

LSST Database

[[LSST Data Management](#)]

Welcome! This page is used for managing internal information related to the LSST Database. For user-oriented information see [mysqlGuide](#). The team working directly on this project include:

- [?Jacek Becla](#) (SLAC)
- Fritz Mueller (SLAC)
- John Gates (SLAC)
- Andy Salnikov (SLAC)
- Serge Monkewitz (Caltech/IPAC)
- [?Andy Hanushevsky](#) (SLAC)
- Fabrice Jammes (IN2P3)
- Brian Van Klaveren (SLAC)

Past:

- [?Daniel Wang](#) (SLAC, 2003-2015)
- Douglas Smith (SLAC, 2013)
- [?Bill Chickering](#) (Stanford, student, 2012-2014)

Past collaborators:

- Vyassa Baratham (Summer Student, 2013)
- Bipin Suresh (Stanford), July - August 2010
- Maria Nieto-Santisteban (JHU), 2005 - 2007
- Ani Thakar (JHU), 2005 - 2007

This database is expected to reach over 50 petabytes.

The LSST baseline database design is described in detailed at [LSST Database Design?](#) ([?http://ls.st/lpj](http://ls.st/lpj)). Please note that it is a live document, periodically updated.

Query Service (qserv)

- [qserv main page](#)

Schema

- [?Baseline schema](#)
- [Naming Conventions](#)

Queries

- [Common Queries - data analysis](#)
- [Queries in Data Production](#)
- [Queries for Performance Tests](#)

- [Queries for testing shared scans](#)
- [Questions about Data Analysis Queries](#)
- [Queries run by PipeQA](#)
- [SQL Programming Style](#)
- [Queries for Validating Schema](#)
- [Queries per Difficulty-level](#)

MySQL-related

- [Optimizing queries](#)
- [MySQL-related hints](#)
- [Large tables in MySQL](#)
- [MySQL related issues](#)

Simulated Data

- [?Simulated data](#) (500 million stars in Milky Way)

Hardware Sizing

- [Model for how database hardware capabilities will evolve](#)

Design Discussions

Note, this section covers higher level design discussions that do not go into qserv specific details. Qserv-specific design discussions are covered in the "Design" section on [Qserv](#) page.

- [Data Placement](#)
- [Fault Tolerance](#)
- [Schema Evolution](#)
- [Provenance](#)
- [Level 3. User tables and external data sets](#)
- [Time series](#)
- [Ingest](#)
- [Correlations](#)
- [Metadata](#)
- [Cutouts](#)
- [Views](#)

Tests

- todo
 - ◆ [Connect engine](#)
 - ◆ [MyISAM concurrency with key cache segments](#)
 - ◆ [Wide tables](#)
 - ◆ [Compression and decompression cost](#)
- work in progress
 - ◆ [Aria vs MyISAM storage engines](#)

- done
 - ◆ Large tables
 - ◆ Merge engine (functionality test)
 - ◆ Merge engine (scalability to 1K tables)
 - ◆ CPU cost for queries
 - ◆ Tests with Solid state disks
 - ◆ Schema Evolution related tests
 - ◆ Db Performance for Forced Photometry
 - ◆ Experiments with Hive Implementation and testing of sample data with Hive
 - ◆ Building index
 - ◆ Cross match performance for DC2
 - ◆ Fine chunks performance for DC2
 - ◆ Course chunks performance for DC2
 - ◆ Stripes with (zoneId.ra) indexes: Performance Summary for DC2
 - ◆ Stripes with ra indexes: Performance Summary for DC2

Databases in Data Challenges

- DC3
 - ◆ Backup and Replication in DC3b
 - ◆ Database Scalability in DC3b
 - ◆ Database Ingest in DC3b
 - ◆ Database setup for DC3a
 - ◆ Storing Reference Catalog
 - ◆ Some discussions on Database setup for DC3 (old)
 - ◆ Provenance in DC3
 - ◆ MySQL Permissions in DC3
 - ◆ Duplicating the DC3a databases
 - ◆ Synthetic data
 - ◆ Calibrated sources
- DC2
 - ◆ Database and Nightly Processing in DC2
 - ◆ DC2 Database Partitioning Tests
 - ◆ DC2 AP in application tests
 - ◆ New, simplified DC2 AP-in-application
 - ◆ Building DC2 Object Catalog - closest neighbor search

All Others

- Winter 2015 planning
- Compression
- MOPS Partitioning
- Science Data Quality Analysis (SDQA)
- Neighbors
- Synchronizing Object Ids
- Views in MySQL
- Object Split
- Transactions
- Spatial Support

- [Computation in database vs. computation in application](#)
- [PixelProcessing](#)
- [Image Access](#)
- [ImageMetadataProposal](#)
- [Persisting Time](#)
- [UniqueIds](#)
- [Query Results](#)
- [Speed issue with innodb](#)

Related publications of interest

- [?A machine Learning Classification Broker for Petascale Mining of Large-scale Astronomy Sky Survey Databases, Kirk Borne](#)
- [?The WFCAM Science](#)
- Organizing the Extremely Large LSST Database for Real-Time Astronomical Processing, ADASS, London, UK, 2007: [?paper](#), [?talk](#)
- Designing a Multi-petabyte Database for LSST, SPIE, Orlando, FL, 2006: [?paper](#), [?talk](#)
- Useful links
 - ◆ [A quick set of links to useful database related documents in Docushare](#)
 - ◆ [External Links](#)
 - ◆ Databases and GPUs
 - ◇ [?Accelerating SQL Database Operations on a GPU with CUDA](#)
 - ◇ [?Low Power Amdahl-Balanced Blades for Data-Intensive Computing](#)
 - ◇ [?Data-intensive computing in sql with cuda and distributed queries](#)
- Related talks
 - ◆ [?LSST Database Architecture](#), presented at LSST Conceptual Design Review, Tucson, AZ, September, 2007
- Some other notes
 - ◆ [Stored Procedures: cross-dbms compatibility](#)
 - ◆ [Amdahl Numbers for LSST](#)
- Various meeting notes
 - ◆ [Meeting notes](#)
 - ◆ [Percona Live MySQL conference 2013](#)
 - ◆ [AWS summit 2014](#)
 - ◆ [Percona Live MySQL conference 2014](#)

Internal laundry lists

- [Issues to Discuss](#)
- [Pipeline Database Setup for DC3+](#)

Obsolete

- [Merging DR Object catalog with Alert Prod Object catalog](#)
- [DataAccWG meeting notes](#)

- Action Items
- MySQL Server Configuration
- CORAL - Relational Abstract Layer
- DBMSStorage Prototype
- Association Pipeline
- DC3 schema:
 - ◆ DC3a Database Schema (MySQL-compatible)
 - ◆ DC3b Database Schema (MySQL-compatible)
 - ◆ DC3bDataProducts
 - ◆ DbDC3bSchemaUpdates
 - ◆ DbDC3bSimSchema
 - ◆ DbSchemaUpdateApr09
- Post-DC3b Schema Changes
- FlexibleSchema
- DC2:
 - ◆ ?DC2 Database Schema Description
 - ◆ DC2 Database Schema (MySQL-compatible)