Big Data Needs High Energy Physics especially the LHC

Richard P Mount
SLAC National Accelerator Laboratory

June 27, 2013

Why so much data?

- Our universe seems to be governed by non-deterministic physics
  - One measurement tells us very little
  - However carefully we set up an experiment, non-deterministic physics decides what we observe
  - If we want to observe something rare, we may have to find a few occurrences (events) hidden in vast numbers of other events
  - The LHC experiments "see" bunch collisions (tens or hundreds of superimposed events) every 25 ns. That is 40 million/second or about 15 trillion bunch collisions per year

Example – Finding the Higgs CMS and ATLAS at the LHC

CMS Higgs Candidate

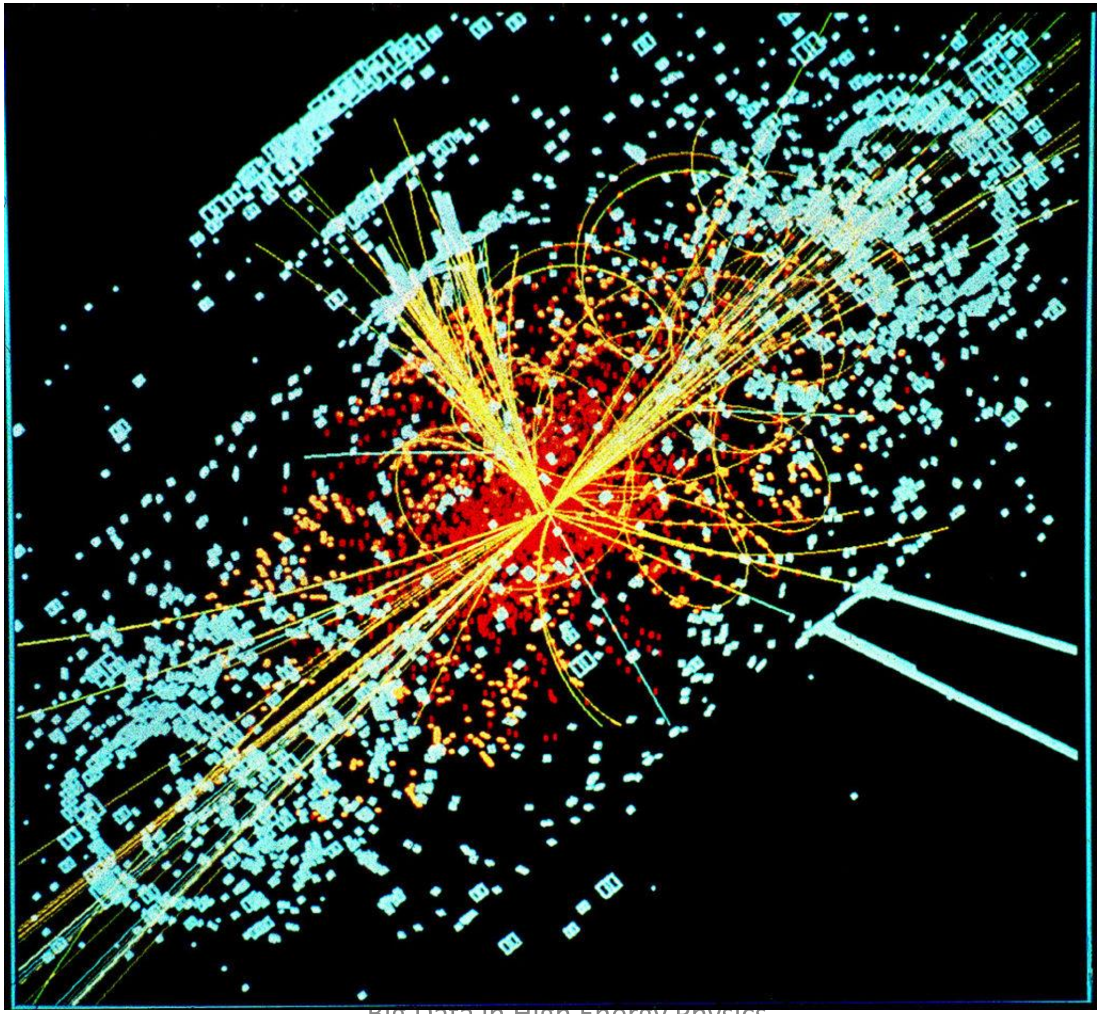

ATLAS
