Data analytics approaches - HEP

Peter Love
2nd November 2016
GridPP/SKA workshop, Manchester

Acknowledgements

- James Walder (Lancaster) - Tutorial material
- James Catmore (Oslo) - Analysis model
- Cedric Serfon (Oslo) - Data management
- Federica Legger (LMU) - Analysis operations
- Hass AbouZeid (UC Santa Cruz) - Data preparation
- Louise Heelan (UT Arlington) - Analysis model
- Nils Krumnack (Iowa State University) - Analysis with ROOT
- Lukas Heinrich (NY University) - Trigger
- Simone Campana (CERN) - Atlas Computing

Overview

Part 1. ATLAS analysis model

- Data reduction from detector to analysis
- ATLAS analysis model during Run 2

Part 2. Analytics overview from CHEP

- Logs, Ops, Physics

The Run 2 Analysis Model

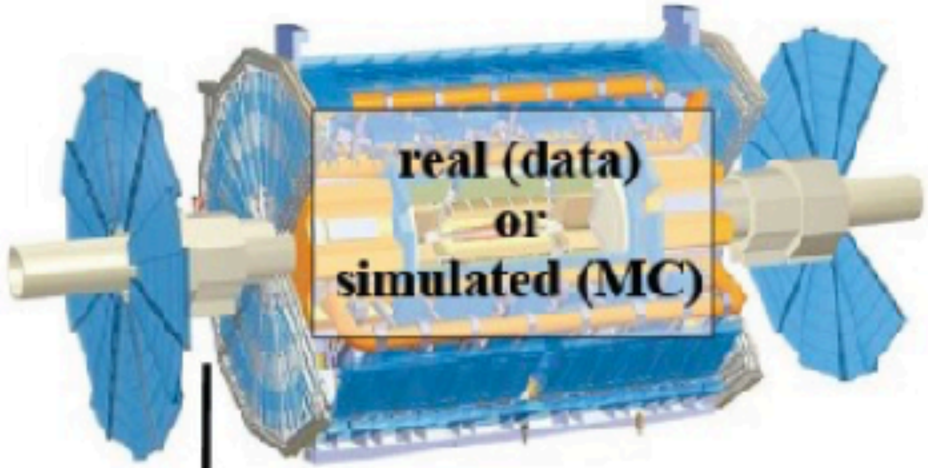

The xAOD format:

- the format is readable by both Athena and ROOT
- it is fast and compact
- it is uniform across all reconstructed object types (electrons, muons, jets, etc.)
- more 'object' oriented

RAW/RDO

Reconstruction (Athena)

xAOD

Derivation reduction framework (Athena)

derived xAOD

Athena or ROOT analysis (dual-use tools)

final analysis ntuple

Will be described in more detail by Attila after lunch

NEW!

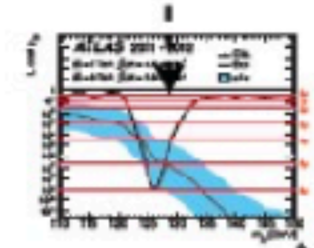

This talk will concentrate on the path from RAW detector data to AOD.

Short Recap Of The Trigger in ATLAS

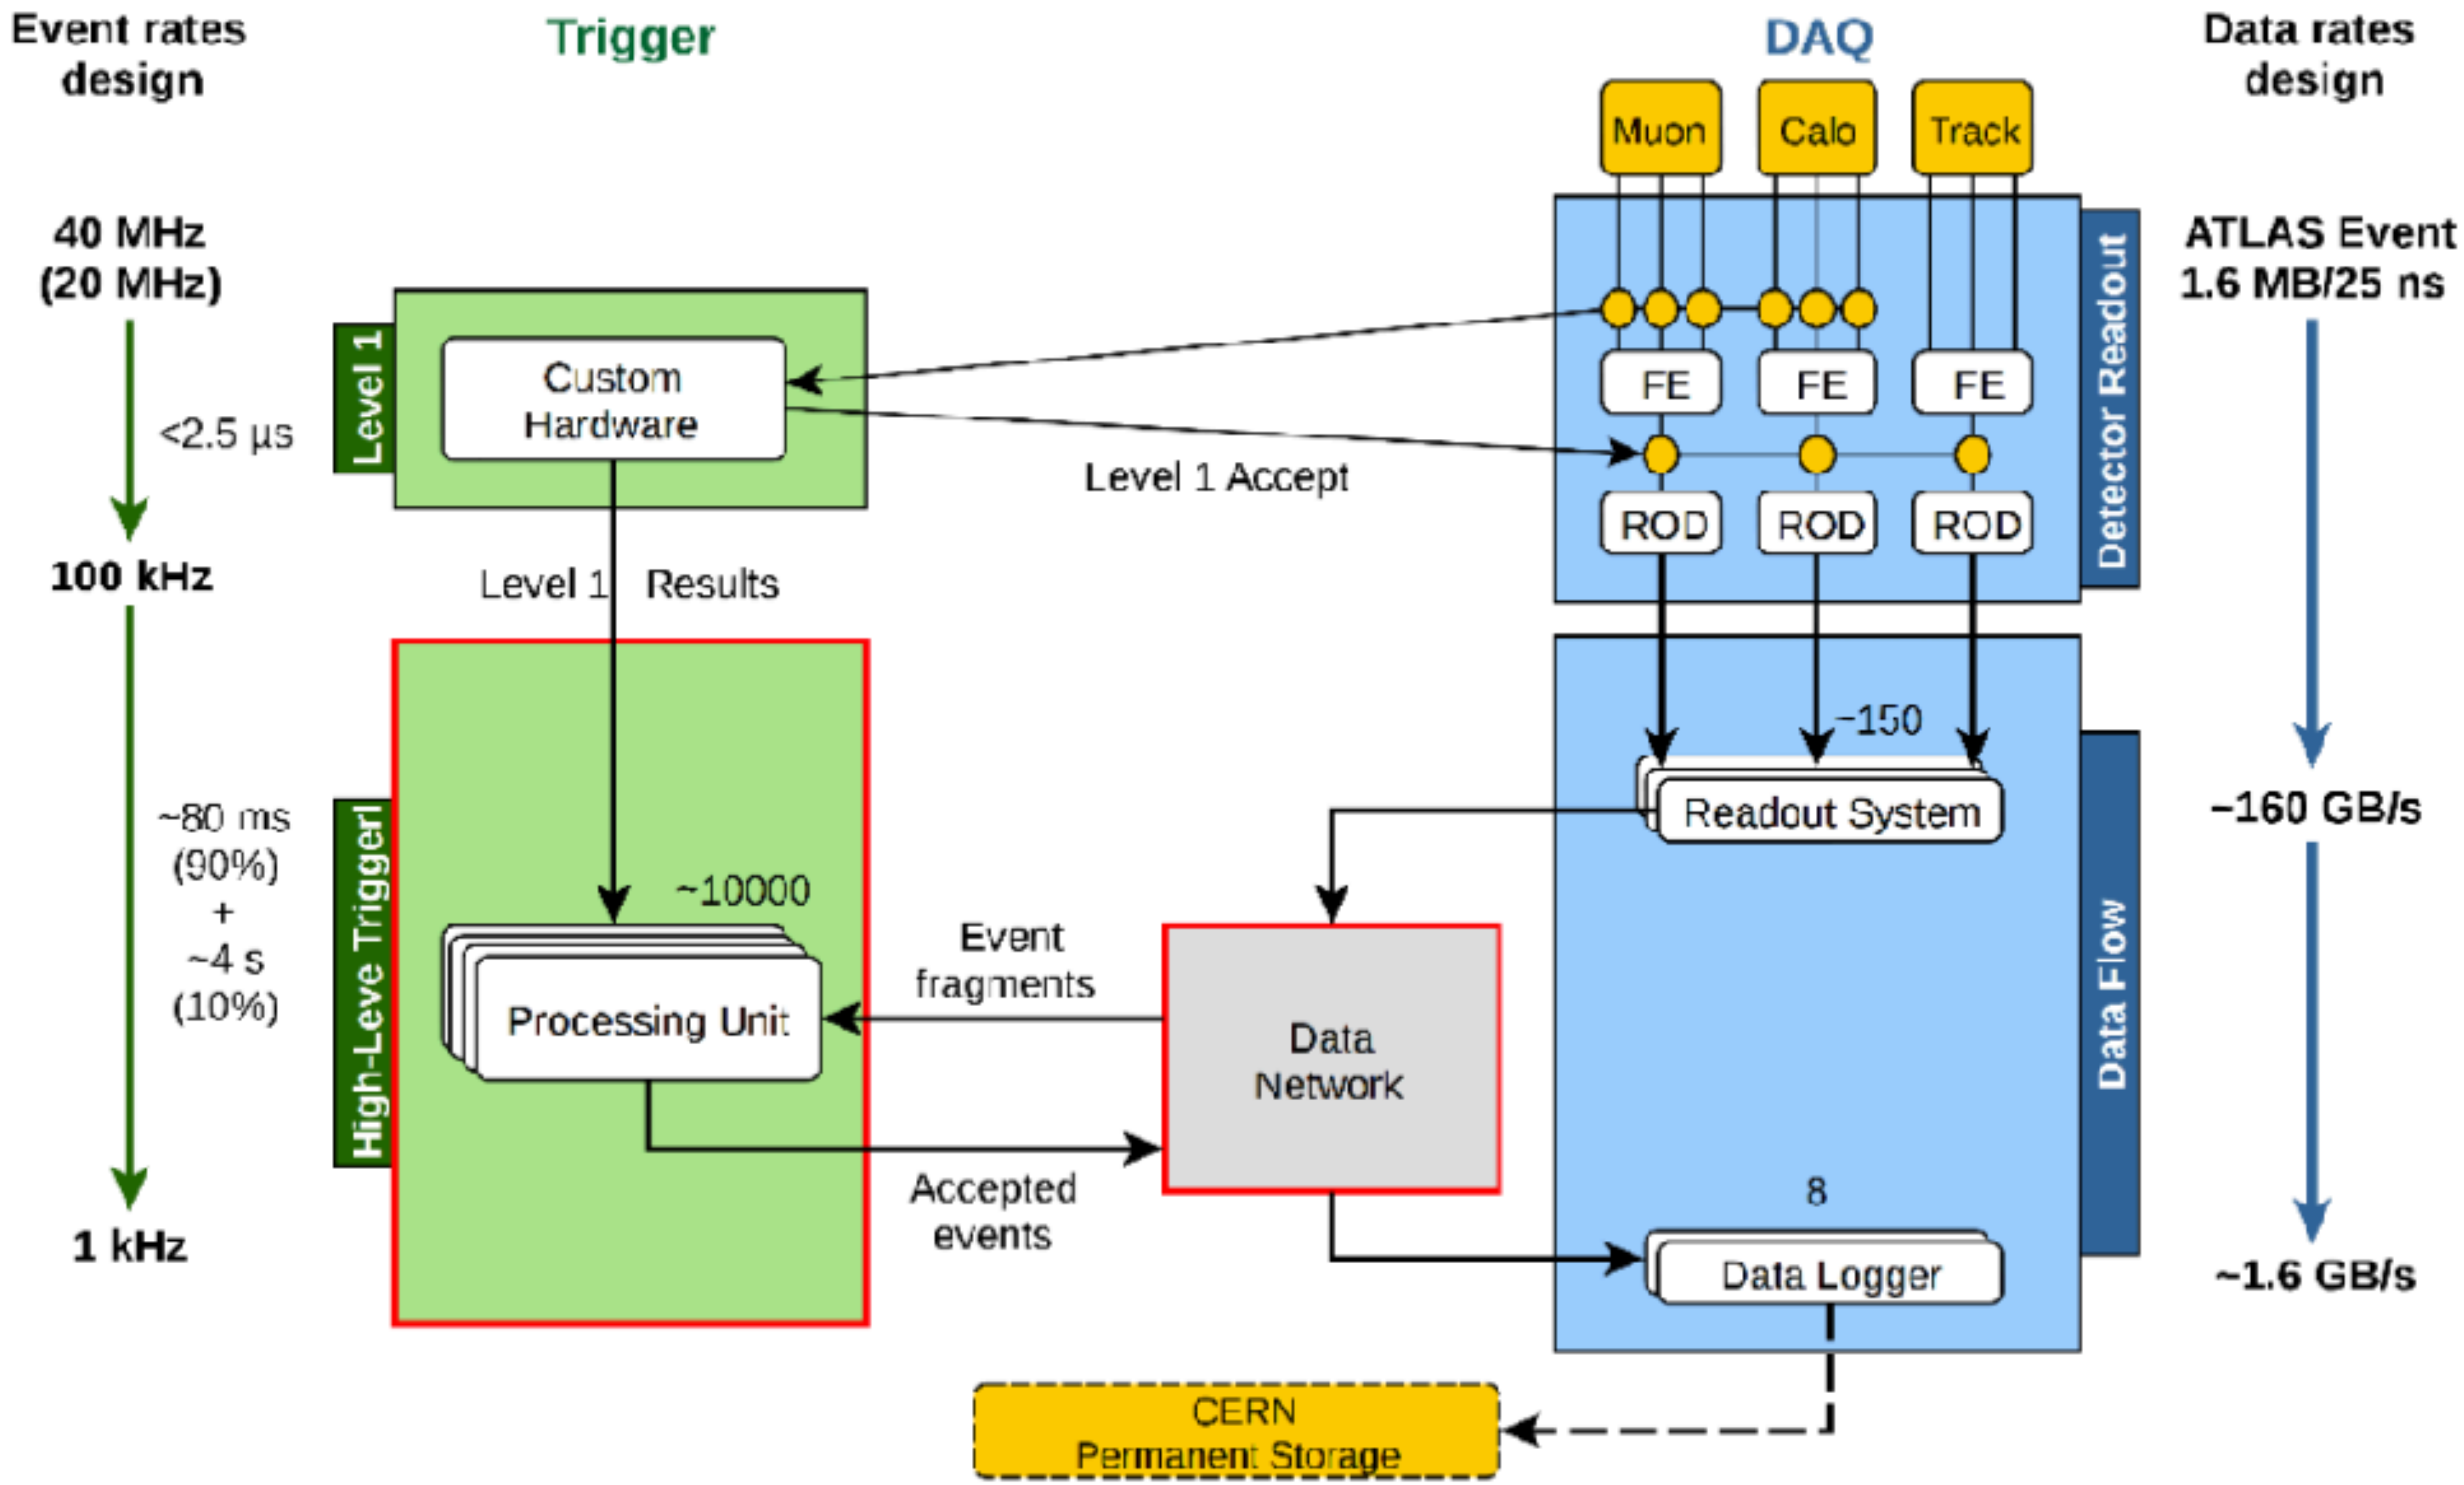

Slide: Lukas Heinrich

Constant pressure on resources

- Limited cash for new resources
- Datasets are growing with time, and increasing rates. Driven by physics ambitions.
- Actual events are more complex
- Hardware and architectures continue to evolve
  - increased core counts
  - constraints on memory bandwidth
  - new architectures (Arm, co-processors)
  - new storage architectures

Pile-up profile

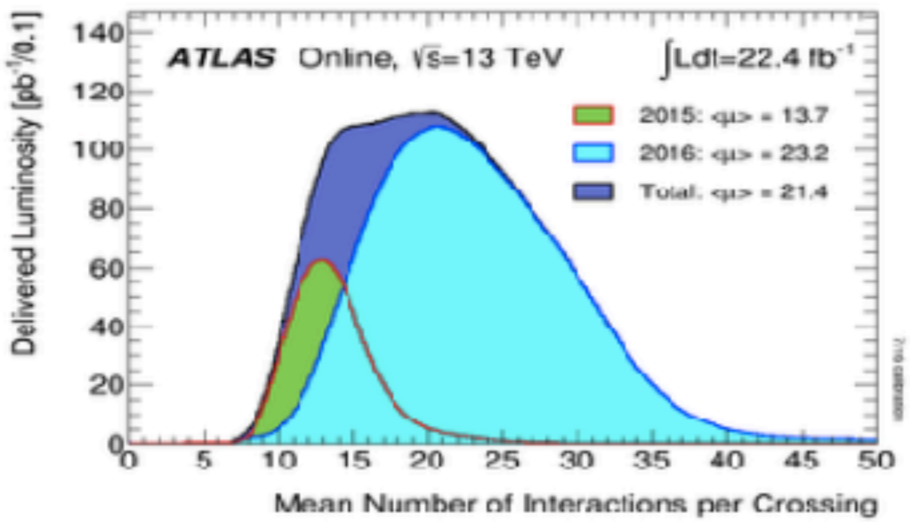

Time in Stable Beams (2016)

