



... for a brighter future

Metadata and Supporting Tools on Day One

David Malon <malon@anl.gov>

Argonne National Laboratory

Argonne ATLAS Analysis Jamboree

Chicago, Illinois

22 May 2009




U.S. Department
of Energy

UChicago ►
Argonne_{LLC}



A U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC



Day One: *“What will this day be like? I wonder.”*



N.B.: from **“I Have Confidence,”** The Sound of Music, Rodgers and Hammerstein



Disclaimer

- Rik asked me to say something about what the day one situation will be.
- Here are some guesses. They are optimistic without being unrealistic--
 - but they do require that people get involved and make things happen.



Finding data

- Identifying which datasets are the ESD, AOD, performance DPD for your stream will be straightforward
 - Metadata in AMI will improve substantially, but
 - Most people's dataset selections will continue to be based upon well-known names, and
 - the early “Good Runs” will be well known, mentioned by number on Twikis and in meetings everywhere



Quality and status information: versions

- Quality information will be updated continually
- Getting the first DQ flags from quasi-real-time monitoring by shifters will be in good shape
- Getting subsequent updates into the database reliably, consistently, timely will not yet be automatic
- Data Preparation/Data Quality coordinators will try to “tag” consistent versions of quality/status updates, but the process may at first look a bit like software release coordination
- Infrastructure for maintaining offline status/quality metadata will be limited and fragile
 - And cause painful problems for some of today’s presenters



Good Run Lists

- Good Run Lists will be a reality, at least for some purposes
 - Not all the flags in the three-tier hierarchy will necessarily be filled
 - Some Good Run Lists may not quite be based upon DQ flags as proposed
 - Volunteers could certainly accelerate progress here
- Standard samples for a physics group, based upon Good Run Lists (explicitly or implicitly), will be okay at the D1PD level, but further derived data products (D2PD, D3PD) will give everyone headaches
 - Particularly when an expert user's dataset needs to be “promoted” to become a production dataset



Tools for luminosity calculation

- Tools like LumiCalc.py myDataFile.py –trigger=myTrigger will work under normal conditions
 - And under certain reasonable error conditions
- Bookkeeping for multi-stream analyses will be fragile and complex
 - And much of the complexity is inherent
- LumiCalc will work even when one begins with a TAG-based selection
 - At least if all the events in the selection can be retrieved, but the prospects for doing better are good
- Getting trustworthy lumi estimates, and getting them into the database, is another matter, but the tools to access the data for analysis will be in place



Analysis tools and metadata

- Simple tools e.g., to cull your sample to Good Run Lists or to the latest quality updates, will be available
 - “Iterate through the events in my file, but skip any events that are (or are not) in this auxiliary list of Bad/Good {run, lumi block}s.”
- *Augmenting* your sample because lumi blocks that were previously “Yellow” have been changed to “Green” will be much harder



Metadata chains and repositories

- Some metadata pipelines are today quite fragile and heterogeneous
- Particularly for globally distributed production:
 - Components to publish metadata from inside Athena
 - Components in the wrapper transforms to marshall metadata from many sources
 - Components in the “production system” to pass file-level, job-level, task-level metadata back to a repository
 - Repository infrastructureAll of these need work.
- On Day One, the situation will be much improved
 - No more “we don’t know how many events are in that file” (?)
 - But a new core transform layer is needed, and it appears that it may not be ready unless new effort is found
 - *Probably already too late given release deadlines*
- On the plus side, more and more physics metadata will be available in AMI



Metadata from online to offline

- Replication from online COOL to offline COOL is relatively reliable
- Need therefore to ensure that what we need offline actually gets into online COOL on a more than a “we’ll take care of it when we have a chance” basis
- My biggest worry (perhaps needless) is getting the necessary trigger counts and deadtime-related information from the HLT processors through the trigger’s IS to the databases reliably
 - People are thinking about this, and *something* will be working
 - *Or maybe you don’t believe the HLT will be a player on Day One*
- Some unmanaged practices that don’t matter in the trigger *do* matter offline
 - Maps of L1 triggers to CTP bits, L2/EF to chain counter numbers
 - Work is in progress to make this as stable as possible



