# LSST Data Management Working Group Report

Data Archiving and Distribution

# Membership

- Alex Szalay(Chair)
- Robyn Allsman
- Scott Anderson
- Bruce Balick
- Todd Boroson
- Kem Cook
- Jim Heasley
- Ho-Il Kim
- Chris Smith
- Glenn Wightwick
- Stephan Luitz

# Scope of the Task

- Create an efficient, scalable environment for PB scale data exploration
- Preserve all legacy information about the project
- Interoperate with the world – VO compliance
- Ensure early community involvement
  - Users group, workshops, summer schools
  - Get NSF involved in supporting archival research

# Where do we need to be at the end of R&D?

- Understand the precise responsibilities of LSST in providing data access
  - cost/benefits analysis
- Have clear idea where archive is in the SYSTEM
- Understand roles of persistent storage vs DB
- 100TB prototype DB built on precursor projects
- Risk analysis wrt rapidly changing technology
- Decision about inclusion of other datasets

# Issues Addressed

- Scope of Archive
- Impacts of technology
- Need geoplexed mirror
- Bottlenecks – moving the data, loading the data
- Usage modes/power users
- Stream subscription vs workbench, multi-tier
- How to deal with rapidly changing environment
- Multiple databases (telescope, provenance, data)
- Scale-up vs **scale out**
- How to capture all relevant information

# Technology Restrictions

- Main obstacle is data transmission
  - From telescope to processing (site selection!)
  - From processing to users
  - Within archive
- IO scaling slower than CPU or disk capacity
  - Random access 10ms? 5ms in 10 years
  - Sequential IO 10MBps? 60MBps in 10 years
  - Needs to be understood more in detail
  - Reliability of large disk farms poorly understood

# Technology Opportunities

- Biggest difference: move computation as close to data as possible
- Multi-Tier Hierarchy (like HEP-LHC model)
  - At the top: involvement of NCSA, SDSC
  - Divides the effort, allows scientific expertise to be co-located with a specific data product
  - Broader community involvement
  - Solves impedance mismatch
  - Replaces BYOC (bring-your-own-computer)

# User Access to Data

- Low latency: Stream subscriptions to alerts
  - Of the order of 10 stream types
- Fast response simple query system
  - Simple object and image extraction, single time series
- High-Latency: Workbench environment
  - Happens outside the mission critical circle
  - Users typically submit workflows
  - Can keep their own temporary data on-site
  - Users can contribute modules
    - Formal acceptance process

# Unresolved Issues

- Multi-Tier vs Single
  - Amount of “outsourcing” – cost-benefit
- Authentication, traffic monitoring, logging
- Fault tolerance (thousands of disks…)
- Data validation (the moment of truth…)
- How stable will the data model be?
- How much redundant capability needs to be built in?
  - As a function of time?
- How many time will we reprocess everything in full?
  - How many versions of the data will be there?
- Everything will change – will switch DB several times
  - How to minimize the impact of such transitions?

# R&D Tasks / Milestones

- Year 1
  - Experiment with data loading/scrubbing strategies
  - Work with DM group to map onto RDB
  - Start process of community buildup
- Year 2
  - First prototype system, maybe several platforms
  - Scalability tests, parallelization
  - Experiment with a prototype streaming system
  - Design workbench prototype, look for code reuse
- Year 3
  - Build a 100TB functional prototype
  - Simple workbench on top

# Tiers and Their Usage

- Tier 0 Usage
  - Mission critical
  - Quality Assurance
- Tier 1 Usage
  - Collaboration
  - Power users
- Tier 2 Usage
  - Astronomers, public

# Data Products

- Raw images
- Calibrated data
- Instrumental and environmental data
  - Alerts and events
- Provenance information
  - pipeline pmts, calibration…
- Value added data
  - Cumulative sky, catalogs…

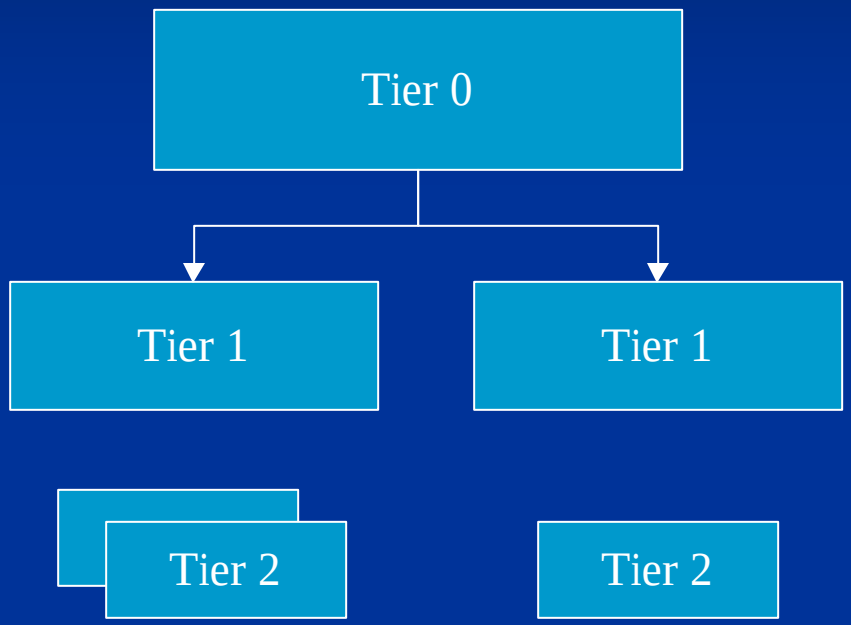
Tier 0
Tier 1
Tier 1
Tier 2
Tier 2