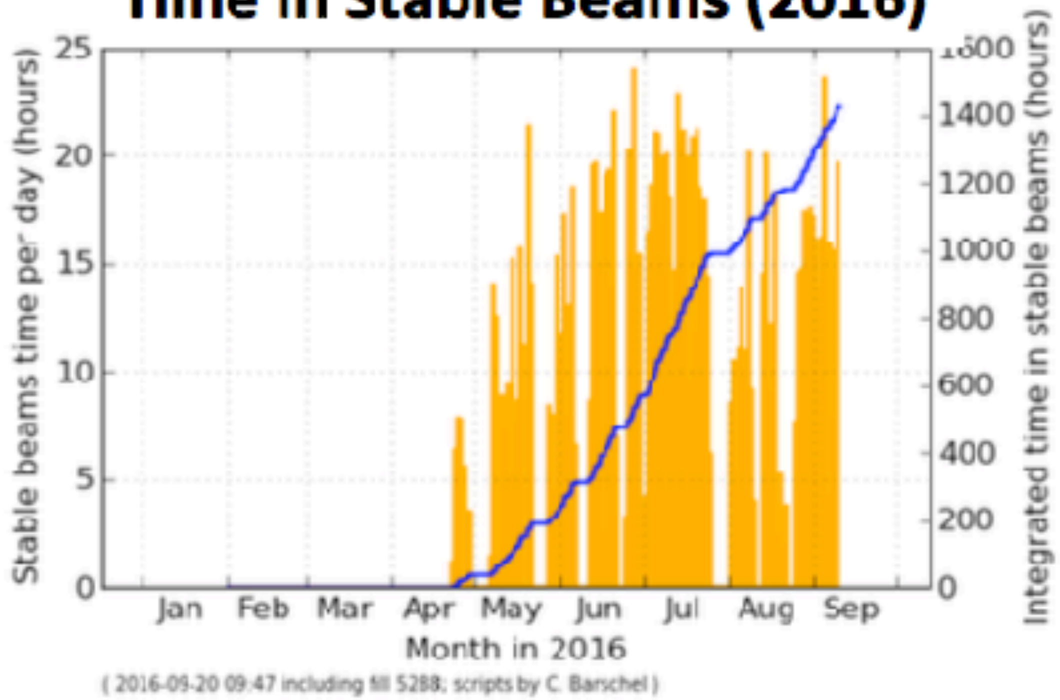

Integrated Luminosity

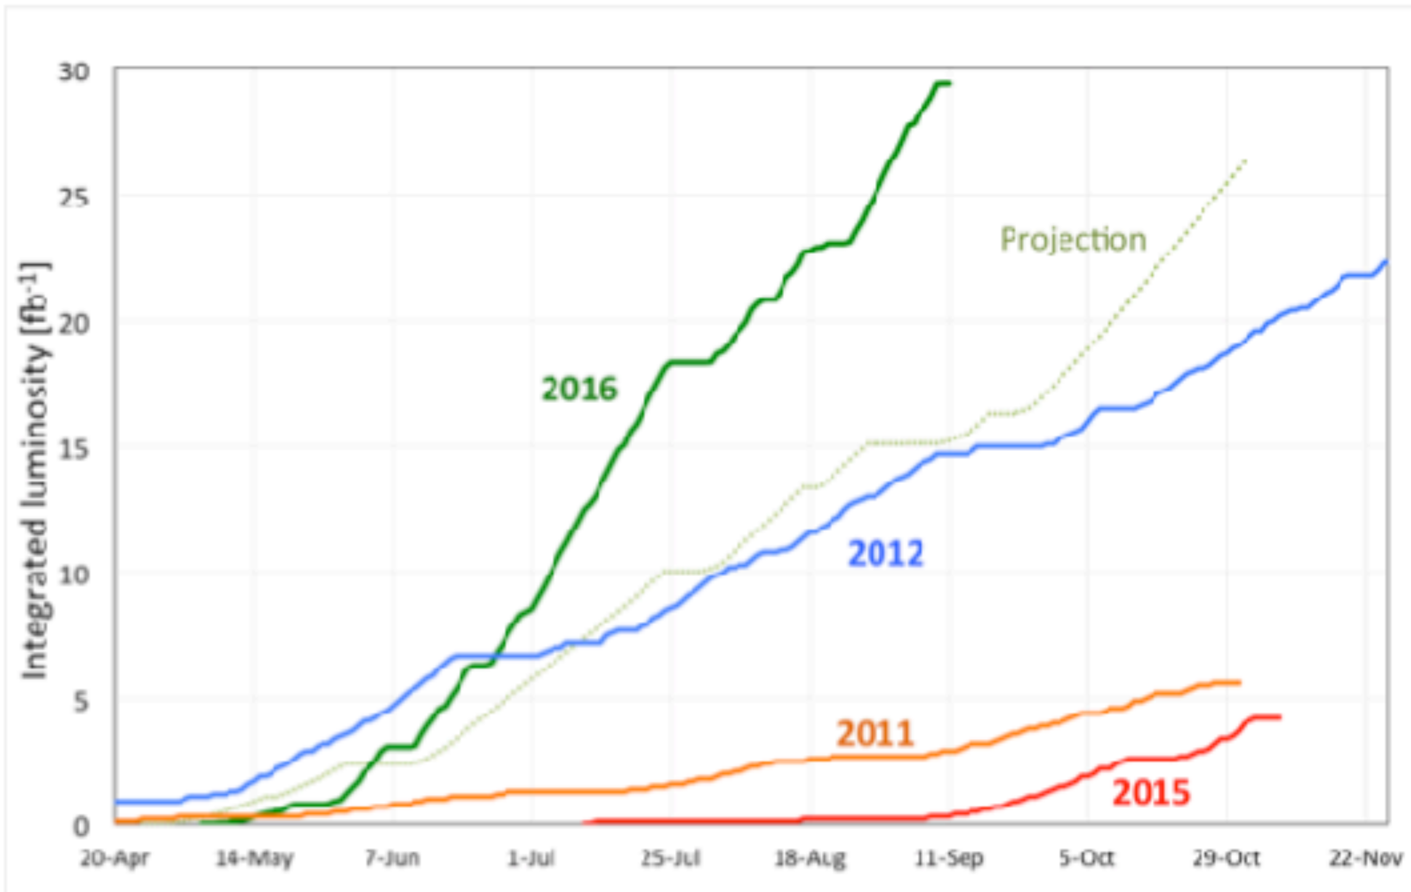

ATLAS analysis model for Run 2

- Derivation framework implemented to alleviate these pressures
- Tighter control on number data formats, number of replicas, and more aggressive deletion policies.

A feature common to most physics analysis

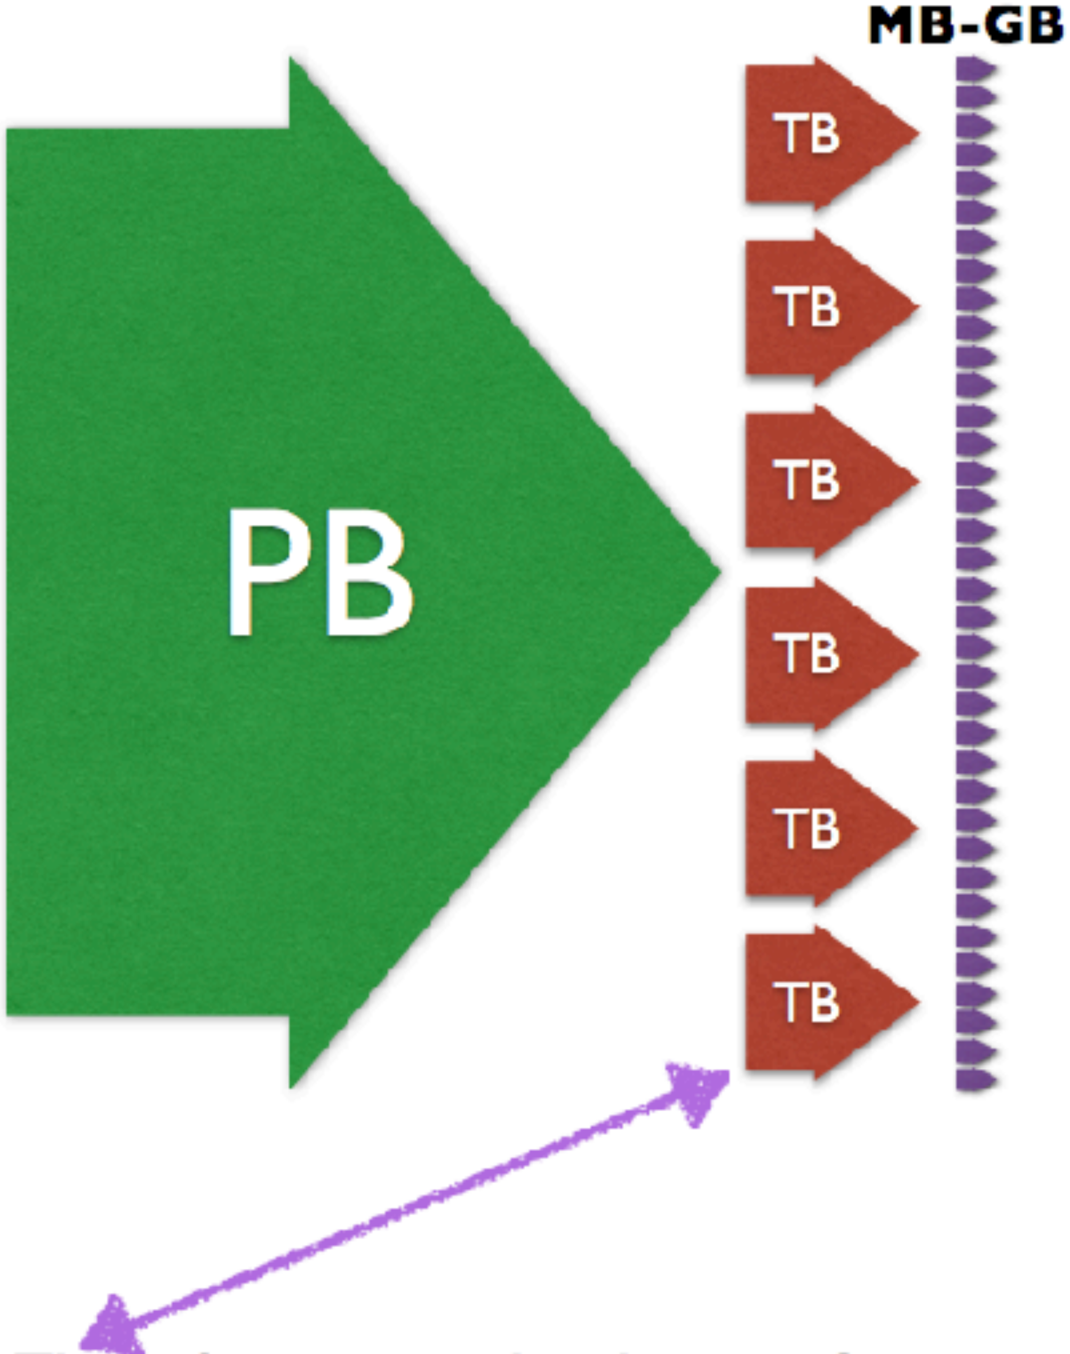

	Full output of reconstruction, ~PB size	One format
	Intermediate analysis format ~TB size	~100 formats
	Final n-tuple ~MB-GB size	~1000 formats

