Data Server Needs from the Science Tools Perspective

[Needs from the Perspective of LAT Scientists Trying to Do Astronomy]

S. W. Digel
SLAC

Outline

- Science tools concept of the LAT data server
- What we need vs.
  - What we really want
  - What we'd like if possible
- Analysis Group Power Users
- Implementation issues

What's in the Data Server?

- **Processing chain [schematic]**

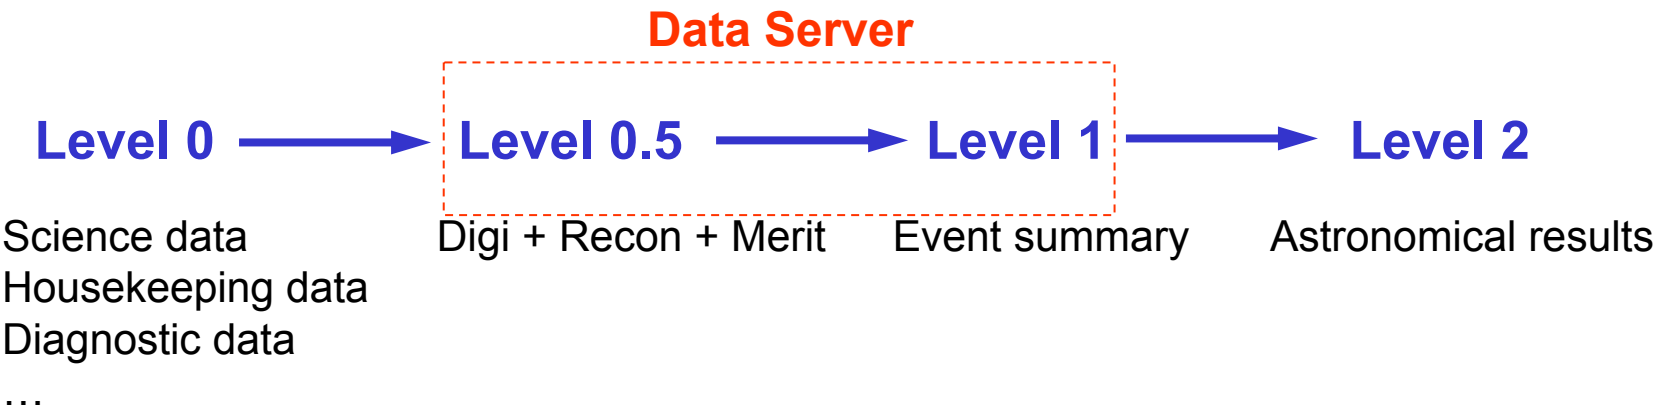

- **Level 1 data are a *subset* of Level 0.5**
  - **The part of Level 0.5 that the science tools care about**
    - **With infinite resources, we wouldn't make a distinction**
  - **Level 1 data don't necessarily* have a separate existence inside the Data Server**
- **Note that handling of pointing history (location, attitude, livetime accumulation, and LAT mode) is not included above as part of the Data Server; see next slide**

* Having them separately stored may be useful for implementation reasons; also, separating cosmic rays from (celestial) γ-rays in the data server might be worthwhile, too

What's Not in the Data Server?

- The answer depends on whether you subscribe to the 'Big Tent' concept of the data server
- The most obvious candidate is the pointing/livetime/mode history (D2 in the SAE), which is unrelated to the events
- Other science tools-related (and ISOC-related) databases including those related to calibration
- What about HK and diagnostic data (and maybe even spacecraft HK)? I suppose this is an ISOC issue.

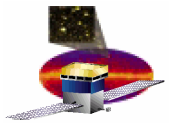

What We Want

- Basic need is to access Level 1 data by region of the sky (and other criteria) in a reasonable amount of time
  - Other selection criteria include: time range, energy range, event class, zenith angle (perhaps energy and angle dependent), maybe even inclination angle range
    - This is a large set, but much more limited than what, say, an Analysis Group Power User might need
  - Access needs to be supported via a Web interface and (for automated analyses) an API
  - And FITS output (FT1) needs to be an option for the Level 1 (gamma-ray) data
- See Bob Schaefer's presentation from the June 2002 Science Tools workshop and corresponding 'requirements' document
  - These were reviewed by Tom today
  - http://lheawww.gsfc.nasa.gov/users/bob/docs/GLAST_Event_DB_Reqs_DR_V1.4.pdf
  - http://www-glast.slac.stanford.edu/sciencetools/workshops/june02/slides/DBReq_summary2.pdf

What We Want (2)

- Regarding performance requirements, we inquired about access rates (instrument team + rest of world) for Chandra data archive to get an idea of performance requirements (queries served per day)
  - Concept of a standard query
    - For high-level analysis, smallest region of the sky and time range of interest is generally not very small
  - My expectation is that rate of queries from LAT team members will be at least as great as for the world at large
  - Also, reasonable real-world requirements on ingest times, availability during ingest/re-ingest, etc. can be specified
- One consideration that the GSSC needed to pay attention to, but we probably won't, is that when you give something away, some people are going to want all of it.
  - We'll have mirror data sites, it has long been stated, but we won't be serving data to the mirror sites so much as sending them whatever gets ingested into the Data Server.