Higgs
Candidates

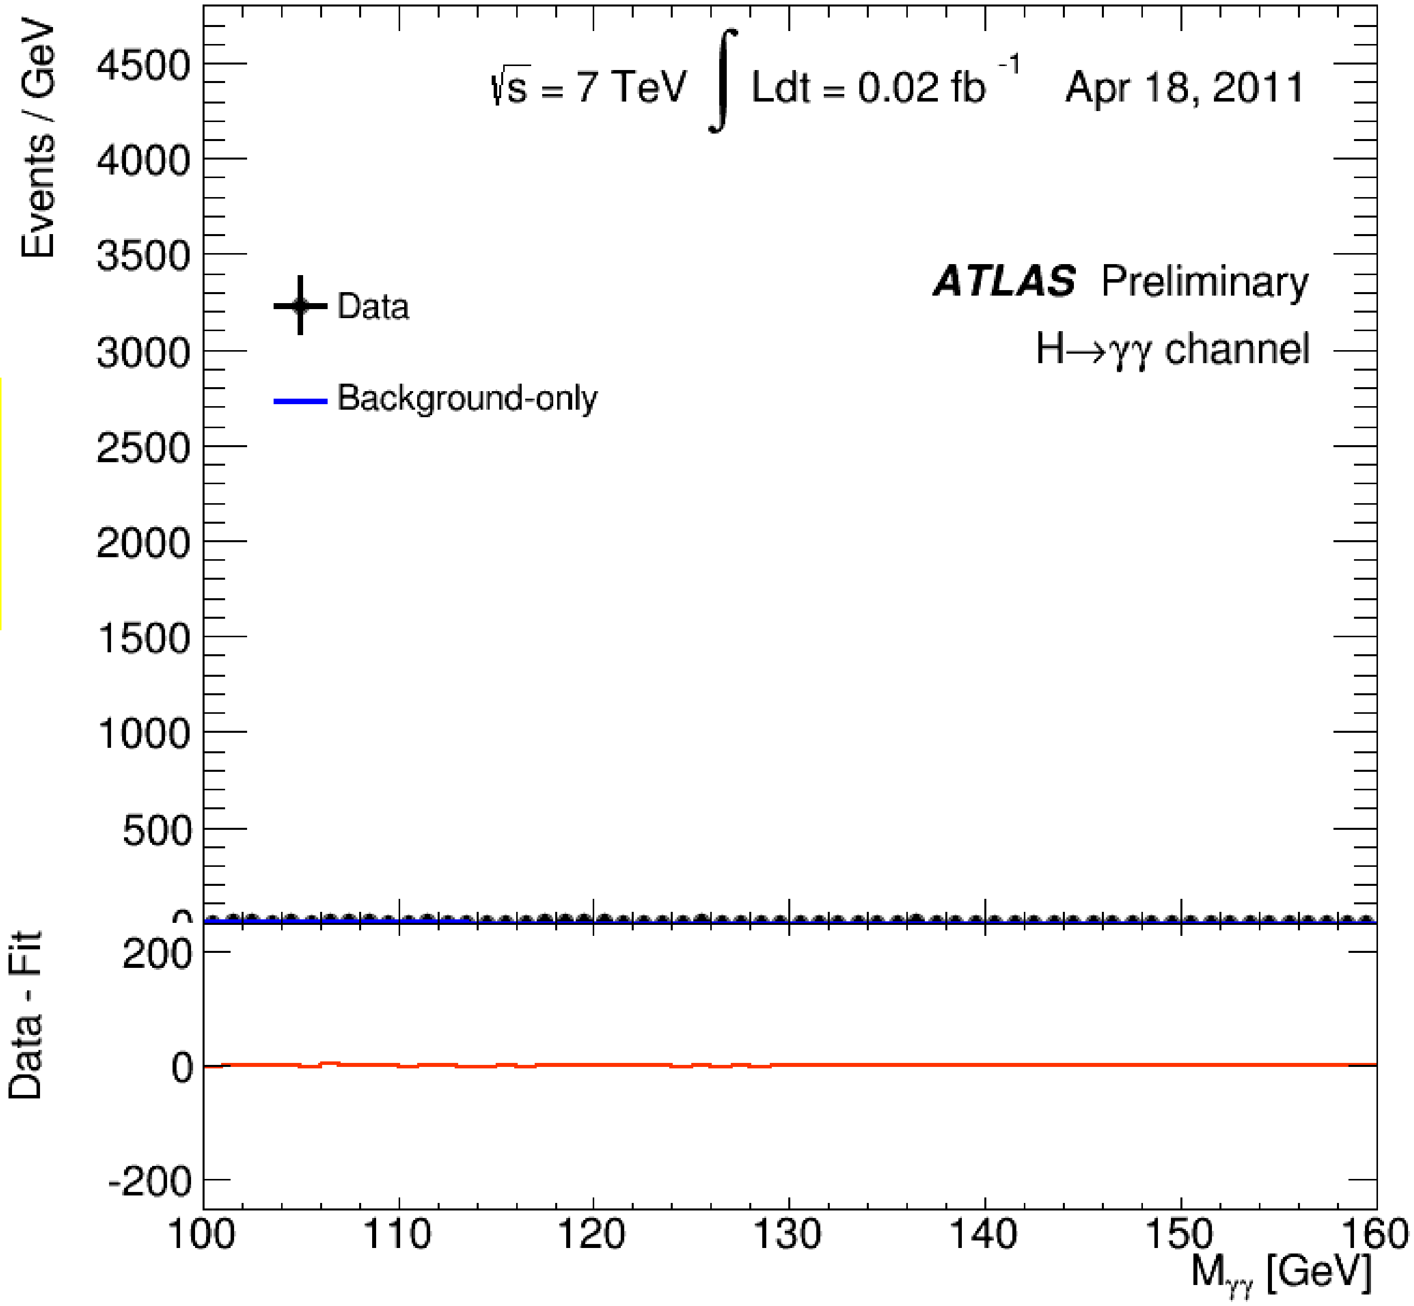

How Much Data?

- Raw analog data rate from an LHC detector
  - About one Petabyte per second
  - One $trillion per year for storage alone

- OK, wrong answer, what’s the right question?

Philosophical Aside

- "What are your data handling and computing requirements?"
  - **WRONG QUESTION**

- "What could the data handling and computing technology that you can afford allow you to do?"
  - **RIGHT QUESTION**

So What Should We Do With One Petabyte/Second?

- In real time, throw away the data not needed to make discoveries
  - But, the discoveries with the greatest impact are those we don’t expect
  - But, we are smart, no?
- We really do throw away 99.9999% of LHC data before writing it to persistent storage

