

A New Petabyte-scale Data Derivation Framework for ATLAS


James Catmore (University of Oslo, Norway)
on behalf of the ATLAS Collaboration

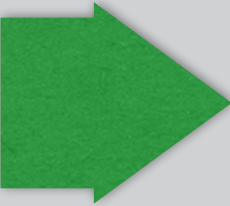
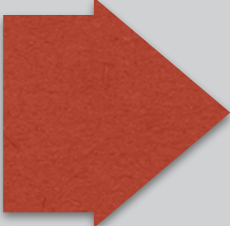

CHEP2015, Okinawa, 13th-17th April 2015

Contribution ID 164

A feature common to most physics analysis


2

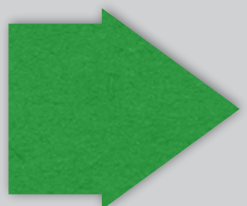
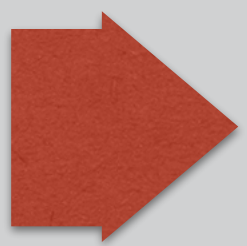
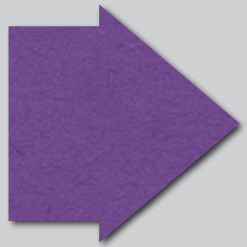


	Full output of reconstruction, ~PB size	One format
	Intermediate analysis format ~TB size	~100 formats
	Final n-tuple ~MB-GB size	~1000 formats

A feature common to most physics analysis

2




	Full output of reconstruction, ~PB size	One format
	Intermediate analysis format ~TB size	~100 formats
	Final n-tuple ~MB-GB size	~1000 formats


- These formats tend to be specific to a single analysis or group of analyses
- Calibrations and common object selections are often applied as they are made
- They generally need to contain all variables needed for calculating systematics
- In ATLAS in Run-I they were created by users; in Run-II we will produce them centrally
 - ➔ This is the purpose of the Derivation Framework and the topic of this talk

