

TAG Status and Plans

Jack Cranshaw

07.04.2011

Atlas Software & Computing Week

TAG Building

- TAG building has been in a bit of a limbo since last October when Tulay retired.
- PAT still working on finding replacement
 - Residual issues building up
 - Container changes in AOD
 - Requests for new information
 - ...
- Otherwise the code is stable.
- But builder code is also a quasi-replica of what is in the D3PD makers.
 - Push by PAT management to rewrite TAG building to use the same code base as D3PD making.
- There was an issue with triggers being defined outside the range usable by TAGS.
 - Minor issue for data10_7TeV, large issue for data10_hi

TAG Distribution

- Files
 - *Tier0*: stable
 - *Reprocessing*: varies by coordinator/campaign
 - Are they produced?
 - Do they get subscribed to CERN?
- Database
 - Uploads stable other than a few glitches, distribution issues.
 - Duplicate events (the sequel)
 - Upload optimizations (to merge or not to merge)
 - Which sites (all sites are equal, but some are more equal than others)

User Support

- Feedback
 - Logging of service activity (*under development*)
 - DAST requests
 - Event pick: about 1 thread/week
 - Generic TAG: about 1 thread/week
 - Roughly half of issues are site or tool issues
 - File catalog incorrectly made or not made
 - Wrong parameters
 - ...
 - Tutorial comments
 - Resulting new features
 - Amitag qualifier for event pick
 - Trigger counts for ELSSI
 - ...
- Integration with group production
 - Currently poor
 - Working with PAT to improve

Performance

- 2010 was a focus on functionality
 - objectives achieved!
- 2011 is a focus on performance.
- Issues
 - ELSSI (query construction)
 - Pages built based on information which changes by run and may be large.
 - Many user criteria trigger full table scans.
 - Users looking at full collections/stream rather than single or limited runs.
 - Event counts can take nearly an hour!
 - Use of functions in queries or complicated arithmetic comparisons.

Performance (II)

- Issues (continued)

- POOL

- POOL utilities basis of Lookup, Extract, Skim
 - Full scan rates
 - CollListFileGUID 200kHz
 - CollAppend 50kHz
 - CollSplitByGUID 30kHz

- AthenaPOOL

- Event data written vertically but read horizontally
 - Large basket sizes make event selection inefficient
 - Changes in data storage of AOD being tested which should speed up athena read speeds by > 2x (~20 -> ~50).

