O'REILLY®
Strata
CONFERENCE
Making Data Work

11–13 Nov 2013
LONDON, ENGLAND

#strataconf
strataconf.com/london

A Physics Approach to "Big Data"

Adam Kocoloski, PhD
CTO
Cloudant


Cloudant

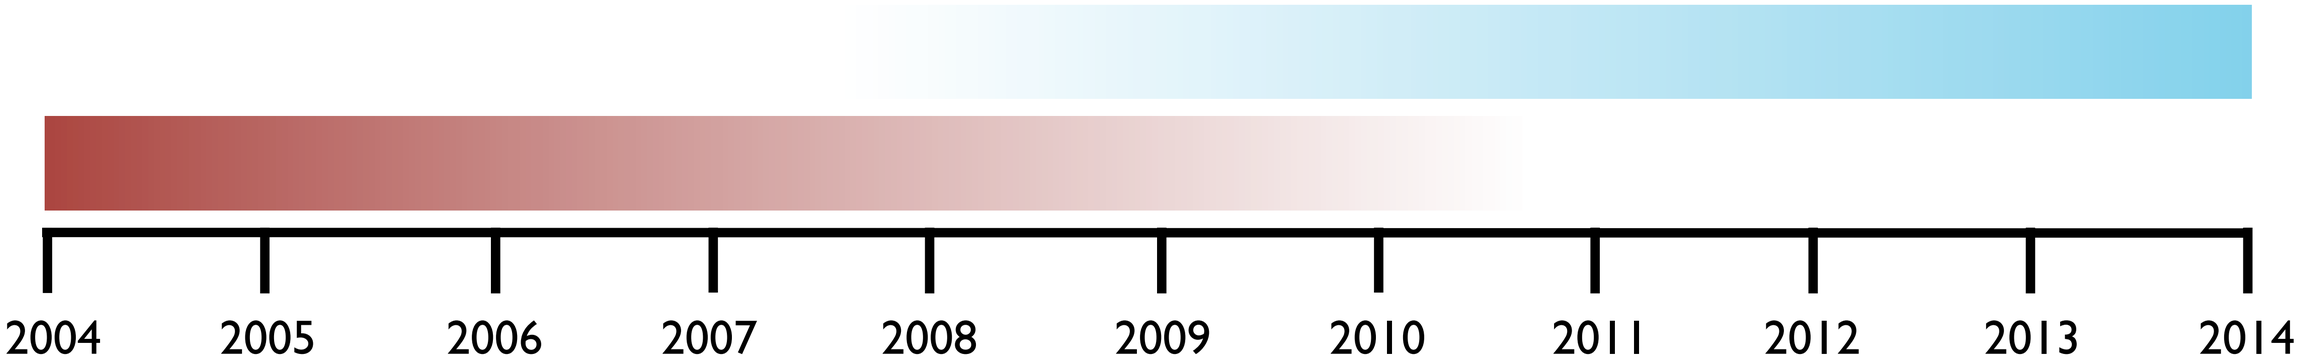
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014

Solenoidal Tracker at RHIC (STAR)

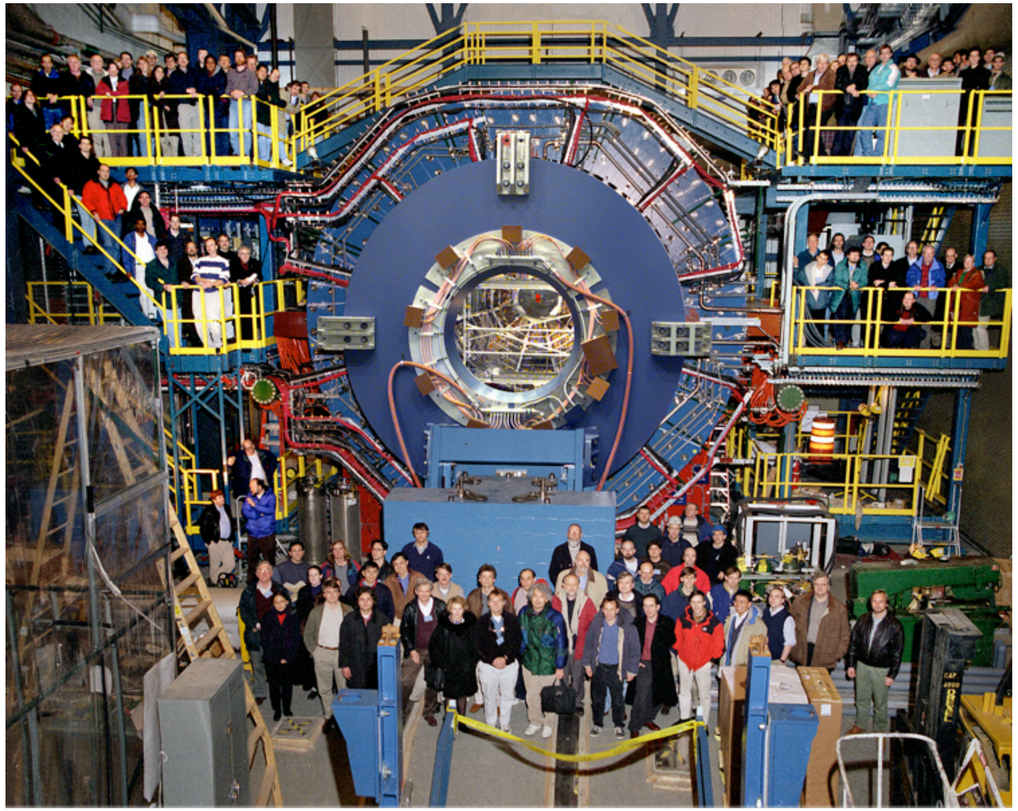

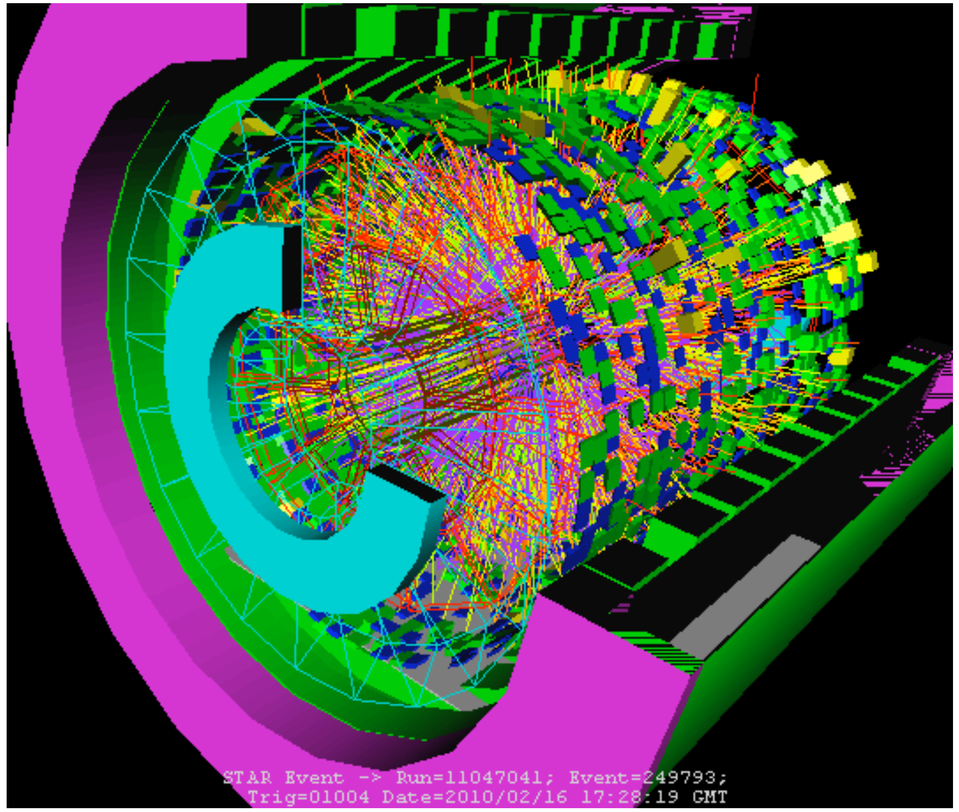

The life of LHC data

- **Detected by experiment**
- “Online” filtering (hardware and software)
- Transferred to CERN main campus, archived & reconstructed
- Transferred to T1 sites, archived, reconstructed & skimmed
- Transferred to T2 sites, reconstructed, skimmed, filtered & analysed
- Written into locally analyzable files, put on laptops
- Turned into a plot in a paper