Data reduction operations


3




Skimming




Skimming:
removal of whole
events based on
pre-set criteria




Thinning



Thinning:
removal of whole
objects within
events based on
pre-set criteria



Slimming

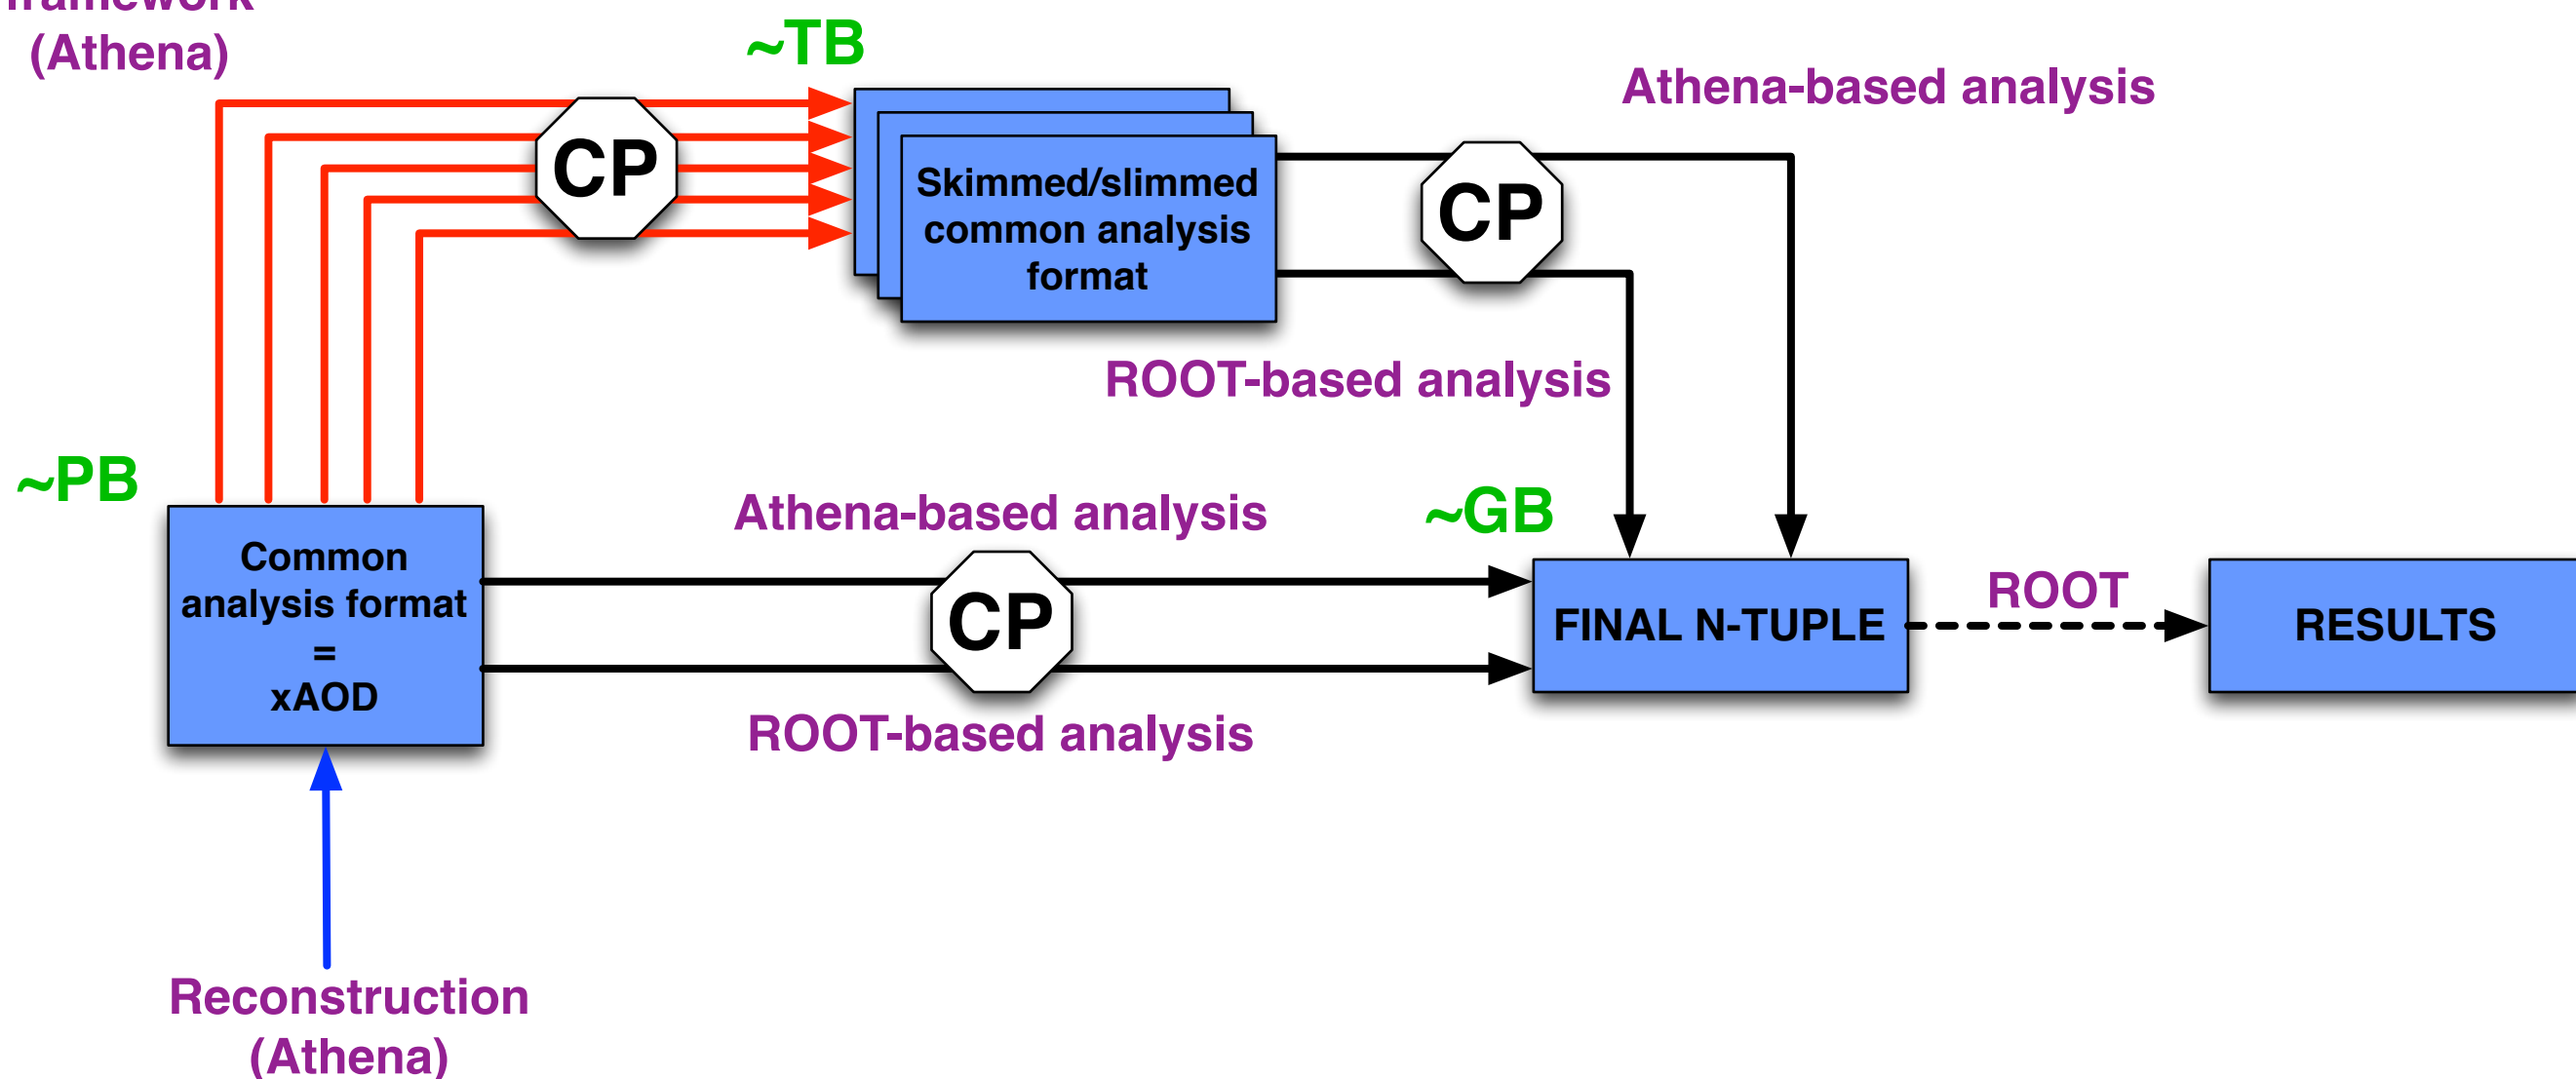


Slimming:
removal of
variables within
objects uniformly
across events

The Run-II analysis model for ATLAS

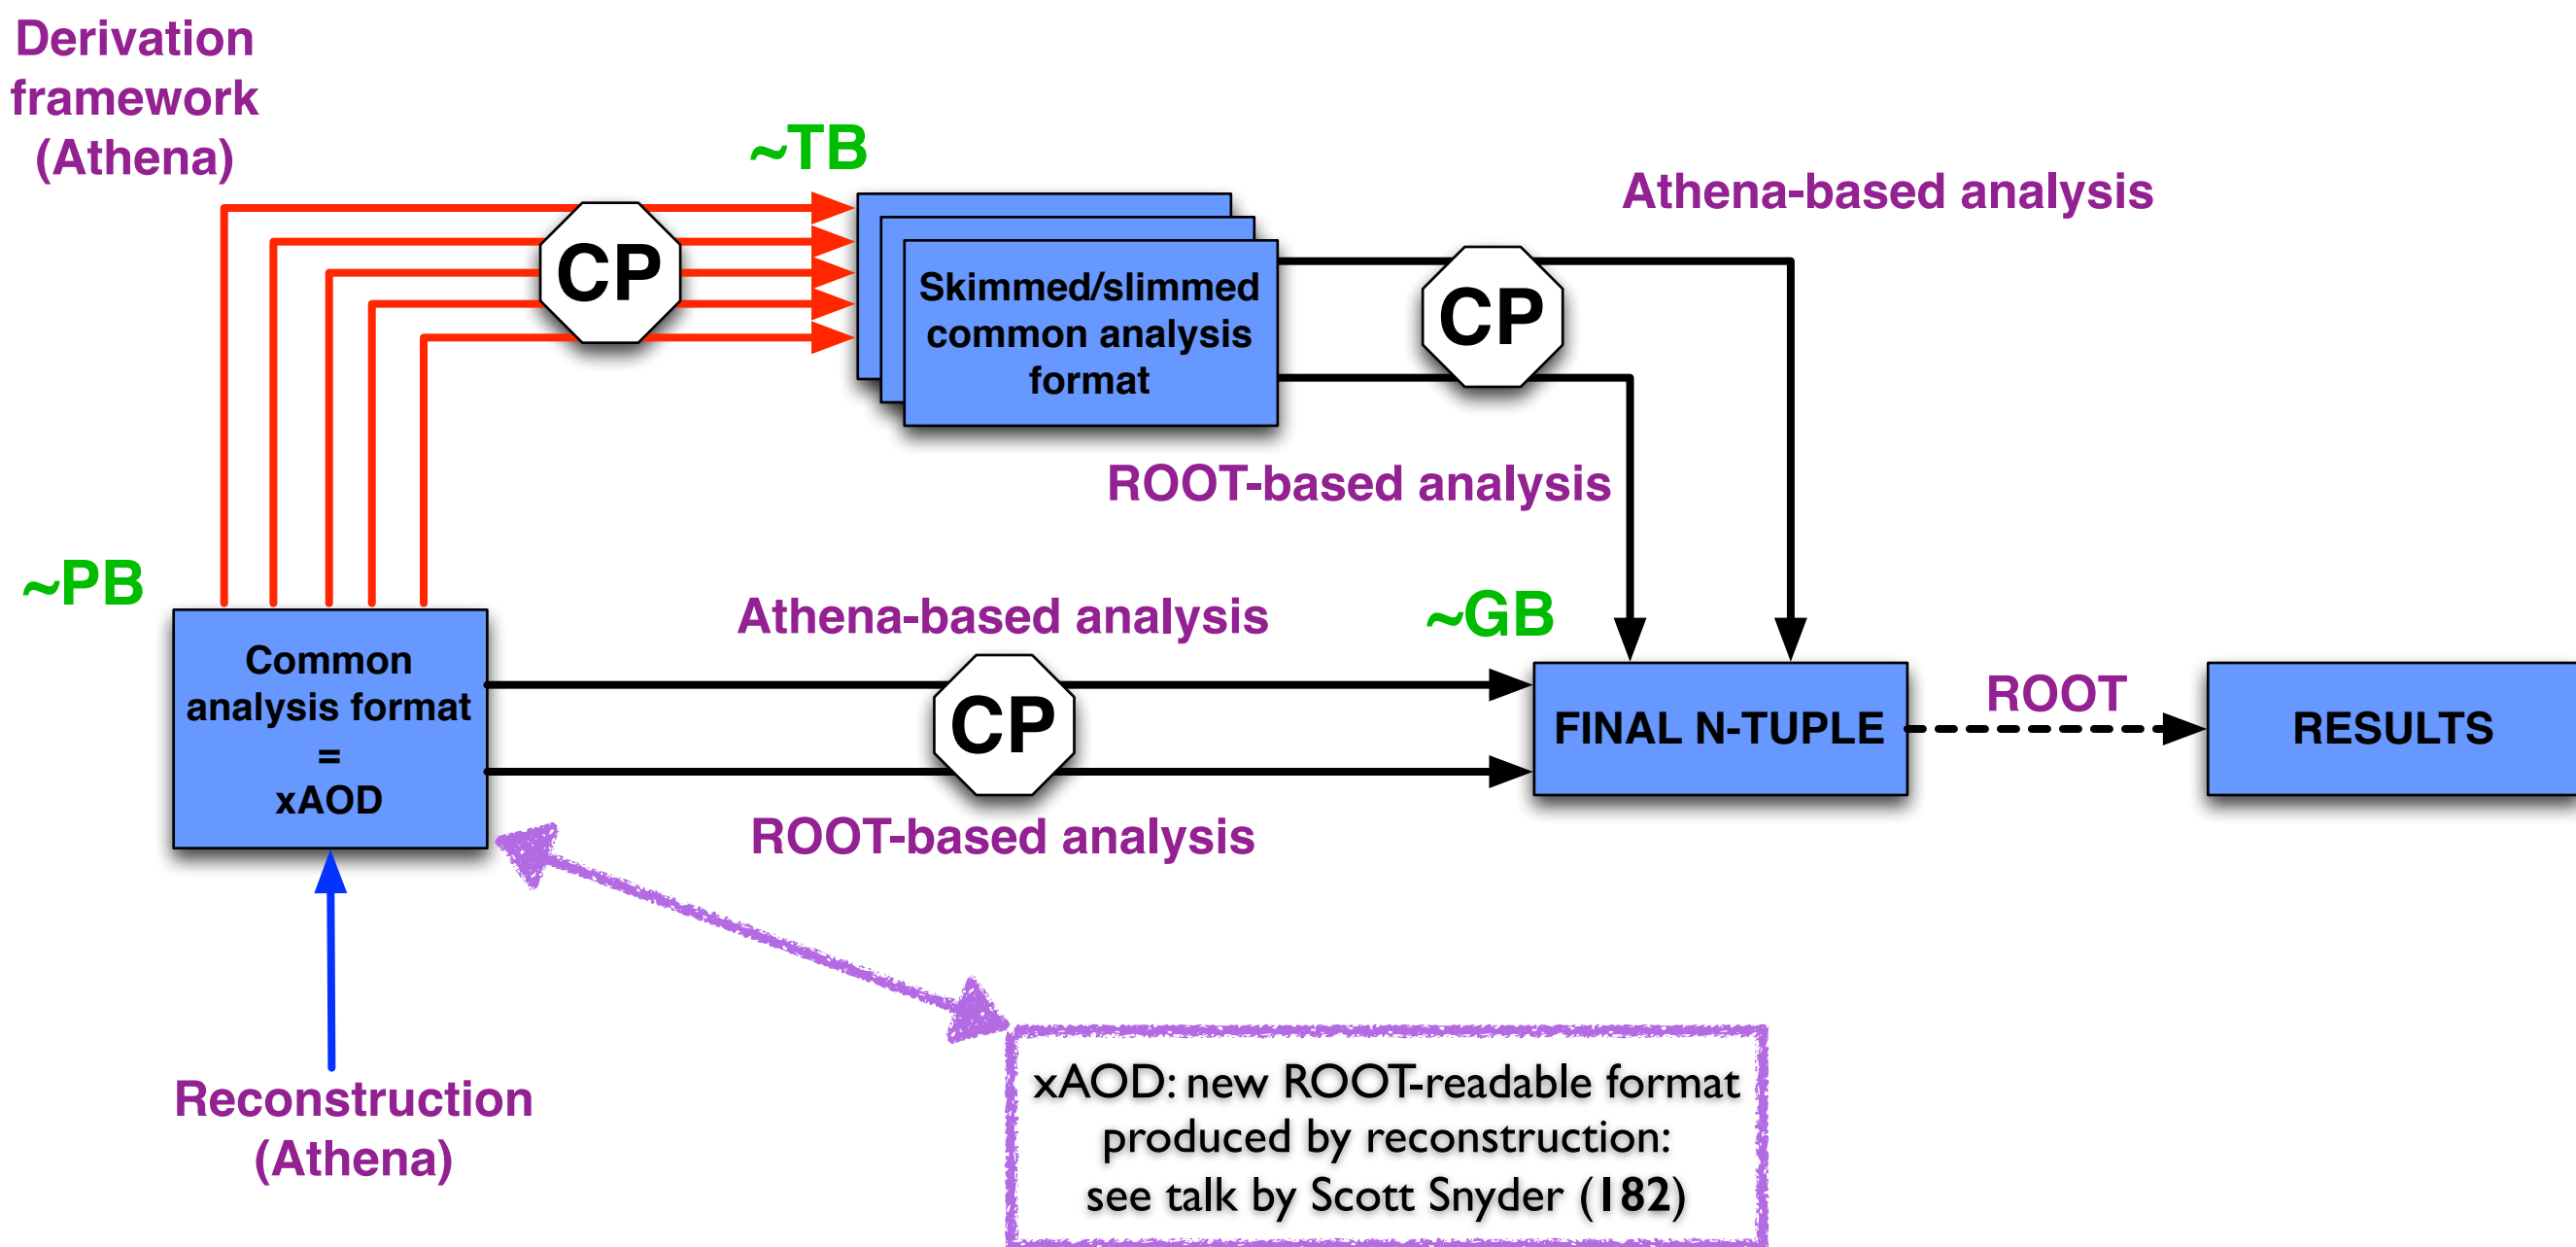
4

Derivation
framework
(Athena)



