

BaBar Databases

Jacek Becla

Stanford Linear Accelerator Center



Outline

- ◆ Bridge federations
- ◆ Other activities
 - Analysis
- ◆ Others
 - Objy releases
 - Communication
 - Manpower
- ◆ Future plans

Bridge Federations

- ✓ Reading (phase I) OK
- ✓ Writing (phase II) almost OK
 - delivered (in 10.4.4 and 11.9.0)
 - Used by skims, backed off last Friday due to internal bug and lock collisions
 - Current status
 - Both problems understood
 - Fixes will be provided within a week

Bridge Federations – cont...

- ◆ Phase III (output fd != input fd)
 - For data export, BdbServer++
 - Due 30 June 02. Good progress already

- ? Phase IV (multiple input fds)
 - For event mixing
 - Estimating benefits vs. implementation cost

OPR

◆ No news

- ...is a good news

- All working

- Redesigning, main effort outside of bdb

◆ Assisting Padova (REP, Linux servers)

- This effort will increase in the next few months

Analysis (1)

- ◆ Bridge federations
 - almost there
- ◆ Data compression (due 1stQ'02)
 - AMS decompresses the data
 - Example compression factors
 - 3.1:1 (evsdhr), 3.6:1 (evt), 1.9:1 (aod), 1.6:1 (esd)
 - On schedule, productizing
 - testing
 - fixing bugs
 - estimating CPU cost
 - working on plugging into production system

Analysis (2)

- ◆ Working on reducing size of event (expected to finish 2nd-3rdQ'02)
 - Current status: developing a tool for finding detailed statistics on sizes, overheads – almost ready
- ◆ AMS redirection and automatic load balance (due 2ndQ'02)
 - Will improve performance
 - distribute load
 - reduce influence of disk failures
 - # disk failures grows (more disks, disks get older)
 - Still far away, coding in progress

