

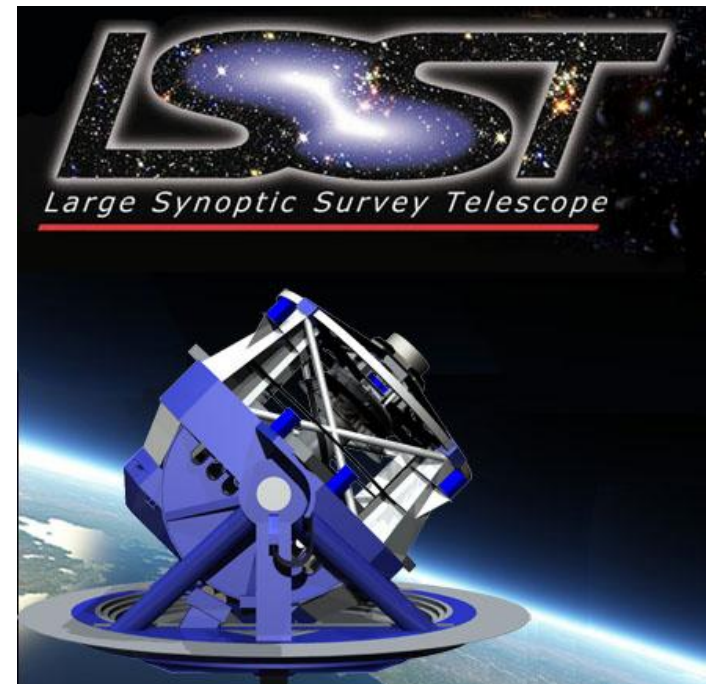
Special Needs for LSST

Jacek Becla

SLAC National Accelerator Laboratory

June 20, 2011

- Schedule
- SLAC role in DM
- Database baseline architecture



Testing

- Most tests on 5 (heterogeneous) nodes
- Have access to 100 node ir2farm
 - but only 70GB storage/node
- 100 node cluster test (LLNL, Jan 2010)
- 100 node cluster test (Dec 2010, boers)
 - Thanks to Fermi, BaBar, Atlas
- 150 node cluster test (May 2011, fells)
 - Thanks to Fermi, BaBar, Atlas

150-node Test

- ~10% of production (1st year)
 - 32 TB, 60 billion rows
- 11 days spent preparing, 4 days testing
- **Hugely** successful
 - Simple queries: 4.1s (requirement < 10s)
 - Full table scan with aggregation: 2m 40s (<1h)
 - Near neighbor across 100 sq deg area: 10m (<10h)
 - Object/source join across 100 sq deg: 2h 6m (<10h)

The Need

- Access to large scale test-bed
 - Few hundred nodes
 - Periodically (2-3 weeks every 3-4 months)
 - Prefer nodes + disks, not cores
 - Many cores welcomed soon
- Lab-wide, multi-experiment, shared, large-scale, data-intensive test-bed?