Metadata component integration

- Interfaces to select which runs and lumi blocks to query will be integrated with interfaces for event selection
 - “Which runs and lumi blocks” is usually a selection from a range of runs the ones that meet your quality and status criteria, AND for which the triggers you care about were part of the menu
 - “Which events” is usually based upon physics quantities and triggers passed and so on (information in the TAGs)
 - You will be able to use the TAG database to select, from the top group’s Good Run List Version 3.1, all the events that satisfy your trigger and physics cuts
 - Will be able to use a single interface to select your runs and lumi blocks, and then optionally to proceed with an event-level selection
 - *You will be able to see integrated luminosity and trigger counts and so on at the end of the first step*



Metadata access

- Access patterns to auxiliary non-event data from Tier 3 users are not well understood
- Already know that a surprising number of tasks require non-event data, for less than obvious reasons
- Most metadata needed for post-AOD analysis will be cached in the data files, but some tasks will need more
 - And you will always want to get the latest lumi estimates and quality information for your sample before hitting the “publish” button—you cannot assume that what is in your file is authoritative
- Some access at Tier 3s to metadata will be required
 - Probably will be extractable by a command-line utility into a local file
 - But this may require that the site allow outbound database connections, and that a remote site allow the incoming connections
 - *This is an occasional connection—one user connects one time for small-volume data even if she is running N jobs*

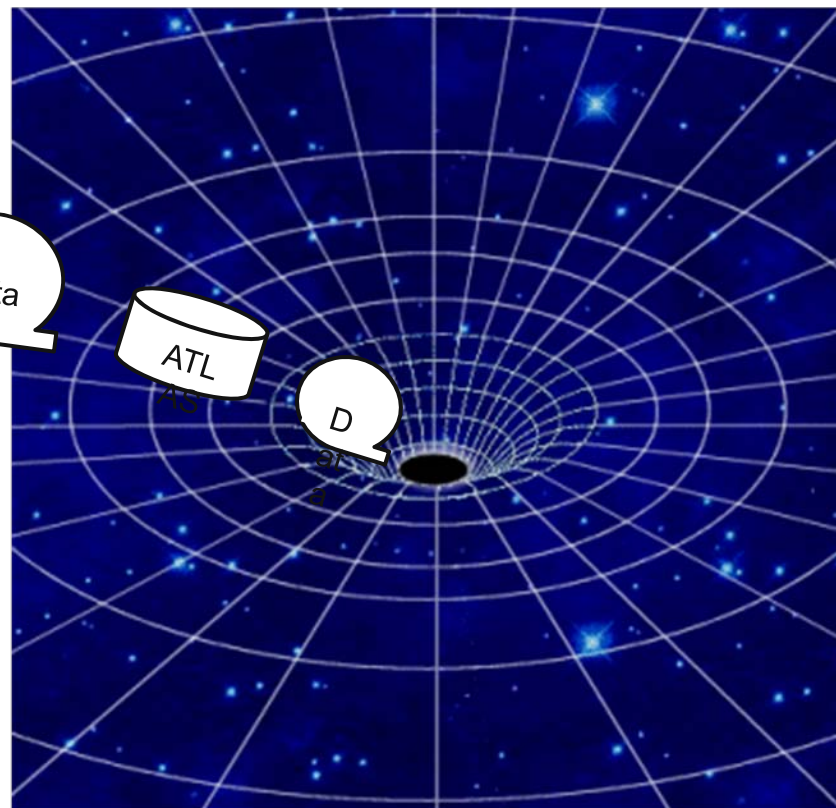




Event-level selection

- Event selections will properly carry and process provenance information
- For efficiency, TAGs will be writable into event data files (AOD, DPD)
- Forward pointing from TAGs (i.e., to data that had not yet been produced when the TAGs were written) will not be possible (yet)
- Replicas of the TAG database at sites others than CERN will not be fully functional
 - E.g., data may be there, but surrounding services may not



Day one in an alternate universe...



"Try not to look worried, David.
The metadata are tracking this."