- 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 produce them centrally

Slide: James Catmore

The Run-II analysis model for ATLAS

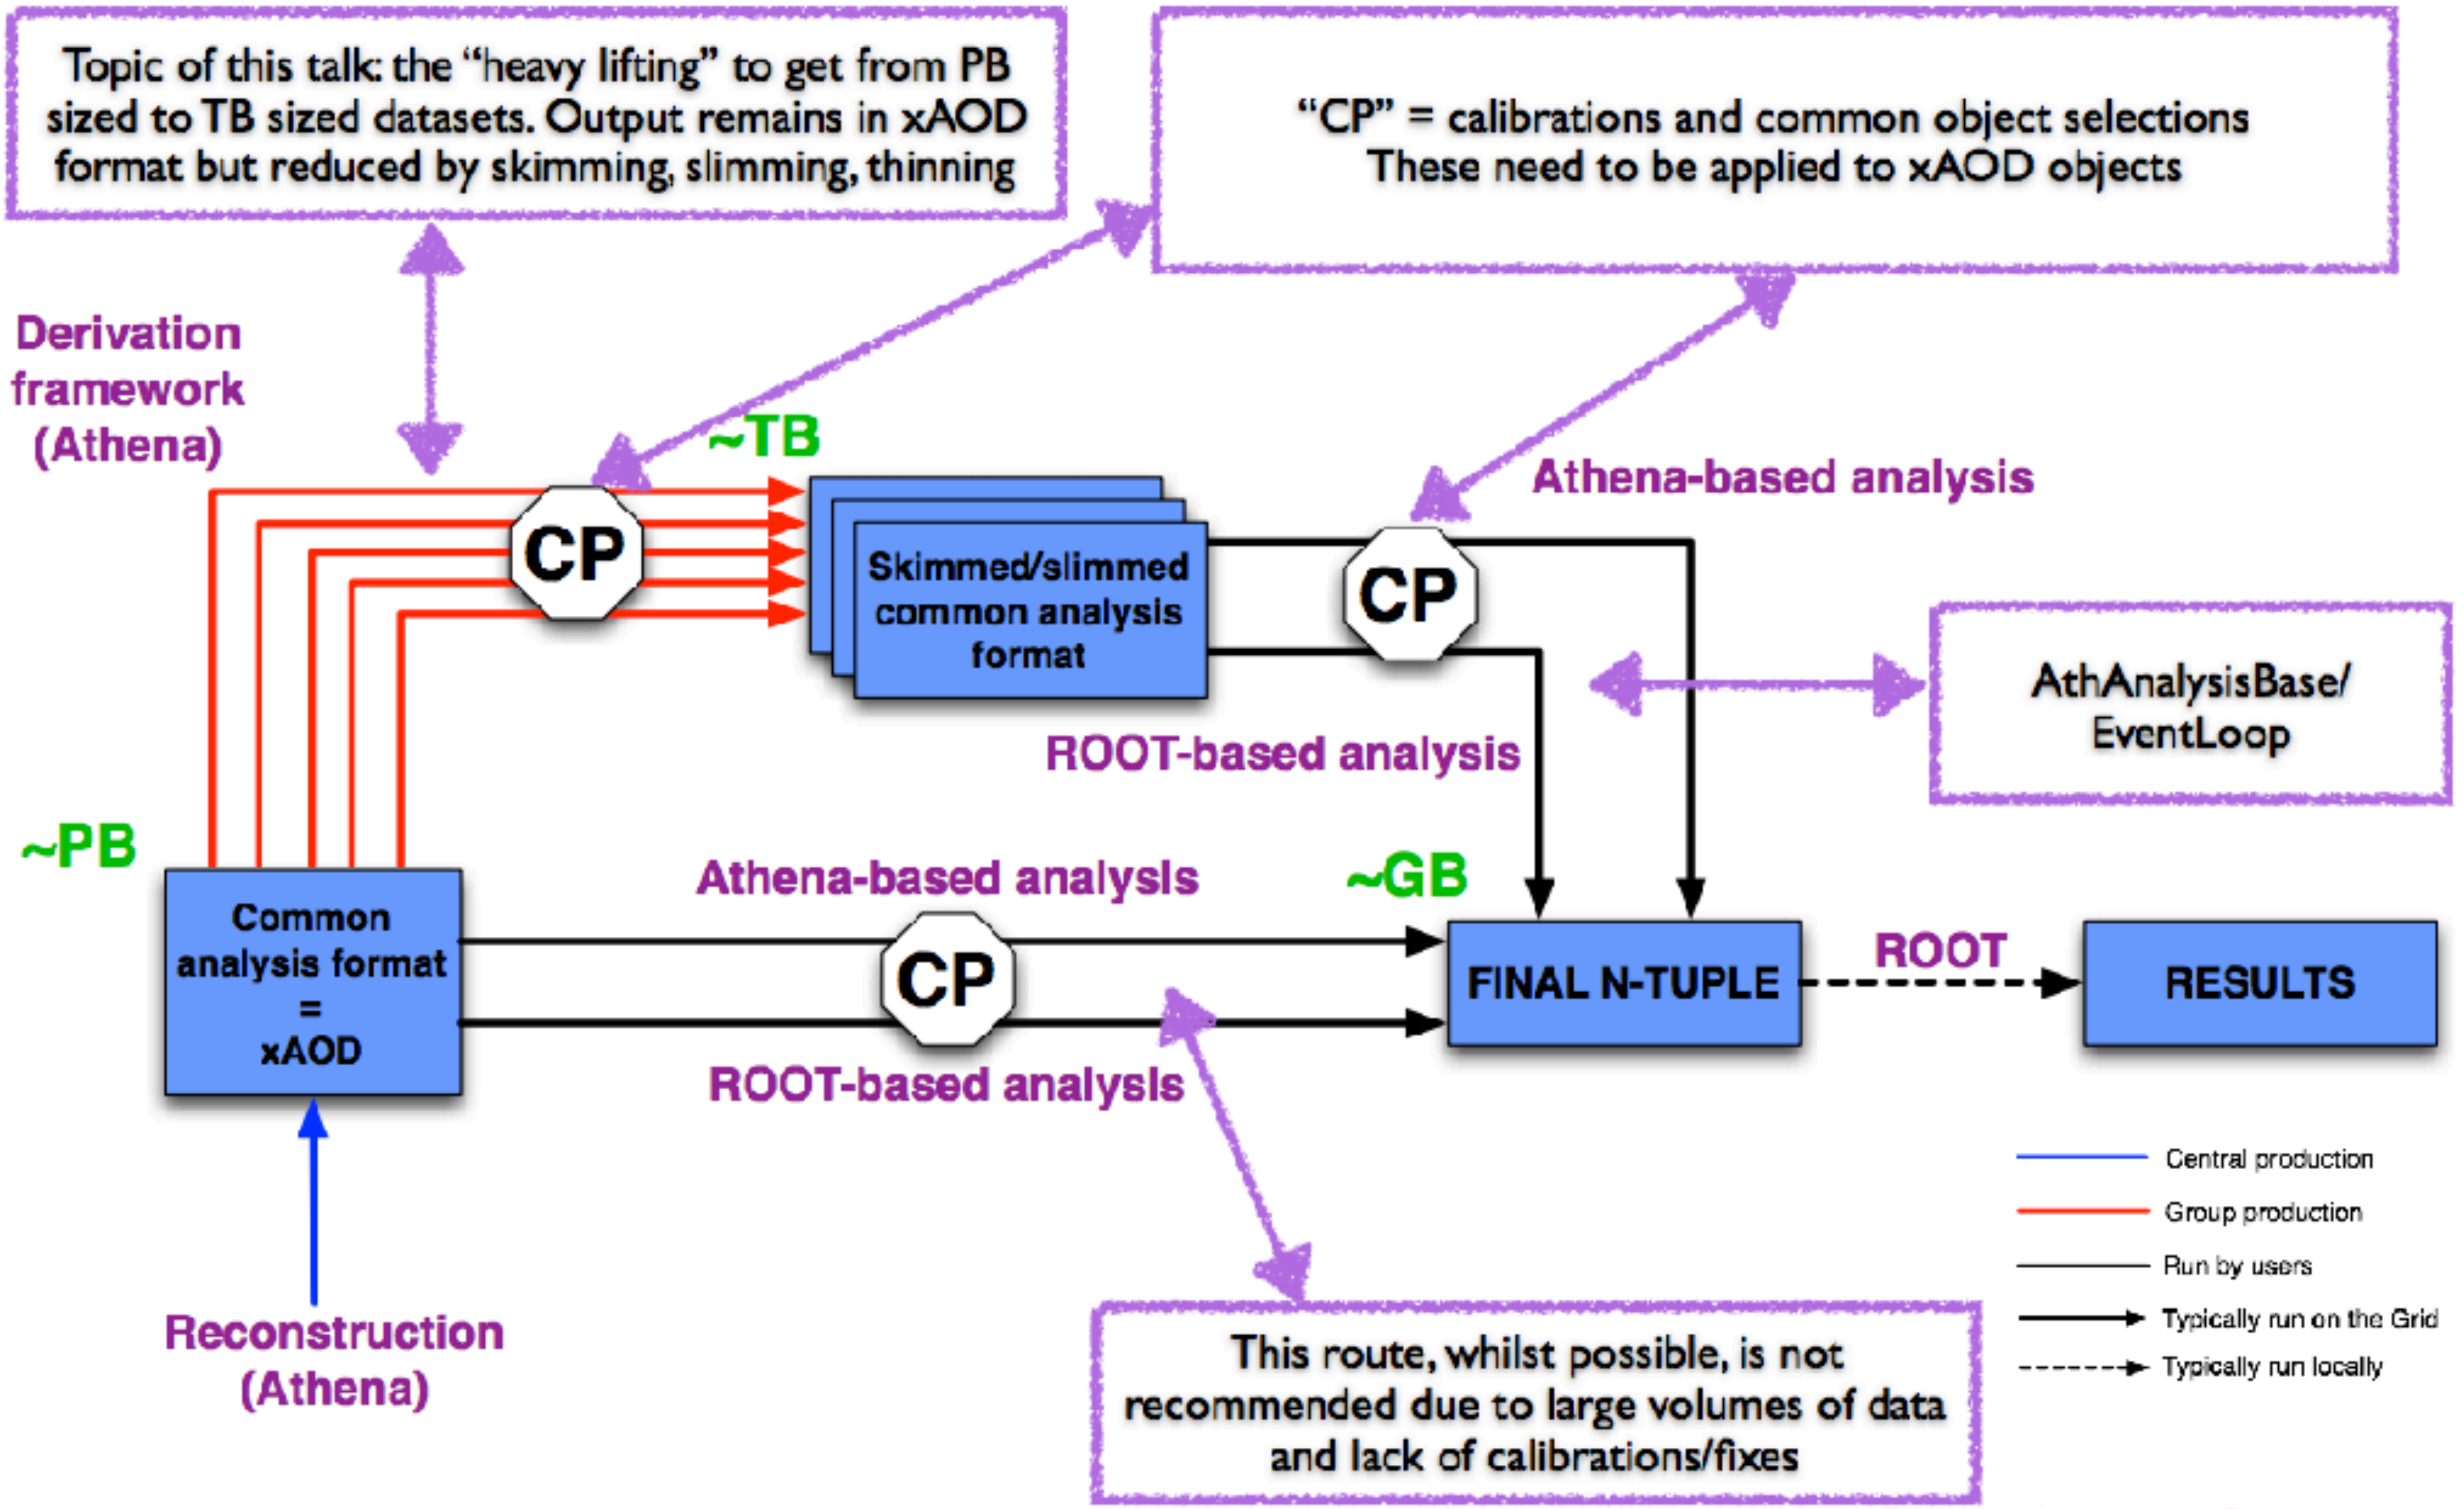

Slide:James Catmore

Data reduction operations

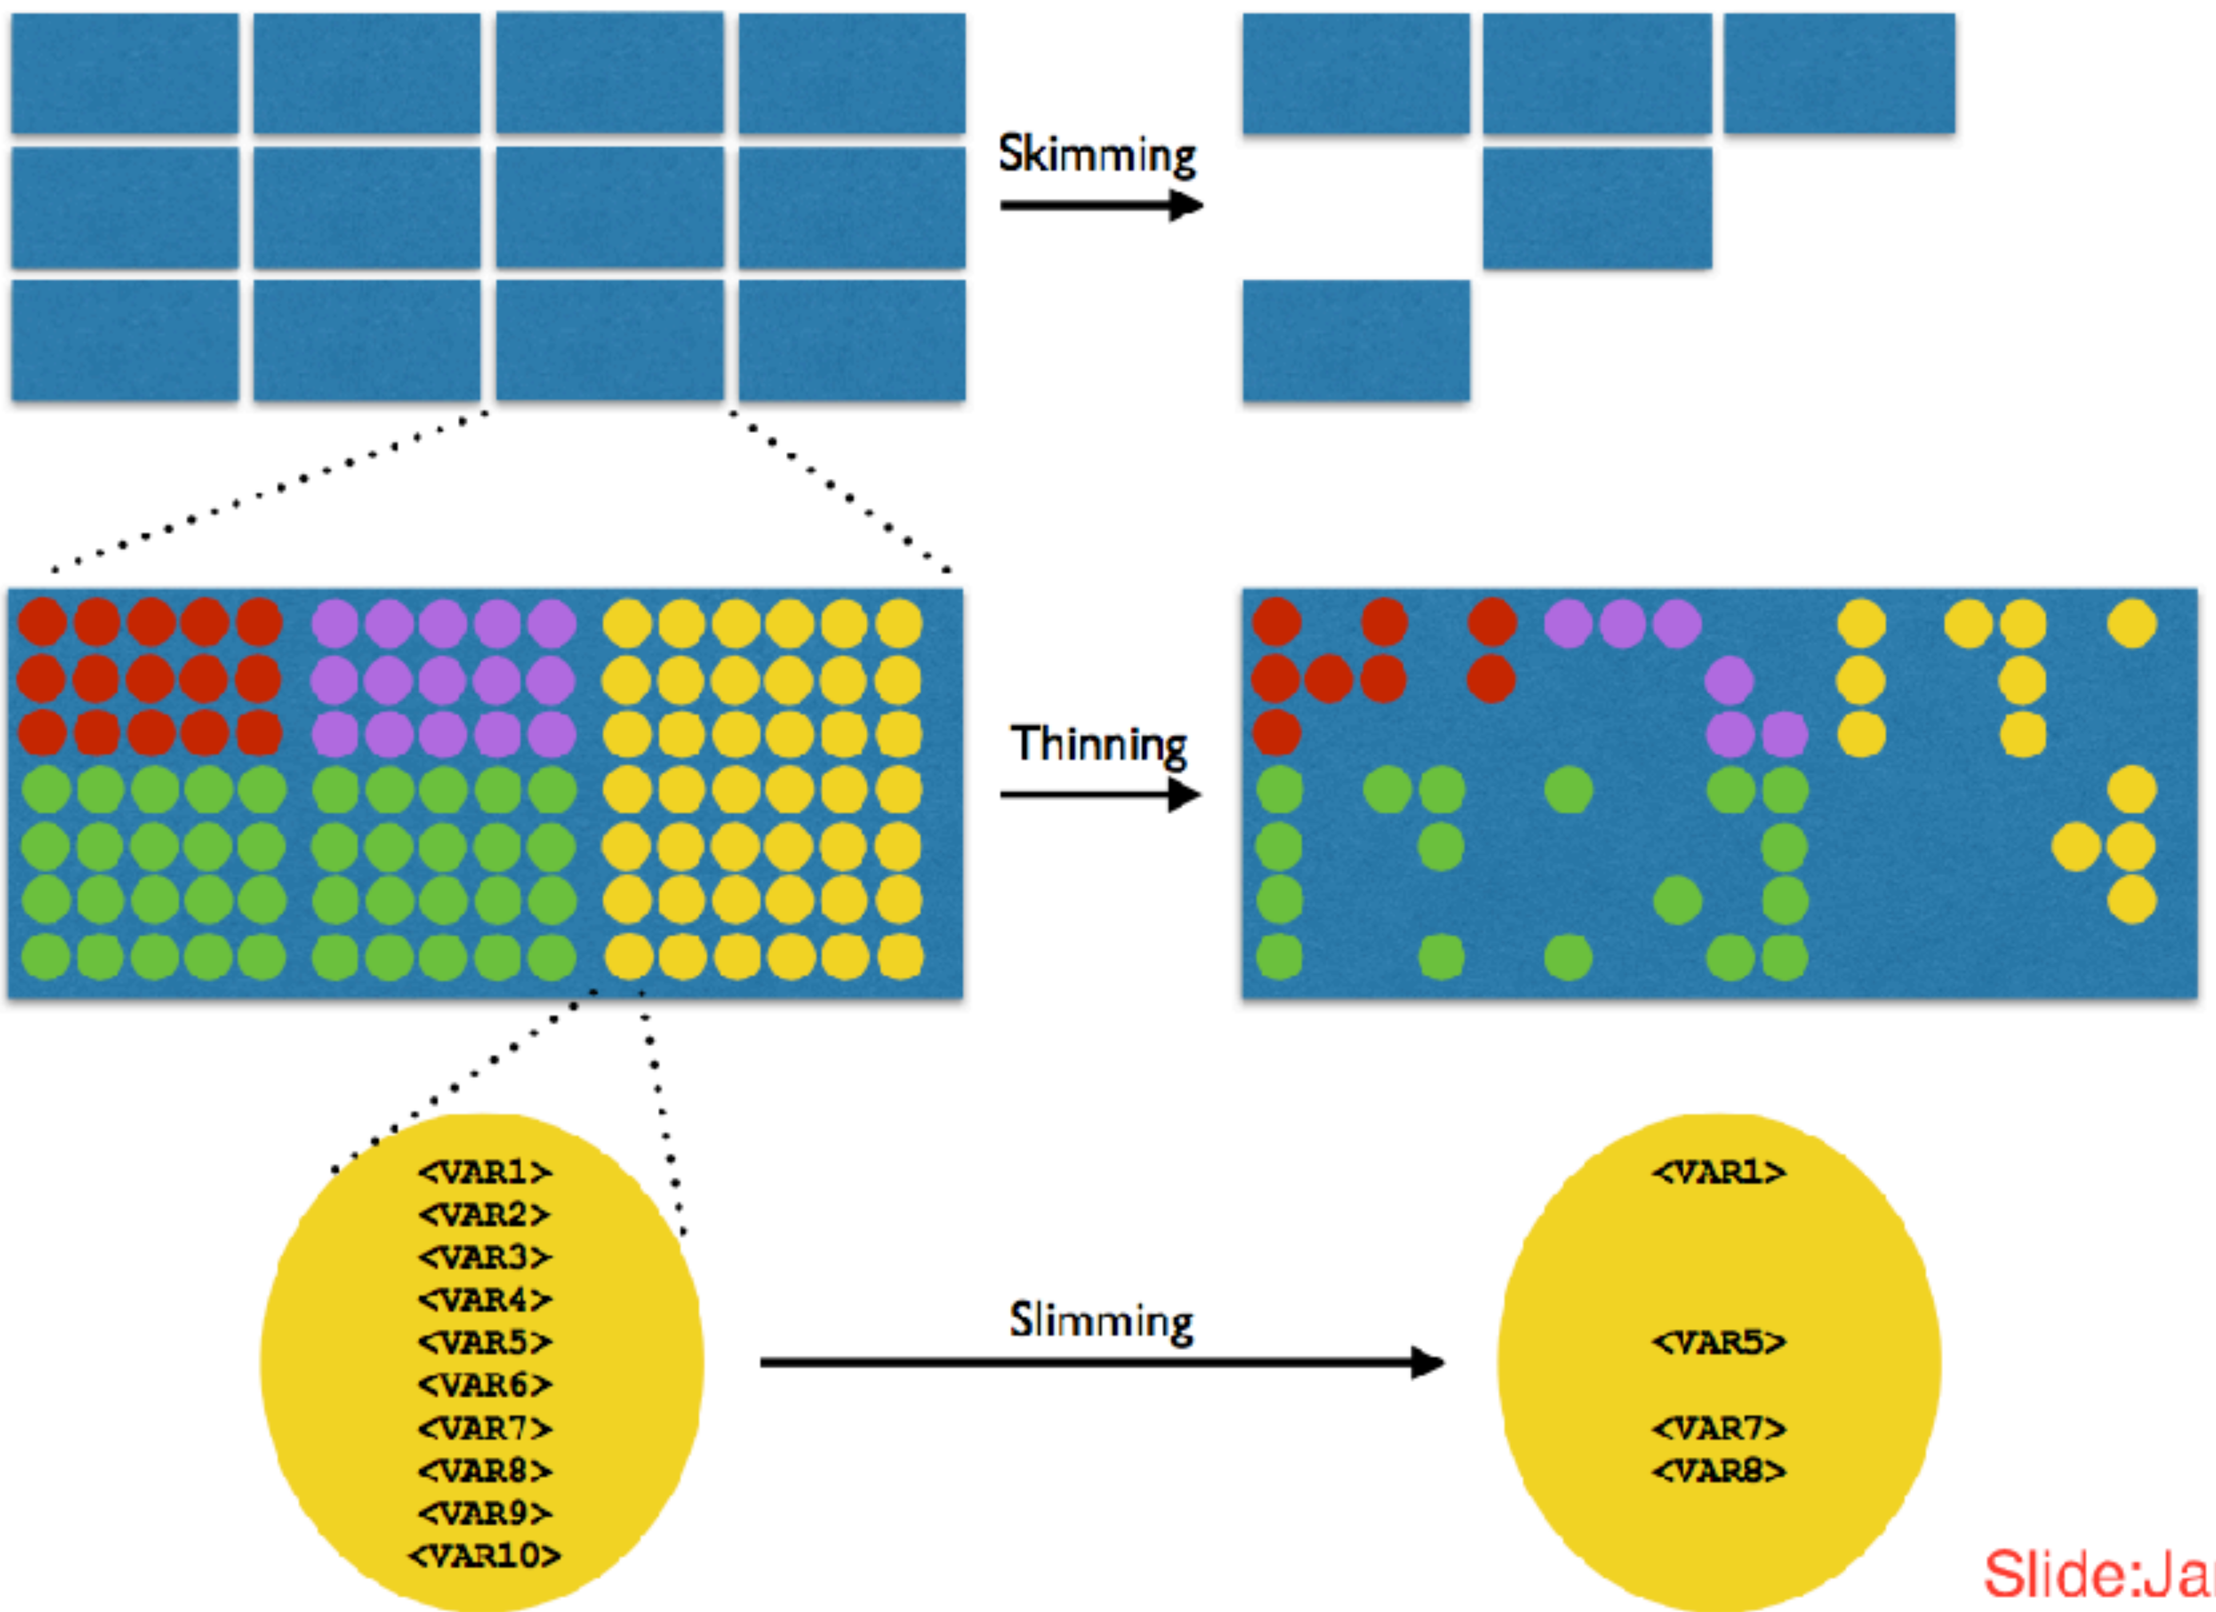

Skimming: removal of whole events based on pre-set criteria

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

Slimming: removal of variables within objects uniformly across events

Slide:James Catmore

Augmentation

- 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")
- Augmentation can be shared across a train, saving CPU