Dig big tunnels

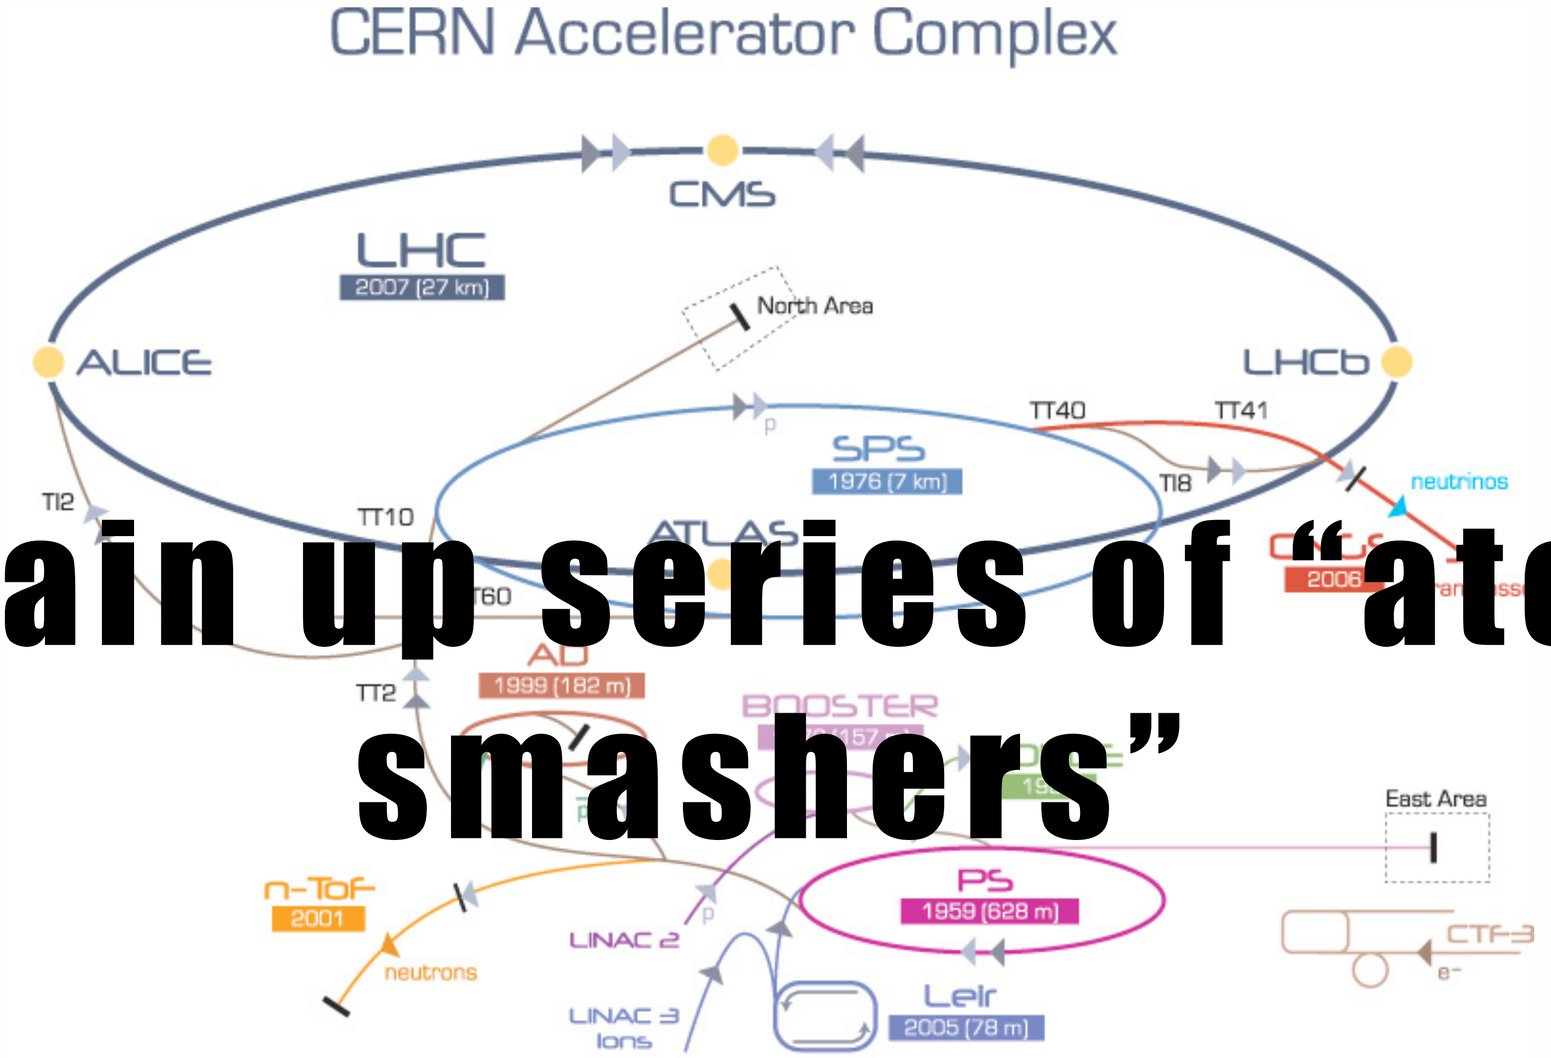

Chain up series of “atom smashers”