LHC Data Acquisition (DAQ) CMS Example

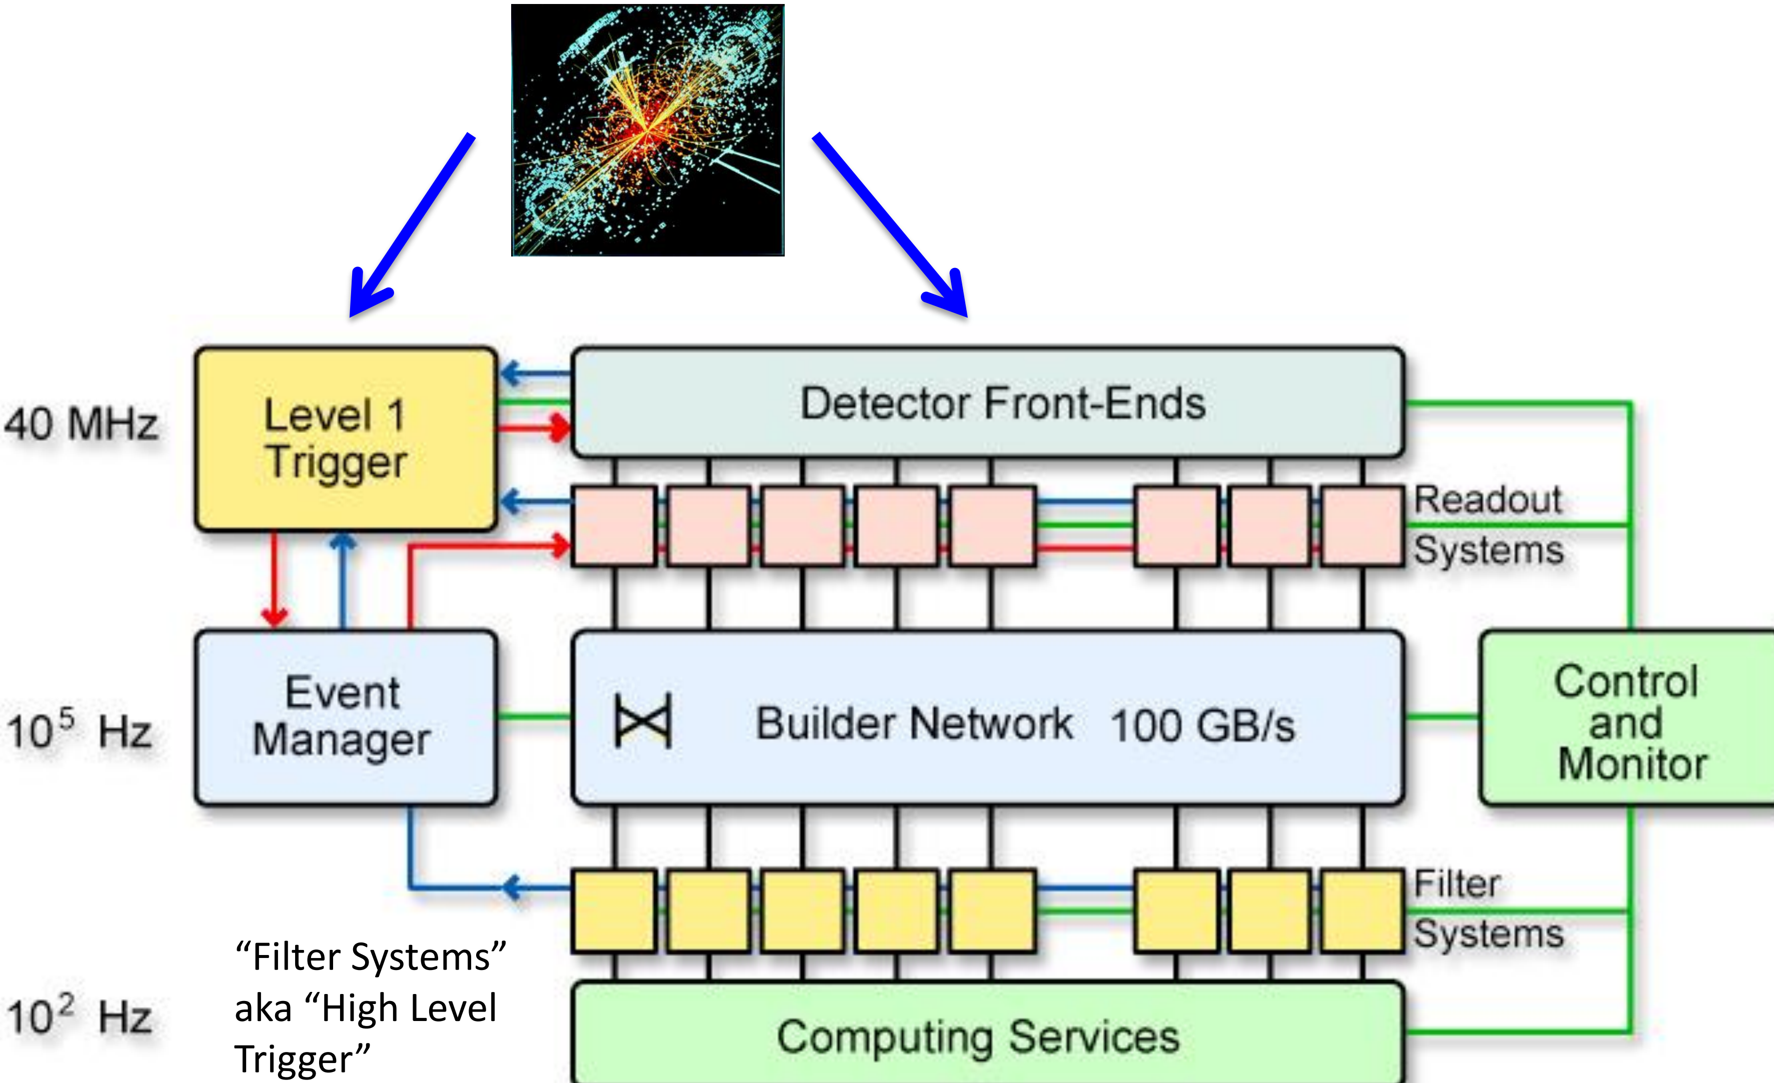

“Filter Systems” aka “High Level Trigger”

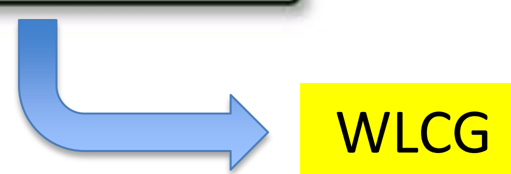

ATLAS High-Level Trigger (Part)