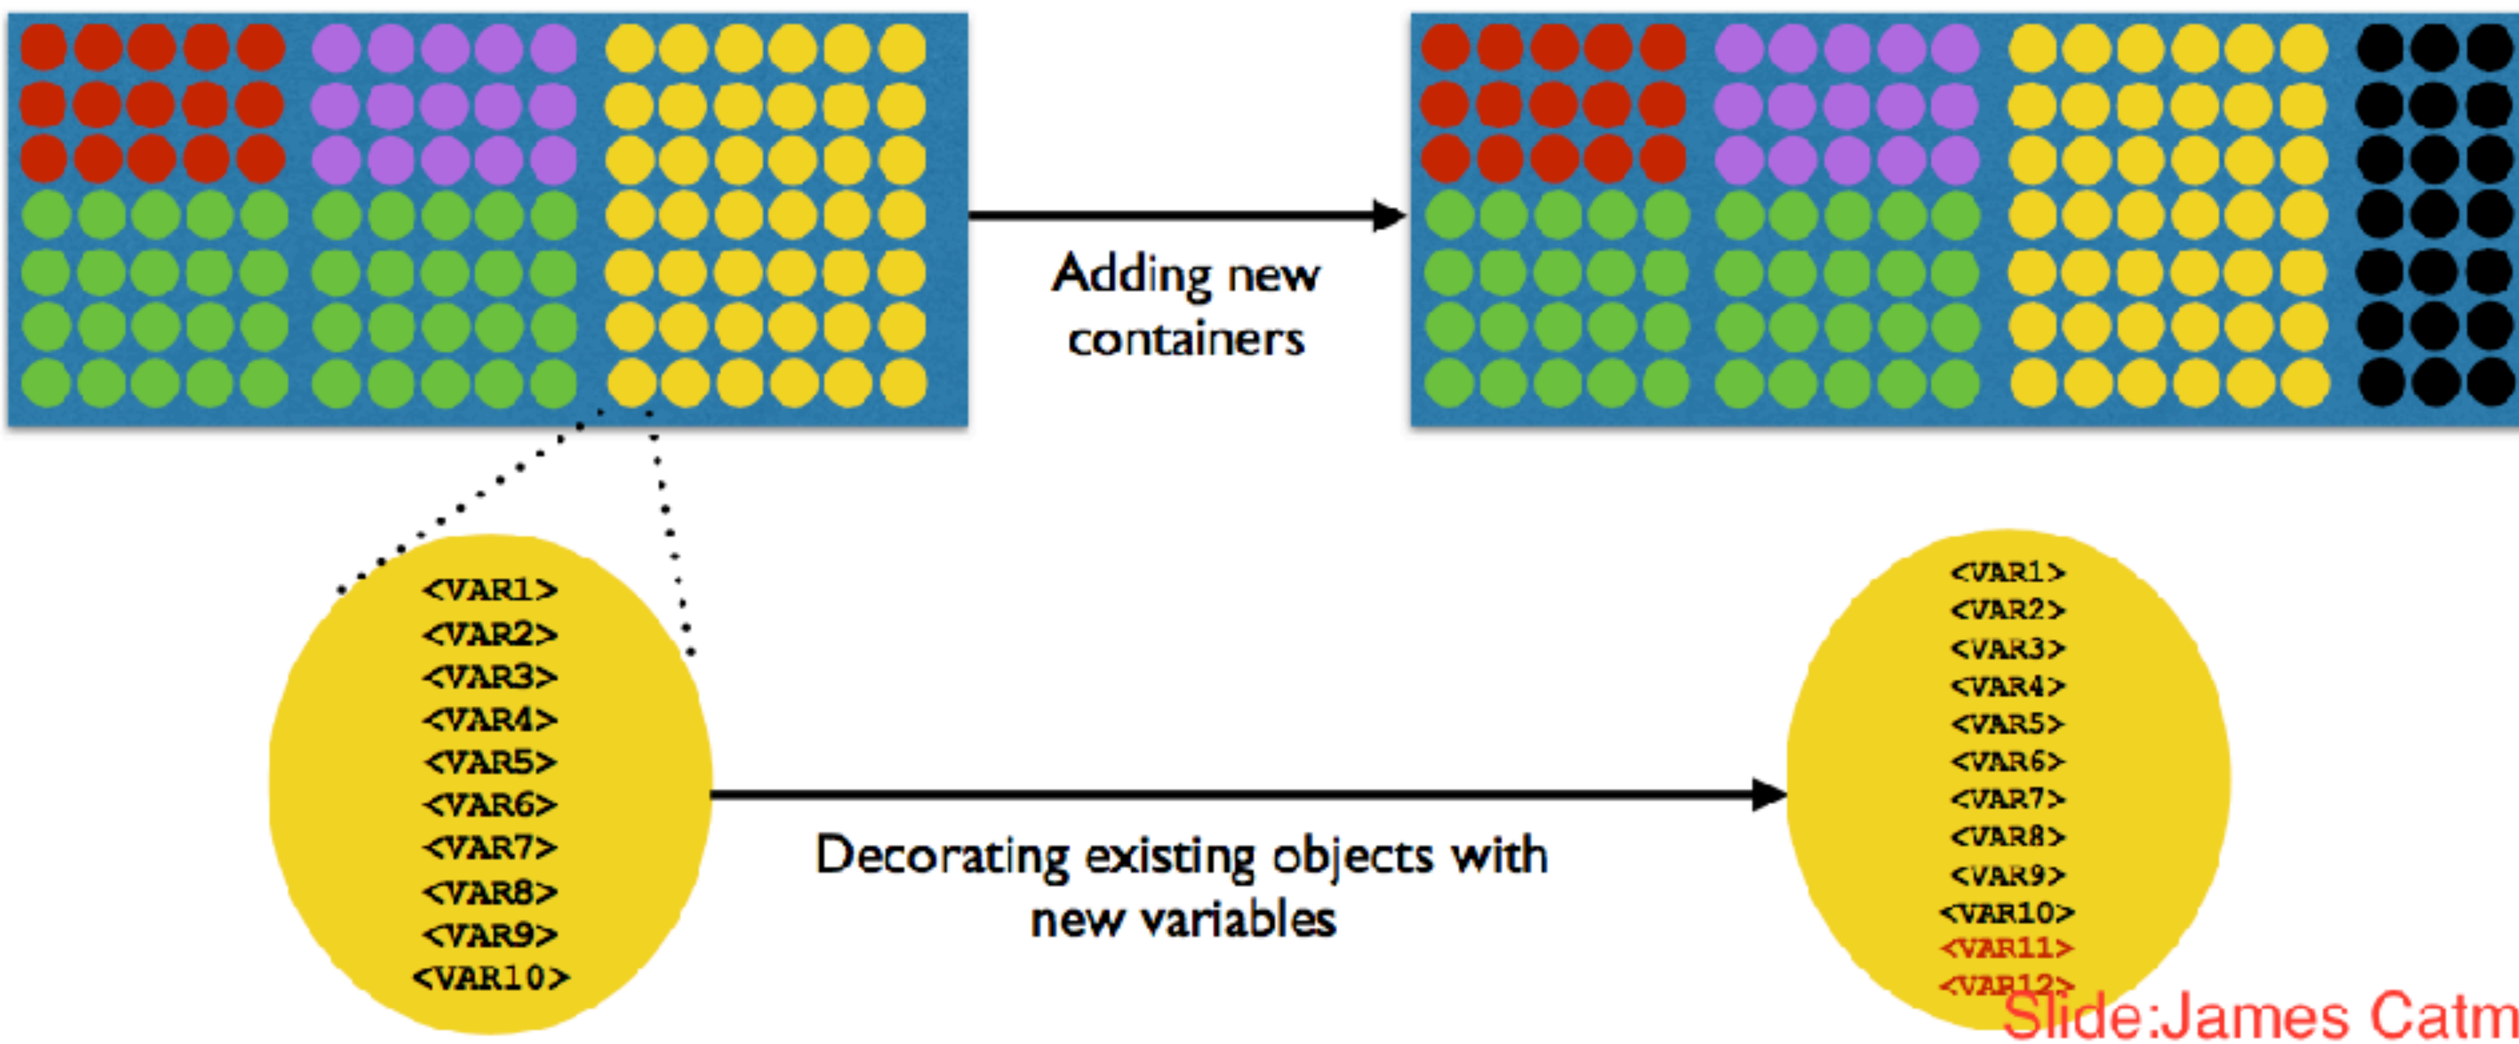

Slide:James Catmore

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 ≤ total xAOD size
  - With ~100 formats, each DAOD should aim to be ~1% of its input xAOD size

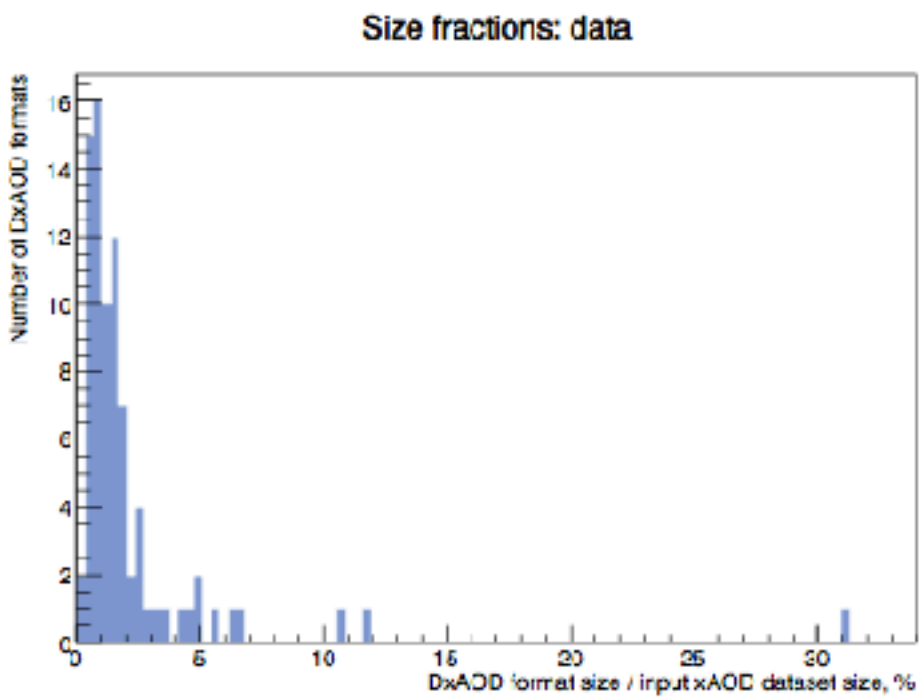

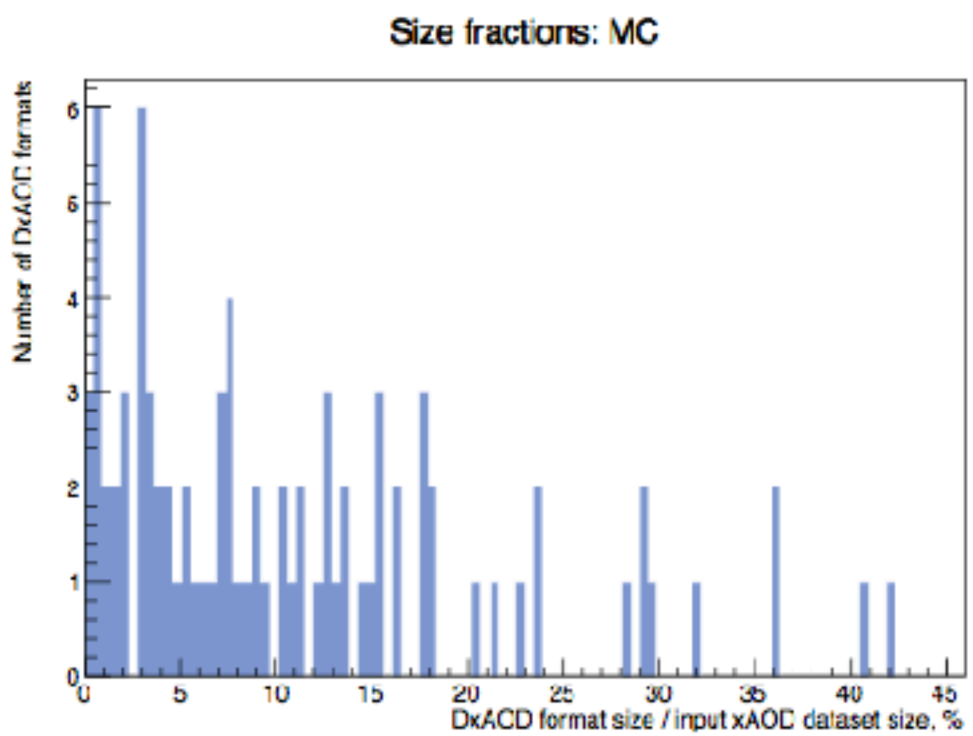

Number of derivations	
B-physics	8
Data preparation	3
E-gamma	9
Exotics	21
Flavour tagging	5
Heavy ions	9
Higgs	22
Jet/missing energy	11
Muons	5
Standard Model	8
Supersymmetry	14
Taus	4
Tile calorimeter	1
Top quarks	5
Tracking	2
[illegible]	6

Slide:James Catmore

Data management

- In addition to the new analysis model was the introduction of a dataset 'lifetime' model.
- Each datatype has a lifetime where the data is 'primary' and not available for deletion.
- Once lifetime has expired the data is 'secondary' and now subject to deletion which is triggered manually when neccessary.
- If a dataset is touched then an extension to the lifetime is allowed.