The Run-II analysis model for ATLAS

4

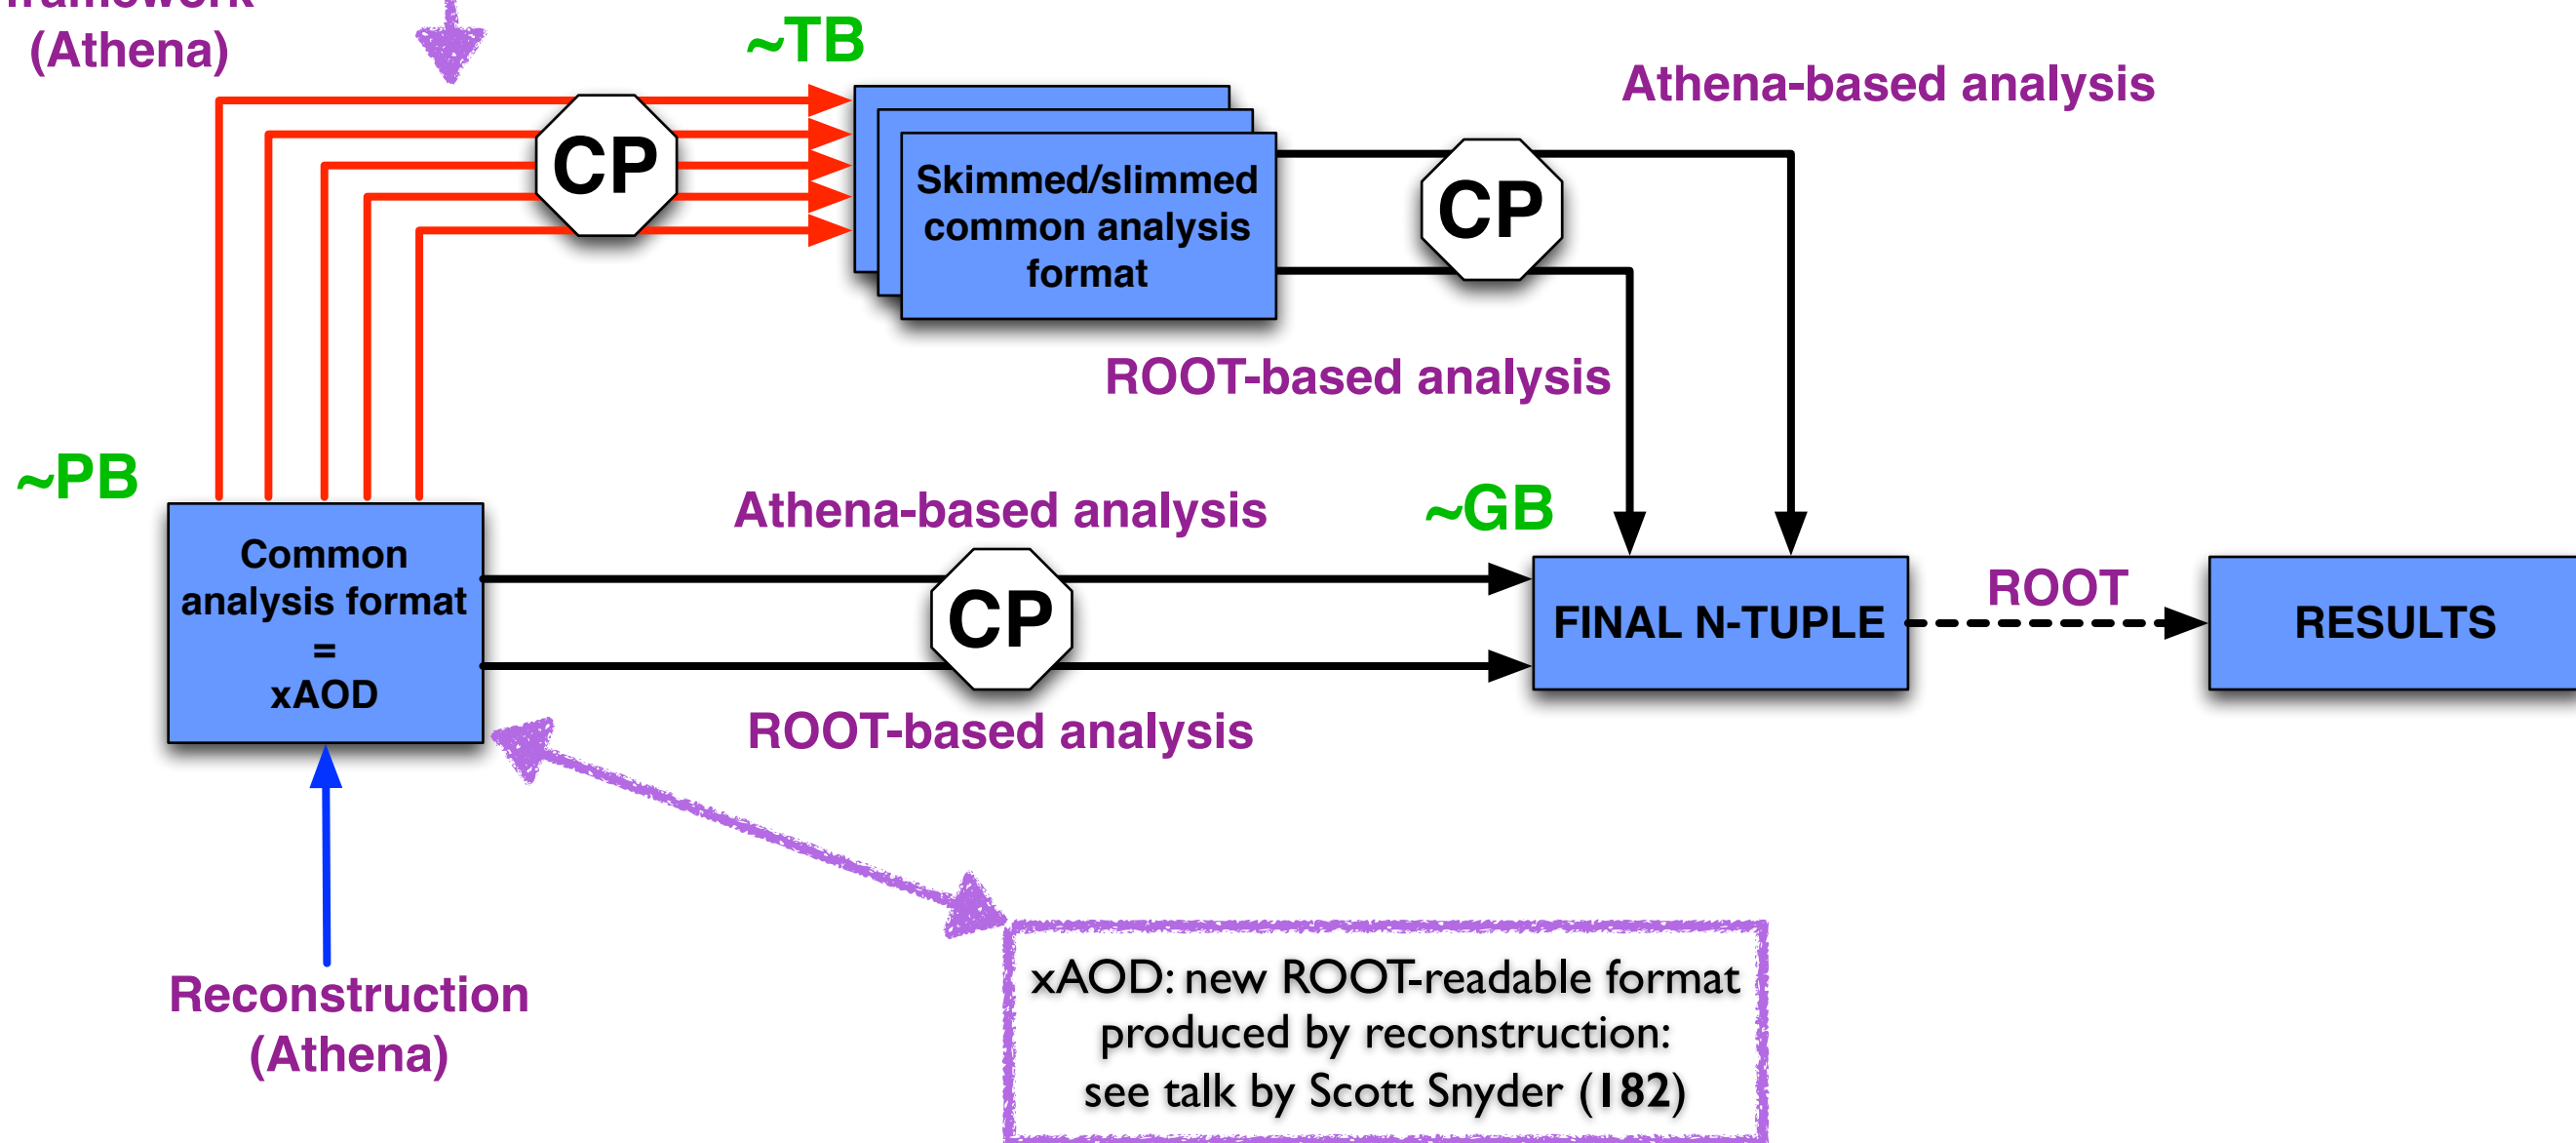


The Run-II analysis model for ATLAS

4

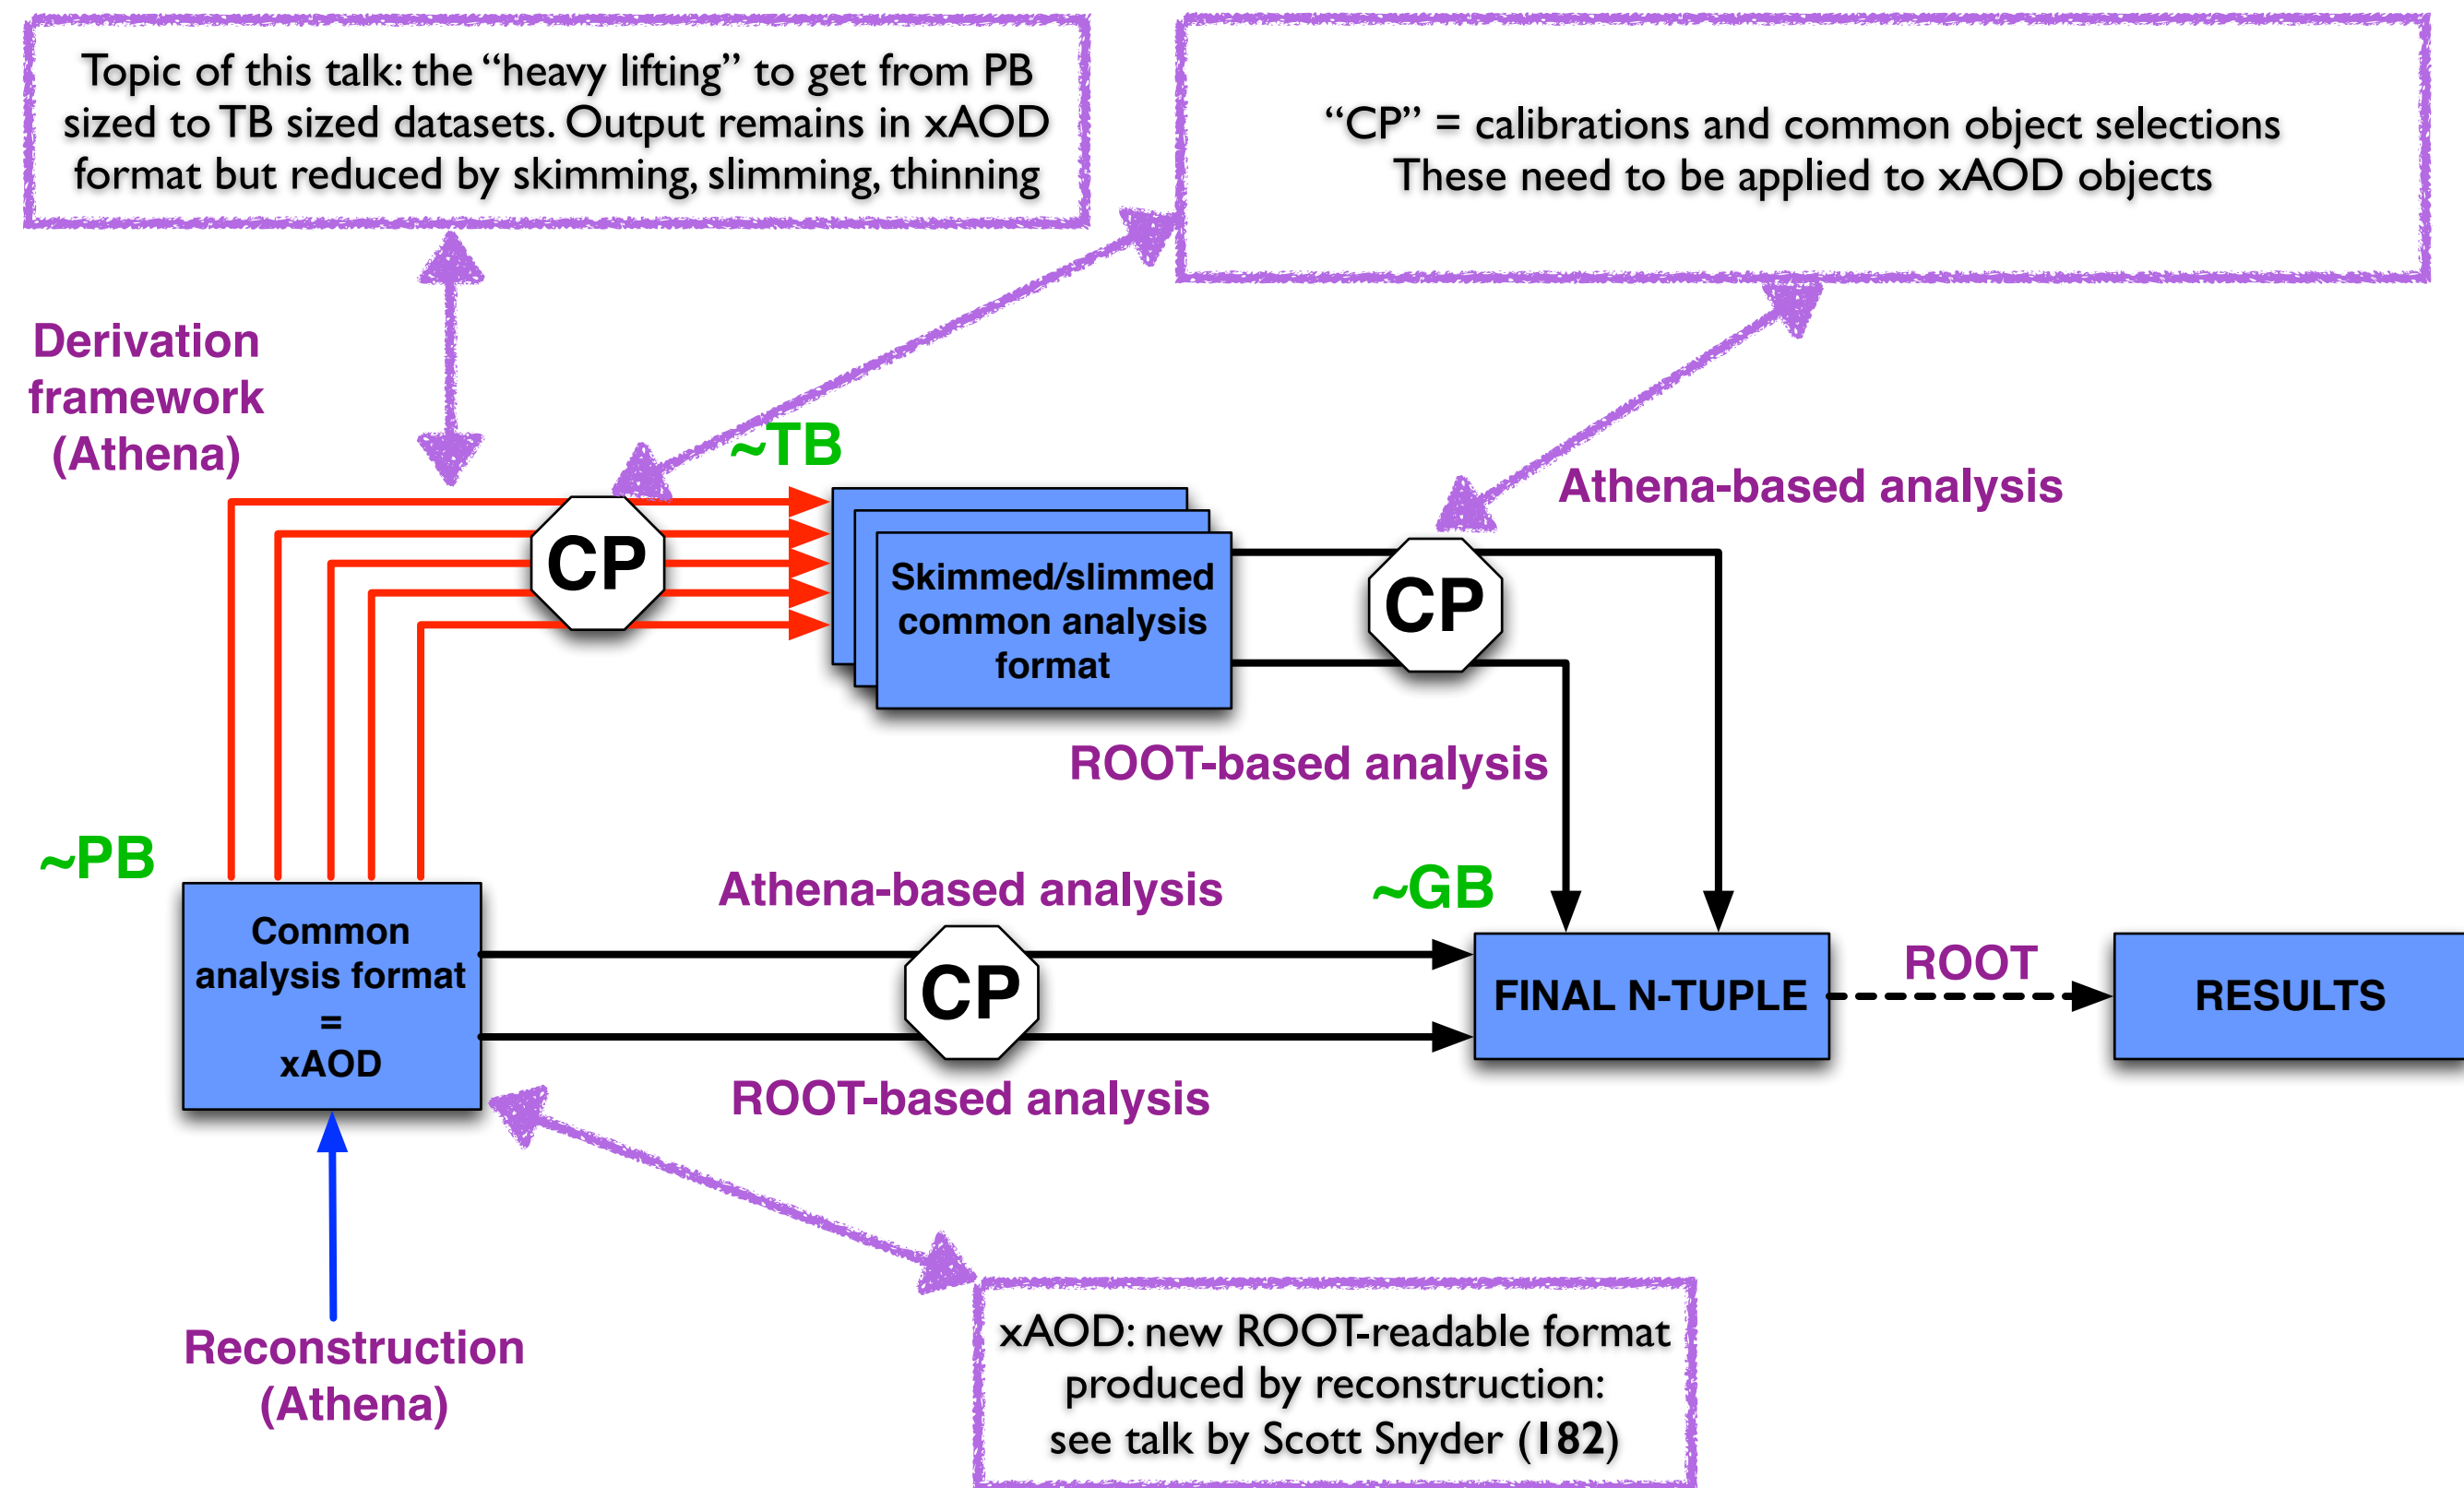
Topic of this talk: the “heavy lifting” to get from PB sized to TB sized datasets. Output remains in xAOD format but reduced by skimming, slimming, thinning

Derivation
framework
(Athena)