Put sensitive cameras in awkward places

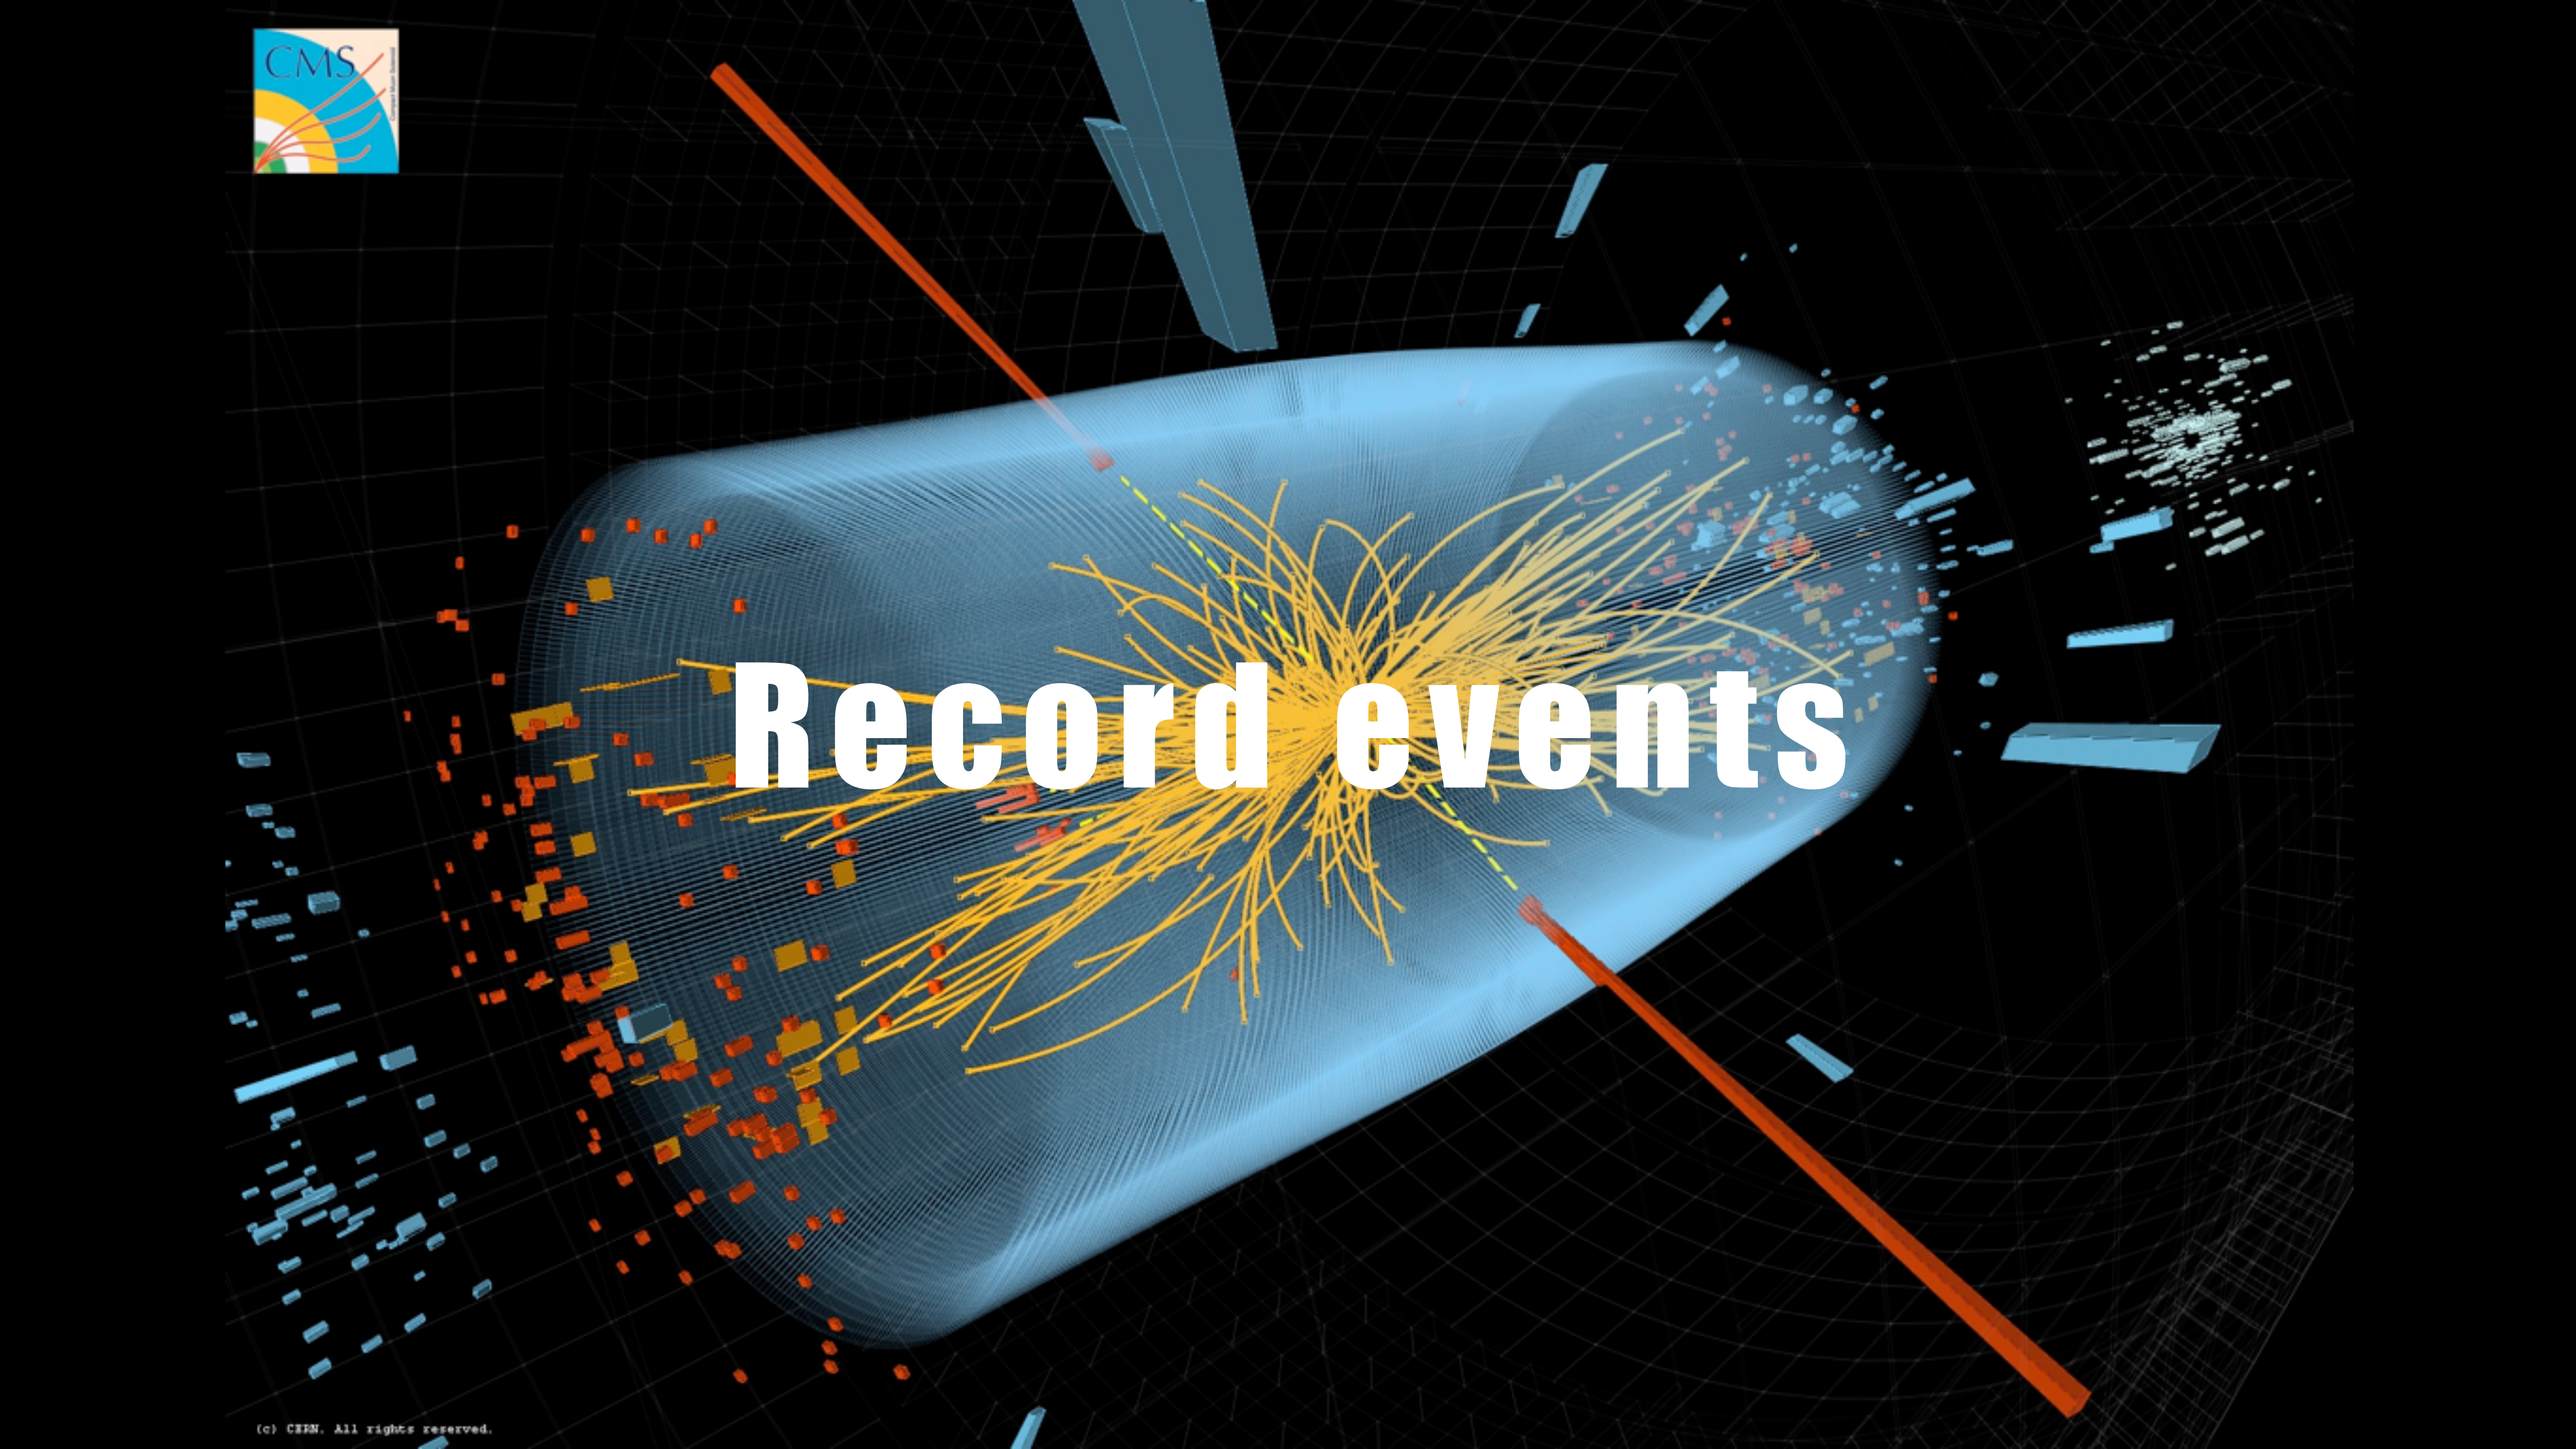

(c) CERN. All rights reserved.

The life of LHC data

- Detected by experiment
- **“Online” filtering (hardware and software)**
- Transferred to CERN main campus, archived & reconstructed
- Transferred to T1 sites, archived, reconstructed & skimmed
- Transferred to T2 sites, reconstructed, skimmed, filtered & analysed
- Written into locally analyzable files, put on laptops
- Turned into a plot in a paper