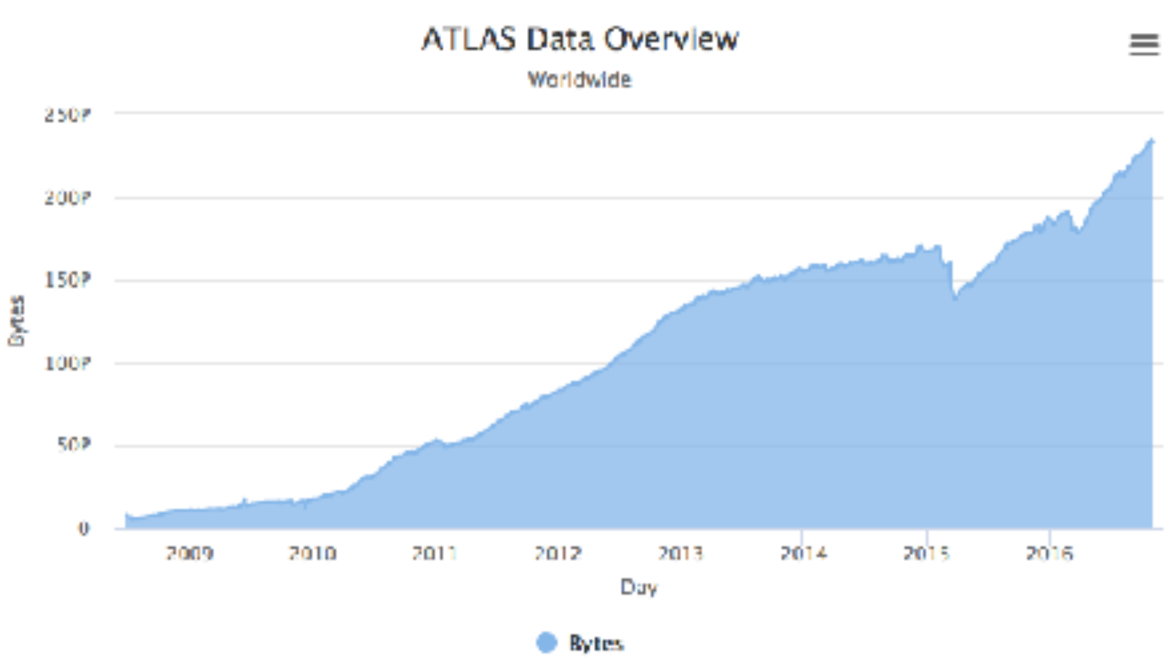

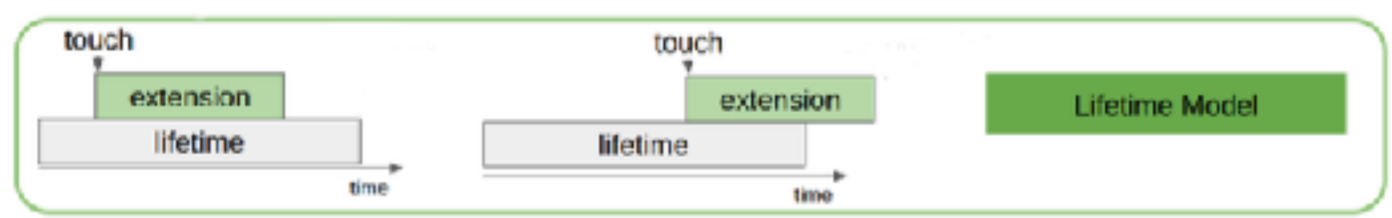

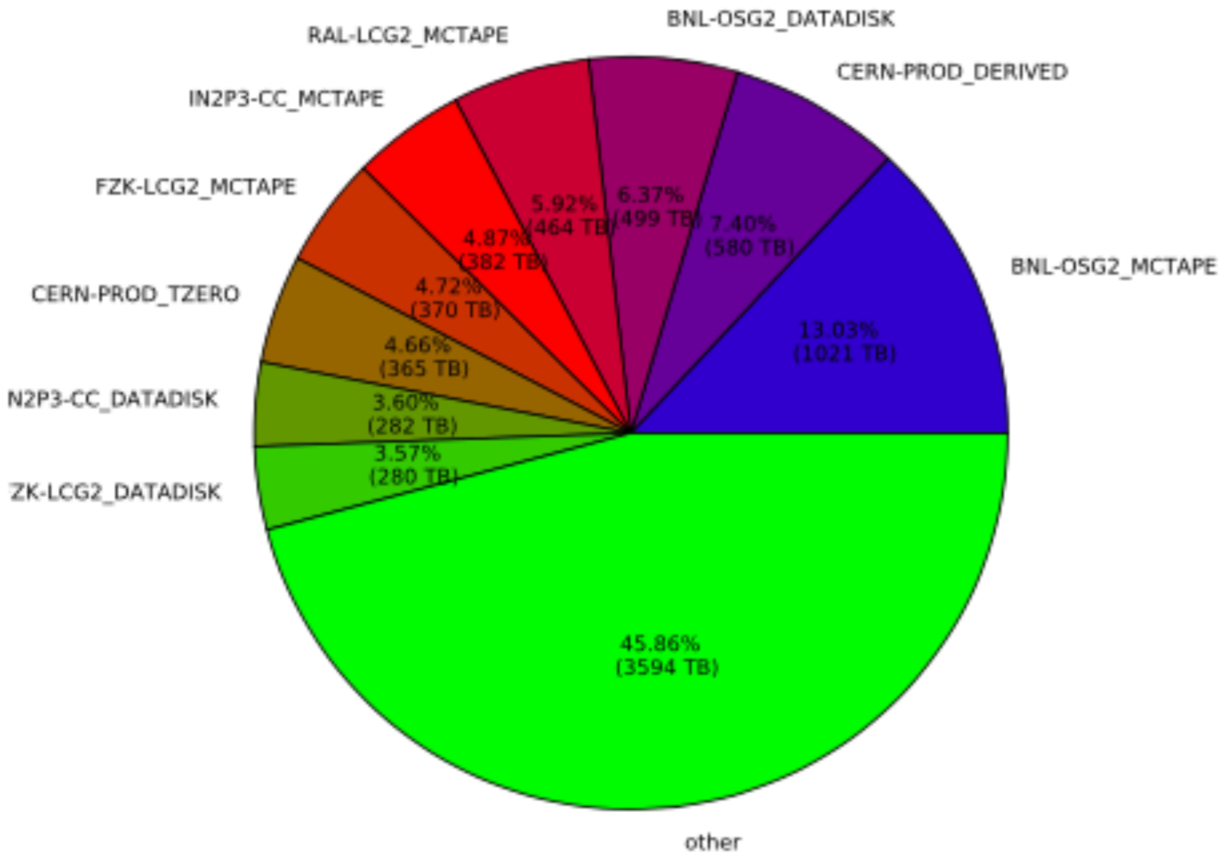

Volume that will be expired on 01/01/16 split by RSEs

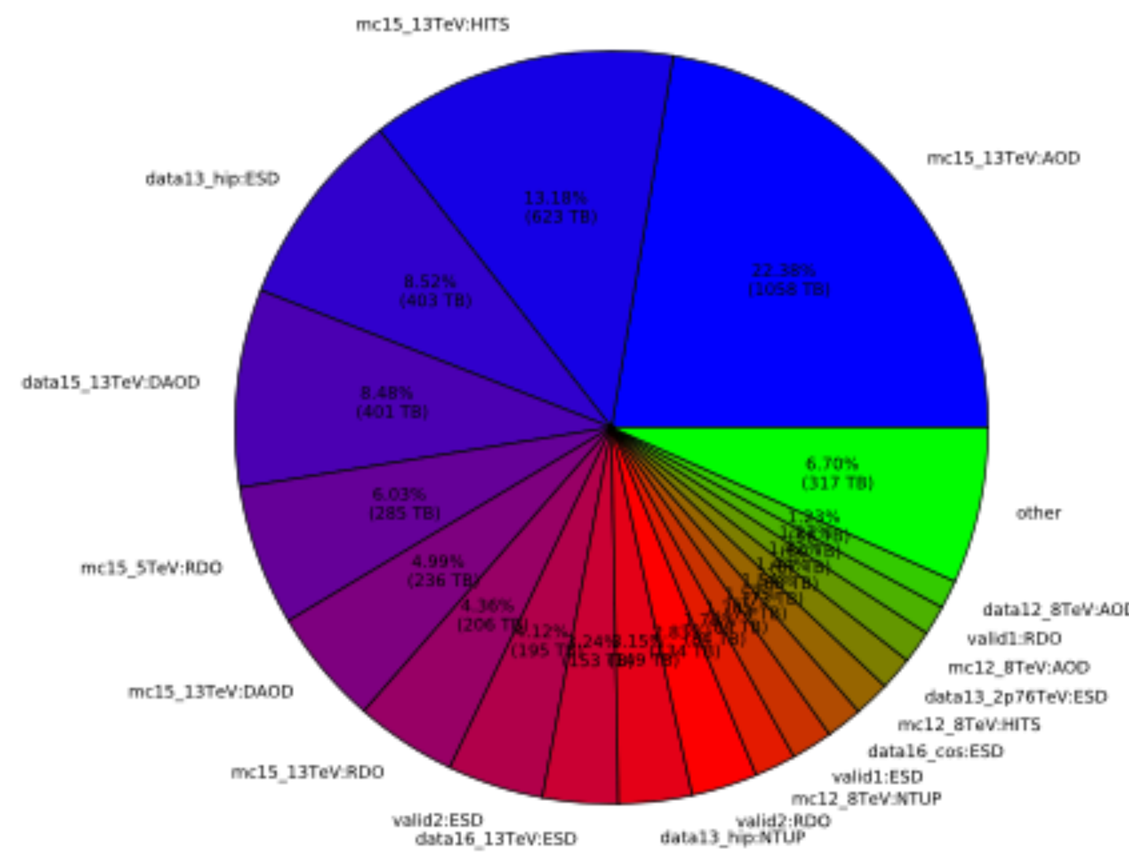

Volume that will expire in September 2016 split by project/datatype

Lifetime Model

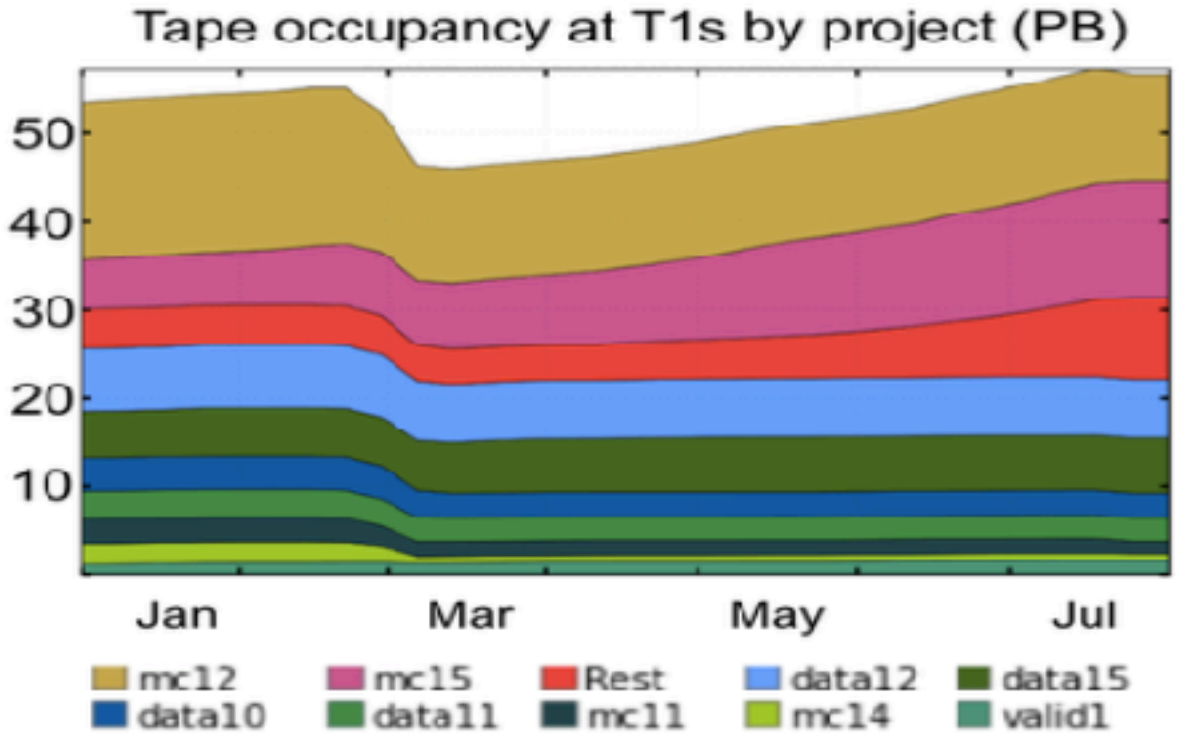

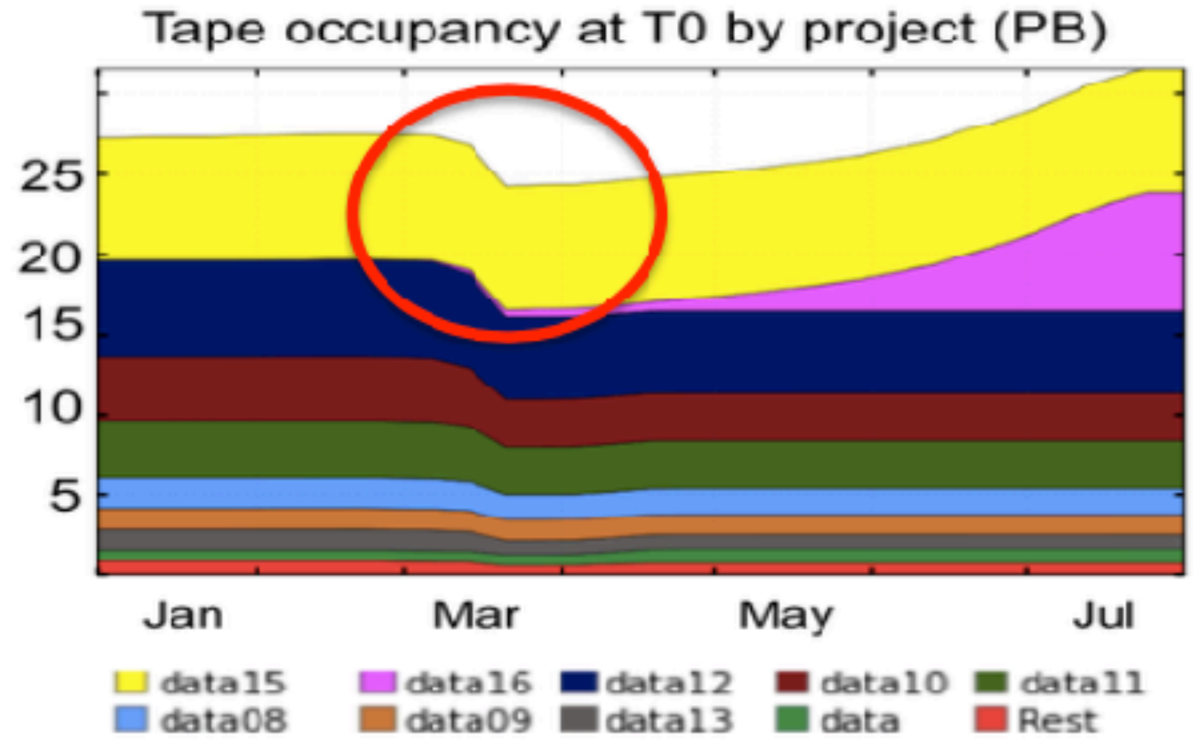

Lifetime Model allows to reduce the unused (useless) data retained on disk AND TAPES
Tuning the lifetimes parameters buys some contingency for disk and tape.
Such contingency proved being crucial to handle the unexpected volume of data in 2016.
Automation.

Run 2 model summary

- xAOD format is readable by ROOT and Athena giving user flexibility for private analysis, preserving provenence of file content. All analysis should be done on DAOD.

- Derivation framework allows production of separate xAOD files tailored to the different physics groups with a large reduction in total size on disk.

- Coherent Analysis Release of recommended tools for users. CP tools operate on xAOD using common interfaces with source-compatiblity between Athena and ROOT.

- The bottom line is these changes enable increased physics output in an environment of static resources (compute, storage, people)

Looking ahead to Run 3 and beyond

HL-LHC baseline resource needs

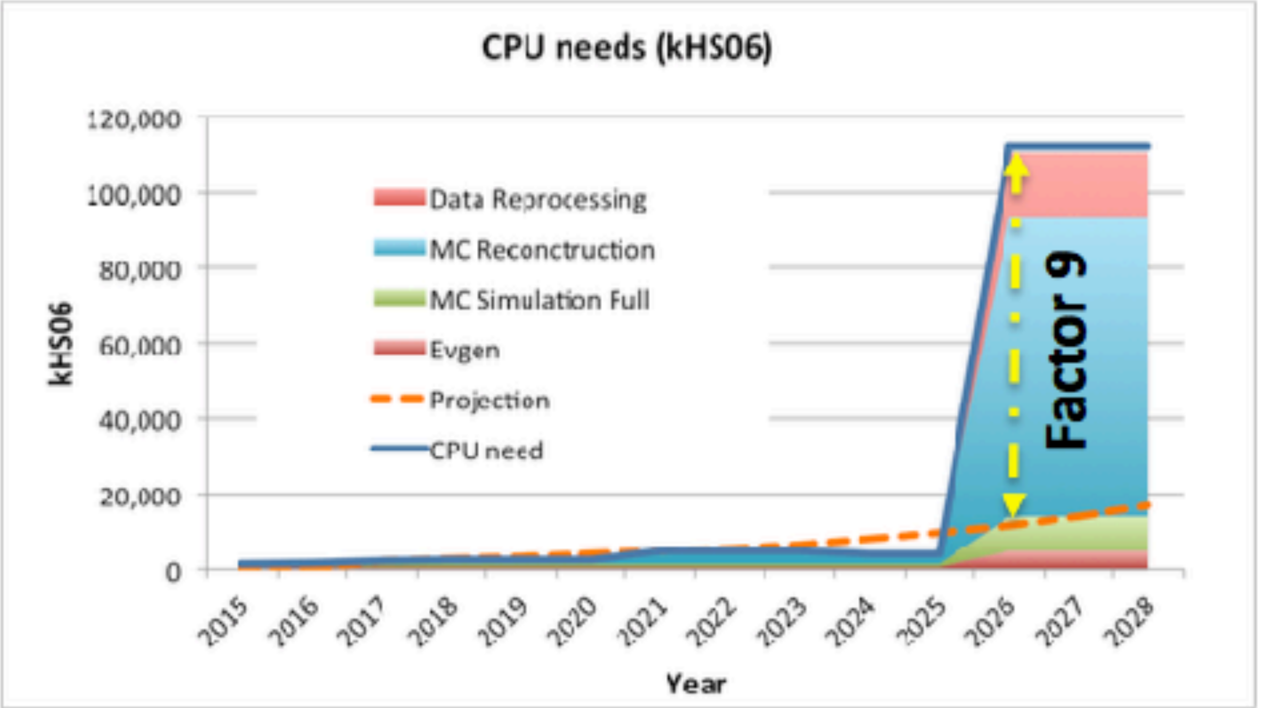

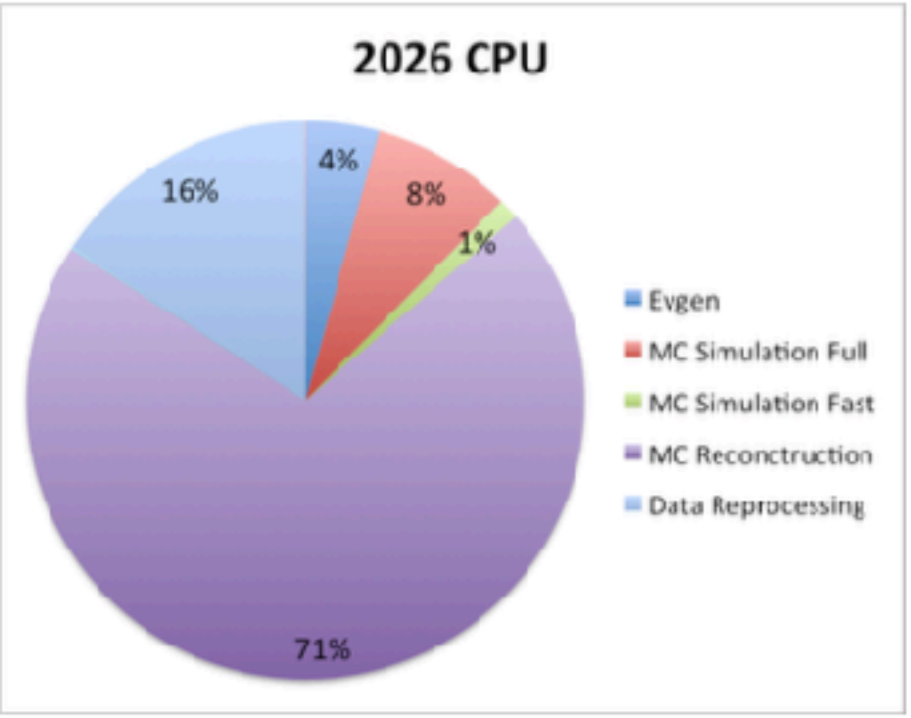

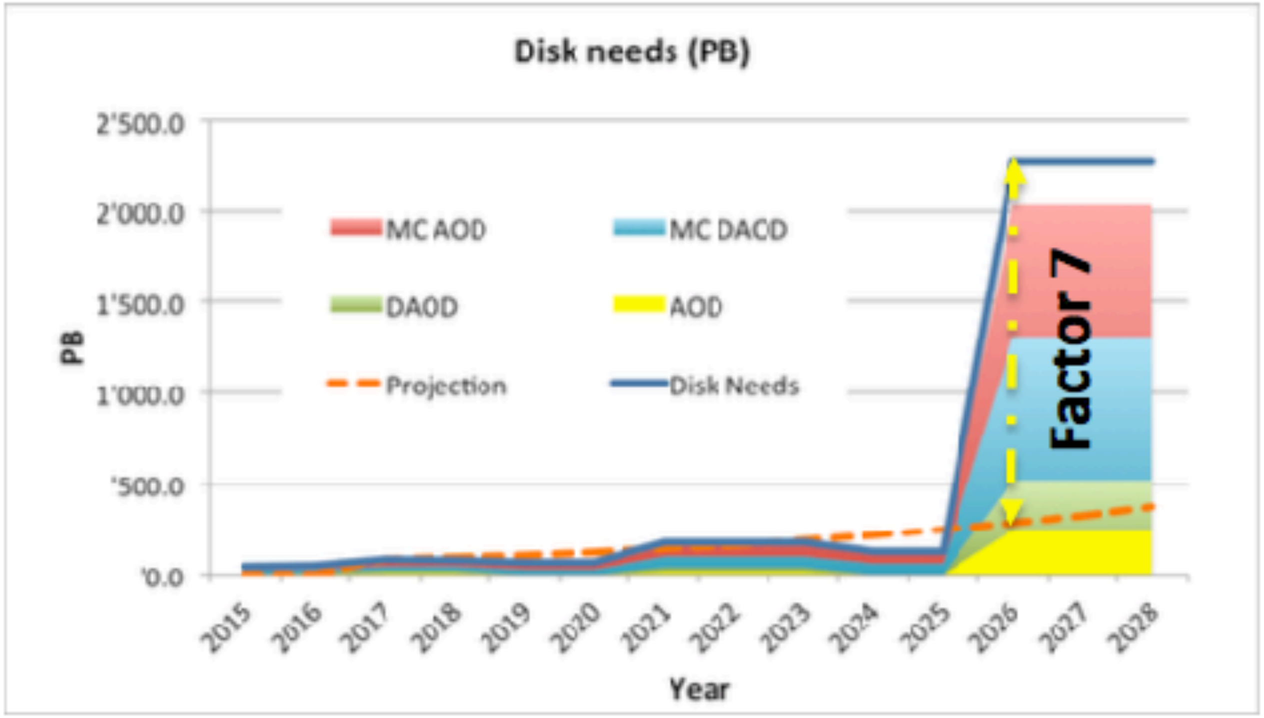

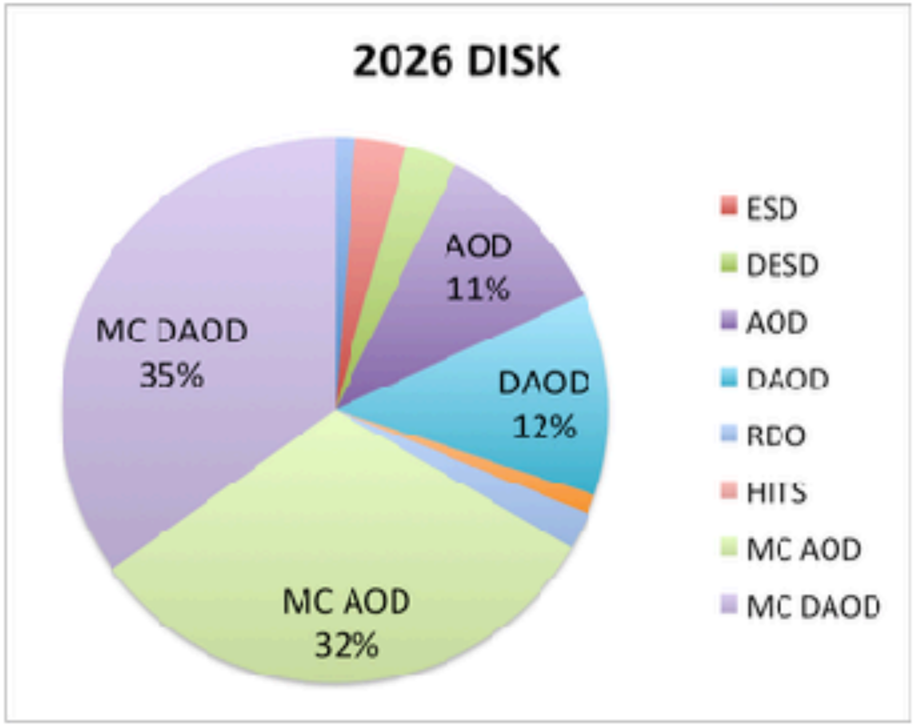

Hardware trend and implications

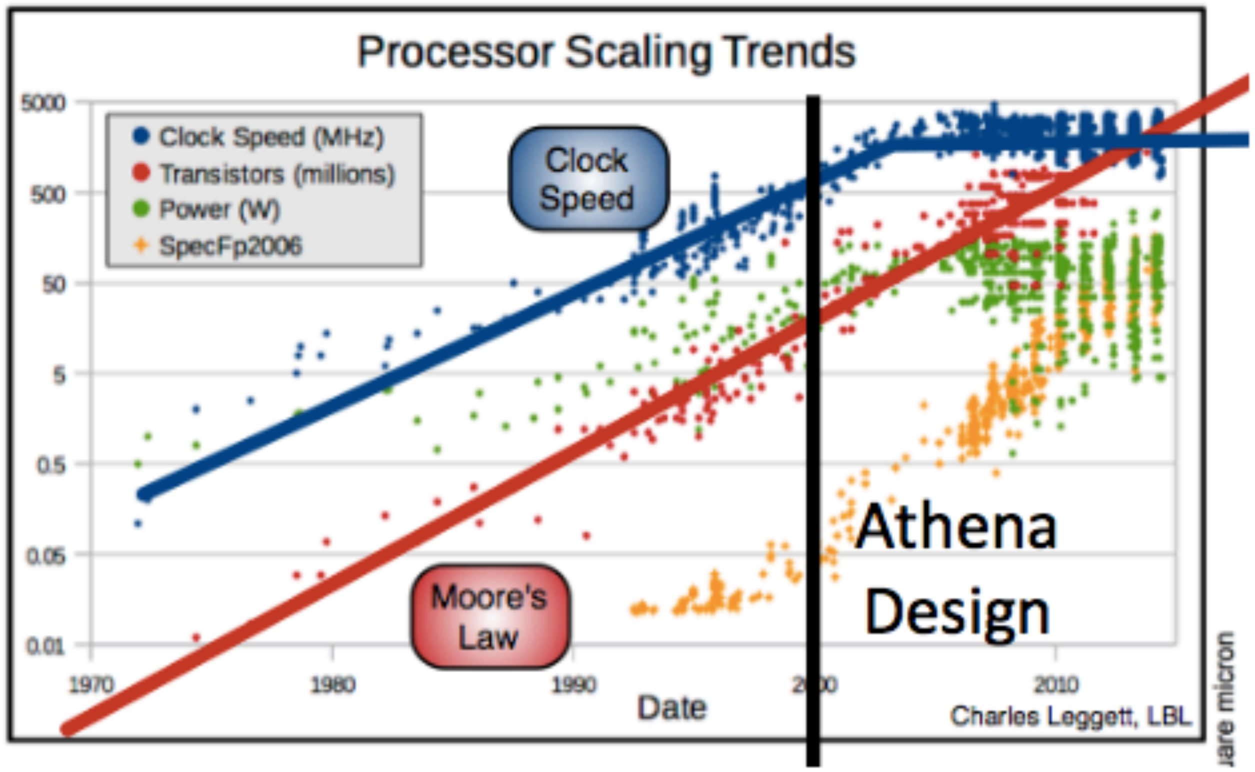

Clock Speed stalled but transistor density keeps increasing. Exploiting hardware becomes more complicated (vectors, memory…)