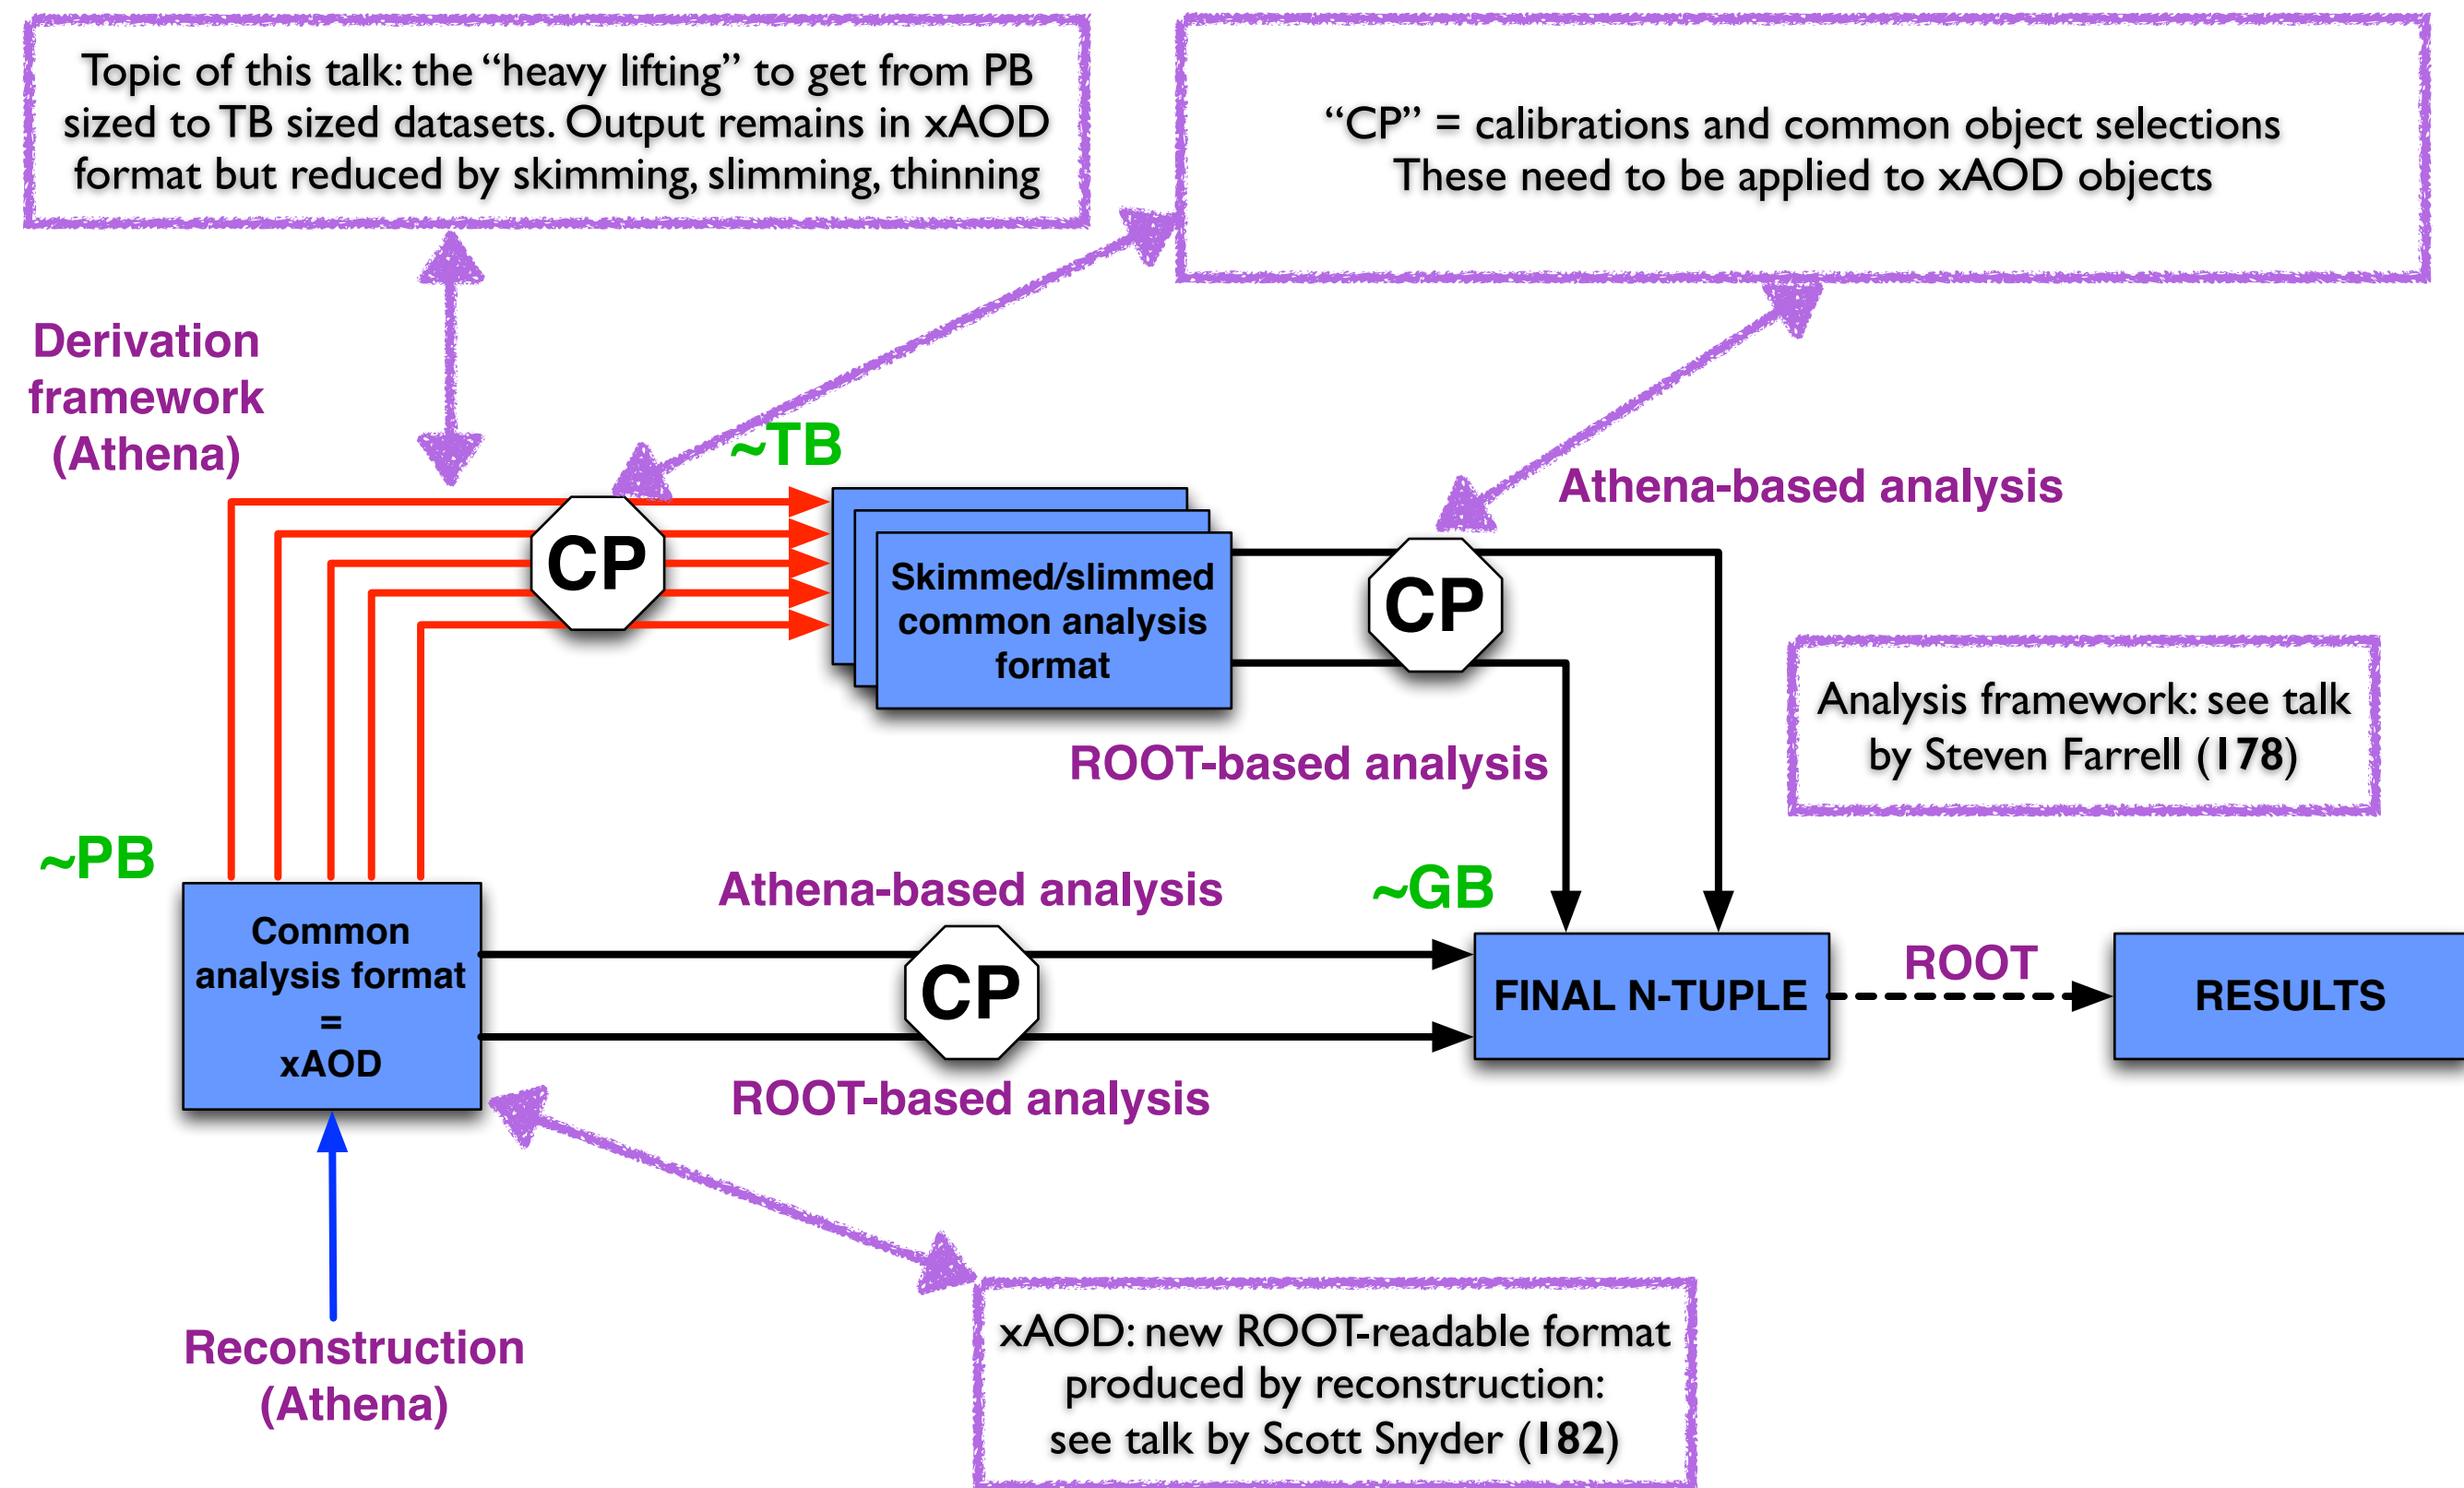
The Run-II analysis model for ATLAS

4

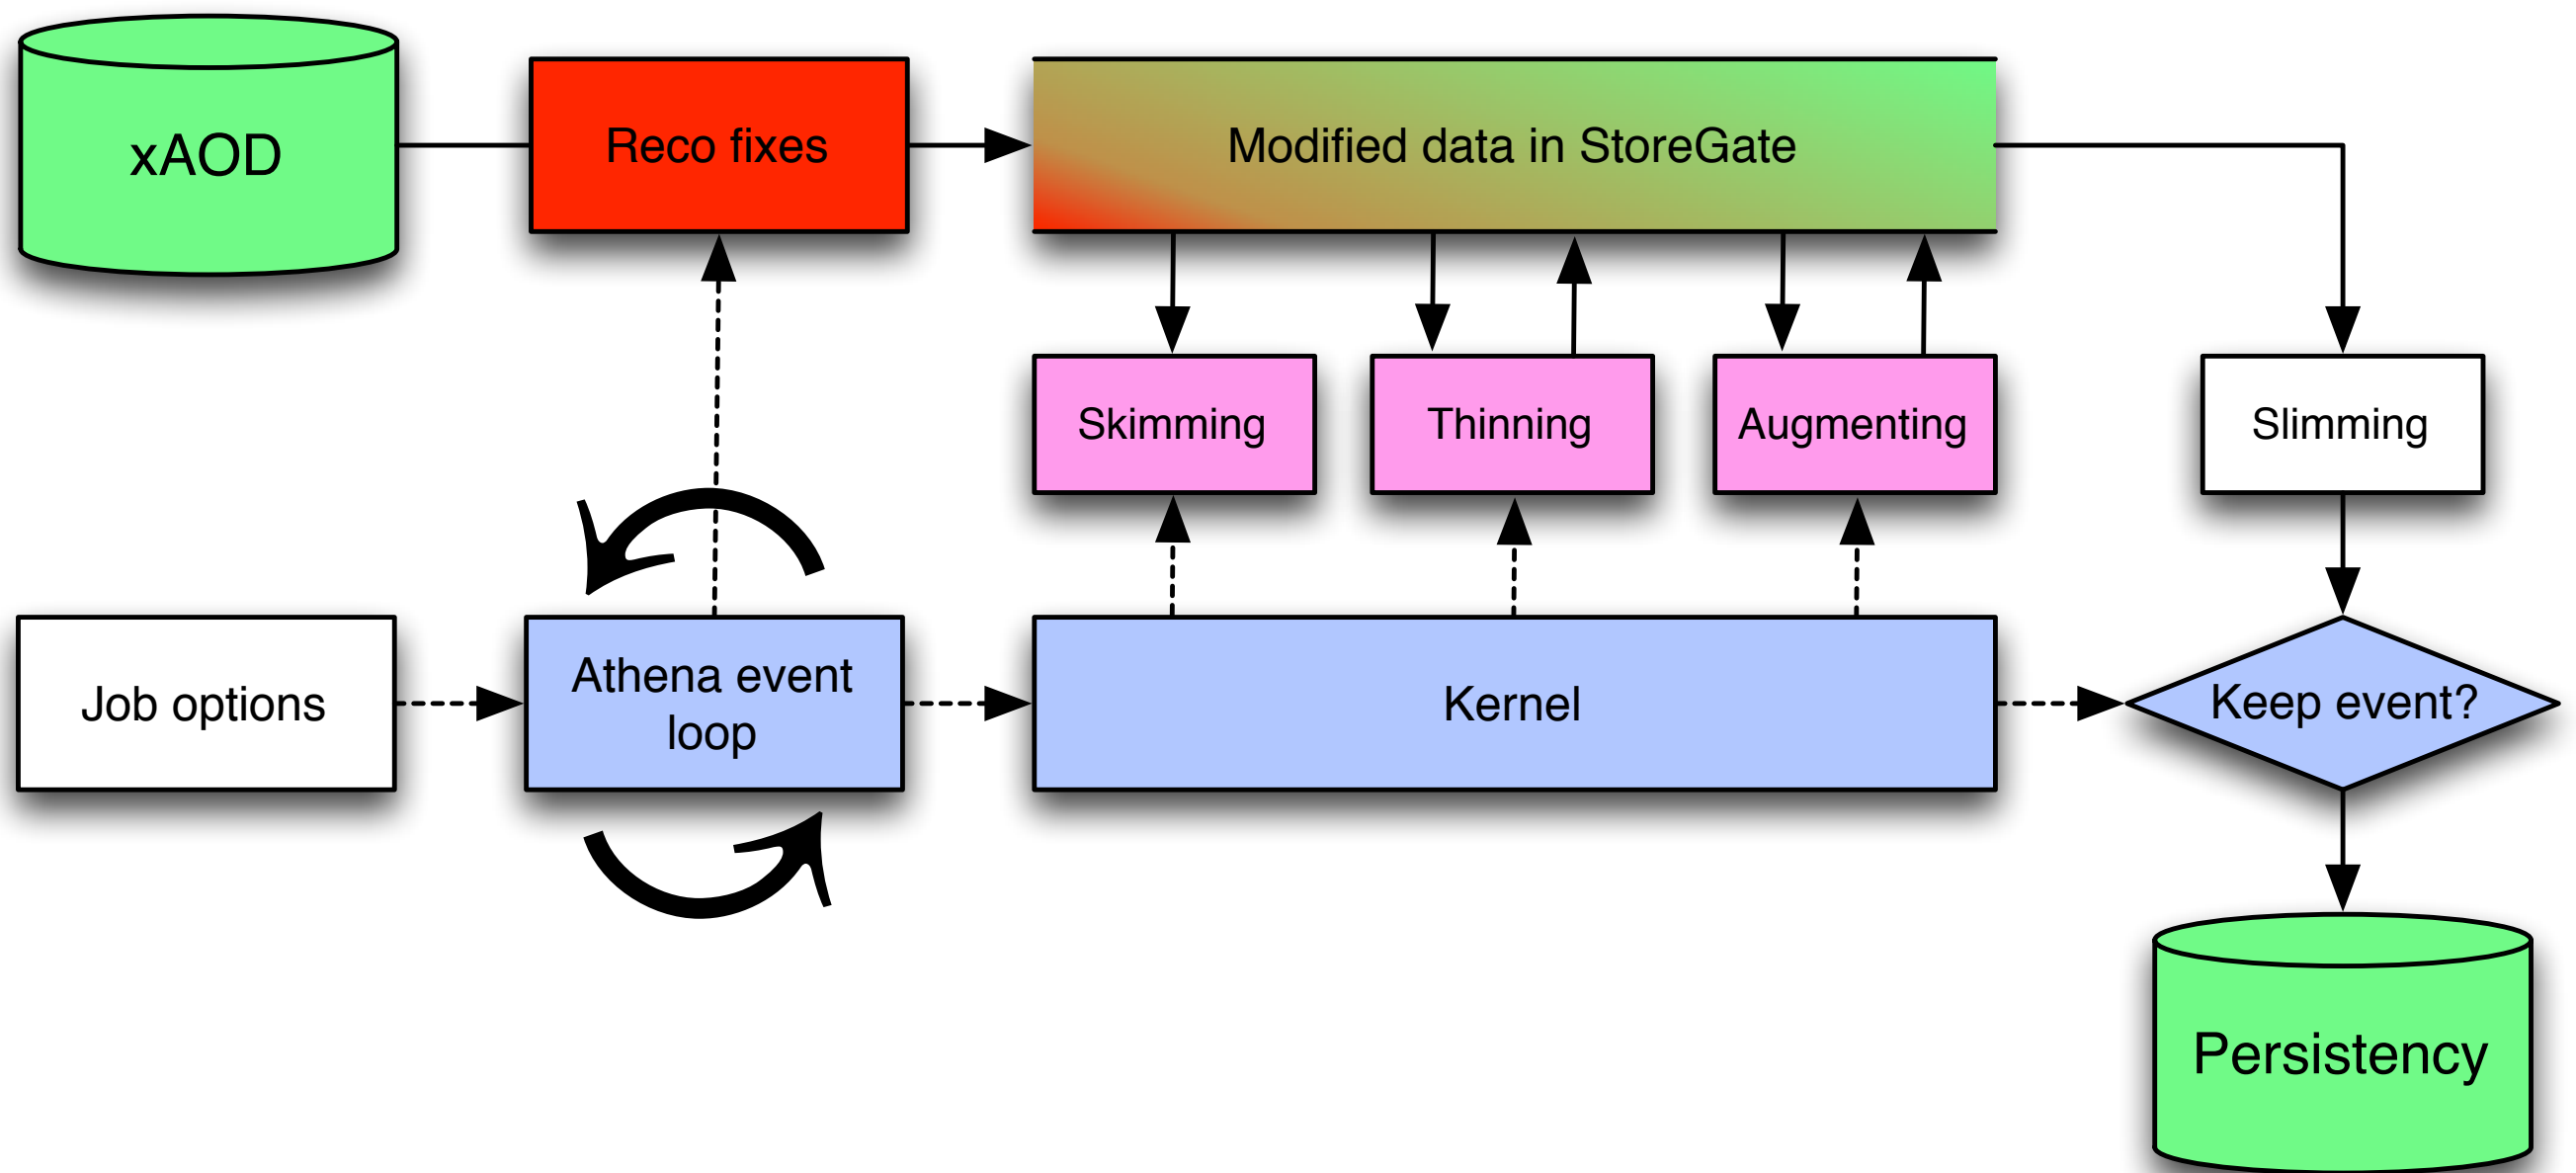


The Run-II analysis model for ATLAS

4

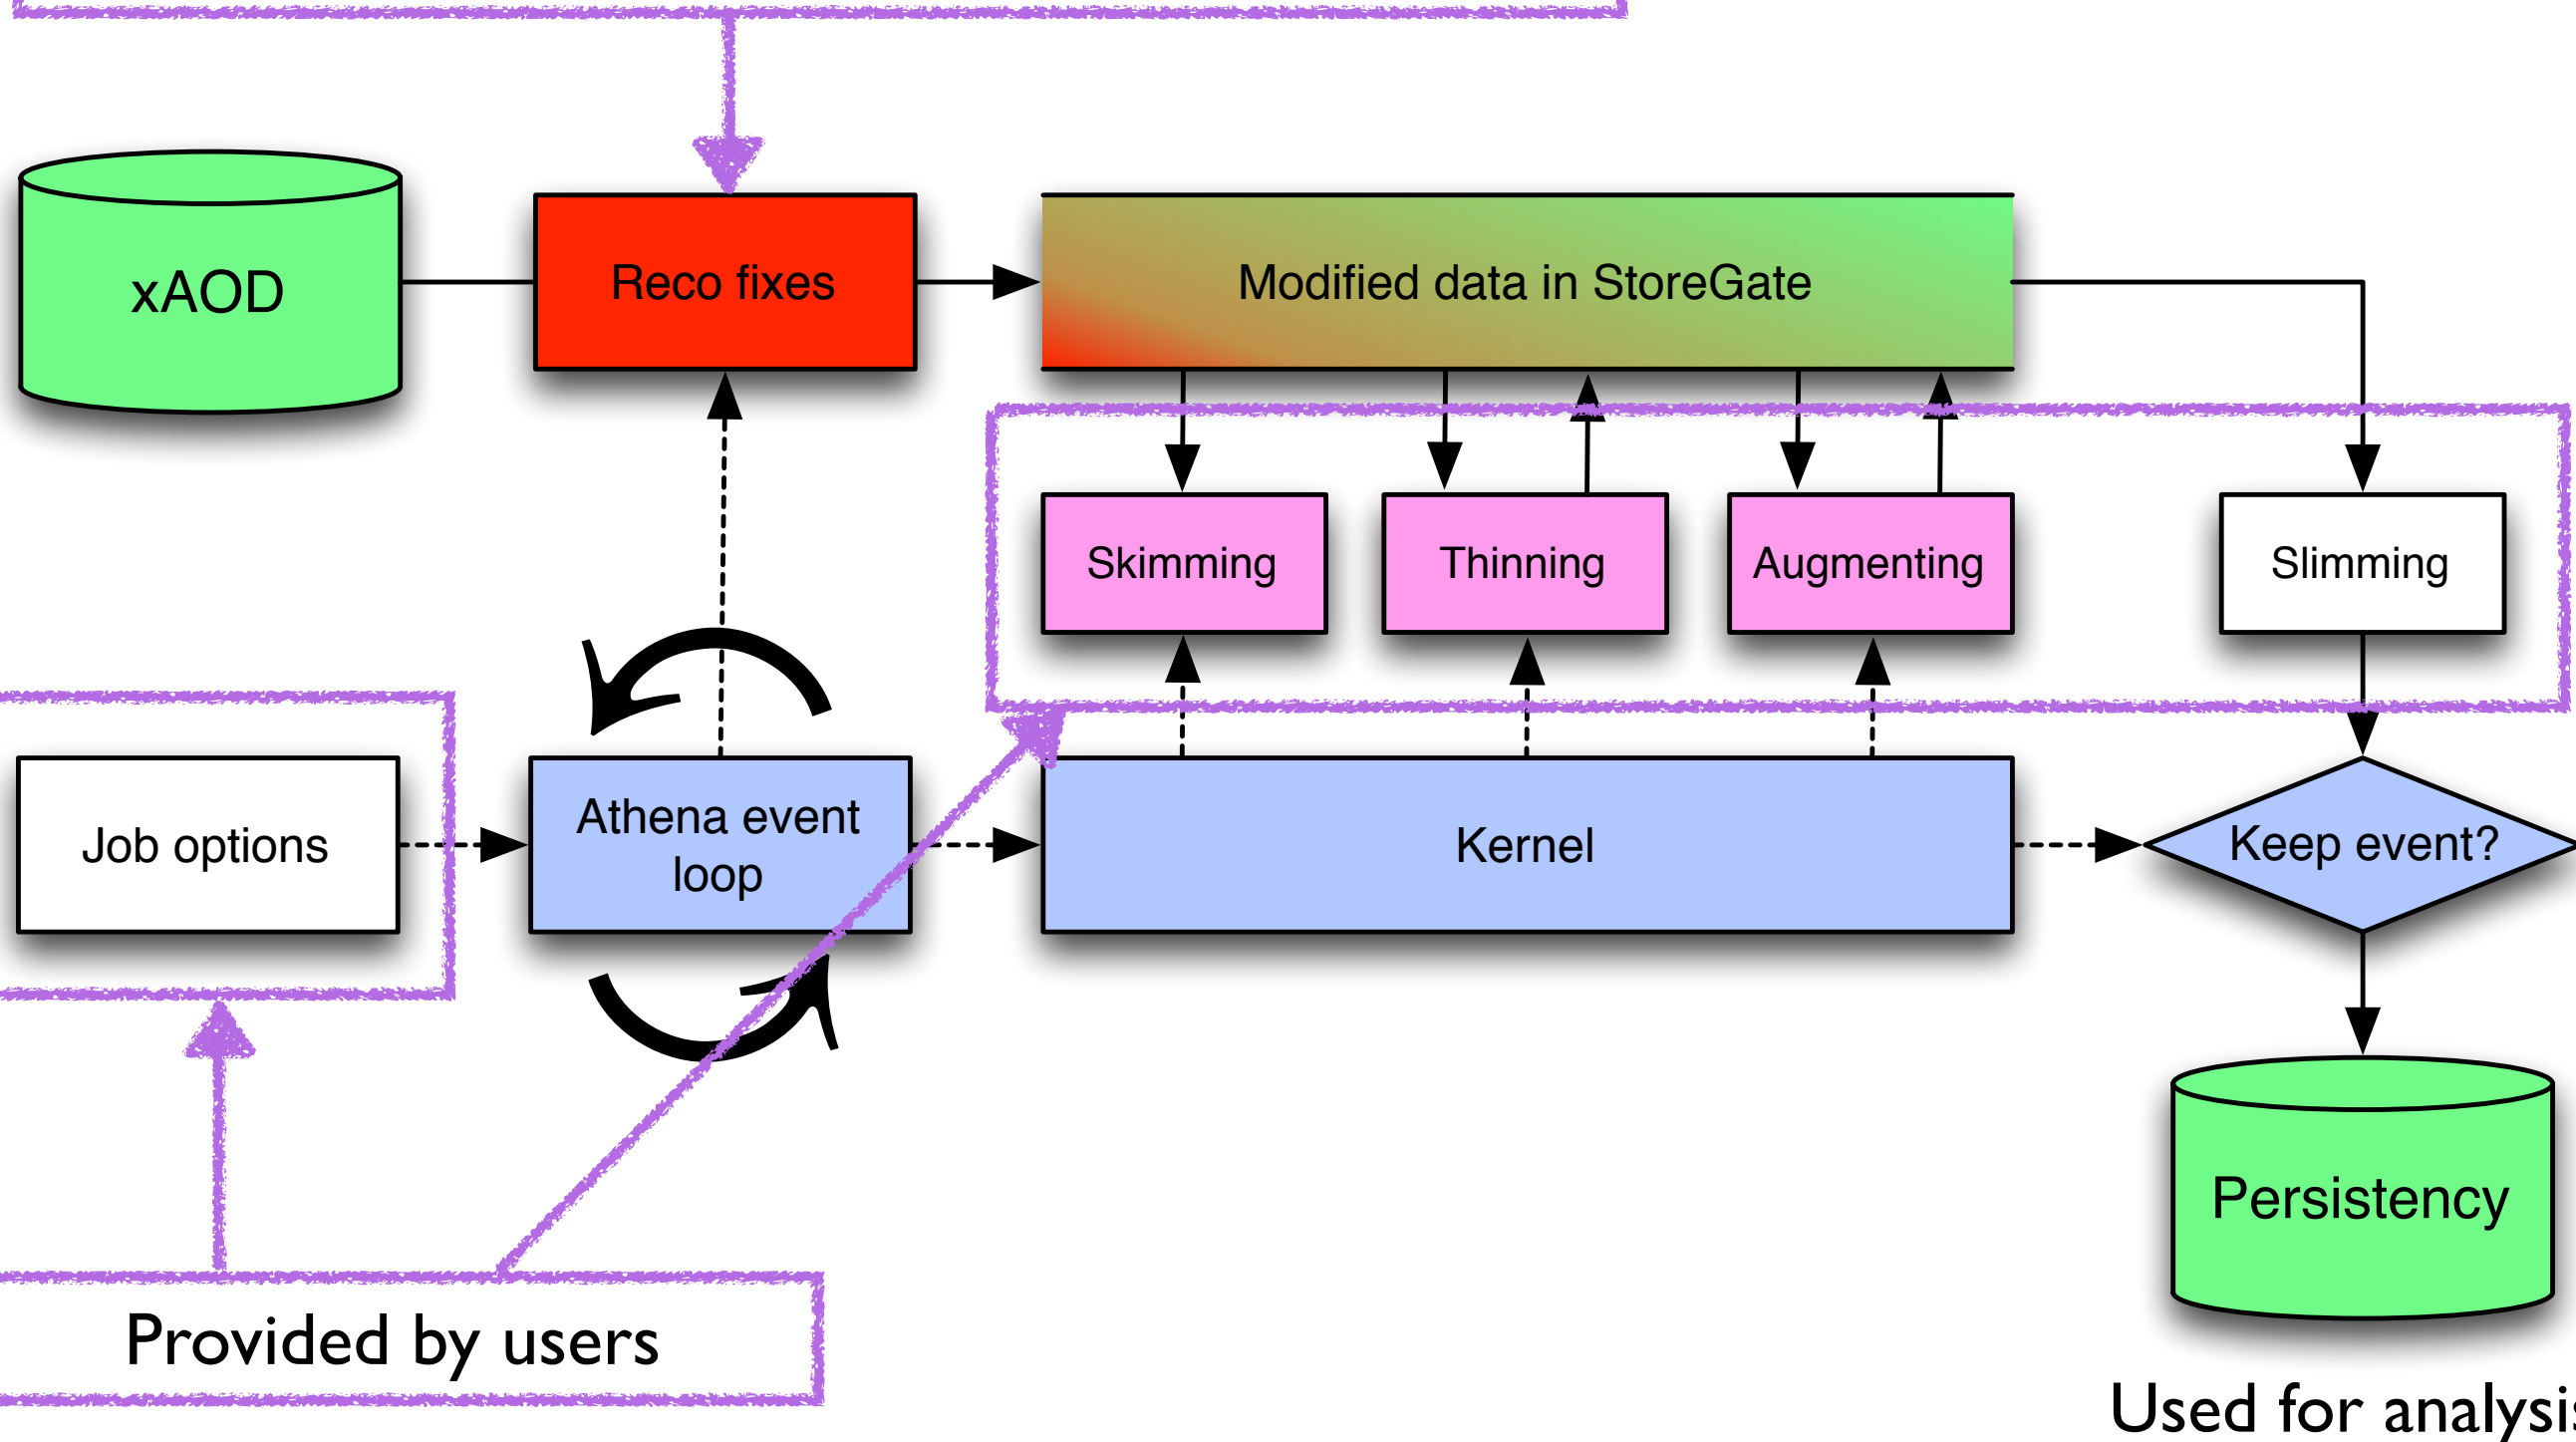


- Built on the main ATLAS data processing framework (Athena)
 - ▶ Enables re-running of parts of the reconstruction
 - ▶ Benefits from the Athena core software (algorithms and tools, whiteboard, thinning service, metadata handling, multiple outputs etc)
- Provides
 - ▶ interfaces for users to implement tools for skimming, thinning and augmenting their data
 - ... and a set of central tools for commonly needed selections