What We Also Want

- After deliberation, the SAE was defined to include an event display tool
- This is problematic for a few reasons
  - Requires Level 0.5 data (which by calling Level 0.5 we signalled that we didn't intend to deliver it right away to GSSC/HEASARC)
  - The existing tool would be hard to distribute and support for GIs
- Current plan
  - Providing a *server of Level 0.5 data at SLAC* – to be accessed by event number (i.e., not for mass downloads) [Richard's suggestion]
  - And a Web-based version of FRED (with limited, but entirely adequate, controls) [Riccardo's suggestion]
  - Full Level 0.5 data to be delivered to GSSC/HEASARC at end of mission in acceptable format [Apparently ok with the GSSC at the time]

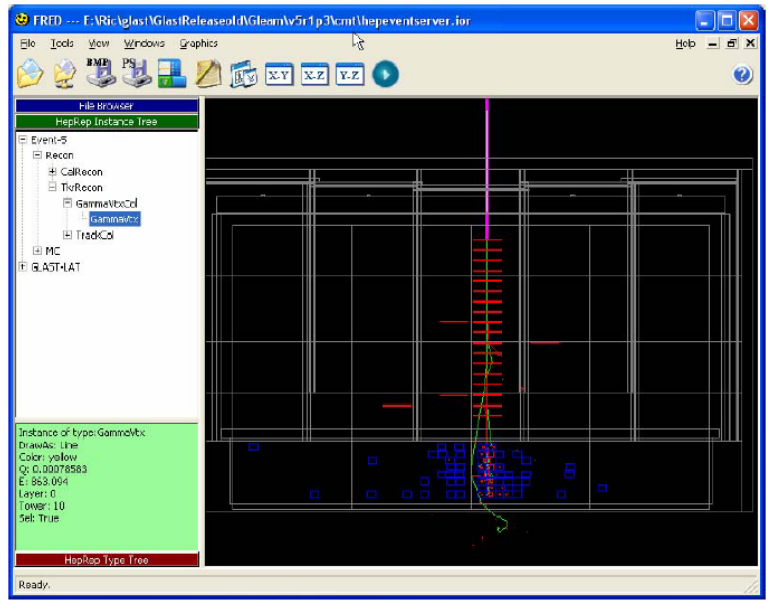

R. Giannitrapani

What We Might Also Want

- **Being able to access all versions of events would prevent analyses from becoming irreproducible after reprocessing.**
  - **More slowly or less conveniently than the current versions would be ok**
  - **In principle, retaining ability to rerun event processing programs could be enough, but Leon's experience last year points out how difficult getting old software to run can be**
- **An implementation that would also allow the Data Server to manage (but not confuse) MC and beam test data with flight data would be desireable**

About Analysis Group Power Users

- From a power user (or ISOC Performance Verification) perspective, one could imagine wanting to access the data in ways that wouldn't be interesting for astronomical analyses
  - For example, by arrival direction in instrument coordinates, by ACD tile, by non-interacting heavy CR interactions in a specific CAL log
- This is not my domain but also may not be covered at this workshop

From Requirements toward Solution

- I don't know the answer, but don't forget the extensive work done by Pat Nolan in 2002
  - Well documented here: http://www-glast.stanford.edu/cgi-prot/wiki?DataBases
  - [Tom Stephens has also documented performance investigations for the GSSC data server, as he has probably presented today]
- Basic findings (see above for details)
  - Region-of-sky searches are hard to make go fast, especially in our circumstance where we are continually adding data (and so cannot sort the database once and for all like, e.g., SDSS)
  - Indexing essentially doesn't help; issue becomes I/O to disk; brute force is faster than RDBMS; ROOT is not magically faster than FITS; local disks are probably better than network accessed

What has *also* been tried or considered or recommended

- Beowulf – [like GSSC Level 1]
  - This is a brute force search with finesse
- ROOT 'Peeler' – DC1 data server implemented by Navid
- ROOT PROOF?
  - Not sure whether this was tried or worked
  - http://www.slac.stanford.edu/BFROOT/www/Computing/Distributed/ROOT2004/files/kelly.pdf
- Storage Area Network
  - Recommended by an SCS database expert; not explored, as far as I know
- All-in-memory (or an index in memory) approach
  - SLAC may become a leader in multi-CPU systems with lots of memory
  - http://daily.stanford.edu/tempo?page=content&id=14302&repository=0001_article
  - http://www.rhic.bnl.gov/hepix/talks/041021am/wachsmann.ppt
  - Some questions include whether we could use such a computer, and time scale for its development

Keep in Mind

- As mentioned above, searching of the γ-rays will go faster if the cosmic rays are stored elsewhere
  - This is part of the approach of the GSSC; the server so far is just for events that are (probably) γ-rays
  - Implicitly, this may be part of the '90%' solution that Richard has discussed – not storing Recon information for events sure to be cosmic rays that aren't otherwise interesting for calibration.
    - At ~300 Hz data rate in the telemetry, <1% of the events reaching the ground will be celestial gamma rays. Richard's filter would increase the fraction to ~10%
- For a brute force search, storing only the essential parts (i.e., the information needed for higher-level analyses) means faster searches
  - This is also part of the GSSC's approach for the data server

Keep in Mind (2)

- Tom has also noted that the Event Summary files generally will be large (100s of Mbyte), and transfer time from the GSSC to wherever they are going can be much greater than the time it took to extract the file from the GSSC data server in the first place.
  - For LAT team use of our data server we can probably assume (to first order) that the files won't have to go anywhere outside of the SLAC computer center (or whatever mirror analysis site is used)