Example: Cori@NERSC (Intel Knights Landing)
1PB of Memory, 9304 nodes
68 cores/node, 4 HW threads/core
=> Approx 300 MB/thread

Heterogeneous Resources

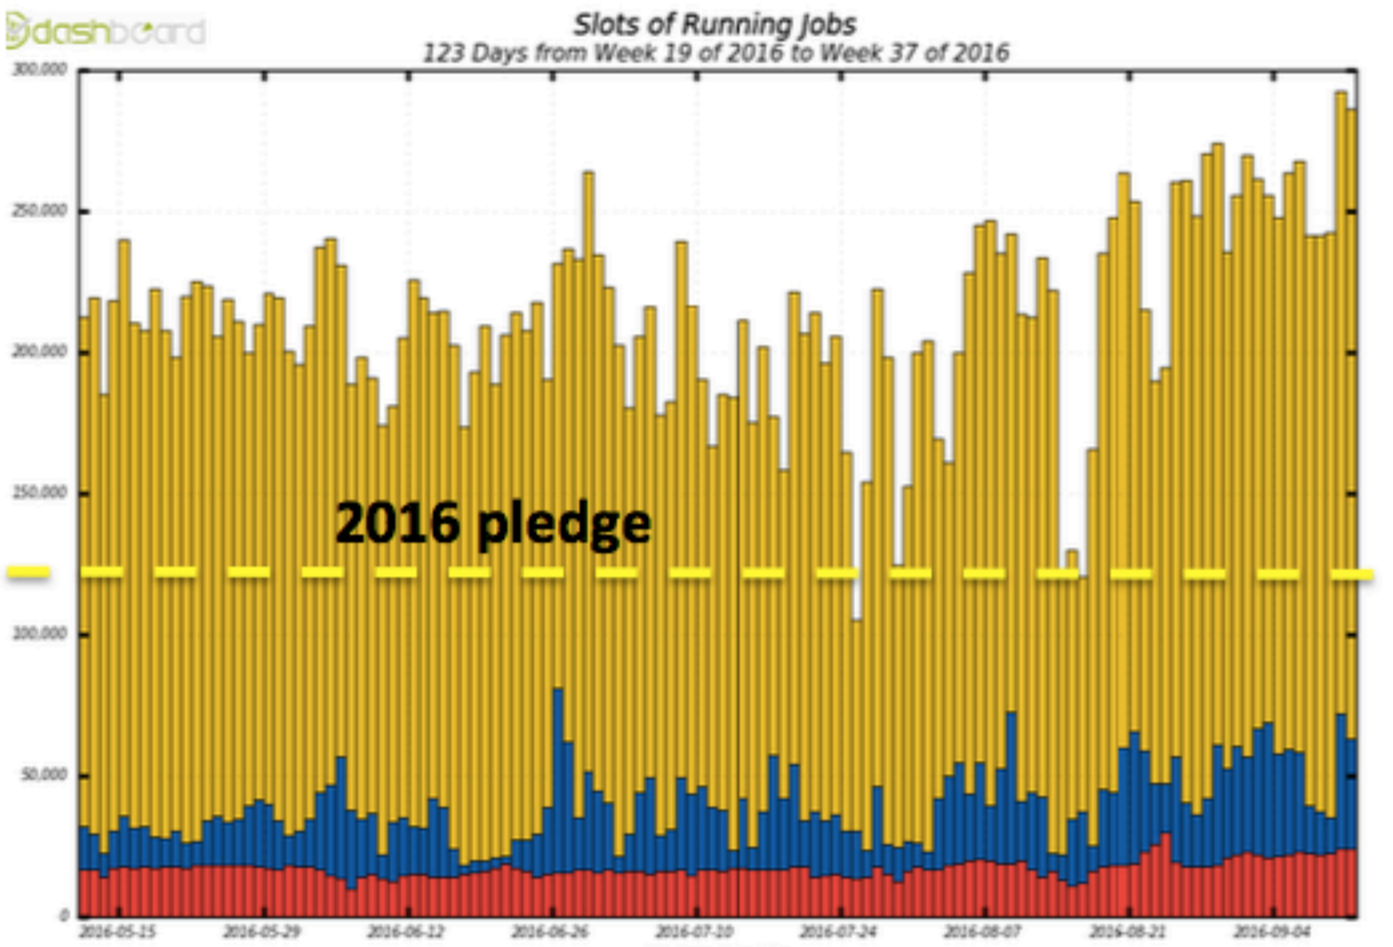

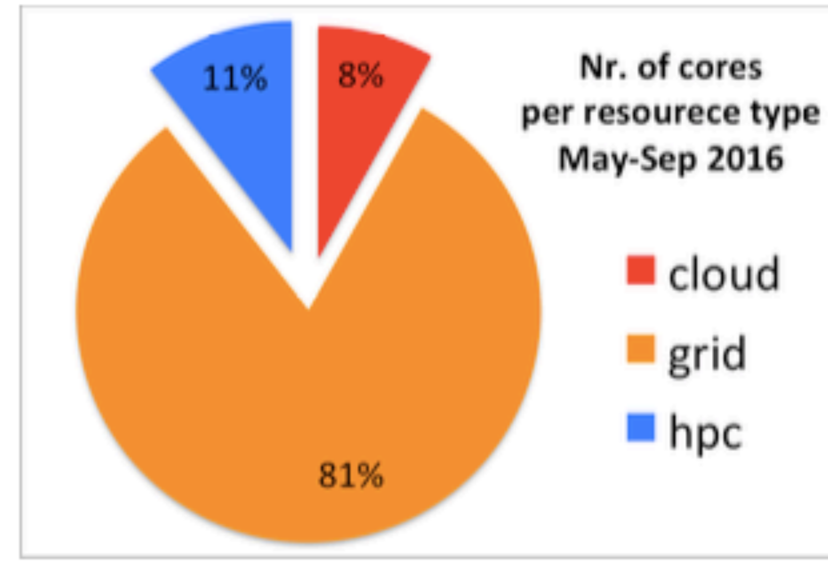

Integration of non Grid resources in ATLAS is a big investment with the potential of a big return

Challenges: resource provisioning, non standard architecture, GPU processing capacity, memory

Computing infrastructure in HL-LHC

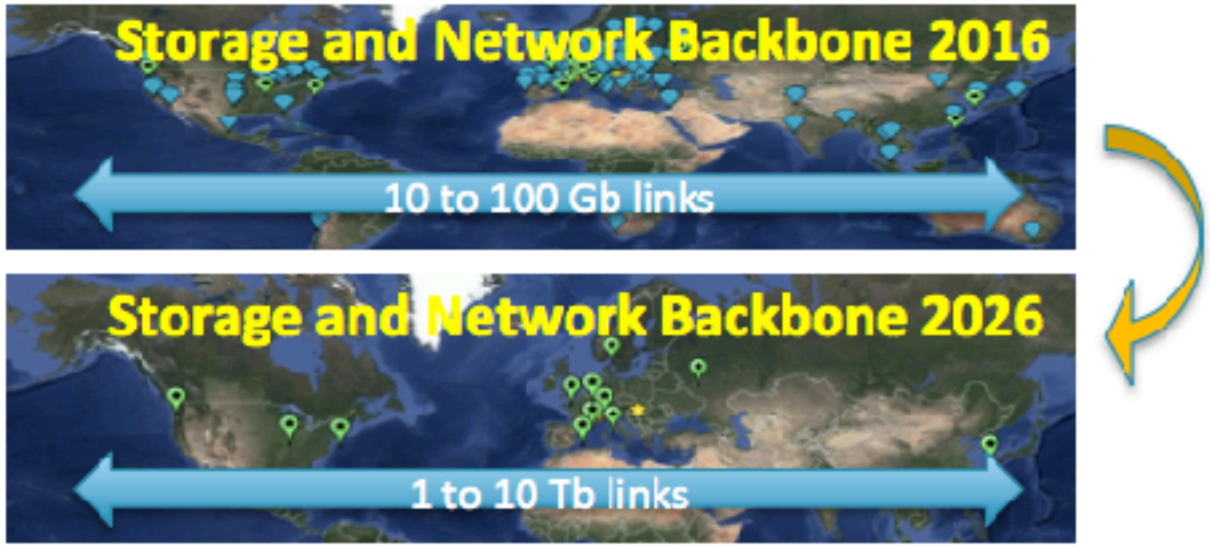

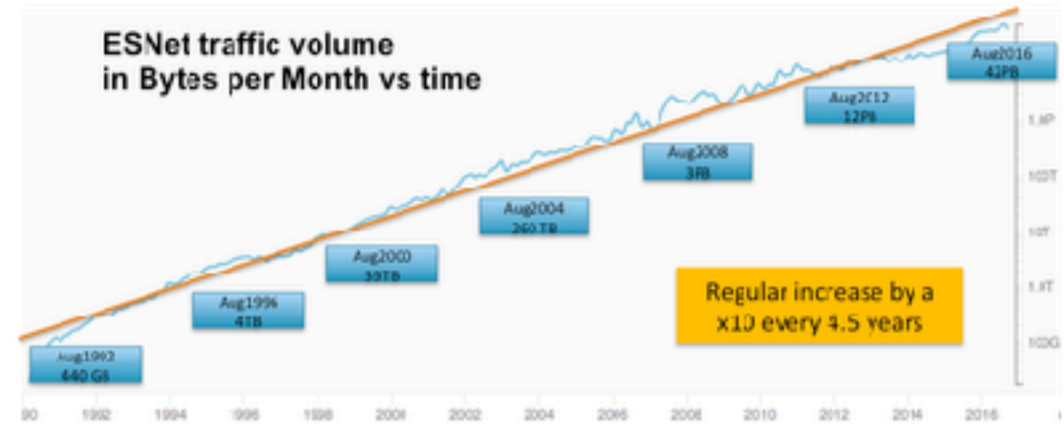

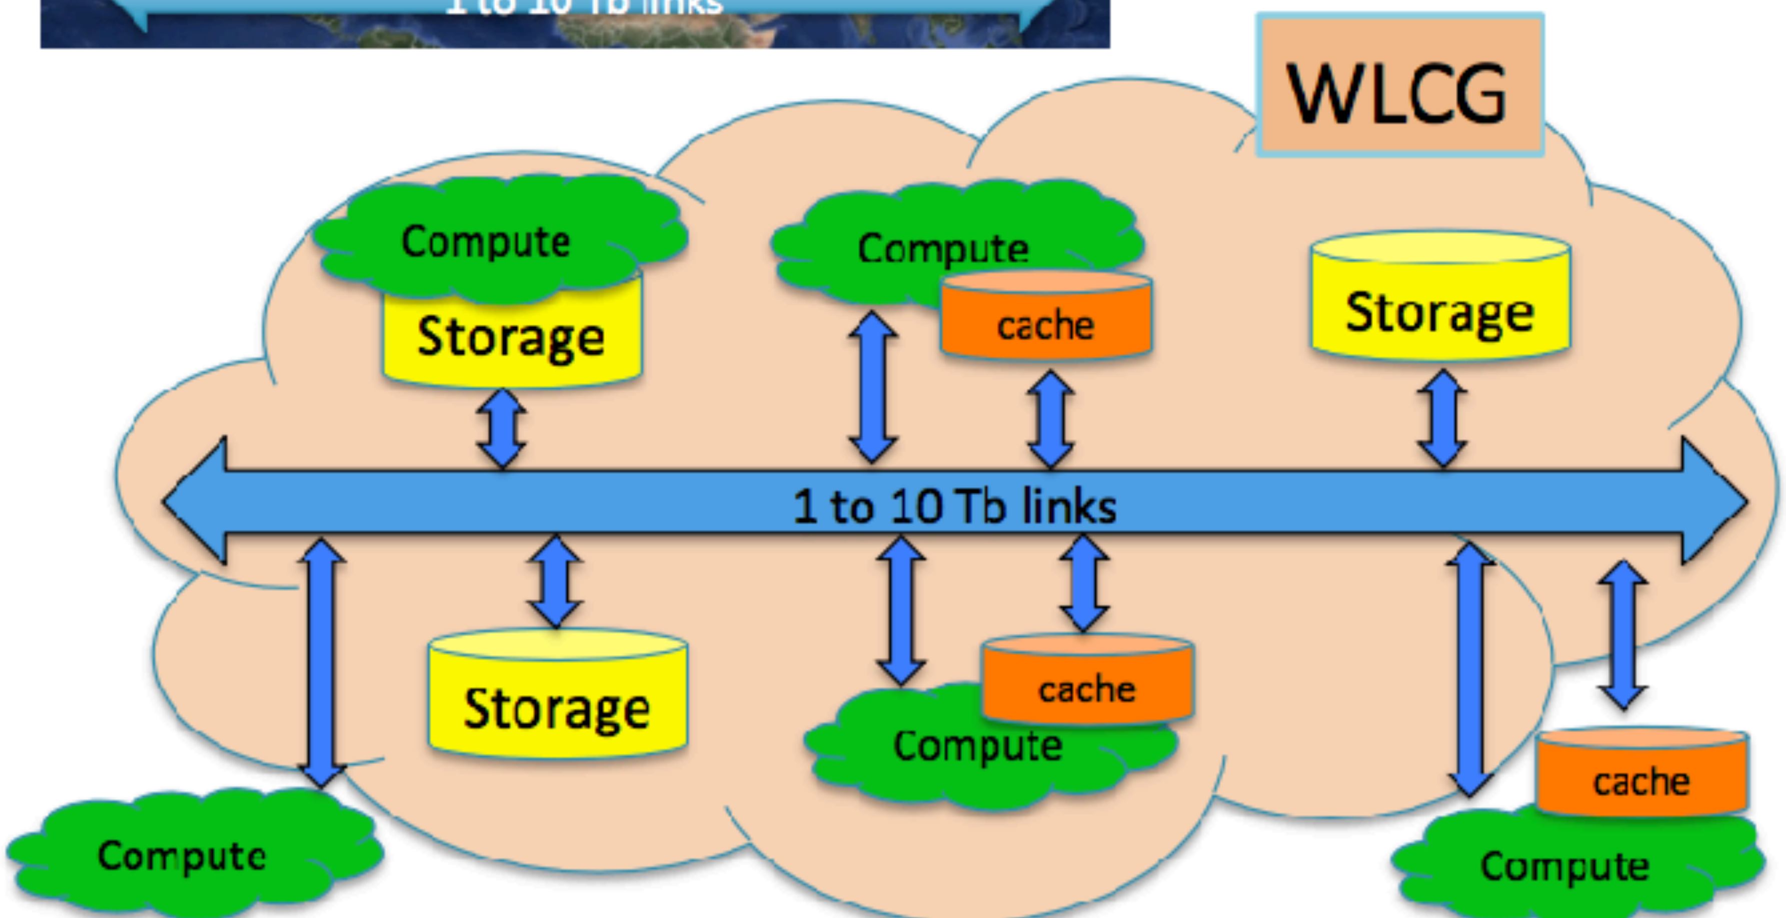

- Network-centric infrastructure
- Storage and Compute loosely coupled
- Storage on fewer data centers in WLCG
- Heterogeneous Computing facilities (Grid/Cloud/HPC/ …) everywhere
- Caching and I/O optimisation to beat RTT

High Luminosity era

- HL-LHC beyond 2026 ten times read-out rate 10kHz
- Cope with the problem from all directions

    1. squeeze all possible resources
    2. adapt software for new hardware architectures
    3. utlilise advances in networking
    4. demand ever more cunning and innovation

Use of generic analytics tools

Analytics in ATLAS et al.

