


LSST Database Scaling Tests

Jacek Becla
SLAC National Accelerator Laboratory

May 27, 2011

LSST - Schedule



R & D


CoDR (Sept 2007)

PRD (~Aug 2011)


Construction

Commissioning

Production




SLAC & LSST Data Mgmt



Data Access (databases and images)
Persistency layer

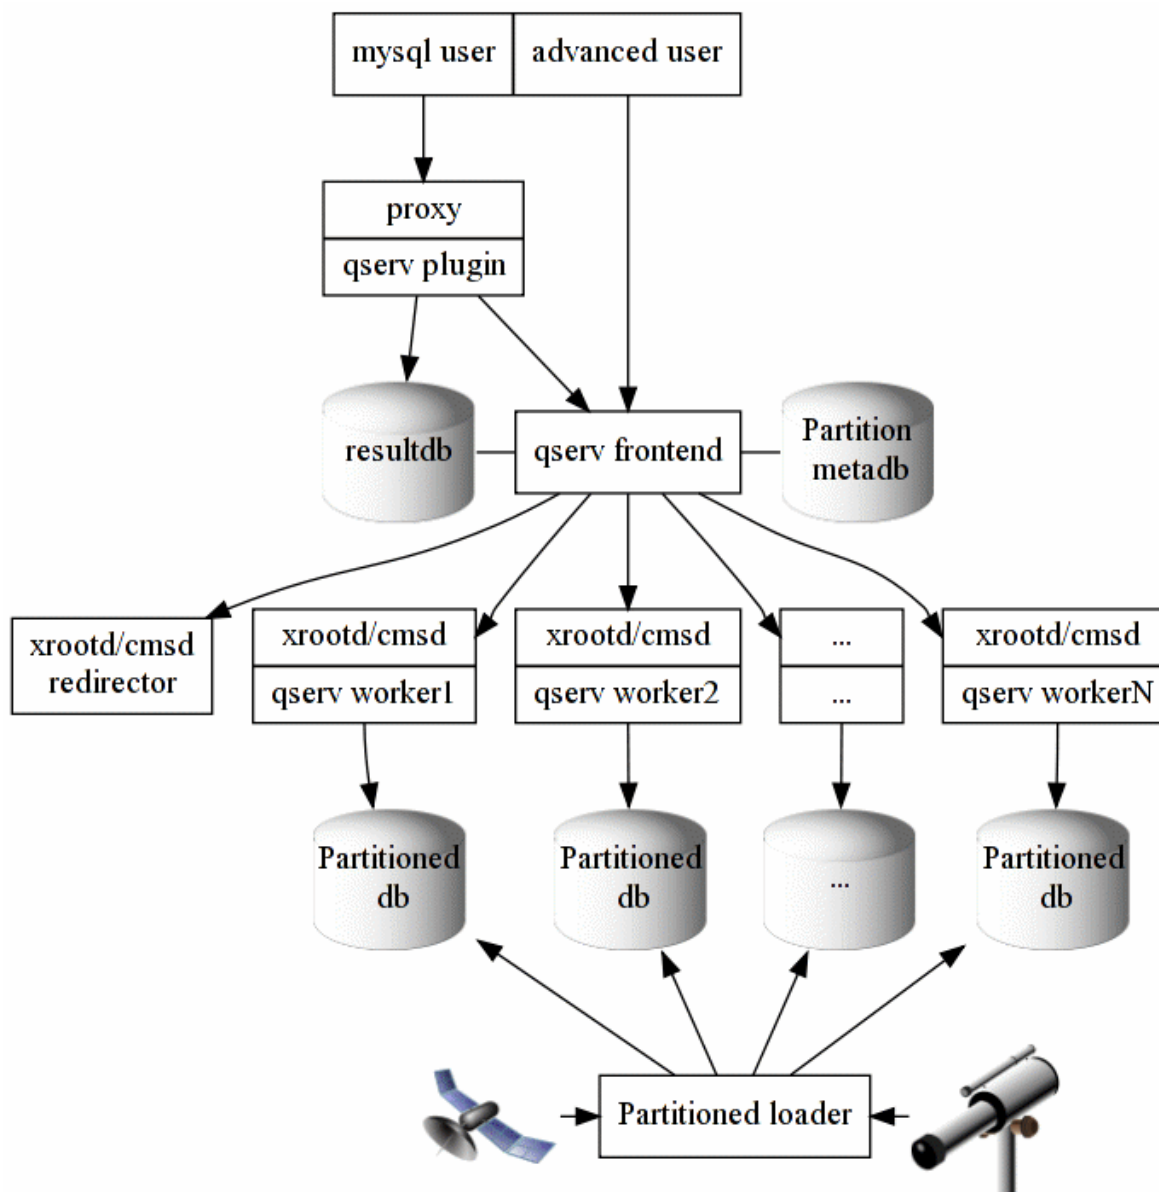
Baseline Database Architecture



- Shared nothing on cluster of commodity nodes
- Parallel execution engine
- Locally attached drives
- Shared scans
- Open source if possible

Qserv Prototype

On top of MySQL
and xrootd



Testing



Lsst-dev02, 03, 04, lsst-db1, 2

100 node cluster test (LLNL, Jan 2010)

Access to

40 node ir2farm (since ~Sept 2010), 64 node memfs cluster, 100 node ir2farm

only 70GB storage/node

100 node cluster test (Dec 2010, boers)

Identified bottlenecks, didn't complete key tests

150 node cluster test (May 2011, fells)

150 node test - Sizes



Object: 1.7B rows, 1.8TB, 9K chunks

Source: 55B rows, 31TB, 9K chunks

Index for Object

3 columns, 1.7B rows in a single table


Expected production numbers

Object: 8-26B

Source 90B-2T

ForcedSource 1T-76T

150 node test - Queries



Trivial

Large area scans, full table scans


Near neighbor

Aggregations

Joins

Mixture of the above under concurrent load

150 node test – data loading



Generating data on the fly
New data set (“pt1.1”), larger
New duplicator / partitioner

~11 days spent preparing

Unexpected

Broken fell 0057

Data purged on fell0294 (2nd day) and substitute node (13th day) due to /scratch retention

Disk failure on fell0304

Security scans

Scheduled power outage knocked xrootd master (lsst-db2)

Recovered and caught up by using ultra-fast ir2srv03

Ran out of space (underestimated density variations)

Fragmentation 55k fragments for one chunk

MySQL misconfiguration

xroot malfunction (misredirection when disks full, min 2% free requirement)

Holes in synthesized data

Slow cluster stabilizations / warm up (~5+ min)

xrootd client bug (mutex missing/corruption)

qserv bug (thread leak under concurrent load)

Results

40,100,150 node configuration tested

All test planned except concurrency beyond 2
successfully completed and timed

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 dev: 2h 6m (<10h)

Overheads for full sky query: <20 sec

Lessons Learned



1 test node prior to big test very useful
Need 1-2 spare stand-by nodes
Using /scratch disruptive

Next Steps



Documentation

Pass PDR

Periodical large scale tests (2-3 weeks every 3-4 months?) very useful

Need to find ways to reduce data loading time

Prefer nodes+disks, not cores

Once we have shared scans, many cores will help

Hugely Successful



Big big thanks to Fermi, BaBar and Atlas and everybody who helped to make this a success