Total of 15,000 cores in 1,500 machines

Disk Buffer → WLCG

Planning For the LHC circa 2005

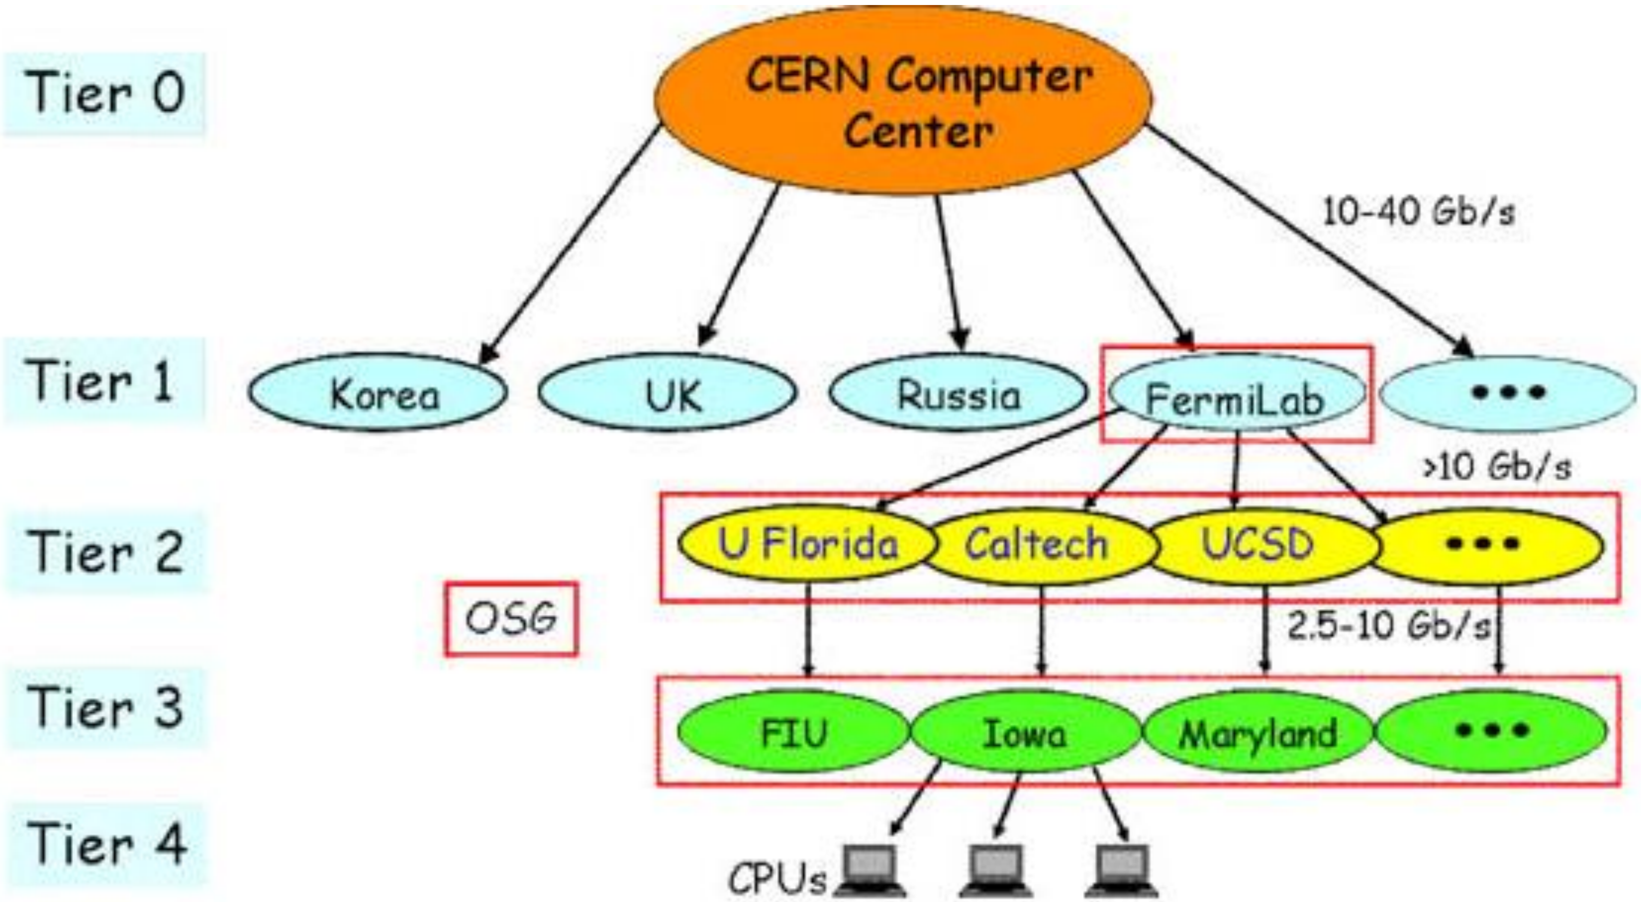

Figure 2: Sketch of the multi-tier Worldwide LHC Computing Grid, where U.S. resources (outlined in red) are incorporated into OSG.
Worldwide, the WLHC contains approximately 11 Tier-1 centers, 100 Tier-2 sites and several hundred Tier-3 institutions.

Simulation in HEP

- The fundamental physics processes are only measurable after non-invertible transformations from
  - Physics (e.g. quarks → observable particles)
