WLCG – Beyond the LHCOPN

Progress and Likely Evolution

Ian Bird

GLIF Workshop
13th October 2010, CERN

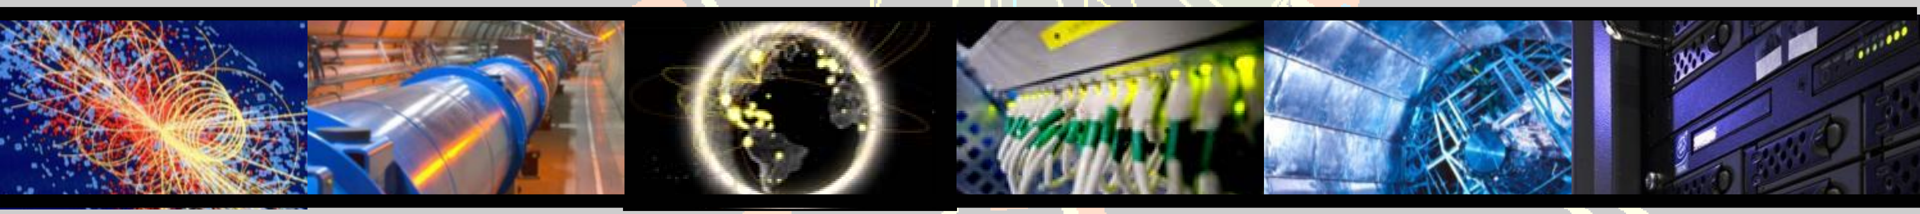

Outline

- LHC and the need for a grid
- Progress in computing for LHC
  – Where we are we today & how we got here
  – What is today's grid compared to early ideas?
- … and the outlook?
  – What are the open questions?
  – Evolution? Sustainability?
  – What next?
  – Network needs