CMS Data Flow

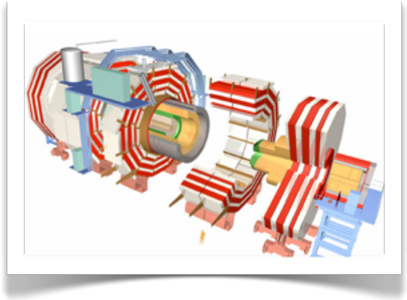

We have a big digital camera

It takes photos of this

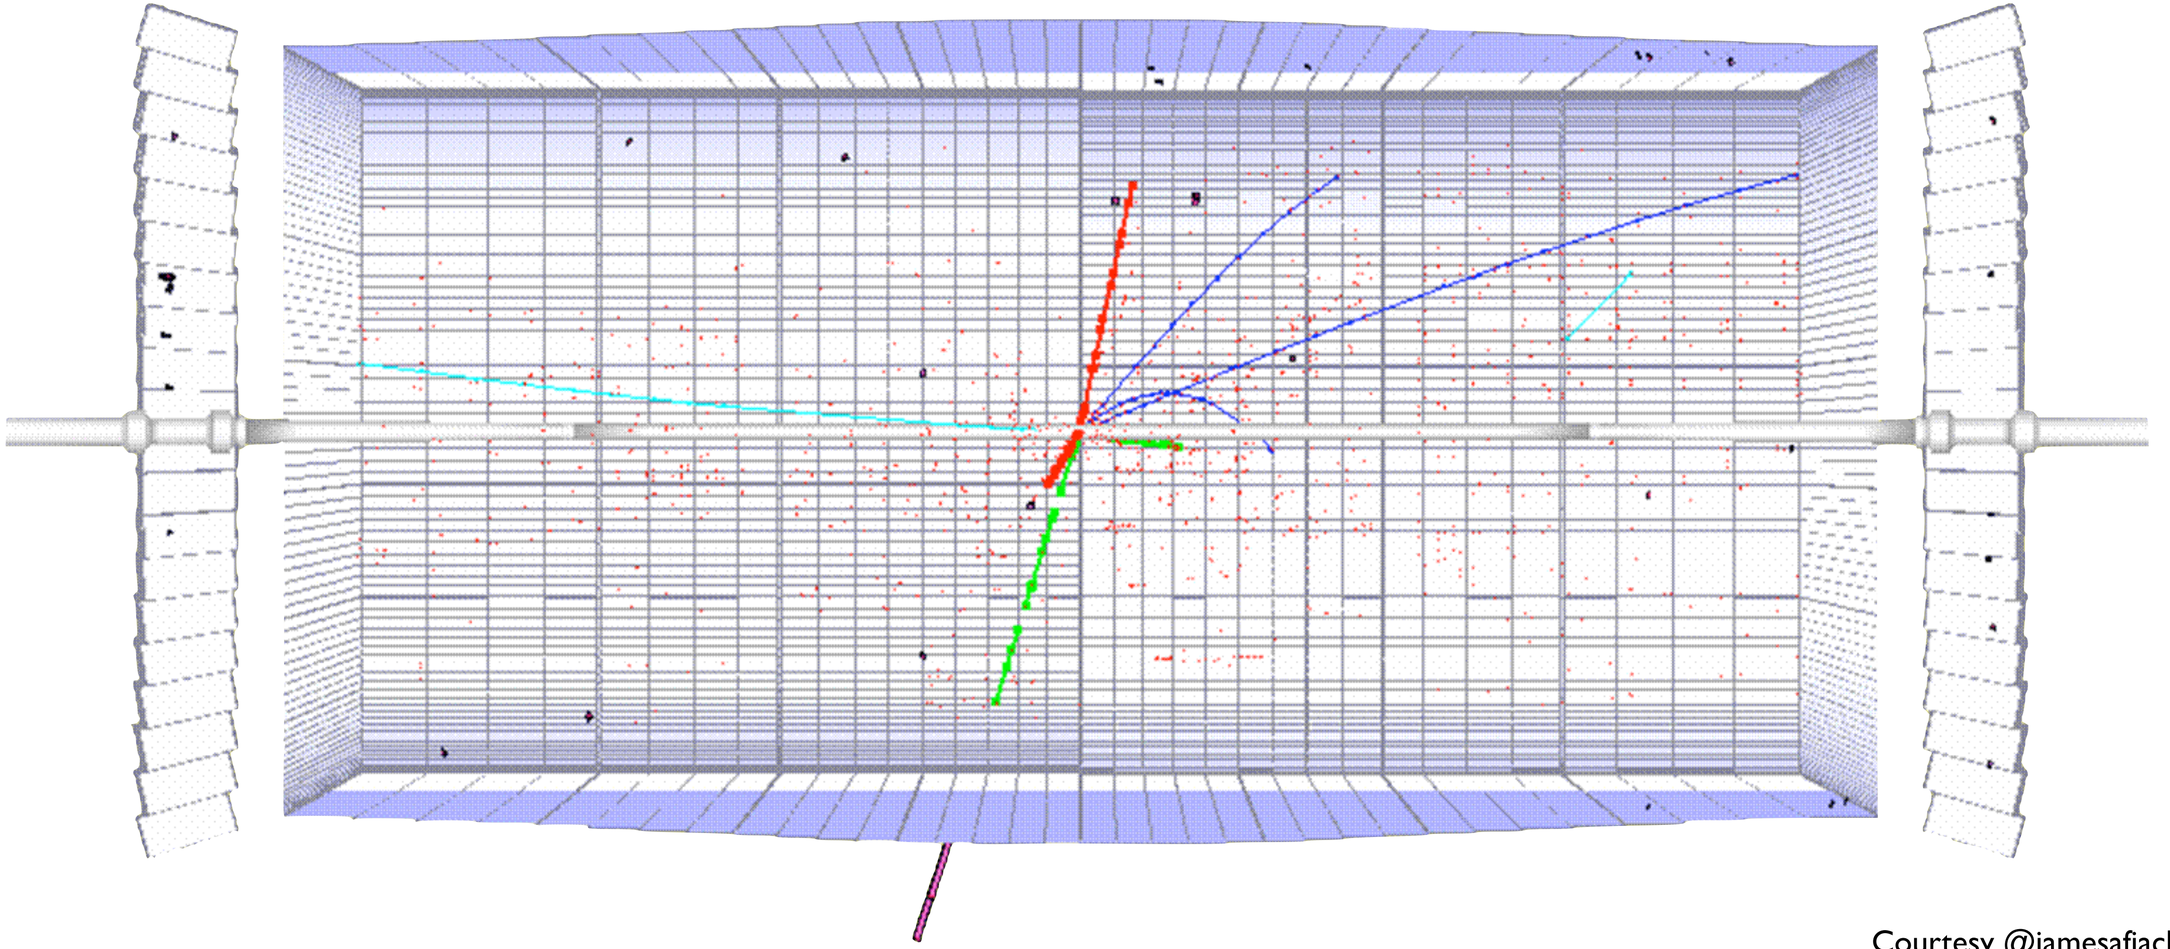

Courtesy @jamesafjackson

Which come out looking like this

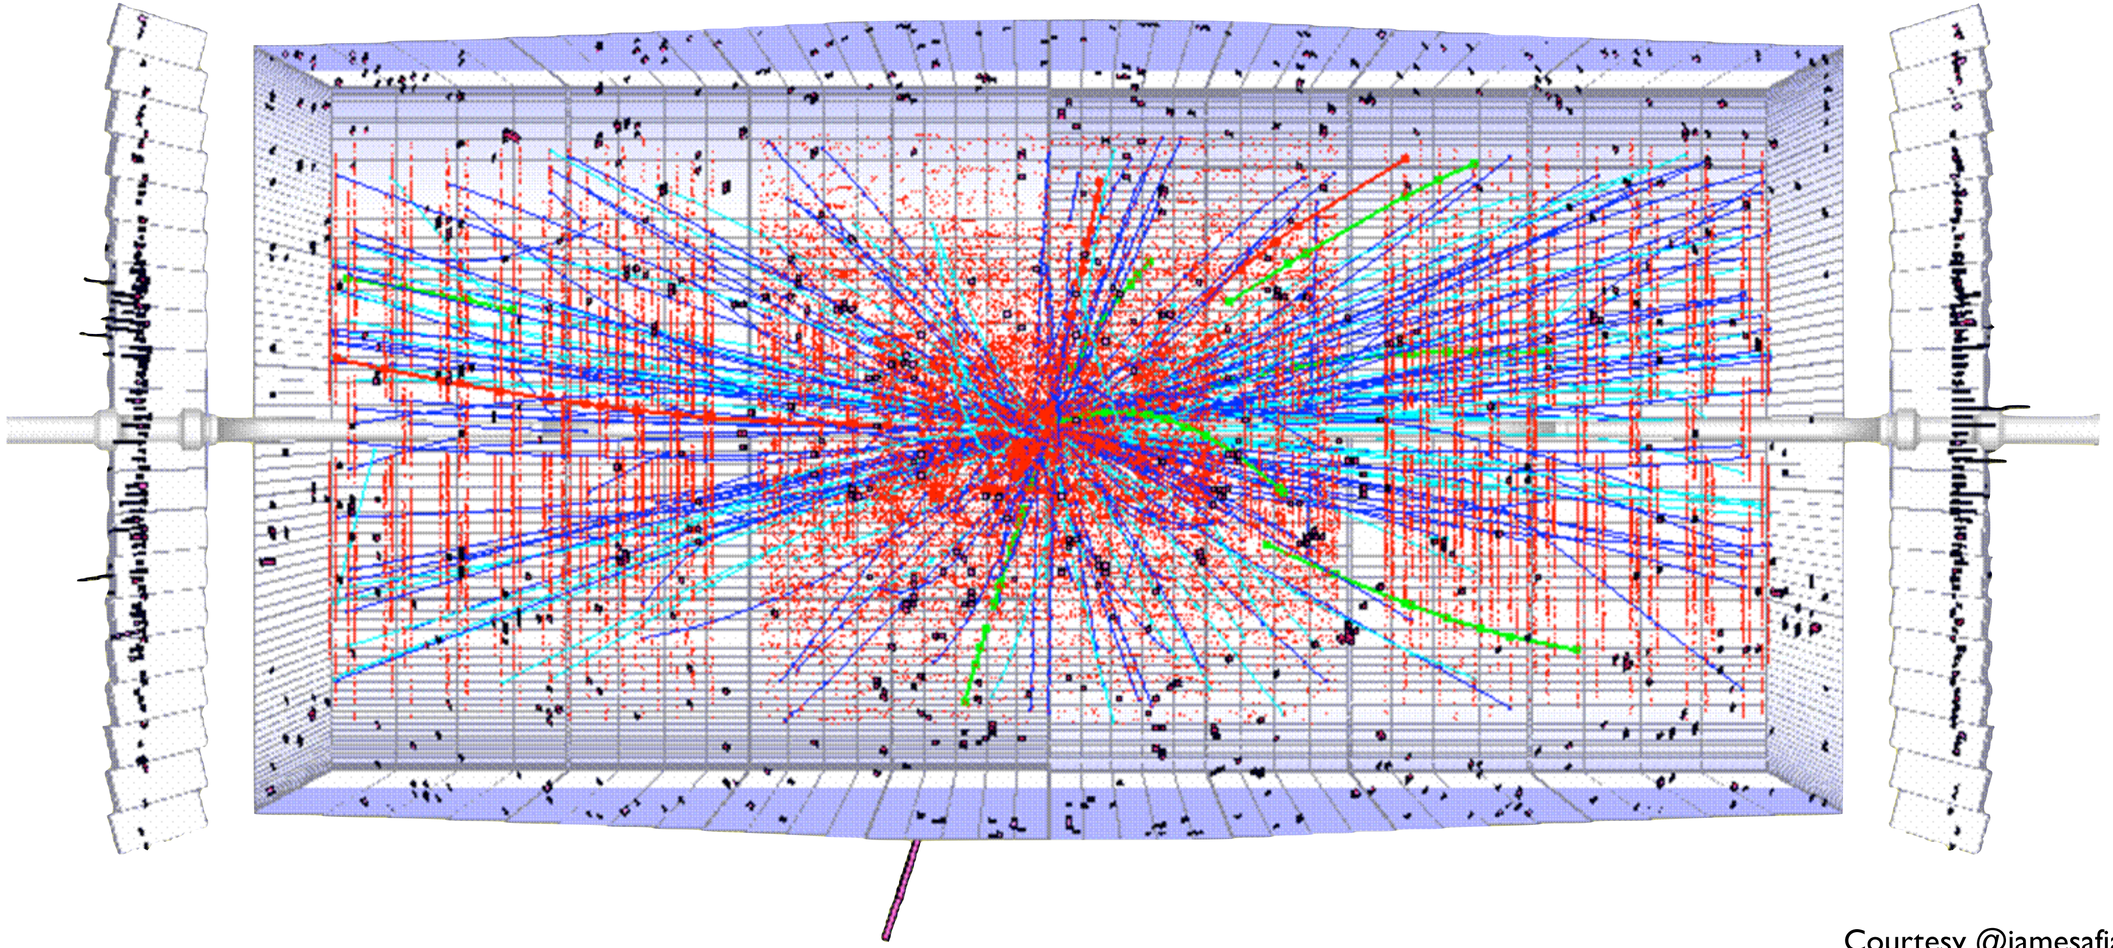

Courtesy @jamesafjackson

CMS Data Flow

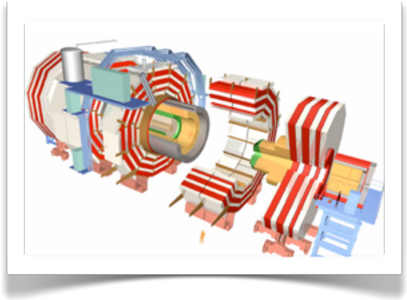

We have a big digital camera

CMS Data Flow

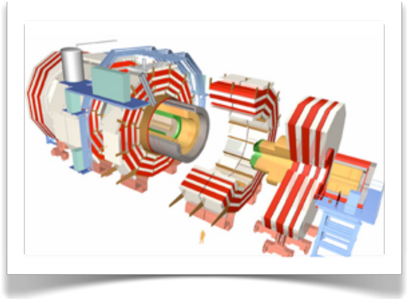

We have a big digital camera

Which goes into lots of computers (HLT)