  - Detector (noise, limited resolution, limited granularity)
  - Software (Imperfect pattern recognition, confusion, bugs)
- Simulate billions of “events”, apply these transformations, and compare with the observed data
- Simulating the needle is much easier than simulating the haystack
- Need much more simulation than data to make the uncertainty contribution from simulation statistics negligible.

And then what happens?

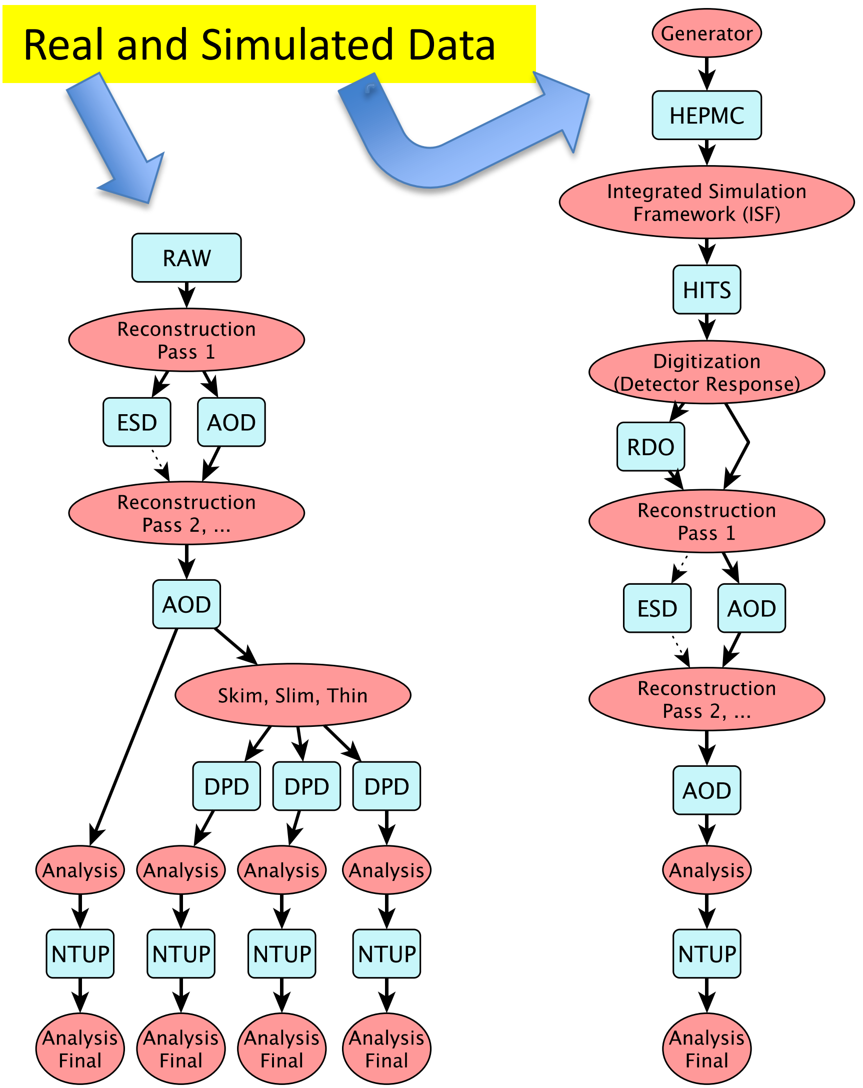

Data-Intensive processing and physics analysis at ~100 computer centers and ~1000 universities worldwide

