



Database Mini-Review - 23rd Sept. 1997

BABAR OO Databases - Issues

David R. Quarrie
Lawrence Berkeley National Laboratory
DRQuarrie@LBL.Gov



Databases Mini-Review - 23rd Sept. 1997

2

Integration into Release Structure

- Successful builds on AIX
 - Regular builds
- DEC Unix problems understood
 - Attempting a build this week
- HP-UX a problem
 - Several attempts, but template problems
 - Decided to defer further work until aCC compiler in use
 - ◆ October?
- This weeks release will have new reconstruction prototype
 - Version 2

9/23/97

David R. Quarrie: BaBar Databases - Issues



Safety & Security

- Lack of AMS protection a real concern
- Andy has suggested using NFS-3 with HPSS
 - Allow file system protection
- Band-aids being put in place
 - Only effective if people don't use direct Objectivity API
- Other band-aids being investigated
 - Making files explicitly read-only?
- Directory creation
 - Databases created by AMS Server
 - Directories created by client?

9/23/97

David R. Quarrie: BaBar Databases - Issues



HPSS integration

- Lack of AMS protection
- Integrate NFS-3 daemon with HPSS API?
- Is a single data mover a good match to our needs?
- Backups?

9/23/97

David R. Quarrie: BaBar Databases - Issues



License Management

- Software distribution mechanism in place
 - Based on AFS ACLs & signed license agreement form
- Need to monitor Objectivity usage at all sites
- Tools?
- Current plan
 - Run oolockmon utility at regular intervals
 - ◆ Lists locks on lock server with node & user names
 - Parse report to compile statistics
 - Need to run on all lockserver machines
- Don't intend to work on this until some indication that we're getting close to 150 limit

9/23/97

David R. Quarrie: BaBar Databases - Issues



Federated Database Strategy

- Protect master federated database
 - Contains schema and database catalog
- *Reference* federated database
 - Contains schema only
 - Used for builds of new *BABAR* software releases
 - Schema evolution disabled by default - managed evolution.
- *Production* federated database
 - Database catalog
- *Developer* federated database
 - Copy per developer, copy of *Reference* FDDB
- Federated DB ID allocation scheme underway

9/23/97

David R. Quarrie: BaBar Databases - Issues



Schema Management

- Incorporating new classes into production
 - Allow validation checks & formal acceptance procedure
 - ◆ Management inertia
 - Accept no data into production federation that wasn't created with production schema.
 - Individual developers can add their own schema to their local federation & generate data using it.
- Sharing of schema & data between developers
 - Don't yet know how to share both data & schema between developers because of possibility of class Type ID clash if they have both added new classes
 - ◆ No problem if willing to regenerate the data



Schema Management (2)

- How do we deal with modifications to our existing schema?
 - Schema evolution or class versioning?
 - Declare that it can't happen & create new classes?
 - Updating of existing data?
- Extensible event structure is designed to minimize schema evolution
 - Also breaks dependencies



Data Distribution

- Active work on this
- Requires significant management tools
 - Allocation of free databases & exporting to remote sites
 - Tracking of databases from free pool
 - Tracking of databases imported back to SLAC
 - etc.
- Goal is to have a functional version by end of 1997.

9/23/97

David R. Quarrie: BaBar Databases - Issues



Partial Reprocessing

- Plan is to use Object Versioning for this
 - e.g. Version the REC Header
 - ◆ Each version has its own set of data objects
- Dealing with partial reprocessing
 - Intermediate “chunk” level between REC & leaf nodes
 - Need versioning at this chunk level
 - How does this propagate downstream?
 - ◆ ESD, AOD, TAG
- Needs more investigation yet
 - Ray Cowan

9/23/97

David R. Quarrie: BaBar Databases - Issues



Data Logging

- Current strategy
 - OEP outputs streamed event data to disk files
 - ◆ *Not* objects
 - Spooled files decouple deadtime generation
 - ◆ Location of spool disk?
 - ◆ Stage to DLT tape if disk becomes full
 - PR inputs from spool files & outputs to database
 - ◆ First contact with HPSS
 - ◆ Alternative could be to have small DB loader jobs first load raw data into database & then have PR access from the DB
 - Each PR node writes to a different output container
 - ◆ Minimizes lock conflicts

9/23/97

David R. Quarrie: BaBar Databases - Issues



Data Logging - Conditions DB

- Both online & offline require access to conditions database
- Proposal is to make online a secondary Objectivity partition
 - SCS is primary partition
- Updates from online automatically propagated to offline
- Updates from offline automatically propagated to online
- Semi-autonomous in case of network failure
- Current implementation of partitioning has some problems
 - Needs more work on understanding limitations
 - Role within data distribution strategy?

9/23/97

David R. Quarrie: BaBar Databases - Issues



Management Tools

- Database management tools
 - Monitor usage & access efficiency
 - Cleanup database
 - Perform replication or re-clustering etc.
 - User Interface rather than programmatic interface
- Large Hole at the moment
 - Long term schedule details tasks but nothing available until next year
 - Hope can use SLAC Database/Web expertise to help us develop user interfaces

9/23/97

David R. Quarrie: BaBar Databases - Issues



Performance Monitoring Tools

- Objectivity gives you some statistics gathering tools
- Not much thought gone into this yet
- I think it's crucial
- Part of the motivation for using BdbRef(T) rather than ooRef(T) *etc.* was to allow for some monitoring to be added
- Intend to spend more time on this in the next 6 months

9/23/97

David R. Quarrie: BaBar Databases - Issues