- Three broad areas:

1. Logs collection, storage, view, insight
2. Ops analysis, insight, steering.
3. Physics analysis, testing the waters.

Treasure trove of activity: CHEP2016 contributions

Log analytics

1. CERN general analytics platform "HADOOP and friends" CHEP-231

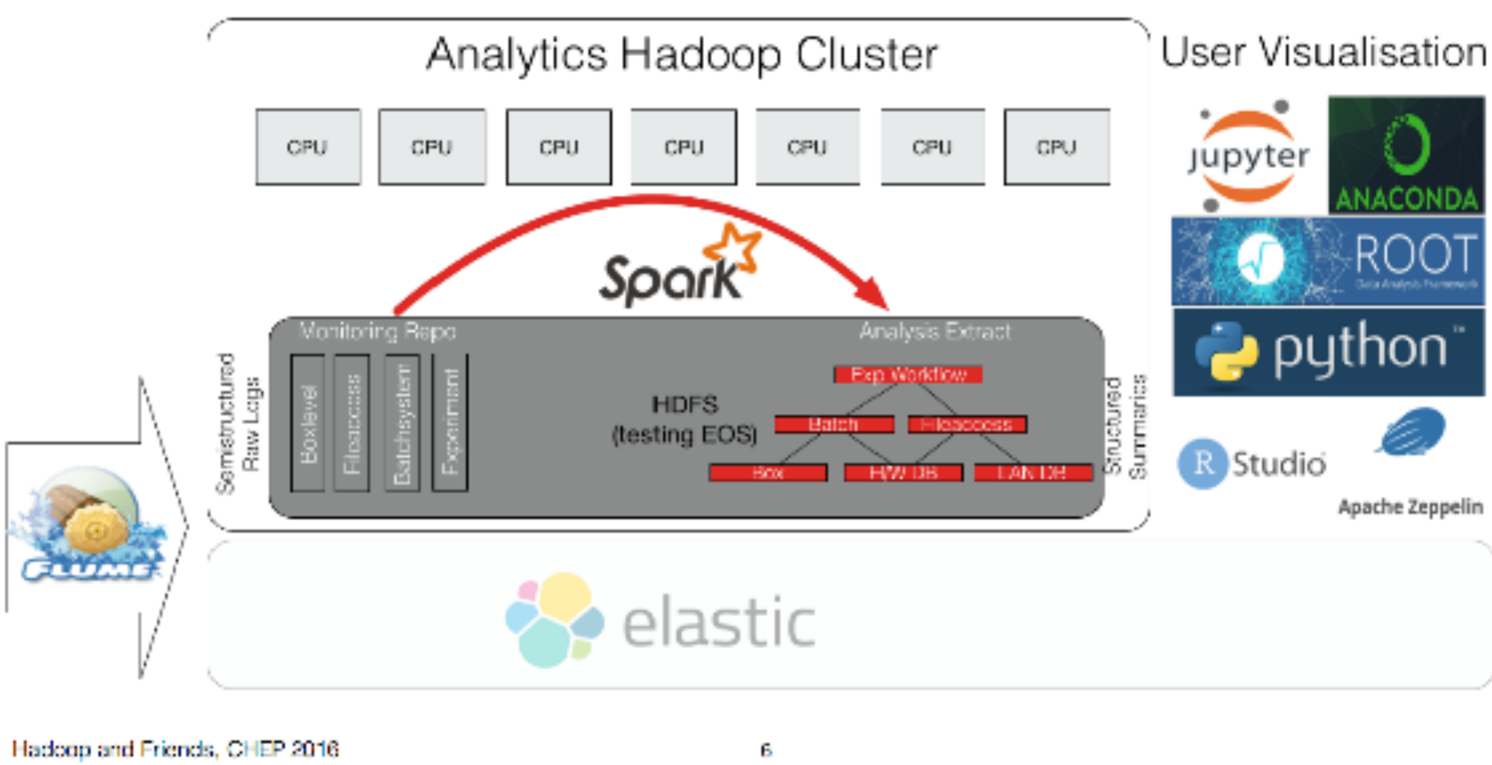

Raw Input Data

Subsystem	Location	Amount	
lemon	hdfs	78 TB	box level
squid	hdfs	110 GB	http cache access
openstack	hdfs	12 TB	agile infrastructure
syslog	hdfs	23 TB	unstructured box logs
eos	hdfs	12 TB	file access metrics
castor	hdfs	55 TB	tape archive access
LANdb	hdfs	small O(100 MB)	host,ip,hypervisor, location
perfsonar	hdfs	small O(10 GB)	network link status
exp. dashboard	hdfs	small (< 1TB)	job summaries
exp. file popularity	hdfs	small O(200GB)	user data access
batch	hdfs	500 GB	accounting & queue config
hw specs	afs	100MB)	h/w rating per model

Log analytics

2. CERN general analytics platform
“First results” CHEP-229

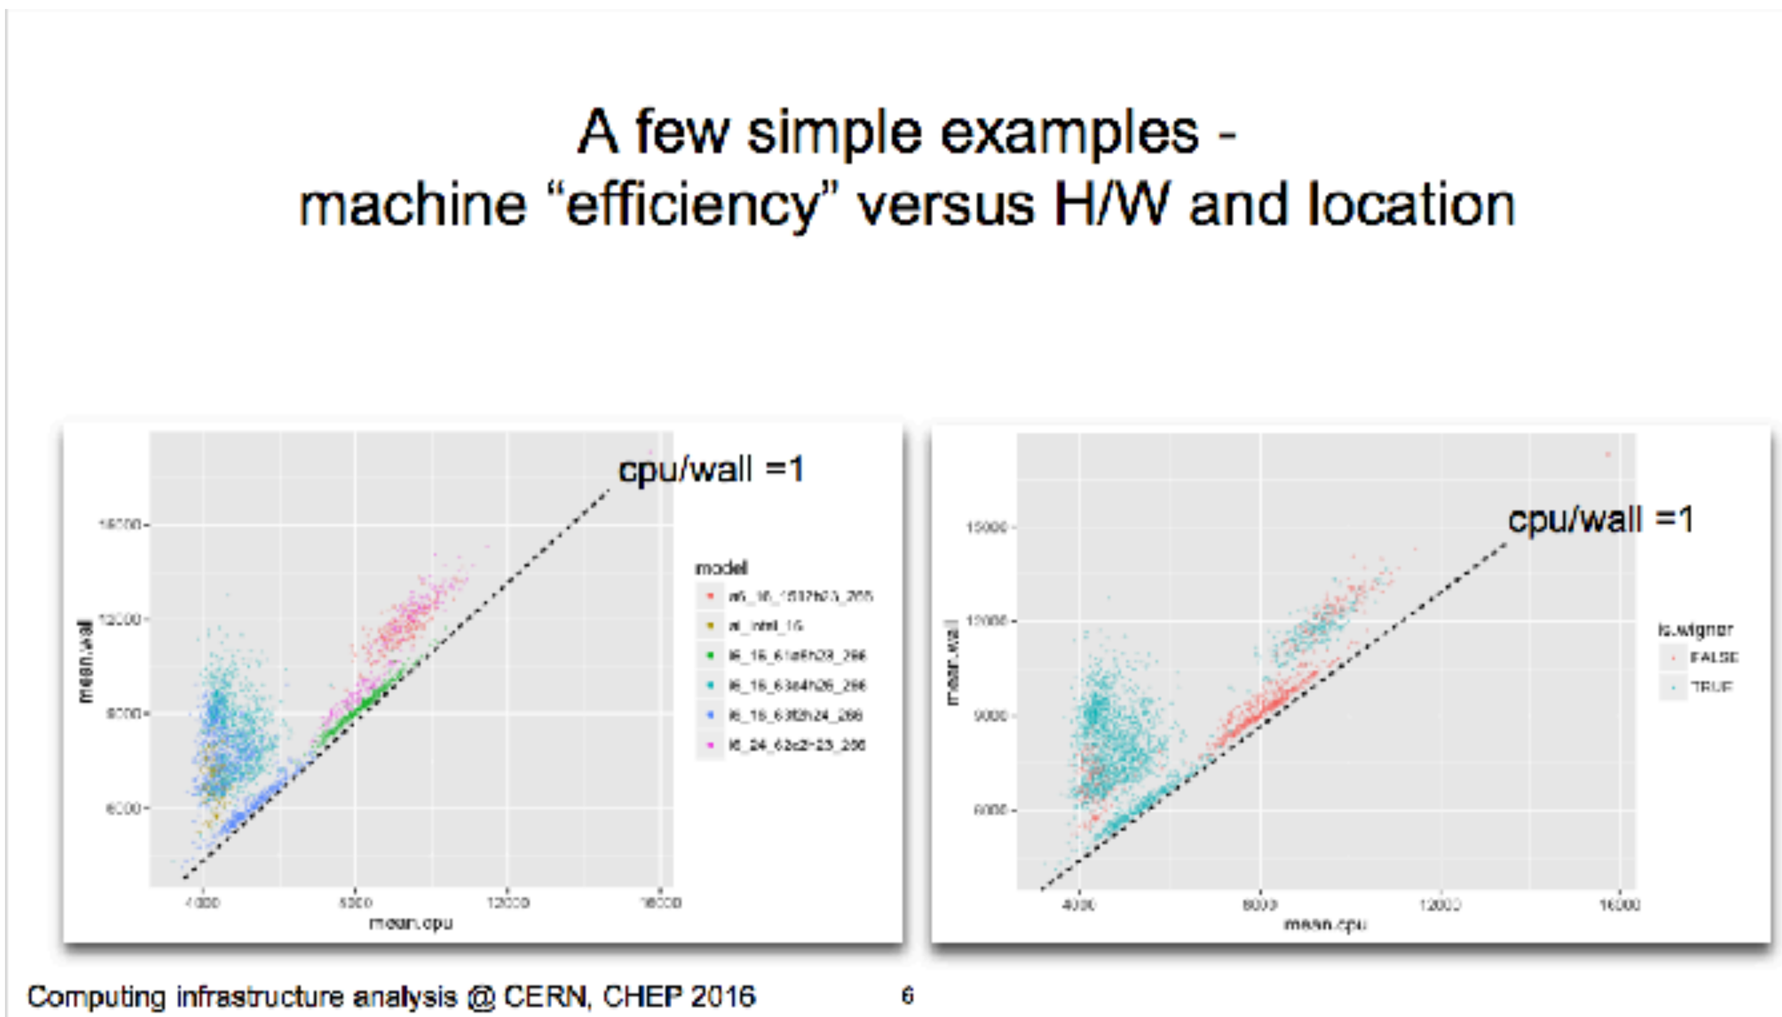

Operations analytics

1. Benchmarking cloud resources CHEP-28

Selection of benchmarks embedded in all ATLAS VMs

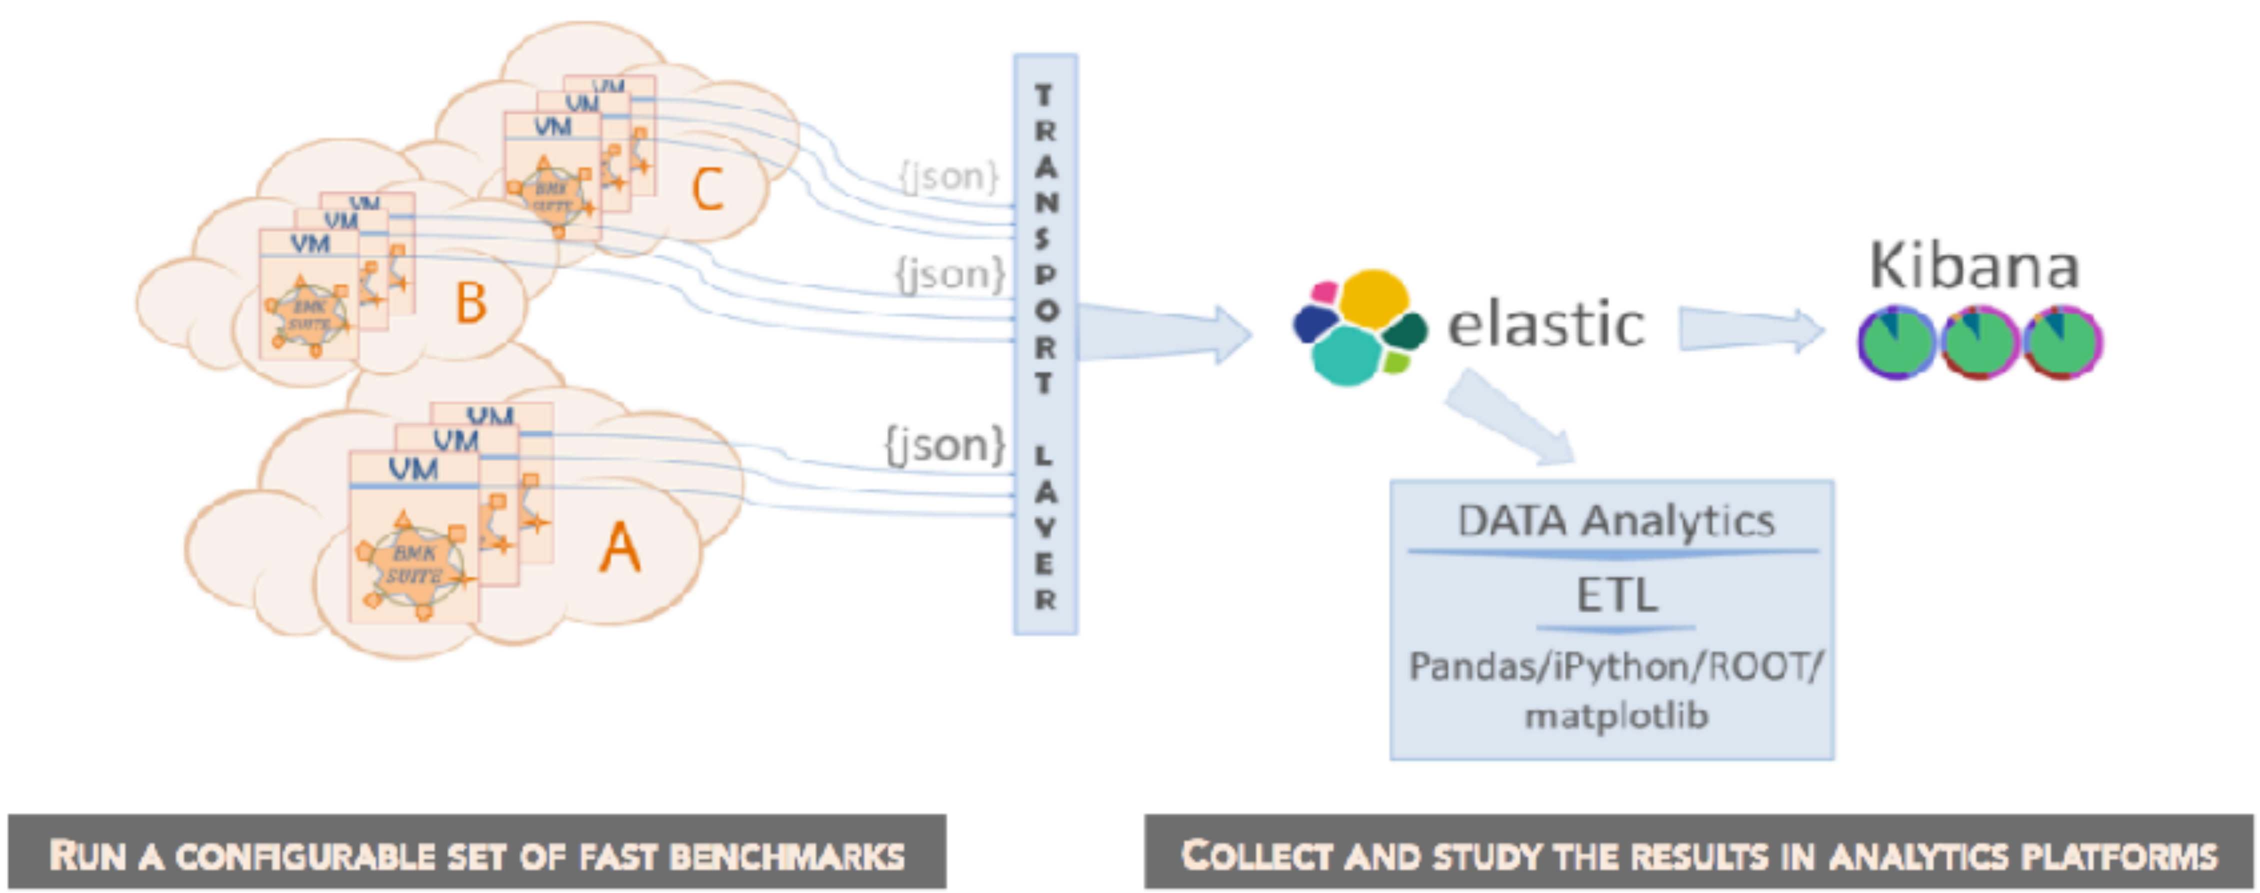

Operations analytics

2. XRootD Popularity on Hadoop Clusters CHEP-176

Analysis of dataset popularity with Hadoop with aim to predict populartity of datasets and proactively place replicas

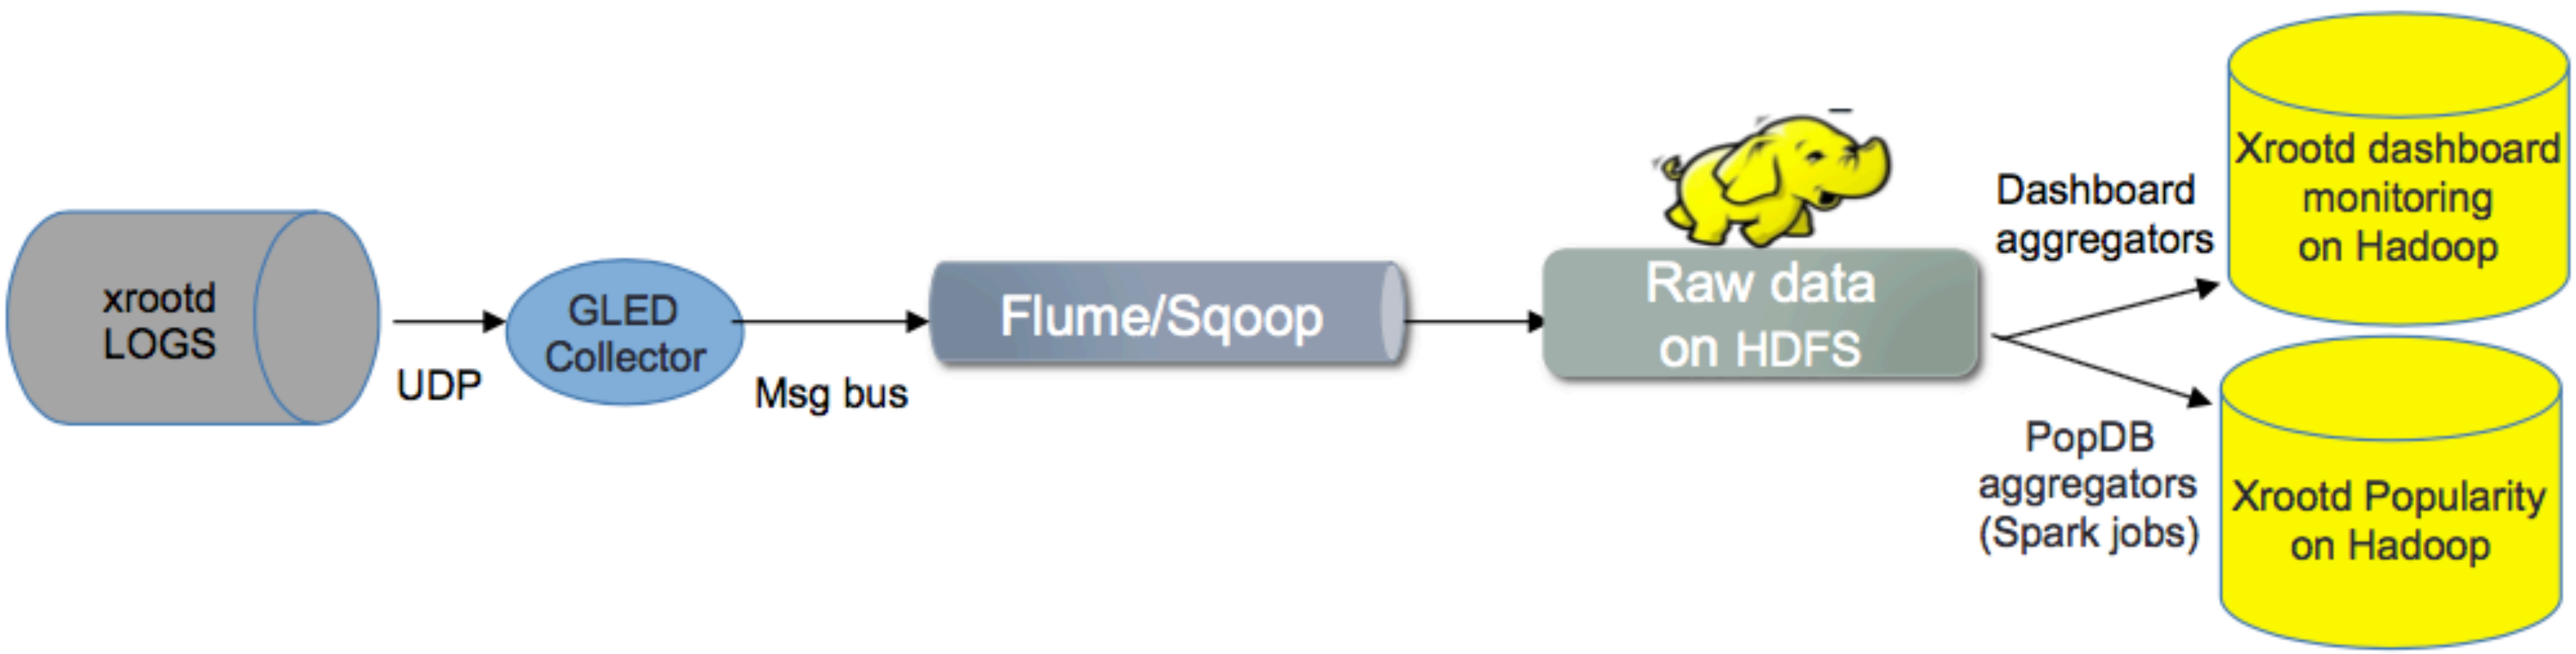

Operations analytics

3. Analysis of empty ATLAS pilot jobs CHEP-89

Quantify the wallclock time used by short empty pilot jobs on a number of WLCG compute resources. Hadoop, Jupyter platforms.