100 Computer Centers are you crazy?

- Political/funding/greed requirement for a distributed system
  - A centralized system might divide unit costs by a factor 2
  - But it might divide funding by a factor 4 at least
- Some real benefits of the distributed system
  - Involves scientists worldwide and distributes expertise
  - Eases access to major opportunistic resources

LHC Reality 2012 (Mainly ATLAS and CMS)

Sites	Number	CPU Cores	Disk Petabytes
Tier 0 (CERN)	1	27,000	22
Tier 1	11	82,000	79
Tier 2	66	215,000	94
TOTAL	**78**	**324,000**	**196**

Many Tier 2s are implemented as multi-site centers

- Conundrum:
  - ATLAS has been taking data at ~300 MB/s, running for typically 2,000 hours/year since November 2009
  - About 3 PB/year maximum of raw data
  - By the Summer of 2011, all the ATLAS disks were full.

How does less than 3 PB become 100 PB

1. Duplicate the raw data (not such a bad idea)
2. Add a similar volume of simulated data (essential)
3. Make a rich set of derived data products (some of them larger than the raw data)
4. Re-create the derived data products whenever the software has been significantly improved (several times a year) and keep the old versions for quite a while
5. Place up to 40[§] copies of the derived data around the world so that when you “send the computing to the data” you can send it almost anywhere
6. Do the math!

§ now far fewer copies due to demand-driven temporary replication – but much more reliance on wide-area networks

Final Aside – Data Rates and Technologies

- I have been deeply involved in many data-intensive high energy physics experiments
- Since the early 80s, I have been personally involved in buying stuff[§] – computers, storage and networks – for these experiments
- What picture emerges?