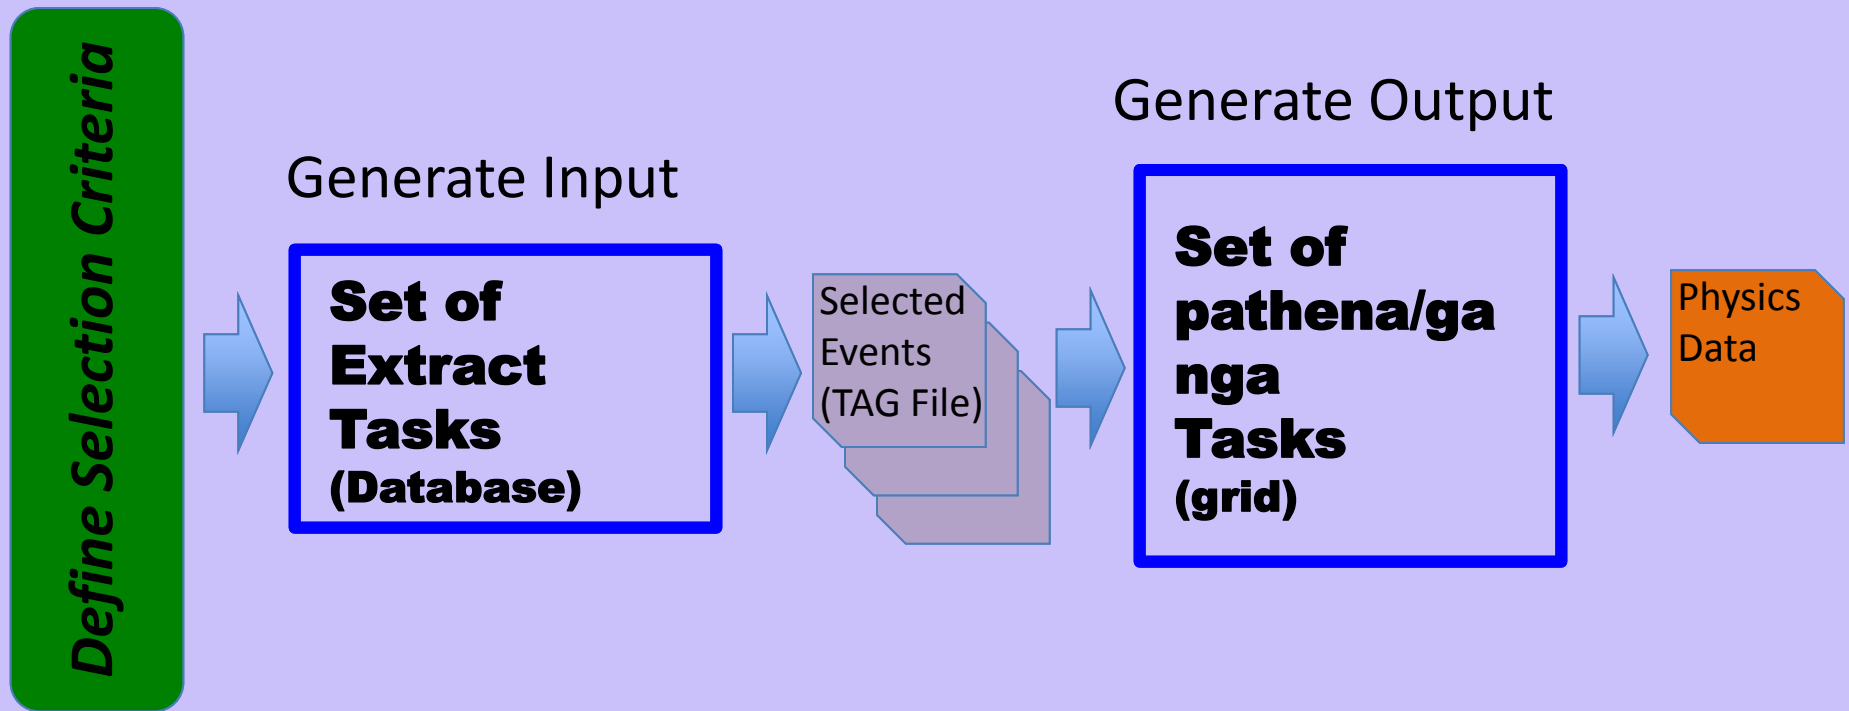
Performance (III)

- Many areas to investigate
 - Look at SQL in detail
 - Change or add indexes where needed
 - Limit user options to avoid bad query plans
 - Optimize looping and fetches
 - Eliminate unnecessary steps
 - Caching of query results or prebuilding summaries
 - ...

