

# Databases Breakout Session Report

Karl Young and Bob Schaefer, Conveners.

Participants: Cathie Meetre, Joane Bogart, Traudl Hansl-Kozanecka, Toby Burnett, Pat Nolan, Dave Davis

## PRESENTATIONS:

There were three presentations in this breakout session, more than in any other breakout session.

1. Traudl Hansl-Kozanecka – Use case of EROS database as an example for structuring the GLAST mirror database. Traudl presented some details about the EROS database, which is similar in size to the level 1 event database. She presented design criteria, architecture, and an overview of operations. This presentation will be valuable for designing the interface to the level 1 database. She also mentioned that there is a French infrastructure already in place for serving out astronomical databases.

2. Bob Schaefer – Beowulf option for the level 1 database architecture. The option to use a Beowulf cluster of PCs to search the level 1 data in parallel has been discussed for some time. Bob presented a working prototype of this database that can do spatial searches. Performance numbers were discussed and it was shown that the prototype exceeds the requirements by roughly an order of magnitude for search speed. Options for improving the performance were also discussed.

3. Pat Nolan: Database Management System Options and constraints. Pat gave us an overview of some of the things he has learned about database management systems under consideration for GLAST. Free databases (PostgreSQL and MySQL) were described as simple and easy to use, but short on features. Oracle's database was discussed at some length. Oracle has an amazing array of customization options for separating the index from data on disk, partitioning the data, and distributing the database. However, the installation package comes with 5 CDs, 2 of which are for the documentation. Oracle is clearly the most powerful database considered, but it is also the most complicated and will probably need skilled care. Pat also suggested looking at Object-relational databases, which would give us more flexibility in the data structure. Pat mentioned that all databases considered can beat the time requirement for ingesting a day's worth of data.

Discussions about the presentations.

Most of the discussion centered around strategies for getting the best performance out of our databases. The topics discussed included:

- Beowulf query “bus” for concurrent searching (searching on multiple queries received in a given time period). Here the server would accumulate queries for a

specified amount of time and then do multiple queries simultaneously as it searches through all the photons. (The next “bus” leaves in 1 minute...)

- Quad tree schemes for pixelizing the data with a rapidly searchable index (HEALPIX, HTM, RTrees, R\*Trees, QuadTrees, KD-Trees).
- Subdivision by usage (Data Warehousing).
- Multiple processors/disks for DBMS.
- All mirror sites should use the same database architecture (which might make use of Oracle problematic).
- Exclusive access of D2 only through U2 was questioned, but we decided that is probably OK.

### **USE CASES:**

It was generally agreed that some query needs to be identified, which goes beyond the point source example used in the requirements document. The fear is that tying the performance numbers to one specific search type leaves us blind to all other types of performance. We will go through all the analysis tools use cases and look for other types of searches (other than single point source) that that may be commonly needed.

### **REQUIREMENTS SUMMARIES:**

With the exception of the dependence of the performance on just the “standard searches”, it was felt that the level 1 event database requirements are fairly well defined. Other databases are relatively straightforward and presumably the requirements are not much beyond what is already in the Science Tools description document.

### **OPEN ISSUES:**

The main level 1 event database architecture is still not defined. While we now know facts about different database options – there has not yet been a head to head comparison of different database options. At a minimum performance comparisons between a Beowulf system and at least one relational database system, somehow deemed to be of roughly equivalent capability, on the set of use cases or “standard searches” described above should be performed. A complete test of Oracle’s capabilities may not be possible by the Fall review as spatial indexing is provided in additional libraries (i.e. not currently available at SLAC) that appear to be fairly expensive (10’s of K range)

### **GOALS:**

- 1) Do a head-to-head comparison of database options and identify the level 1 event database architecture by the Fall review.
- 2) Produce all database requirements documents by the Fall review.

## TASKS:

The tasks listed below have been compiled to accomplish our goals. We have ~2 FTEs until Fall (sum of Pat, Karl, Bob, and Dave and other attendees) to complete these tasks.

- 1) Identify key performance tests – Bob volunteered to try to get more search use cases by studying the searches that would be used in the analysis tools use cases. Any outside help would be appreciated.
- 2) Draft database requirements documents – lead by Bob and Karl.
- 3) Testing Requirements
  - a) Need to identify hardware capable of doing performance tests.
    - i) Bob & Dave will identify hardware at GSFC
    - ii) Karl will find hardware to do testing at SLAC.
  - b) People & Time for testing (Pat, Karl, Dave, & Bob). Sum of time is ~1 FTE until Fall.
- 4) Do measurements of key performance parameters by architecture
  - a) Beowulf – Bob & Karl
  - b) PostgreSQL, MySQL, Sybase – Pat, Bob, & Dave
  - c) Oracle – Pat and Karl
- 5) Identify elements for easy operation to further distinguish databases (to be done by all attendees).
  - a) Data management - how data is to be managed within the database (including making sure that tables contain necessary information with minimal redundancy)
  - b) Maintenance issues for different databases (hardware, software, & data).
  - c) Extensibility questions – how difficult is it to change data formats and sizes?
  - d) Identify risks (where are weak points and what mitigation can be done)