Operations analytics

4. Evolution of User Analysis on the Grid CHEP-190

Insight into analysis performance by mining data from job and data records. Elasticsearch etc.

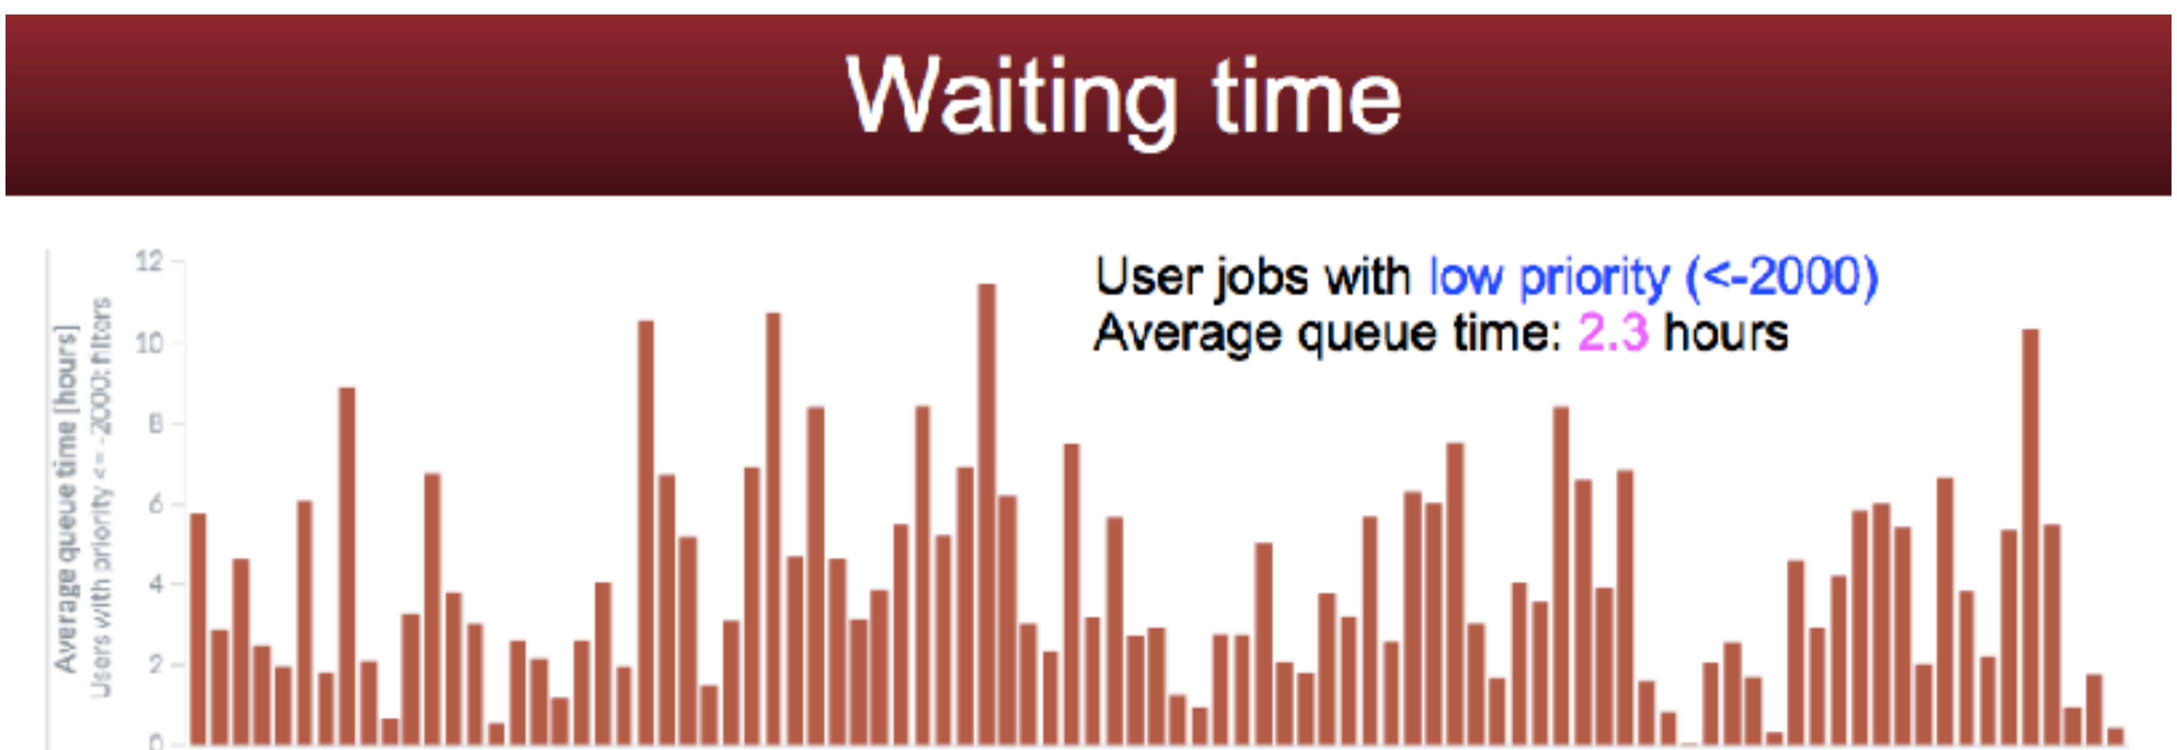

Physics analytics

1. Big Data Analytics Tools as Applied to ATLAS Event Data - CHEP-215

Mapped DxAOD event data into Elasticsearch. Looked at performance of certain operations/filtering. Visualisation useful but simplistic.

Skimming & slimming performance

ATLAS EXPERIMENT

Read test	DAOD data15_13TeV		AOD data15_13TeV		DAOD mc15_13TeV		AOD mc15_13TeV	
	Rate [Hz]	CPU* [%]	Rate [Hz]	CPU [%]	Rate [Hz]	CPU [%]	Rate [Hz]	CPU [%]
1 variable from all events	486	1	264	1	520	1	225	1
10 variables from all events	453	3	263	1	491	3	227	1
10 variables from events passing cut (−2 % events)	606 / 32 kHz	22	251 / 13 kHz	10	971 / 50 kHz	15	226 / 11 kHz	20
Full events passing cut (−2% events)	161 / 6.3 kHZ	77	35 / 1.8 kHz	84	121 / 6.2 kHz	80	32 / 1.5 kHz	82
Streaming all full events	116	86	39	87	120	88	31	86

Selections are basically free, cost is in data transfer back to client and client parsing JSON.

*client CPU usage.

12

Physics analytics

2. New Machine Learning Developments in ROOT CHEP-321

Machine Learning features in ROOT with deep learning toolkit, multithreading suitable for many-core and GPU architectures.

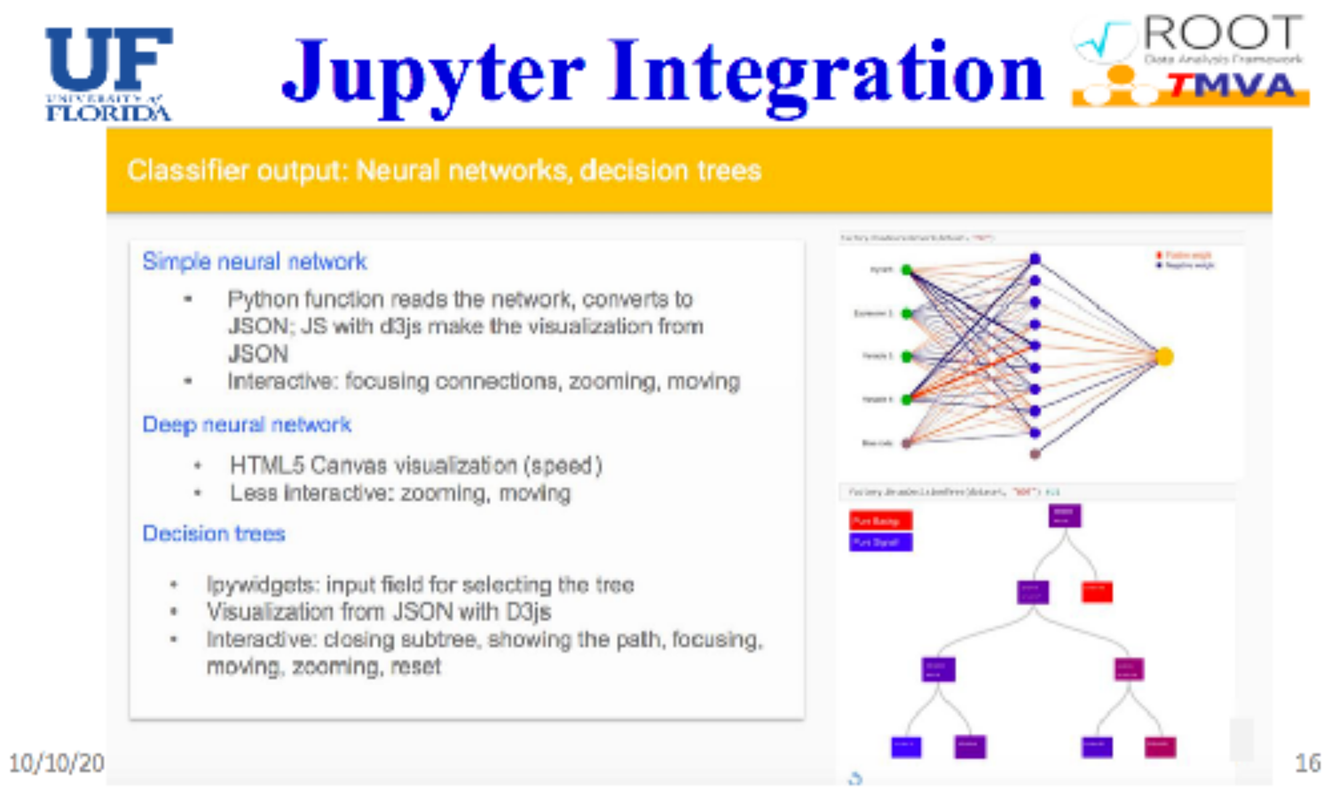

Physics analytics

3. Deep-Learning Analysis Pipelines on Raw HEP Data from the Daya Bay Neutrino Experiment at NERSC CHEP-554

Machine-learning analysis pipeline developed and operating at NERSC.

t-SNE plot: color-coded

5120 data points shown

Each point represents the neural net's encoding of one IBD or accidental event

Nearby points on this plot have similar encodings

Color represents which data set the point belongs to (IBD or accidental)

NN *was not* given this information!

Separation of red and blue indicates *neural net discovered different features* for IBD and accidentals events

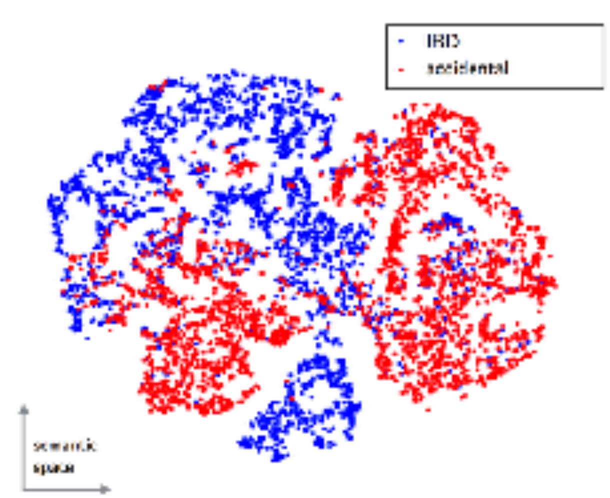

Physics analytics

4. Giving pandas ROOT to chew on: experiences with the XENON1T Dark Matter experiment CHEP-96

Modified computing model to move between ROOT and non-HEP analysis tools (e.g.,NumPy, SciPy, scikit-learn, matplotlib)

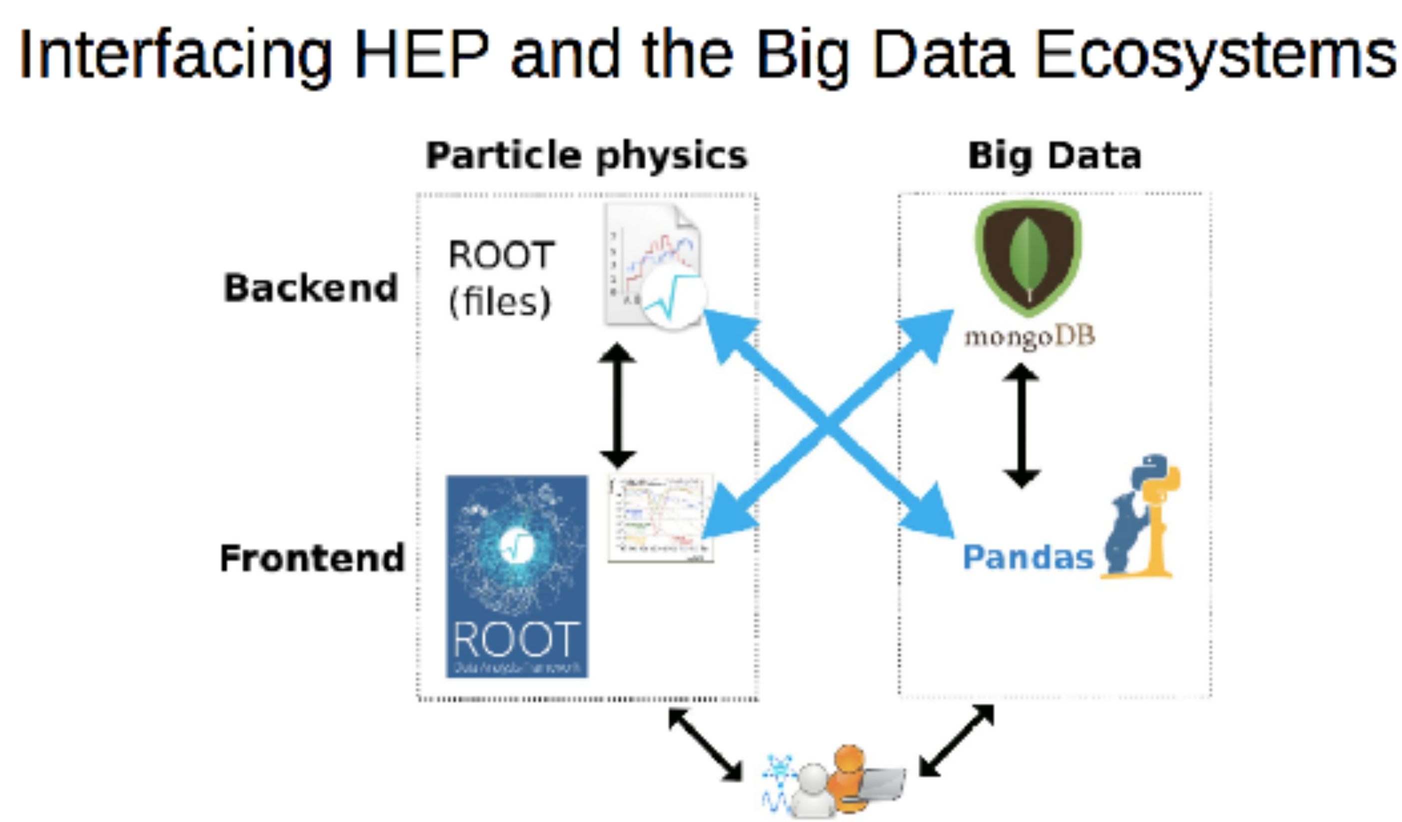

Conclusions

- On the face of it the gritty details of ATLAS analysis model is of little interest to SKA. Treat as a story from the coal face.

- ATLAS was prepared to replace large chunks of its analysis model to ensure a successful Run 2 physics programme given available resources.

- New analytics tools will help shape strategies by providing valuable insight about operational matters. Namely, how to squeeze resources and avoid wastage.

- Robust platforms will be needed, open to all, allowing researchers to explore new ideas and techniques in both computing and physics domains.

- Ask us how we’d do things now, from scratch.

We've got no money, so we've got to think. -ER

Lancaster University