observing window. In addition, ingest for data indexing and re-processing brings the total ingest rate across all Centers to 6 Gigabits/s.

The LSST database system supports multi-petabyte scalability, good transactional insert performance, ad-hoc queries with ability to scan data at tens of gigabytes per second, and high data availability—all at a reasonable cost. The proposed LSST database system design was developed by a group of engineers and scientists with extensive hands-on experience in building very large scientific database systems (e.g., BaBar, SDSS) and in-depth understanding of astronomical problems.

The proposed design assumes relational DBMS technology, which is much more widely used than object-oriented DBMS technology. The design features several independent server farms, each dedicated to serving a specific task or a specific group of users and tuned accordingly for the expected access patterns. Key design features providing high performance for query access include:

— Caching major indexes in RAM
— Partitioning tables, indices, and queries
— Executing queries at the “best” place

— Combining disk arrays with virtual file system to provide load-balancing
— Replicating most frequently used parts of data.
— Separating ingest and query servers

Ingest servers are mirrored to provide fault tolerance. Newly ingested data are available for query together with previously taken (archived) data in real time. To minimize competition between readers and writers, ingested data are transferred to archive data servers in real time, i.e., through replication. Archived data are partitioned to reduce the amount of searched data.

4.5.6 Data Management Subsystem Integration

The DMS subsystem integration approach is to create the foundation of infrastructure and middleware necessary for the Nightly Pipelines and Data Products first. The Archive Center will be used as a development, integration, and test center for these capabilities while the Base Facility is being constructed. Once the Base Facility is constructed, the tested components are moved to the Base Facility and replaced at the Archive Center. Next, the Data Release Pipelines and Data are integrated into the Archive Center. Note that the Data Access Centers and their corresponding networks are not required for Observatory System Integration and Test or Commissioning. They will be deployed and tested during Science Center Commissioning. Figure 4-69 depicts the integration sequence against the construction timeline.

4.5.7 Results of Data Challenge 1

This section provides a summary of results of the recently completed Data Challenge 1 (Plante 2006a). The technical objective in DC 1 was to validate the DMS infrastructure and middleware scalability to 5% of that required for the DMS (see Table 4-6). The results significantly exceeded our targeted performance levels.

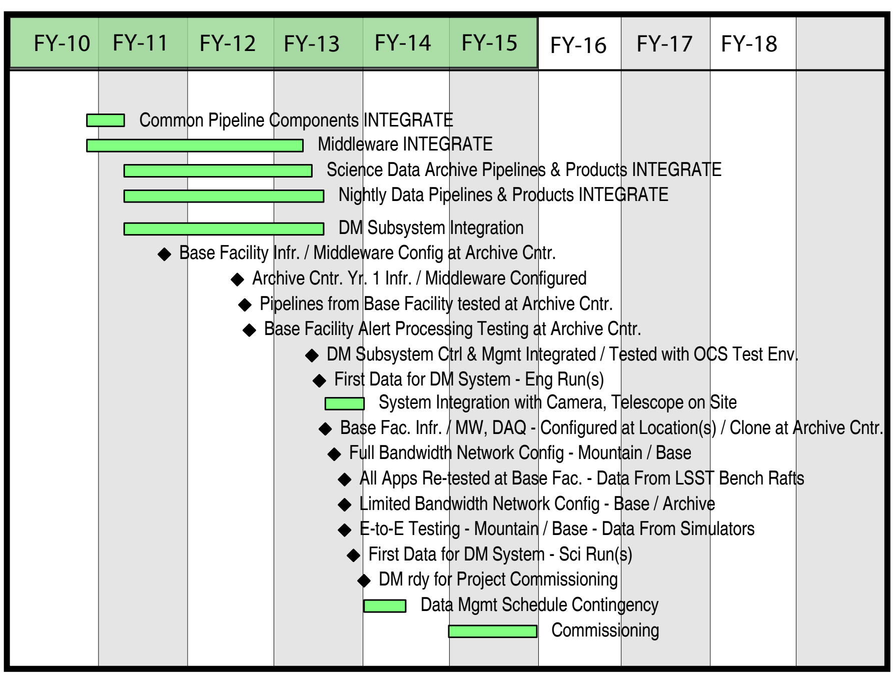

FIGURE 4-69 The data management integration timeline

In addition to meeting the 5% performance target, the team wanted to understand how the system scales toward the full required data volumes and number of processors. While not explicitly part of the Data Challenge goals, the team also wanted to prototype techniques for automated software deployment. Finally, by actually building a working prototype system as a collaboration, DC 1 also prototyped the design and development process to ensure that it is best suited for this project.

Three sites of the NSF TeraGrid represented the mountain, base. and archive facilities as illustrated in Table 4-6.The mountain site fed data files to the base facility where it was entered into an input queue for processing through a simulated nightly pipeline. Simultaneously, the data files were forwarded on to the archive facility and fed into an identical pipeline. Timing was measured for the data transfers as well as for various intervals in the pipeline processing. The pipeline itself was composed of tunable "consumers"—components that consumed computational, memory and I/O resources but otherwise performed no real processing. Instead, the consumers were configured to match the load estimated from our requirements and sizing models. The last stage of the pipeline created simulated catalogs, which were then fed into the database ingest service. This service stored them in the database.

Planning and development for DC 1 began in January, 2006. Execution nominally began July 1, 2006. The original plan called for completion by July 31; actual completion of the demonstration test meeting the 5% goal was completed on September 30, 2006. This represented an approximately 25% schedule and 15% cost overrun. Identified sources of the overrun included slower than planned startup, staff turn-over, lack of team experience on the TeraGrid, test execution delays due to the latency between TeraGrid requests and time allocations, and insufficient UML modeling for certain portions of the software. These lessons learned are being incorporated into the DC 2 plans and process. While the DC 1 goals were achieved on September 30, the team continued to exercise the DC 1 system until October 30 to see how far they could extend the performance. The following are the results:

- — Inter-site transfer rate: 85.0 ± 0.5 MB/s: 17% of required for full-scale DMS
- — Time to process one image: 76.3 ± 3.1 s: 25% of required for full-scale DMS (i.e., time to complete estimated compute cycles for one image)

— Database ingest rate: 6.0 ± 0.2 MB/s: 100% of required for full-scale DMS

These results indicate that in every aspect tested, the design is consistent with achieving the LSST DMS requirements.