Collisions at the LHC: summary

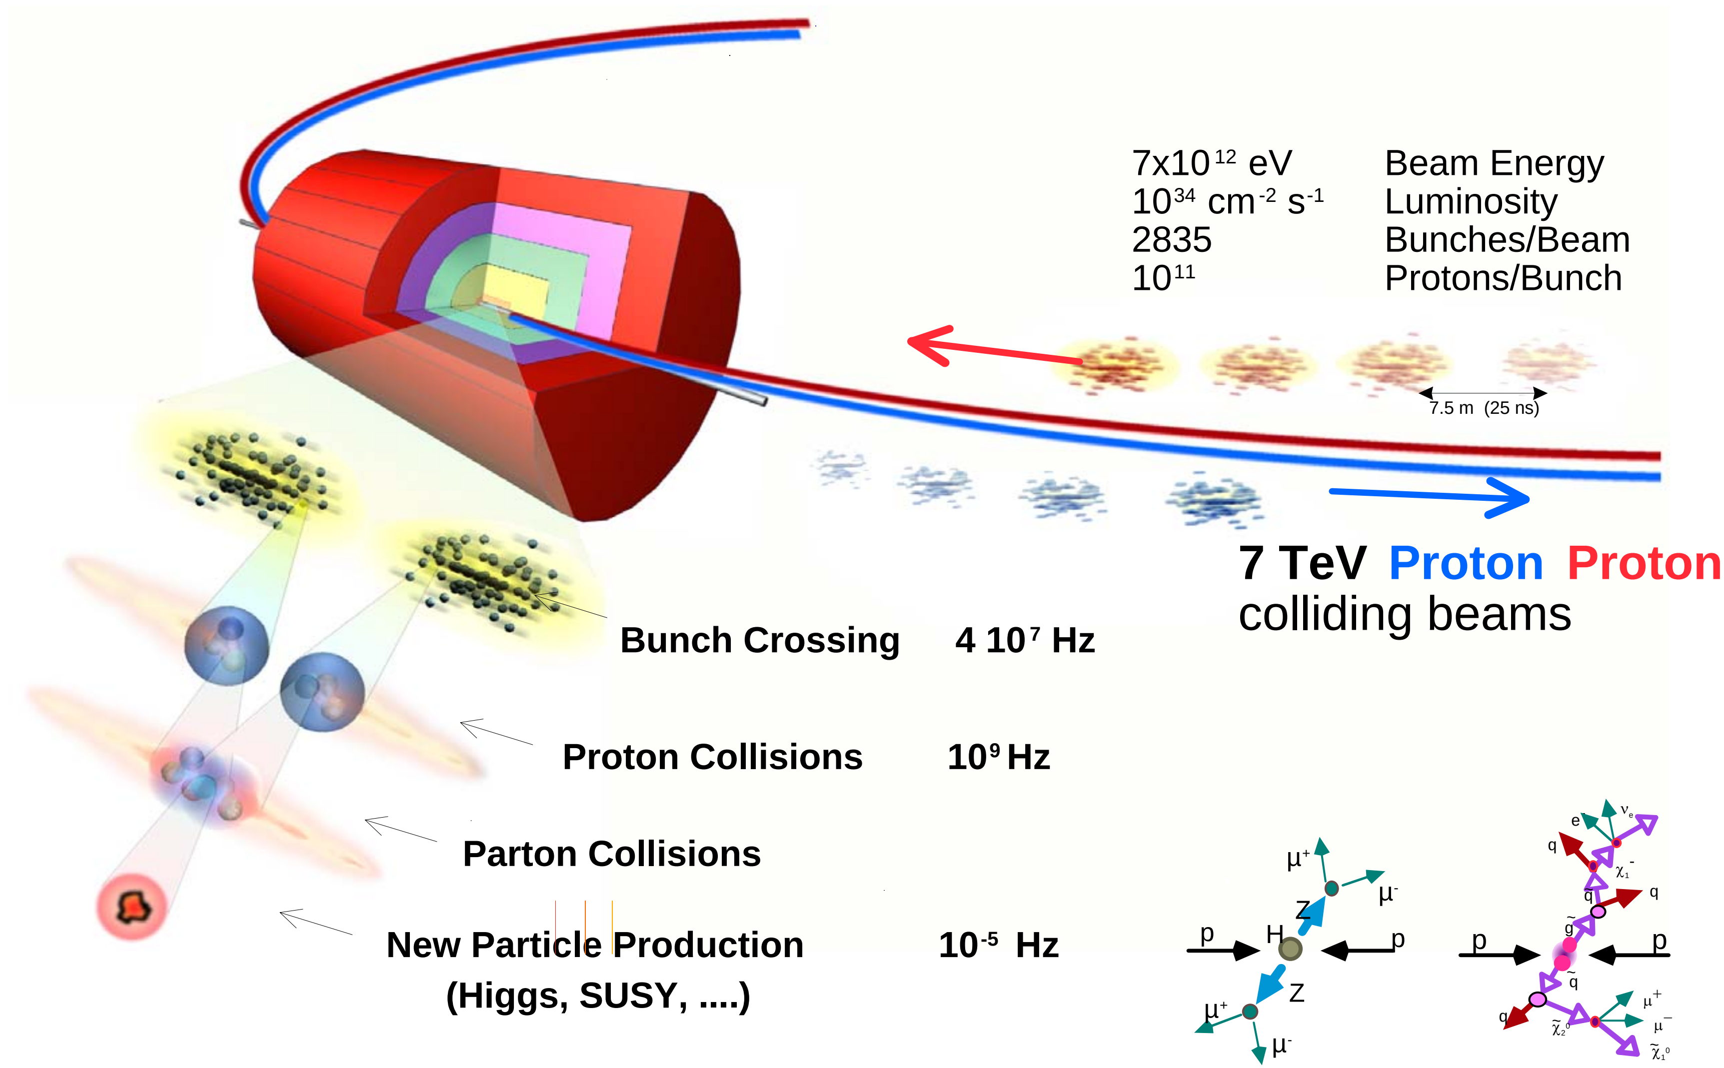

Selection of 1 event in 10,000,000,000,000

pp collisions at 14 TeV at 10^{34} $cm^{-2}s^{-1}$

A very difficult environment …

How to extract this:
(Higgs → 4 muons)