CMS Data Flow

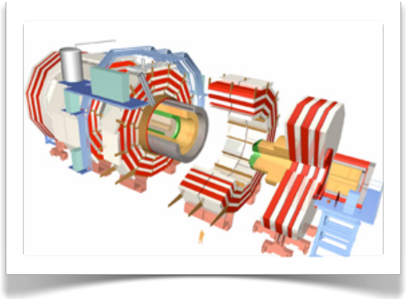

We have a big digital camera

Which goes into lots of computers (HLT)

And then into lots of disk (Storage Manager)

CMS Data Flow

We have a big digital camera

~200 GB/s

Which goes into lots of computers (HLT)

~2 GB/s

And then into lots of disk (Storage Manager)

Throw away data ASAP

The life of LHC data

- Detected by experiment
- “Online” filtering (hardware and software)
- **Transferred to CERN main campus, archived & reconstructed**
- Transferred to T1 sites, archived, reconstructed & skimmed
- Transferred to T2 sites, reconstructed, skimmed, filtered & analysed
- Written into locally analyzable files, put on laptops
- Turned into a plot in a paper

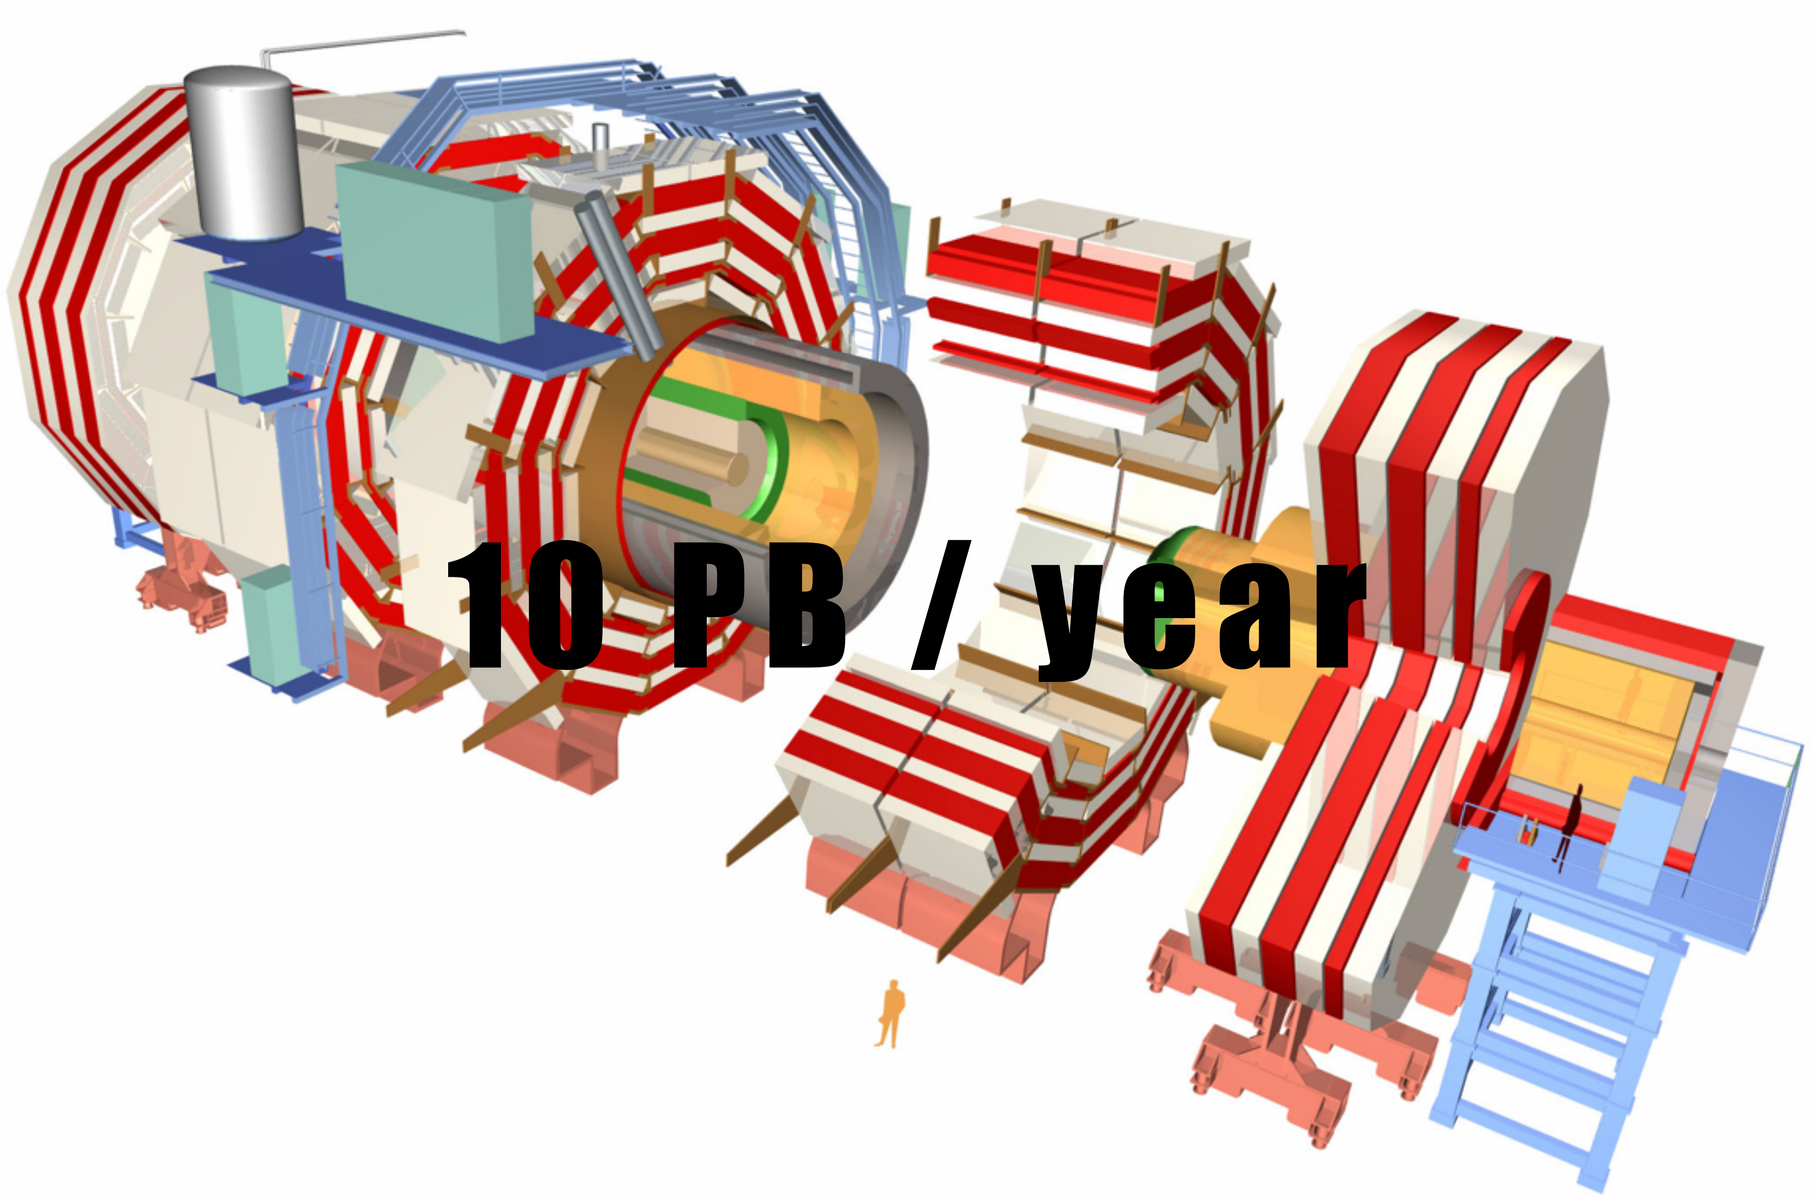
10 PB / year

The life of LHC data

- Detected by experiment
- “Online” filtering (hardware and software)
- Transferred to CERN main campus, archived & reconstructed
- **Transferred to T1 sites, archived, reconstructed & skimmed**
- **Transferred to T2 sites, reconstructed, skimmed, filtered & analysed**
- Written into locally analyzable files, put on laptops
- Turned into a plot in a paper

Tiers

- Maybe a bit different
- Many (>100) sites with 100's TB storage, 10000's worker nodes
- Why so many? Politics, power budget, cost