 - ▶ a text-based event/object selection mechanism to minimise user-developed C++
 - ▶ built-in lists of variables required for each calibration/object selection/systematic tool used in the analyses (“smart slimming”)
 - ▶ detailed monitoring of CPU, skimming rates, overlaps between formats



Used for analysis

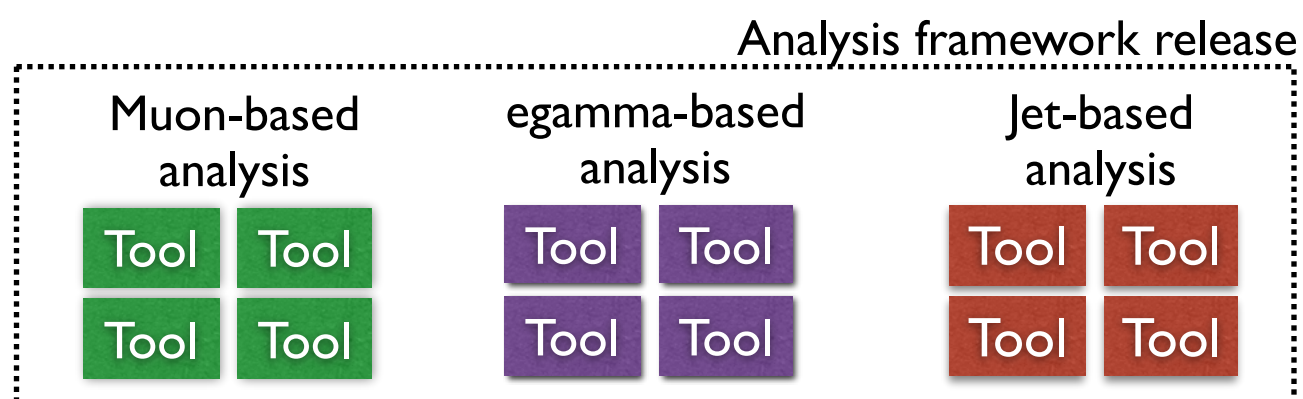
Allows corrections to be made to the reconstruction via “AODFix” (configured centrally)



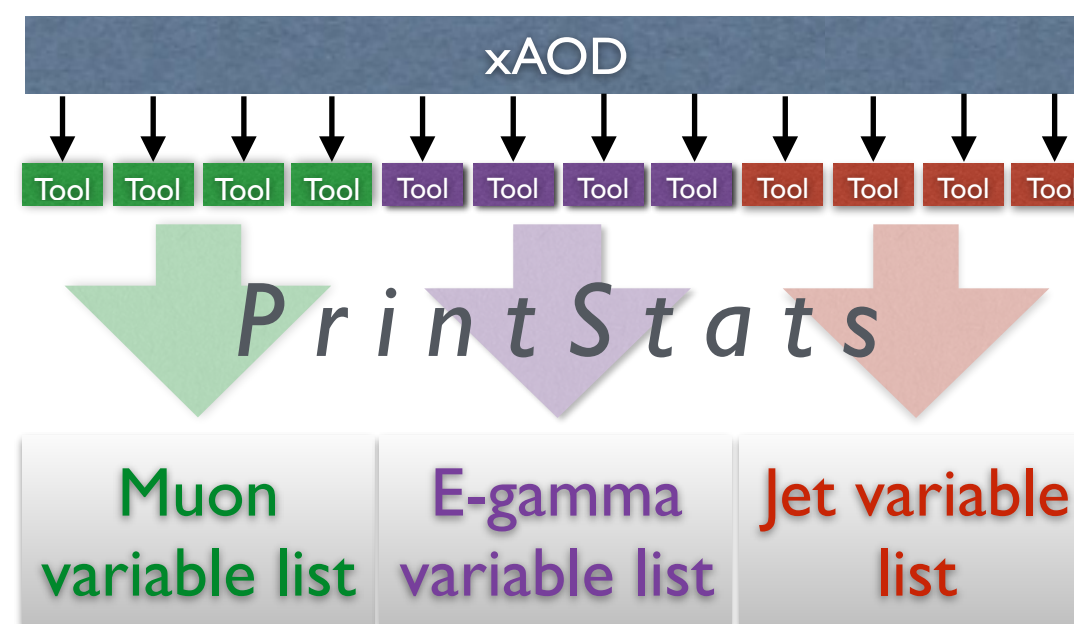
- To avoid large numbers of C++ tools being written the event and object selection is configured where possible from the Python job options alone
 - ▶ Execution in C++ is performed by a single “expression evaluation” tool
- Allows arbitrarily complex selections of the following type to be made:
 - ▶ Events (slimming):
 - `count (Muons.pt > 25.0*GeV && Muons.eta < 2.5) >= 4`
 - ▶ Objects (thinning):
 - `InDetTrackParticles.pt > 5.0*GeV`
- Can access any variables in StoreGate (either from the input file or added by other tools) - also supports operators, unary mathematical functions, constants
- Text parsing is done in the initialise step (once per job) minimising the CPU overhead