From

W

20 proto
collision

And this
every

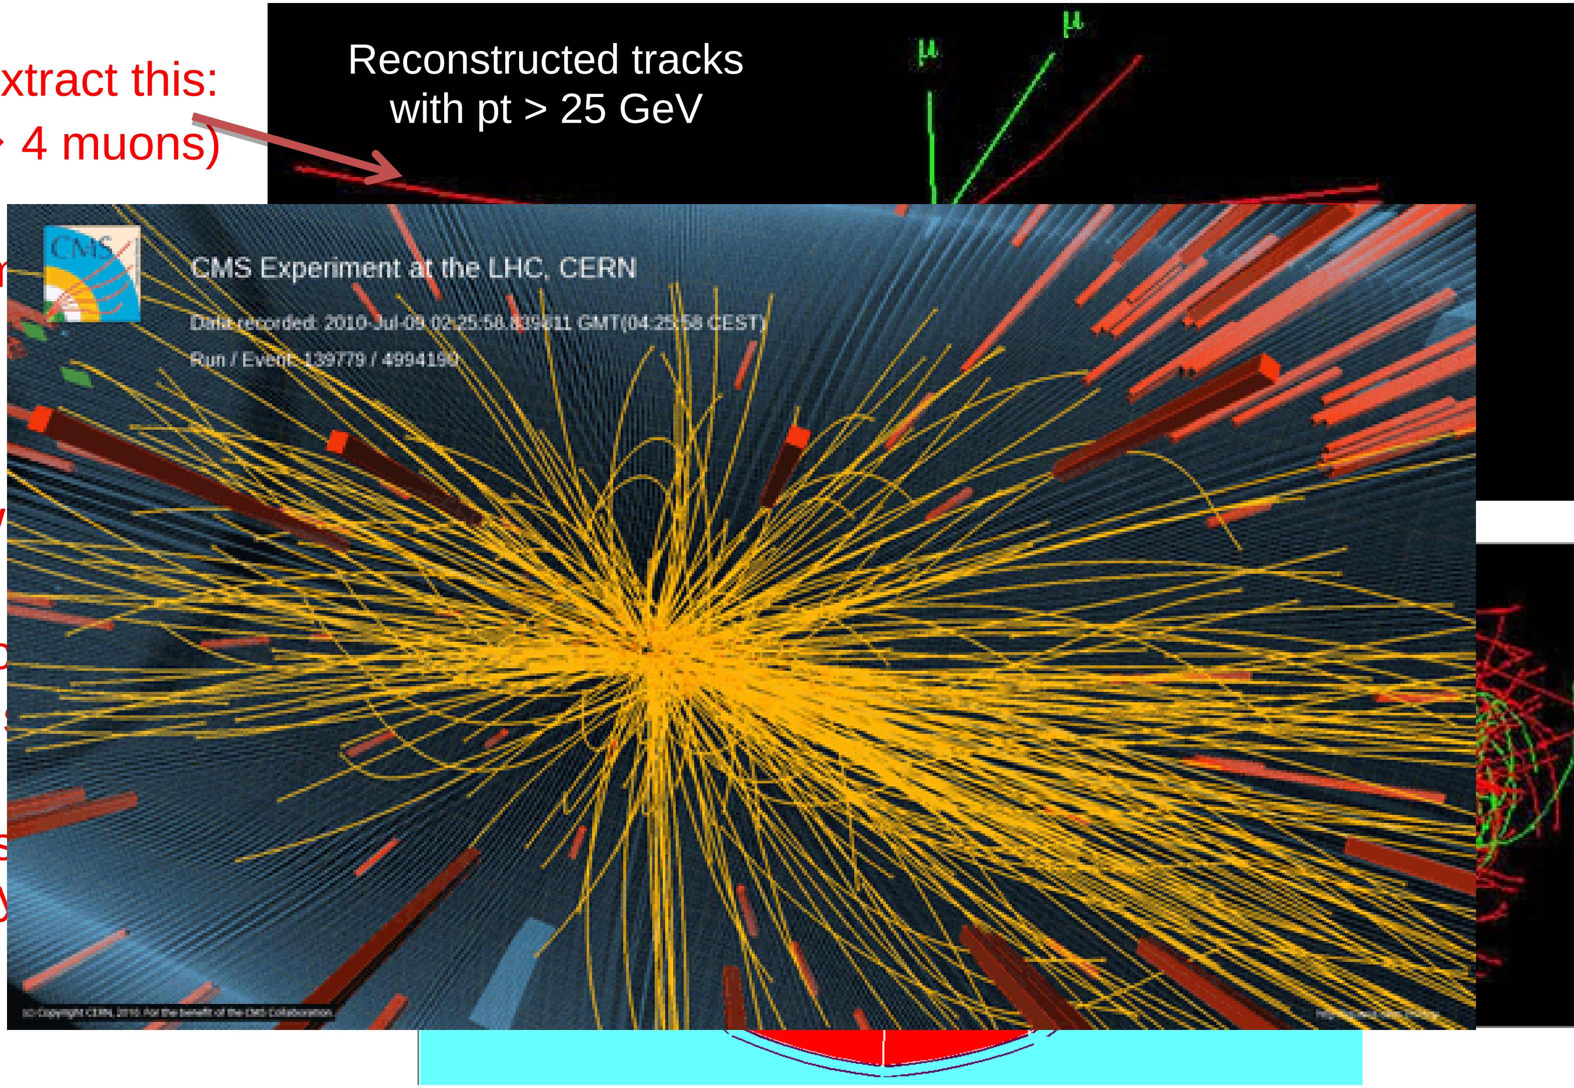

The LHC Computing Challenge

- Signal/Noise: 10^{-13} (10^{-9} offline)
- Data volume
  - High rate * large number of channels * 4 experiments
  - ➔ **15 PetaBytes of new data each year**
- Compute power
  - Event com… * Nb. events * thous… …ers
  - ➔ **200 k o… (today's) fast… …Us**
  - ➔ **45 PB of disk storag…**

>250 k cores today

100 PB disk today!!!

- Worldwide analysis & funding
  - Computing funding locally in major regions & countries
  - Efficient analysis everywhere
  - ➔ **GRID technology**