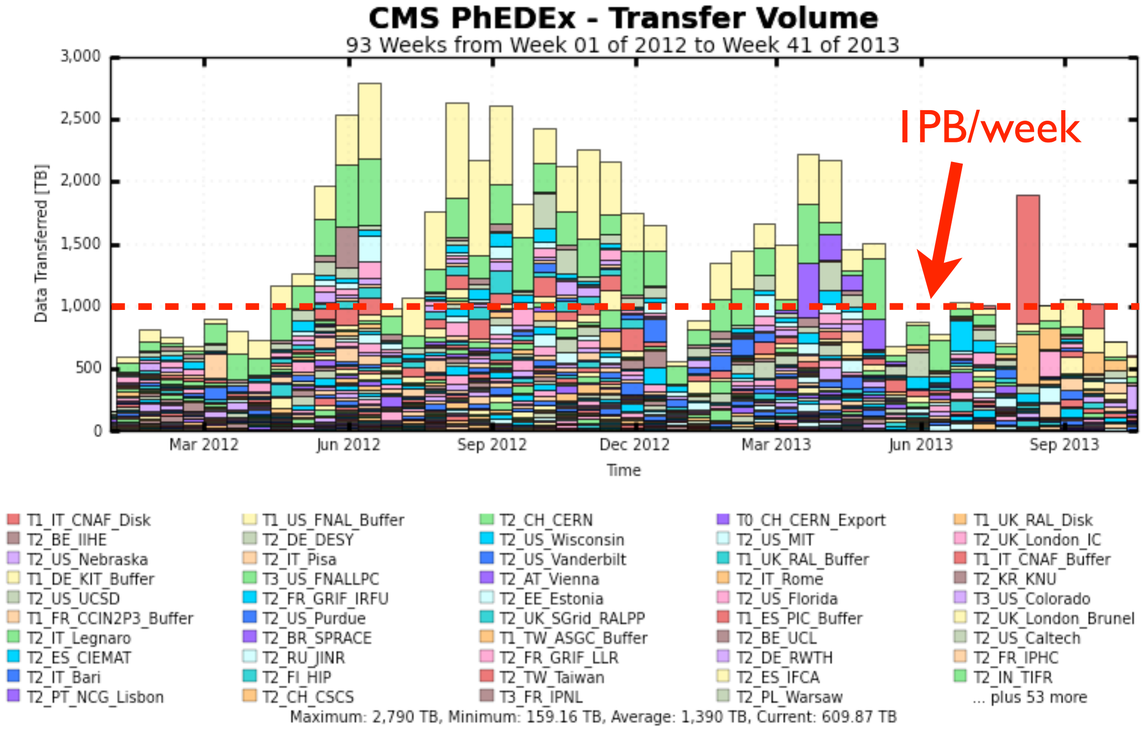

The life of LHC data

- Detected by experiment
- “Online” filtering (hardware and software)
- Transferred to CERN main campus, archived & reconstructed
- Transferred to T1 sites, archived, reconstructed & skimmed
- Transferred to T2 sites, reconstructed, skimmed, filtered & analysed
- **Written into locally analyzable files, put on laptops**
- **Turned into a plot in a paper**

Analysis Workflow

- Each analysis is ~unique
- Query language is C++
- Runs on distributed and local resources
- Myriad “cut” selections to identify interesting events
- Data in final plot substantially reduced from the original dataset