- Aim is to keep the variables needed for analysis and nothing more

1. We need to keep variables required by any of the tools in the analysis framework

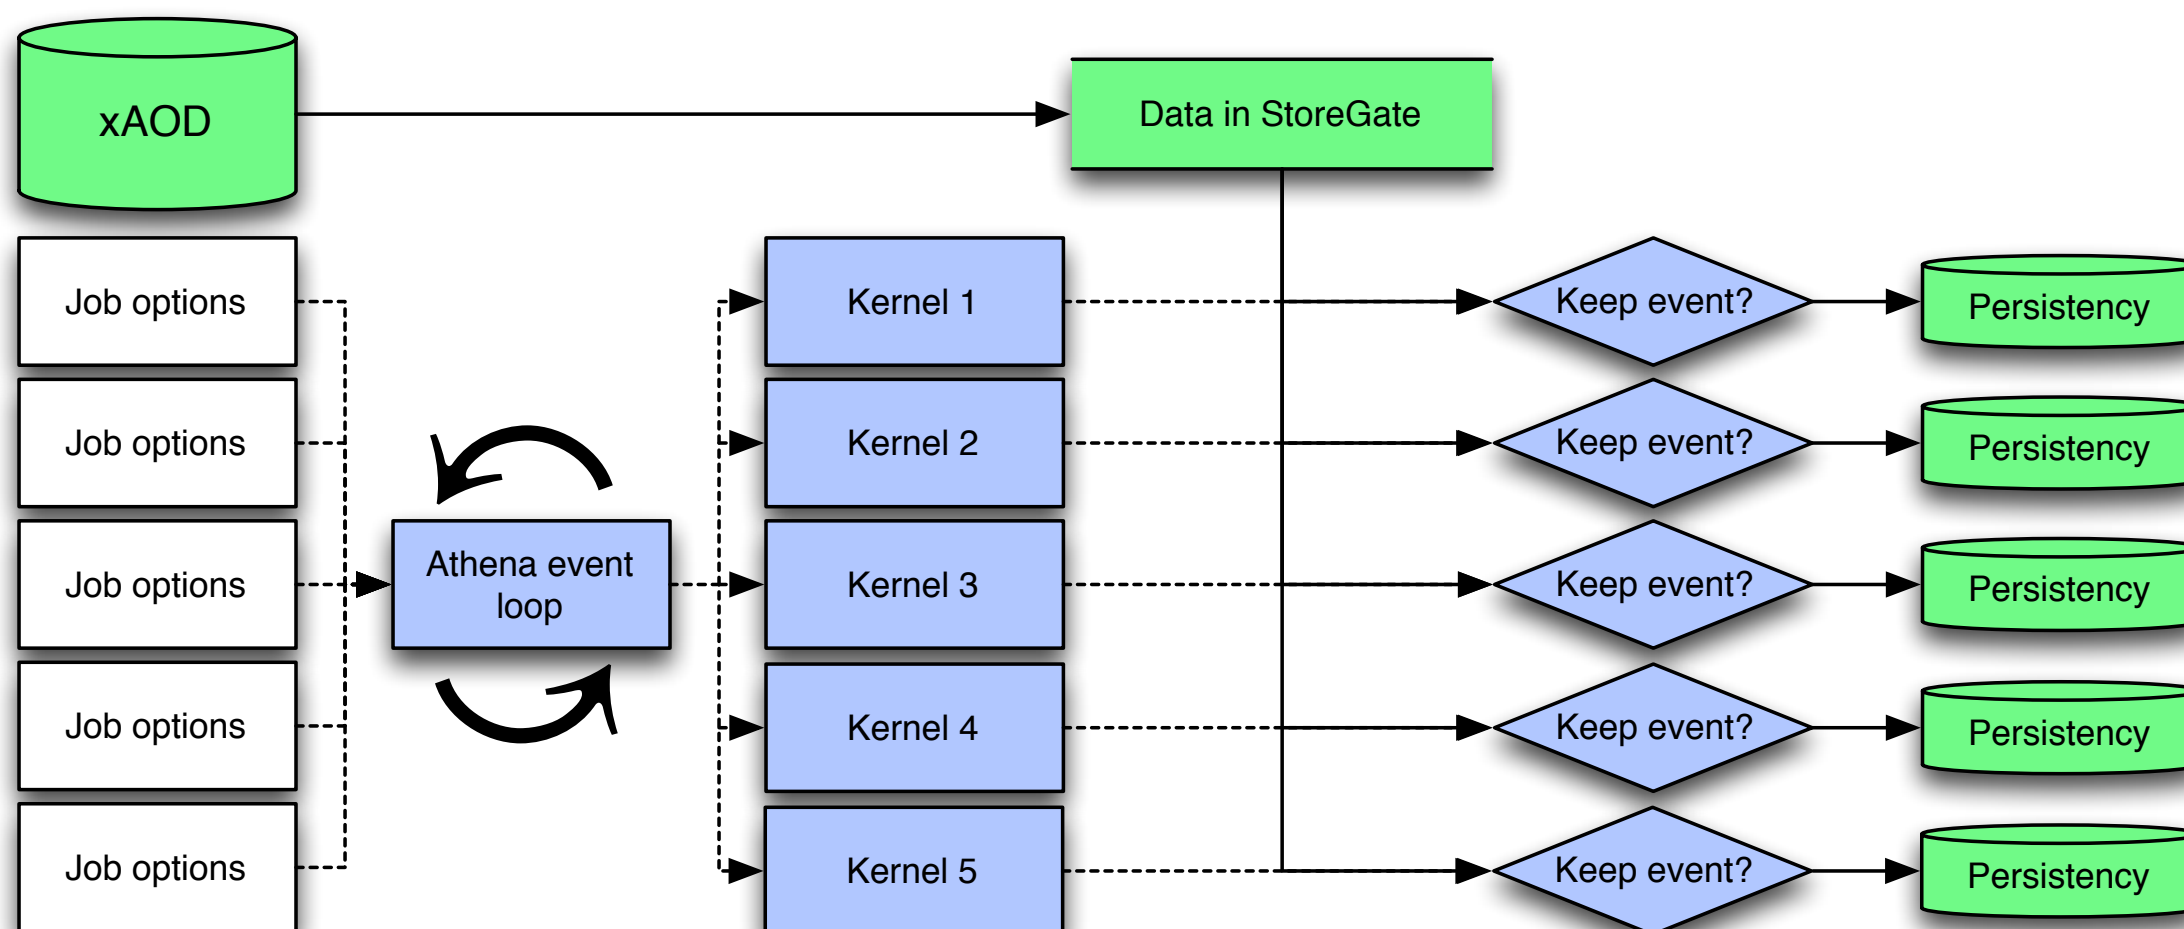


2. Run over an xAOD, calling each of the tools in turn. The PrintStats service generates lists of all accessed variables.




3. The variable lists are installed in the Derivation Framework release and automatically included. Extra variables can be added as needed.

- Multiple outputs per input file produced in a single job
 - Steered by the existing multiple stream manager in Athena
 - Each output stream is independent of the others
- This allows a “train model” of central production: several outputs per input
- Each train typically handles 5-10 output formats from a single physics group




- New information (augmentation) is typically done in two ways:
 - Adding new reconstructed object containers: typically jets made with a modified algorithm.
 - Decorating existing objects with extra variables: typically the results of object selection by combined performance tools (e.g. “this is a good muon”)
- See talk of Scott Snyder (182) for more information on the decoration mechanism
- Augmentation can be shared across a train, saving CPU



Implemented derivations

- No limit on the number of derivations: only on the total size
- It should be possible to analyse a derivation dataset on the grid with normal user privileges in approximately 1 day
- Budget:
 - ▶ total derivations size \leq total xAOD size
 - ▶ Each derivation should aim to be $\sim 1\%$ of its input xAOD size
 - ▶ Each physics/performance group should aim to write not more than 4% of the xAOD as derivations



Number of derivations	
E-gamma	5
Exotics	13
Flavour tagging	4
Higgs	14
Jet/missing energy	8
Muons	4
Standard Model	5
Supersymmetry	5
Taus	2
Top quarks	2

Overlap monitoring


12

- We do not want to write out the same selections in multiple formats

- ▶ If two formats strongly overlap, and they are large, we should merge them

- Event-wise overlaps can be monitored both in production using the EventIndex database and offline by Athena

- ▶ See talk on EventIndex by Dario Barberis (208)
- ▶ Content-wise overlaps can also be monitored




An event-wise overlap plot: enables us to determine whether a pair of formats might be so similar that they can be merged

- The Run-II analysis model for ATLAS includes the centralised production of analysis specific data formats “derivations” containing less information than the reconstruction output, but in the same format
- Of order 100 derivations are foreseen, each with a size of approximately 1% of the input
- A software framework, built on Athena, has been developed to produce these formats in bulk
- The framework is in use already and the output data products are within the resource limits
- It will be deployed from the start of data taking next month and will reduce the computing workload on individual physicists, who should not have to run large private productions any more

Supplementary slides

- Thinning: object removal
 - Performed by a pre-existing service within Athena
- Sophistication is not in the removal of the objects themselves but in the re-setting of *ElementLinks* between objects
 - *ElementLink* = persistent representation of pointer links between objects = an index number
- Example: consider links between inner detector tracks and muons:

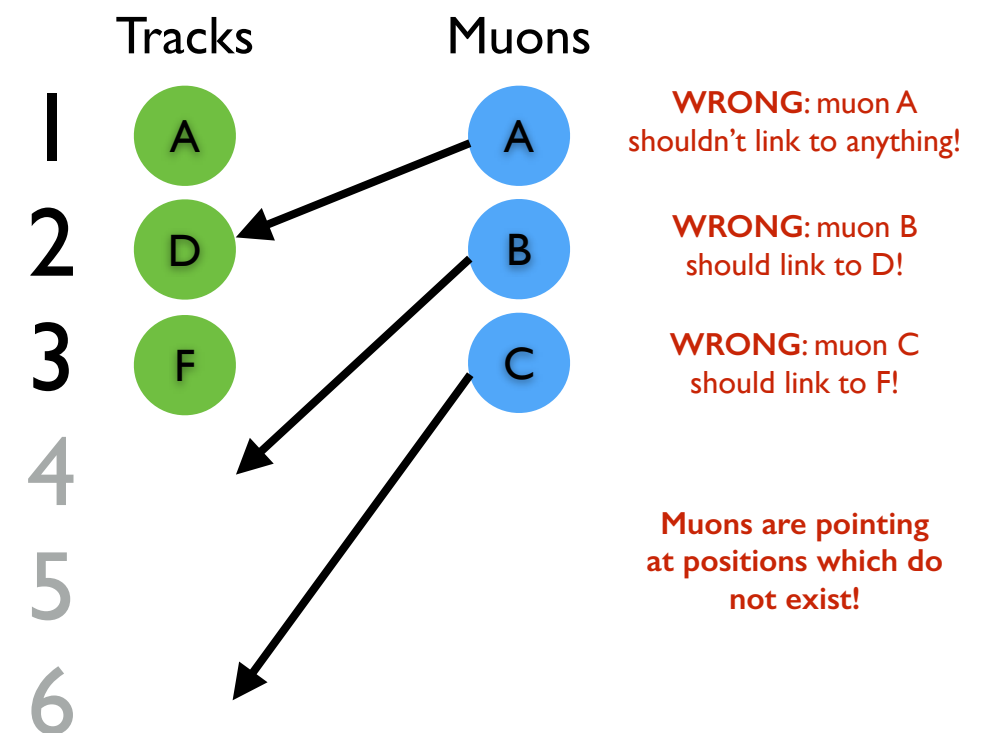
Before thinning



Muon A links to 2nd track in the container
 Muon B links to 4th track in the container
 Muon C links to 6th track in the container

Tracks B, C and E
 naively removed
 by user


After thinning



Muon A links to 2nd track in the container
 Muon B links to 4th track in the container
 Muon C links to 6th track in the container

- Thinning: object removal
 - Performed by a pre-existing service within Athena
- Sophistication is not in the removal of the objects themselves but in the re-setting of *ElementLinks* between objects
 - *ElementLink* = persistent representation of pointer links between objects = an index number
- Example: consider links between inner detector tracks and muons:


Before thinning



Muon A links to 2nd track in the container
Muon B links to 4th track in the container
Muon C links to 6th track in the container

Tracks B, C and E
removed by
ThinningSvc

After thinning



Muon A links to nothing
Muon B links to 2nd track in the container
Muon C links to 3rd track in the container

- Design: text parsing done by the Boost::spirit library
- “Compiled” at initialise method into a “virtual machine”: time consuming but only done once per job
- during each event the only action necessary is
 - ▶ the loading of the numerical quantities into the virtual machine
 - ▶ the actual decision
 - ▶ this is fast