Remaining Functionality

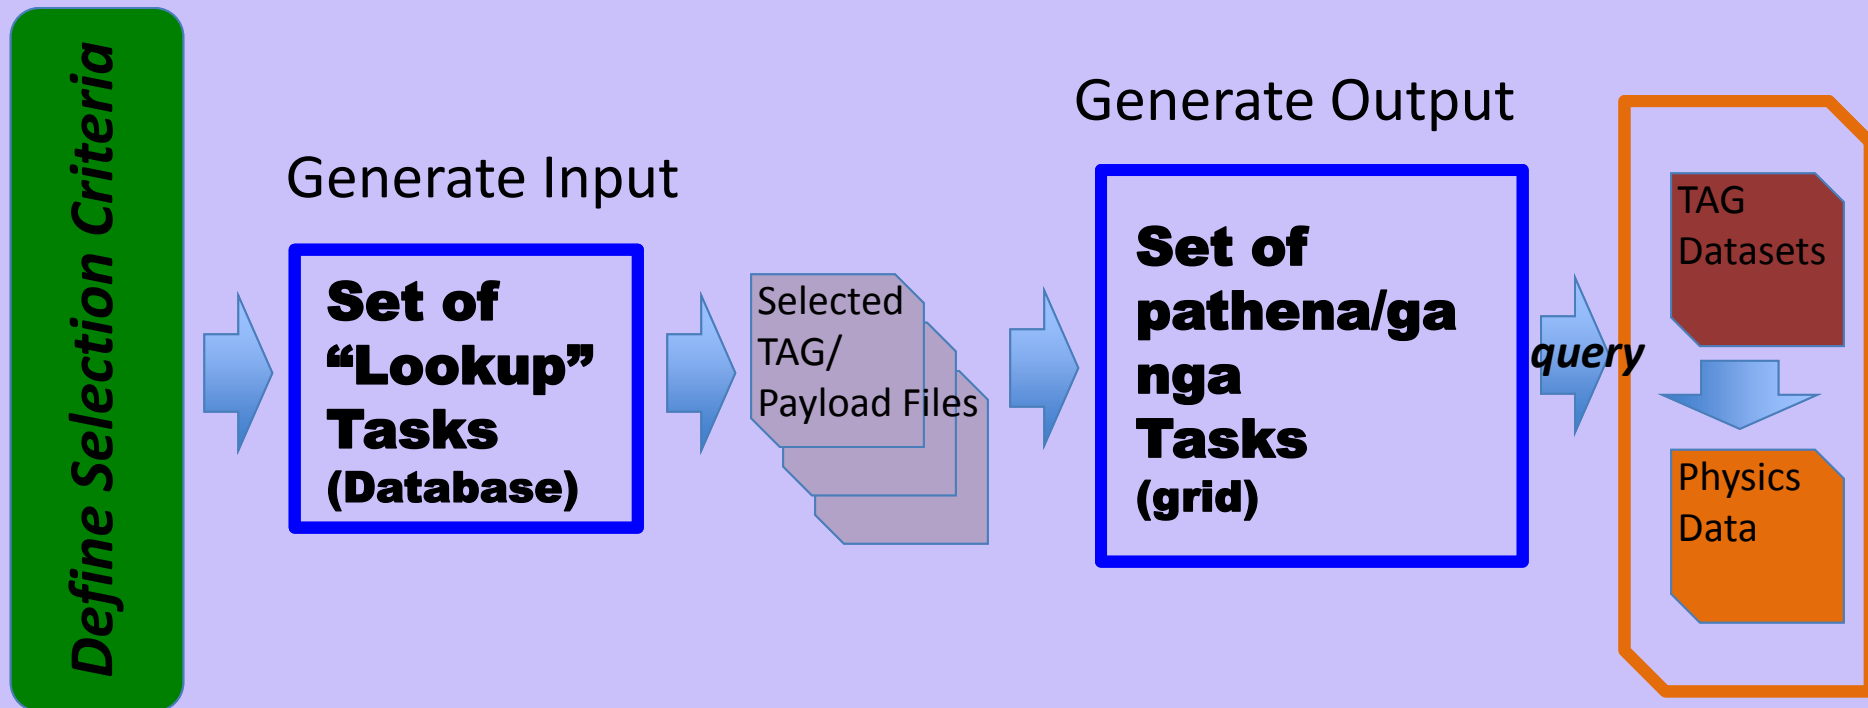
- In Progress
 - Skimming is much improved in capabilities and robustness.
 - Using panda backend with ganga (more robust)
 - Nightly testing added
 - Close to getting large, full stream, skims working
 - New student at Glasgow added to work on optimizations and integration with Higgs group
- In Development
 - Workflow Extensions
 - Can I go from the results of a skim back to an extract or to iELSSI?
 - Further factorizations of ELSSI innovations
 - Generate luminosity range record for queries.
 - Split query across sites for other services.
 - Simplifying service deployment procedures

Generalized Skim (Extract-based)



- Database activities separated from grid activities.
- Exchange format with full navigational capabilities.
- ADC has complained that this has overhead which has large effects on performance

Generalized Skim (Deferred Query)



- Database role is to provide guid/dataset information to job splitting.
 - Bookkeeping information also needed but not part of job submission.
- Actual “event” selection is done at the site.
- Being implemented by Tadashi and Mark Slater.
- TAG ‘file’ distribution model may need improvement/changes.

Documentation

- Florbela's departure generated another round of documentation review.
- EventTagInfrastructureAndTesting wiki had become a mess.
- Entry point is now on the EventStore wiki
 - [https://twiki.cern.ch/twiki/bin/viewauth/Atlas/EventStore#Navigational Metadata](https://twiki.cern.ch/twiki/bin/viewauth/Atlas/EventStore#Navigational_Metadata)
 - Improved, but still under development (*will that change?*)
 - Use cases being collected on the Overview page.

Personnel Issues

- Gaps appearing and not getting filled
 - PAT liaison (since last October)
 - TAG Database Administration (since last week)
 - 1 Florbela = some Gancho, some QZ, some Jack (shaken not stirred)
 - Tier0 TAG Coordinator (after October 2011)
- External budget pressures
- *Coordinating work across the Atlantic can supplement but not replace personal contact.*

Summary

- 2010 established a set of services with core functionality that are deployed and being maintained.
- Extensions of this functionality are being developed.
- Monitoring and user feedback are being incorporated into the development.
 - Better logging and testing.
 - Better integration with user groups: DAST, tutorials, physics groups.
- Focus is now on *performance* and response to user needs.