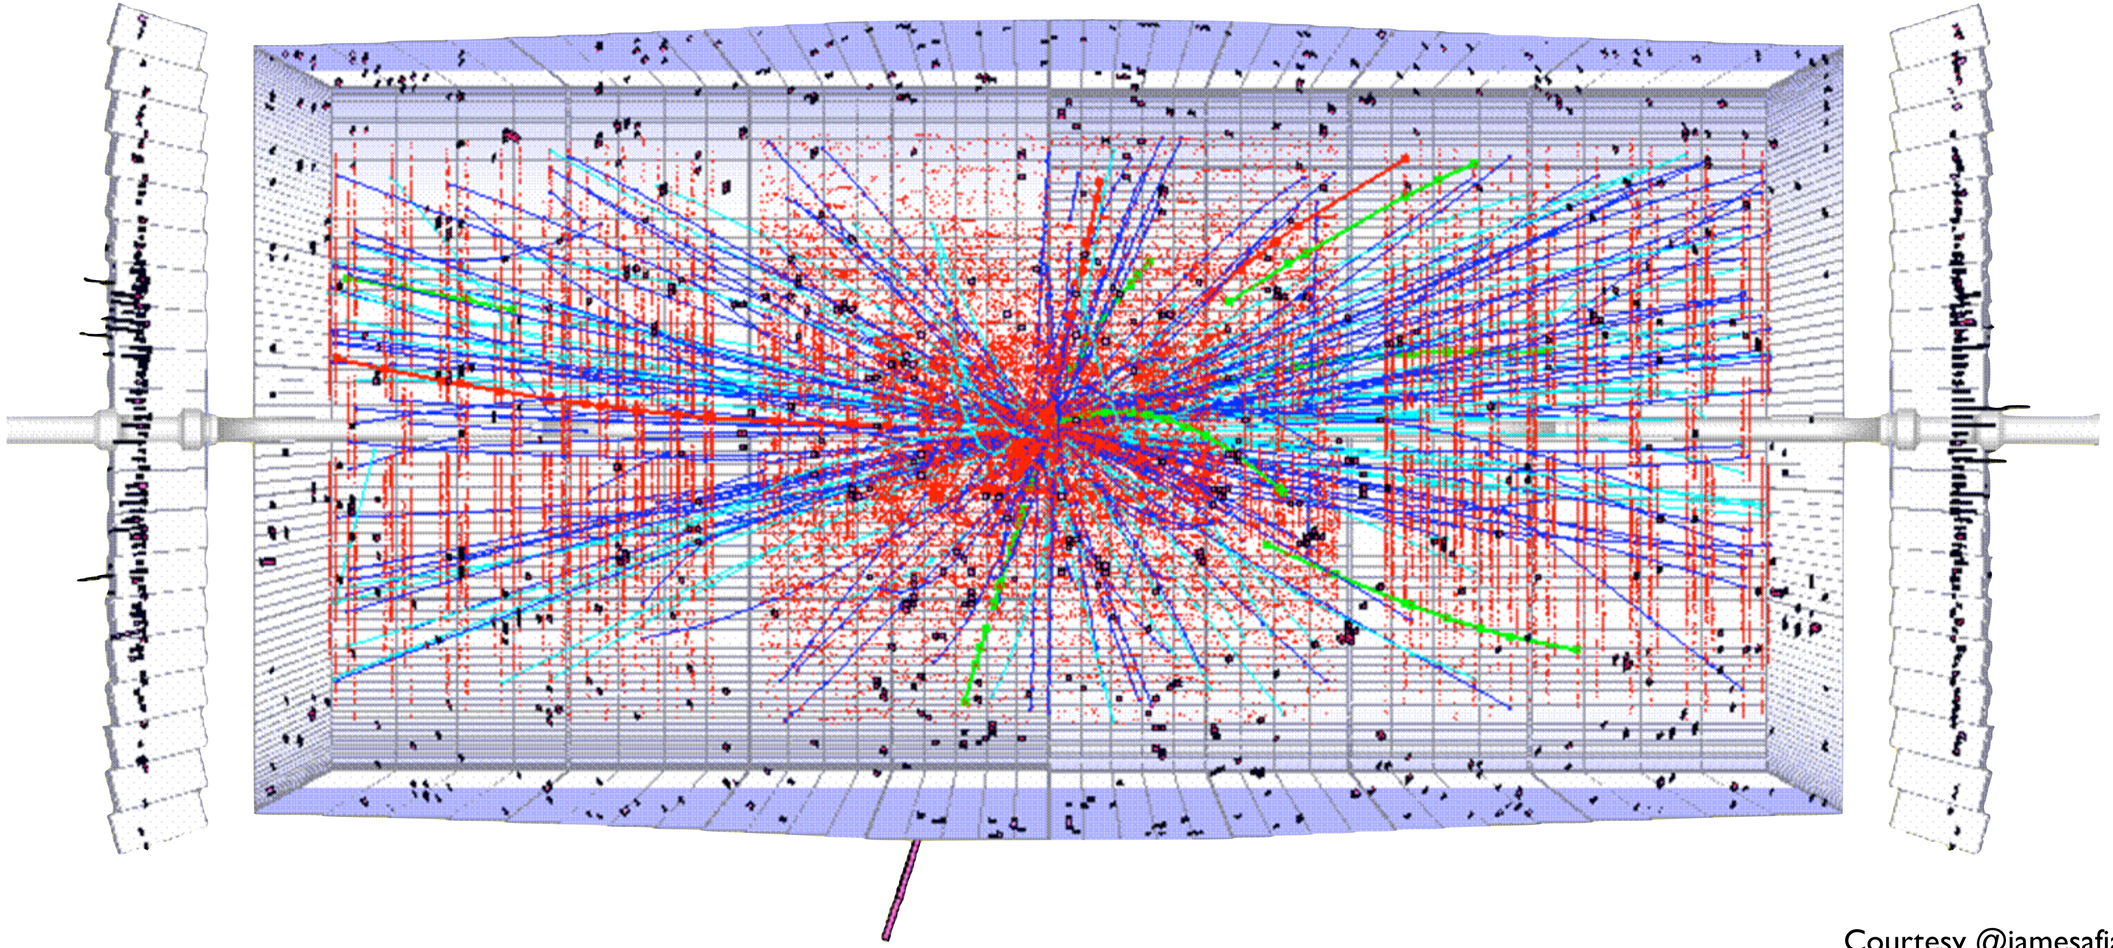

Courtesy @jamesafjackson

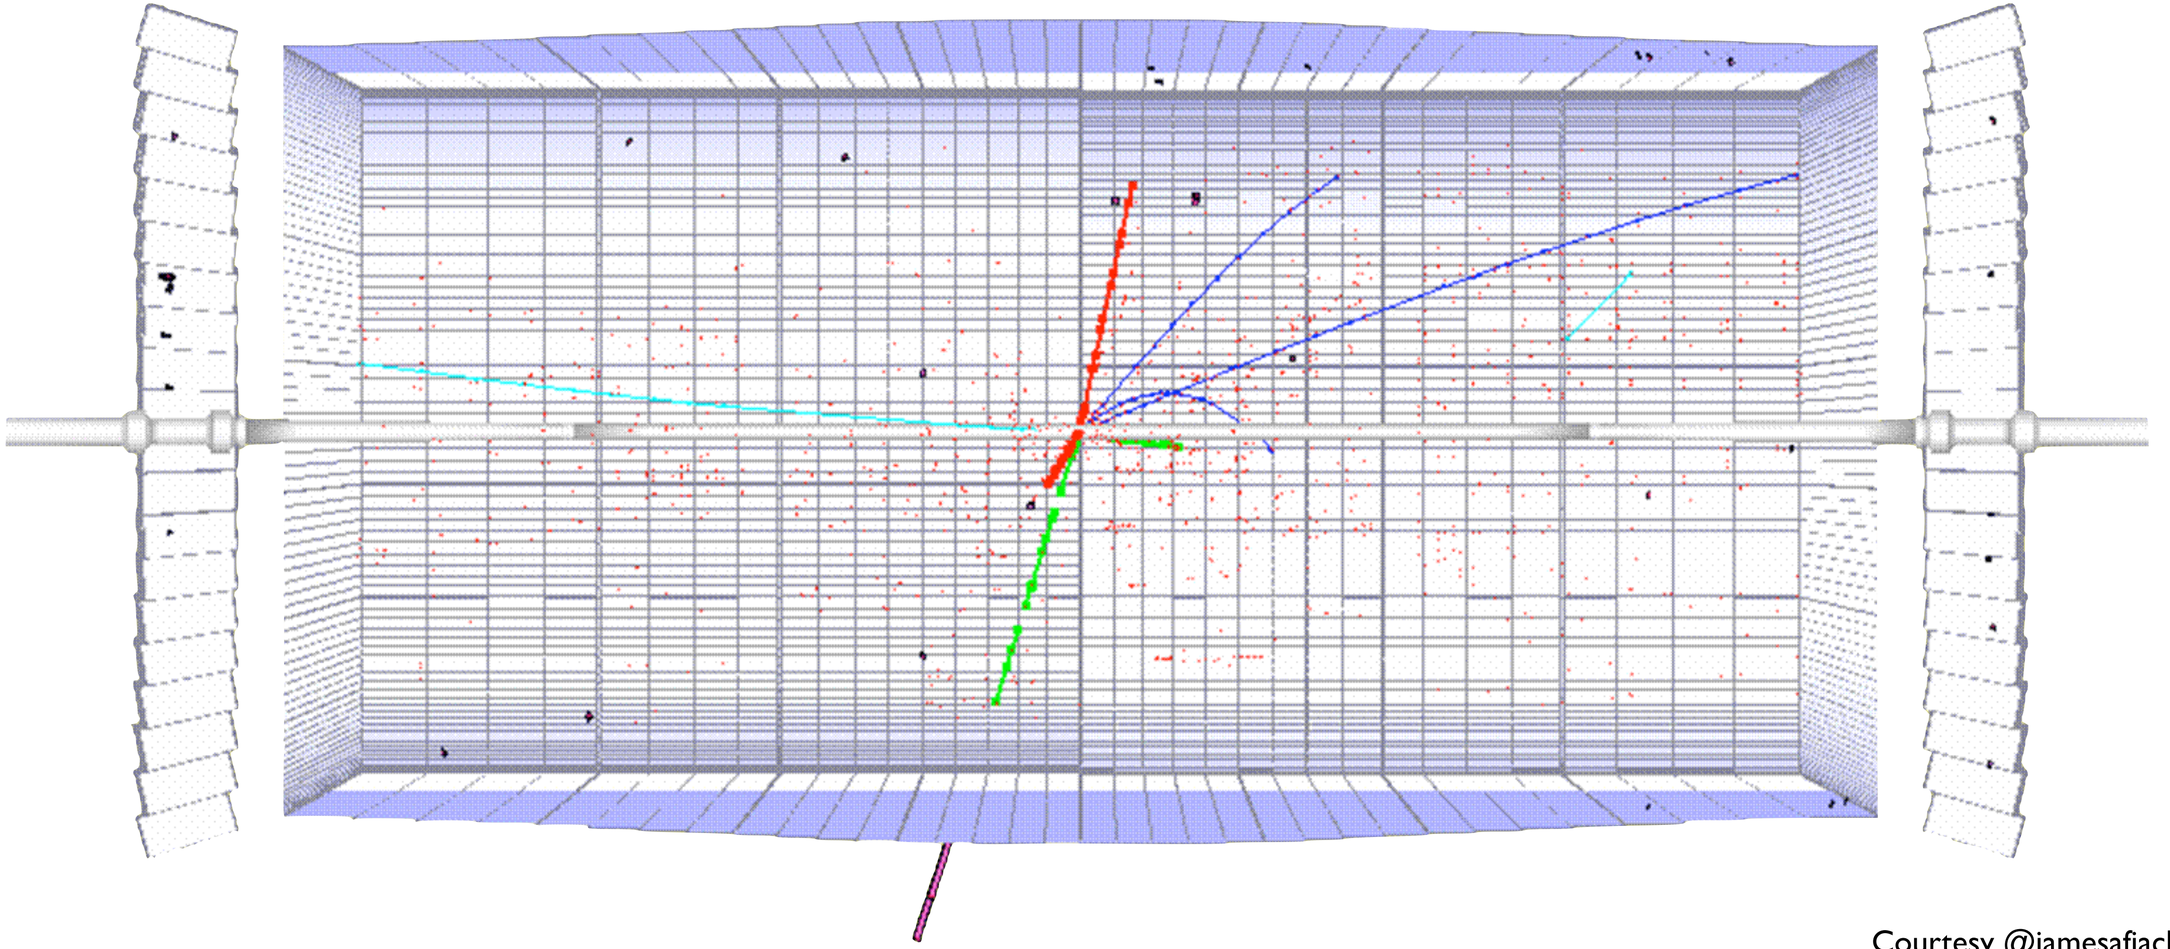

Courtesy @jamesafjackson

Multivariate analyses extract signals that are impossible to find via “hand-drawn” cuts