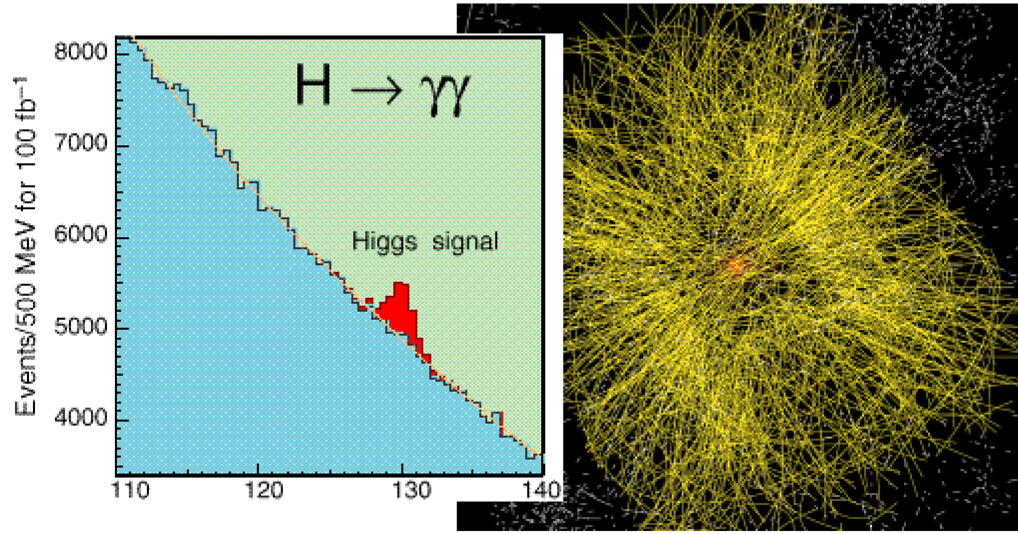

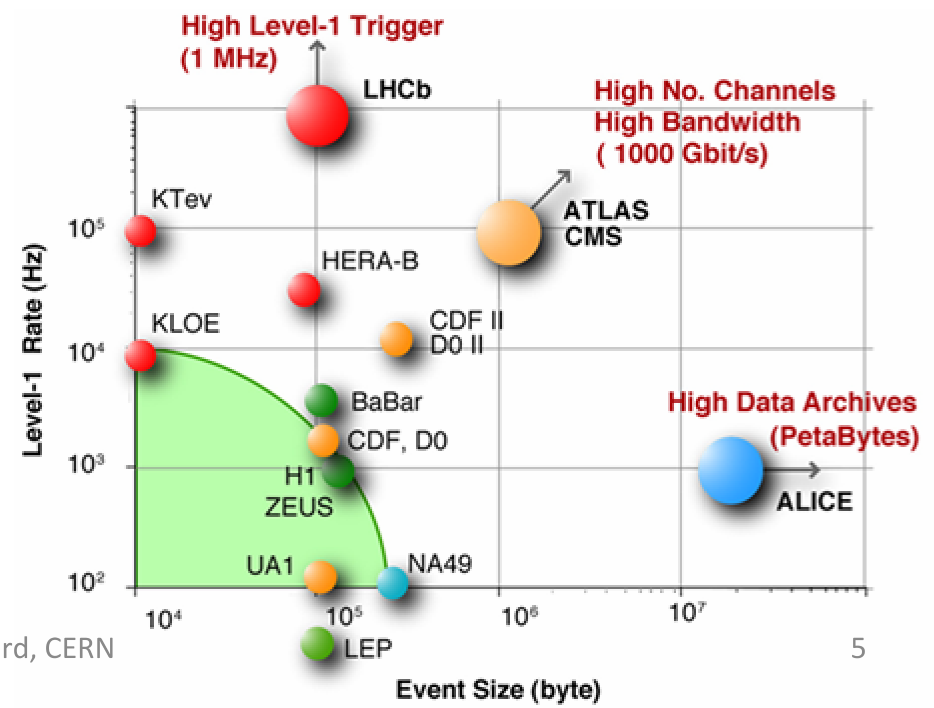

WLCG – what and why?

- A distributed computing infrastructure to provide the production and analysis environments for the LHC experiments
- Managed and operated by a worldwide collaboration between the experiments and the participating computer centres

- The resources are distributed – for funding and sociological reasons

- Our task is to make use of the resources available to us – no matter where they are located
  – We know it would be simpler to put all the resources in 1 or 2 large centres
  – This is not an option ... today

WLCG Collaboration Status

Tier 0; 11 Tier 1s; 68 Tier 2 federations

Today we have 49 MoU signatories, representing 34 countries:

Australia, Austria, Belgium, Brazil, Canada, China, Czech Rep, Denmark, Estonia, Finland, France, Germany, Hungary, Italy, India, Israel, Japan, Rep. Korea, Netherlands, Norway, Pakistan, Poland, Portugal, Romania, Russia, Slovenia, Spain, Sweden, Switzerland, Taipei, Turkey, UK, Ukraine, USA.