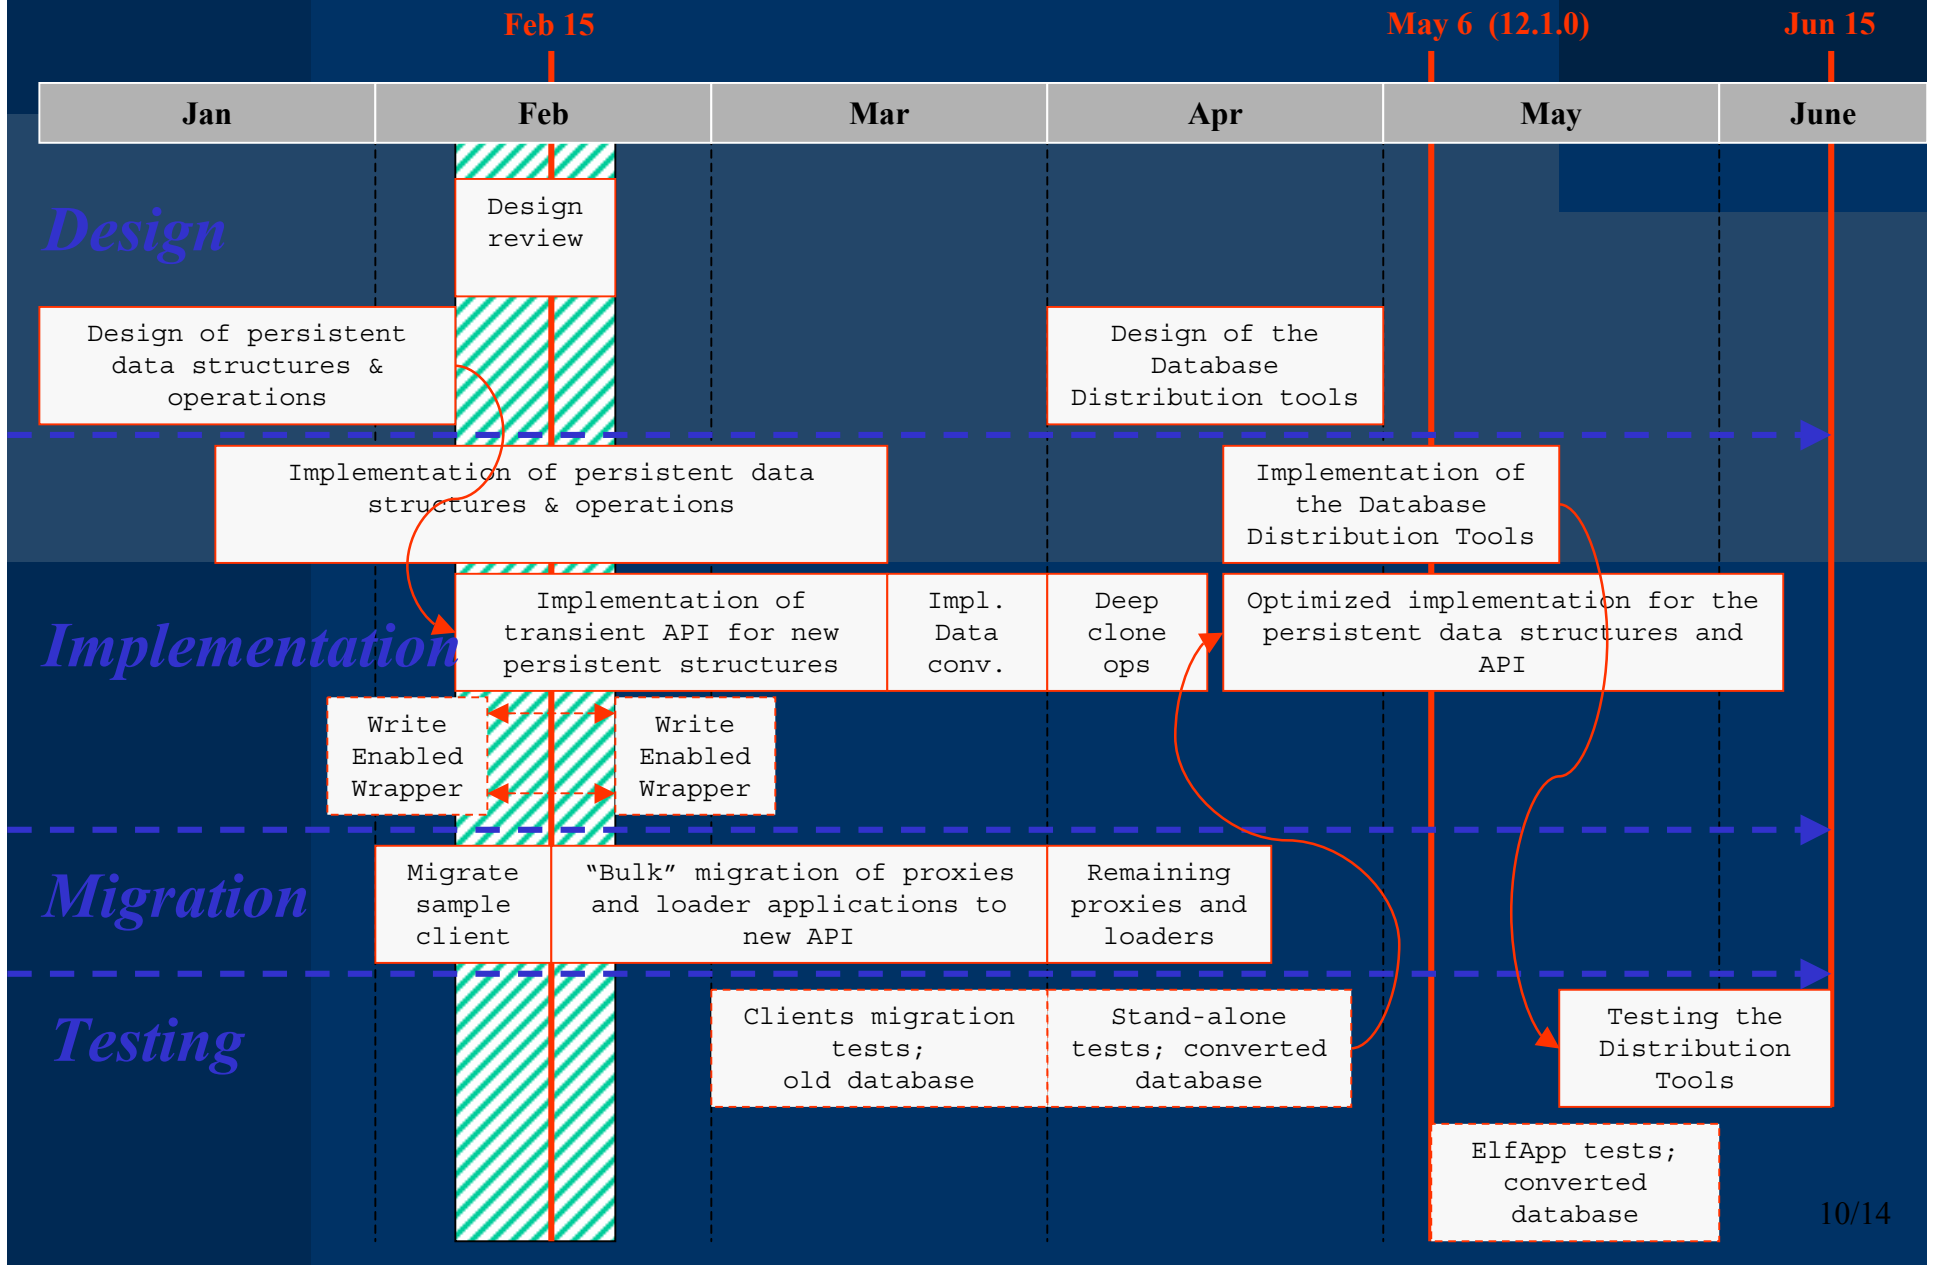
Analysis (3)