Humans find signals even when none exist

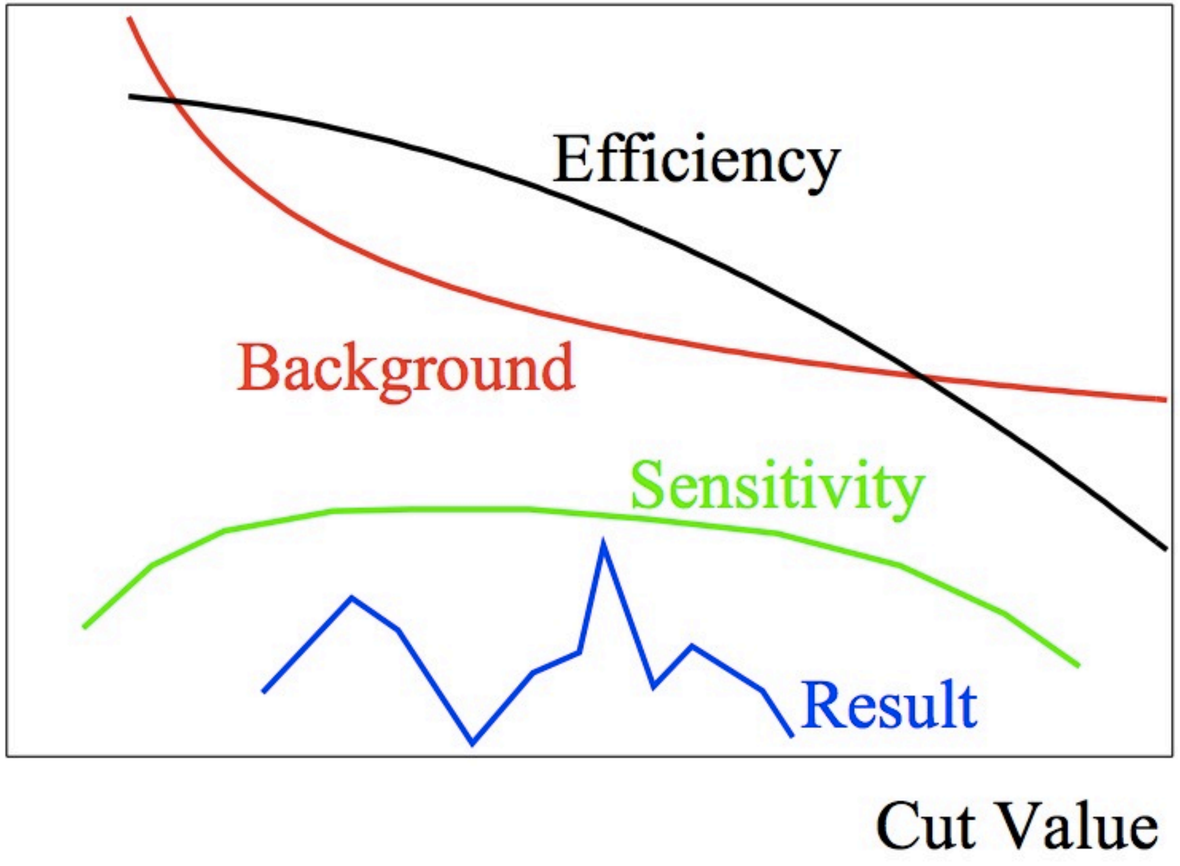

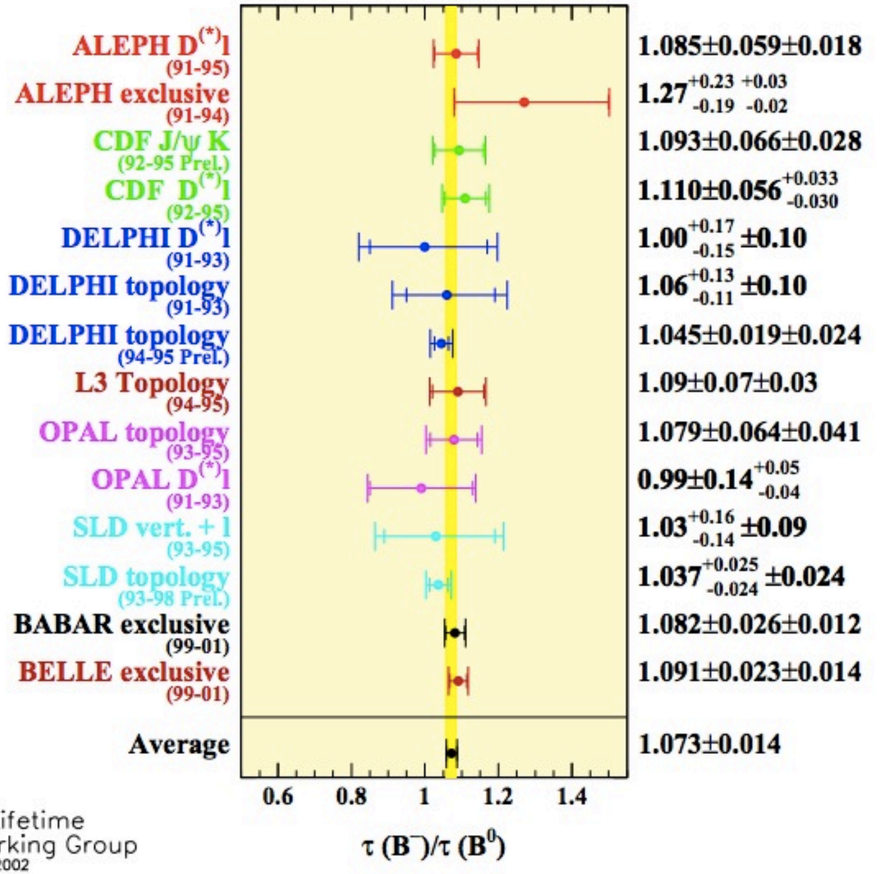

http://www.slac.stanford.edu/econf/C030908/papers/TUIT001.pdf

Workflow Ladder

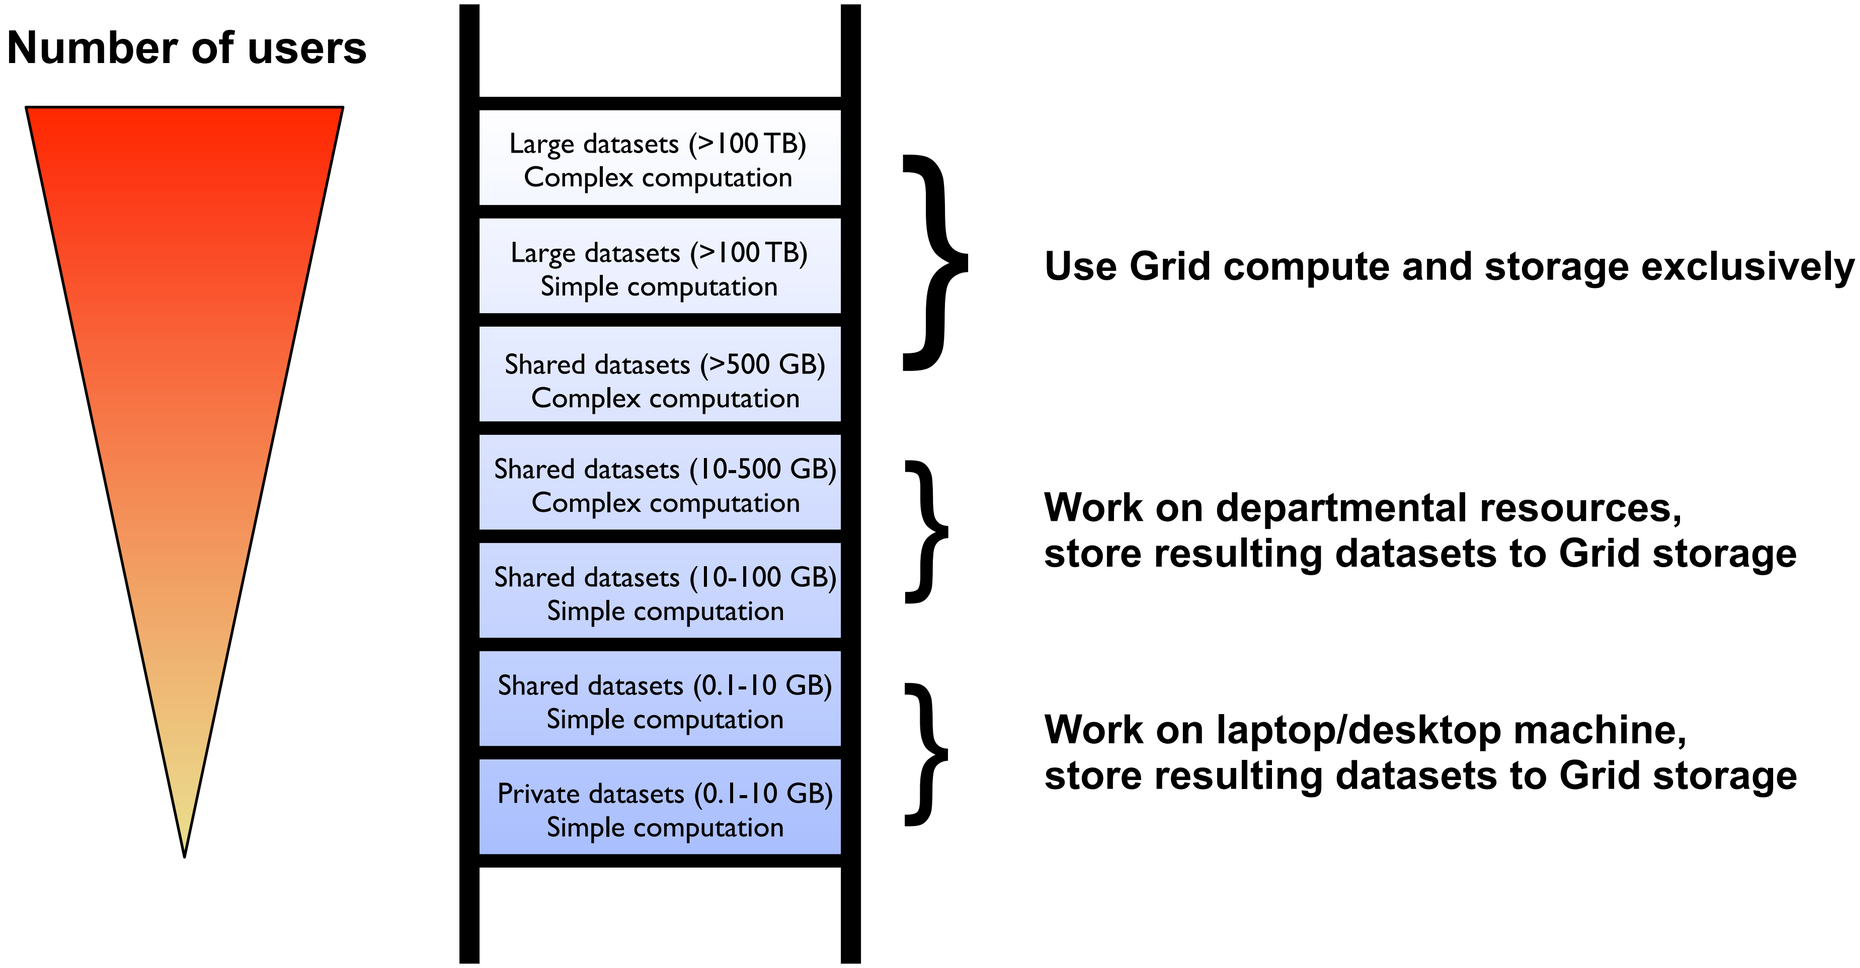

Collaborative “skimming” remains valuable

Let’s Review

Let’s Review

- Particle physics is special

Let’s Review

- Particle physics is special
- But not as special as it used to be vis a vis data operations

Let’s Review

- Particle physics is special
- But not as special as it used to be vis a vis data operations
- And that’s a good thing

Let’s Review

- Particle physics is special
- But not as special as it used to be vis a vis data operations
- And that’s a good thing
- Keep only the data you can analyze

Let’s Review

- Particle physics is special
- But not as special as it used to be vis a vis data operations
- And that’s a good thing
- Keep only the data you can analyze
- Remember that we’re biased to find patterns, and lots of data means lots of patterns

Let’s Review

- Particle physics is special
- But not as special as it used to be vis a vis data operations
- And that’s a good thing
- Keep only the data you can analyze
- Remember that we’re biased to find patterns, and lots of data means lots of patterns
- Multivariate analyses win in the end

Let’s Review

- Particle physics is special
- But not as special as it used to be vis a vis data operations
- And that’s a good thing
- Keep only the data you can analyze
- Remember that we’re biased to find patterns, and lots of data means lots of patterns
- Multivariate analyses win in the end
- Skim datasets for fun and profit

Thanks!

@kocolosk

adam@cloudant.com