§ Some of the network purchases were made by Harvey Newman/Caltech

HEP Data versus Technologies

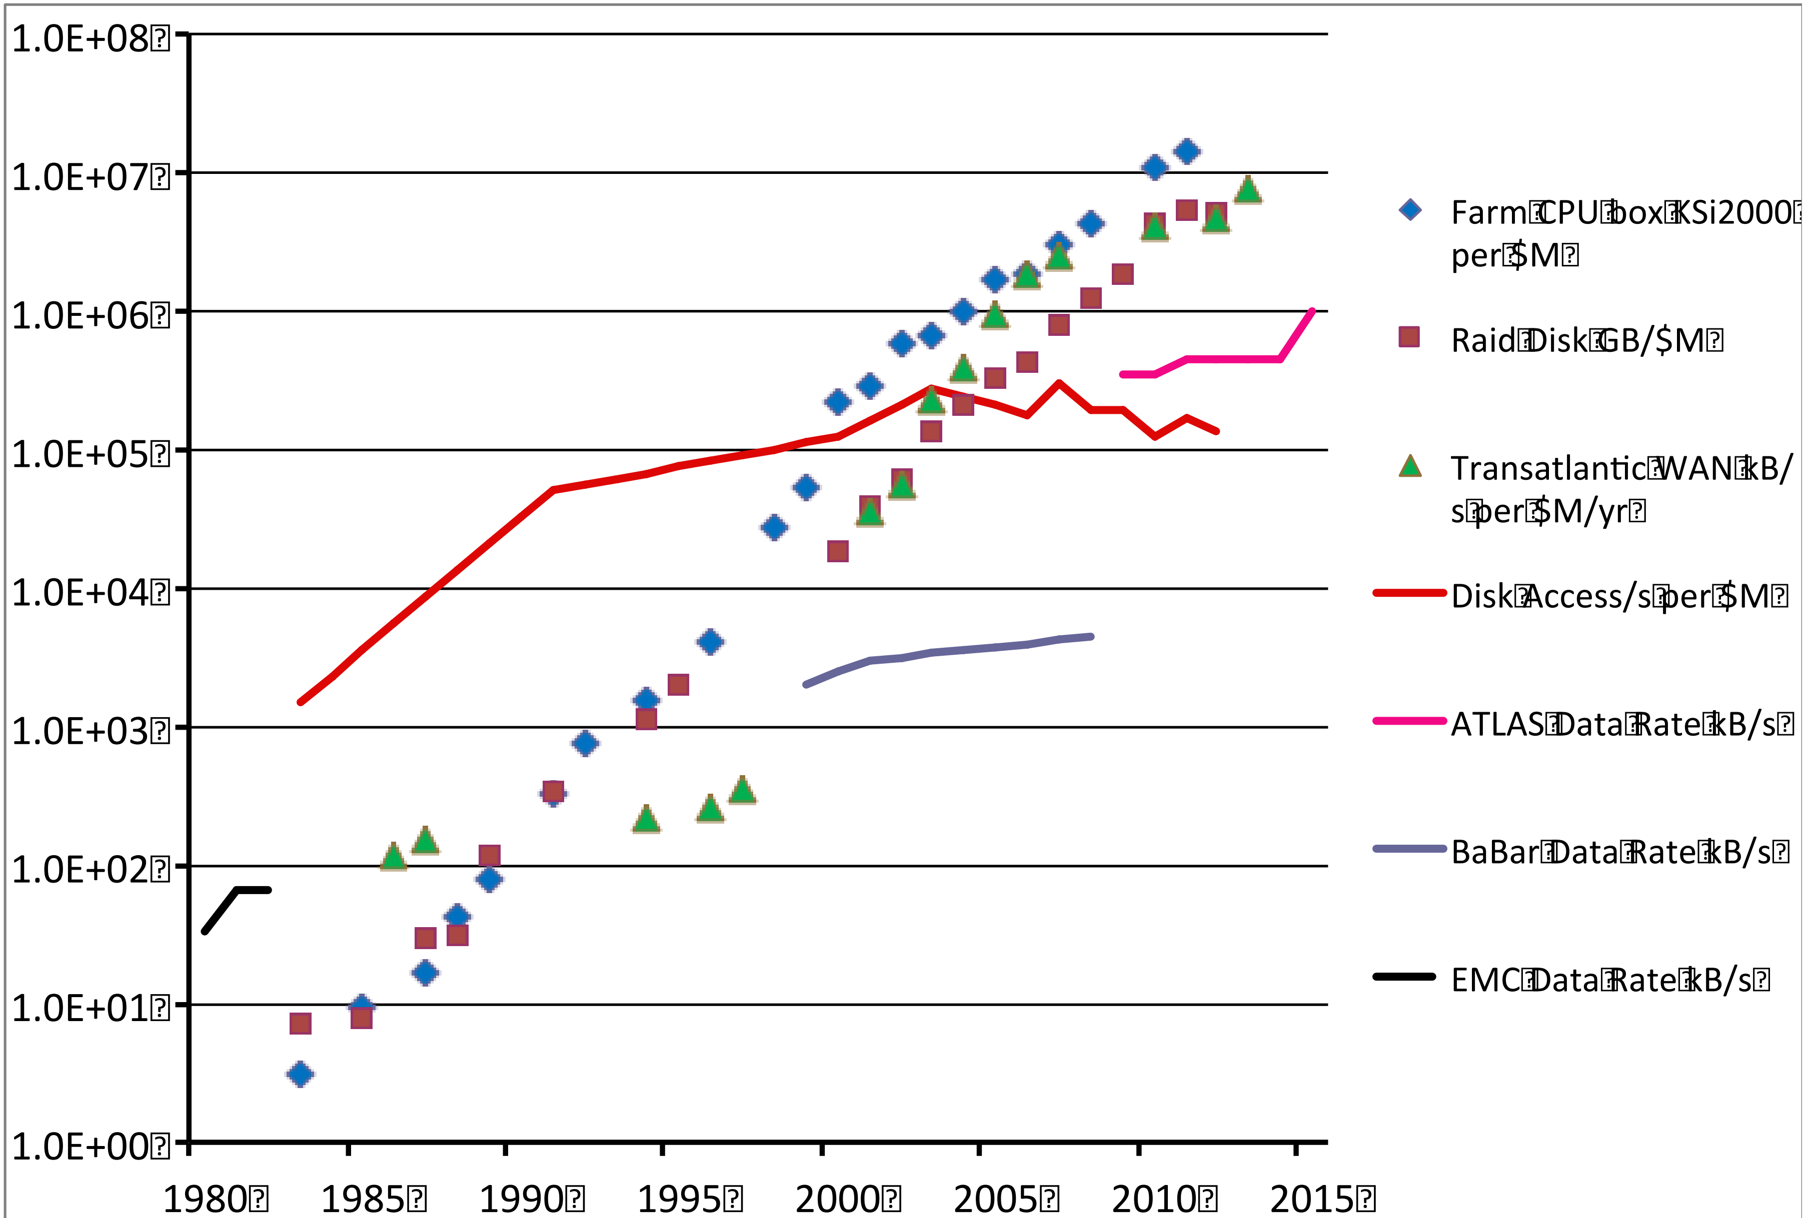

HEP Summary (almost)

- Hundreds of Petabytes of data
- Worldwide automated processing and data flow
- But, fortunately, we are losing our leadership position in the data-intensive world