- ◆ Starting to seriously focus on performance
 - “testbed” to measure sizes & performance
 - Improving collection code
 - Fixed problems with crashes due to lock collisions
 - *All* lock collision problems will be fully addressed in ~ a week
 - Working on mapping collection → set of associated files
 - Much more difficult than expected due to “borrowing” events
 - No good time estimate yet
 - Other changes if manpower permits
 - Plans to reduce disk & network traffic (2nd-4thQ'02)

Analysis (4)

- ◆ Reworking Conditions DB
 - Will address all current problems
 - Lightweight and flexible
 - Redesign in progress, due 15 Jun'02
 - Tight schedule

CondDB: Implementation & Deployment



Objy

- ◆ 7.0.1 available
 - In 11.x.x series, starting from 11.9.0
 - Lock traffic reduced 33-50%
- ◆ Linux support
 - Still RH 6.1, kernel 2.2.12 and 2.2.14, gcc 2.95.2
 - Official certification from Objy coming end of Feb:
RH 7.2, kernel 2.4.9-21, gcc 2.95.3
- ◆ Next major release 8: ~1st Q'03
 - Will get some useful features much sooner

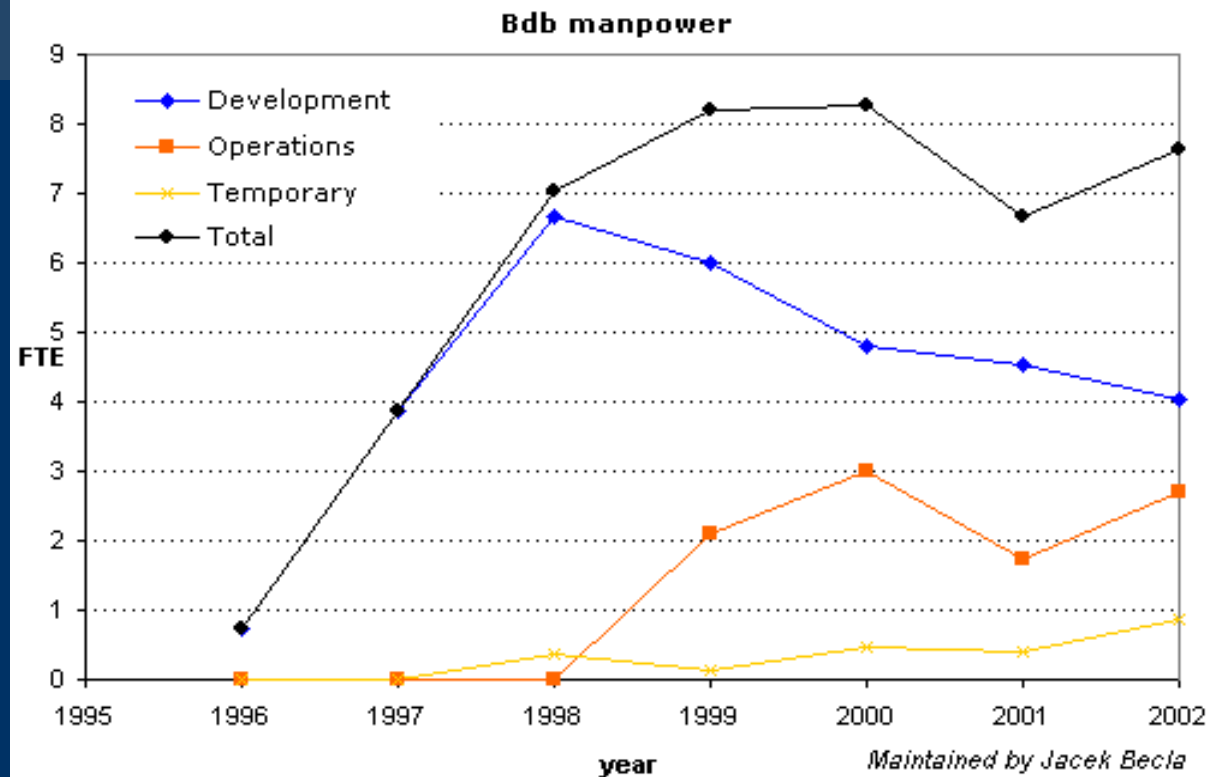
Communication

- ◆ Bdb website (www.slac.stanford.edu/bdb)
 - Statistics show good interest
 - Especially Real Time Monitoring
- ◆ Error reporting & response
 - New HN: BdbSOS
 - Works well, high traffic, focused discussions
- ◆ Feedback, suggestions always welcome

Did you know...? You can track status of most major Bdb projects. See website for details

Manpower

- ◆ No changes since Dec'01
 - Development ~4 FTE
 - Working on hiring one person
 - Operations ~3 FTE
- ◆ Total ~ 43 FTE



Summary

- ◆ Do you remember OPR state ~2 years ago?
(~8 Hz, 50 nodes, <math><10\text{ pb}^{-1}/\text{day}</math>?)
- ◆ ~18 months of tuning \rightarrow $>300\text{ pb}^{-1}/\text{day}$
 - + Objy no longer perceived as a bottleneck
 - + *“Keeping up with data, still room to increase rate”*
 - + Gained a lot of experience!
- ◆ We’re just starting to focus on Analysis...