Worldwide resources

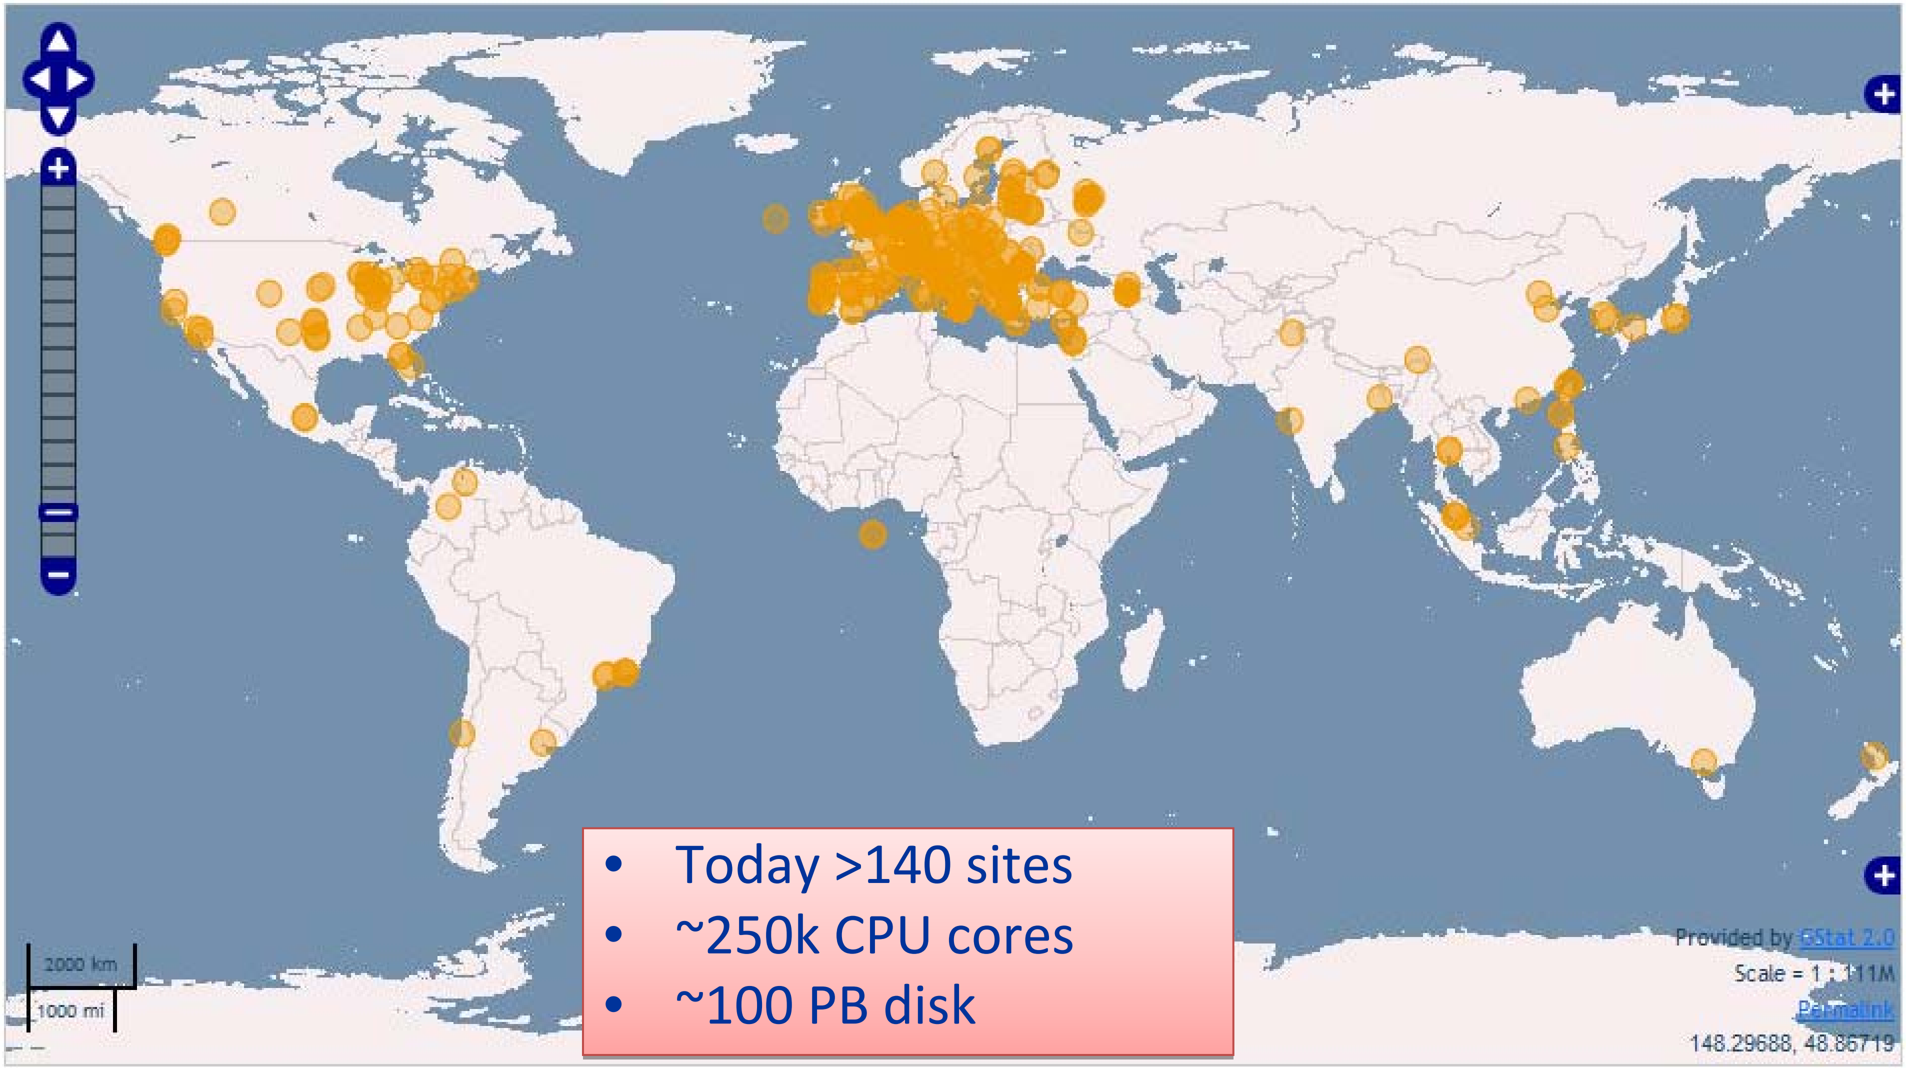

- Today >140 sites
- ~250k CPU cores
- ~100 PB disk

Computing models have evolved

- Models all ~based on the MONARC tiered model of 10 years ago
- Several significant variations, however

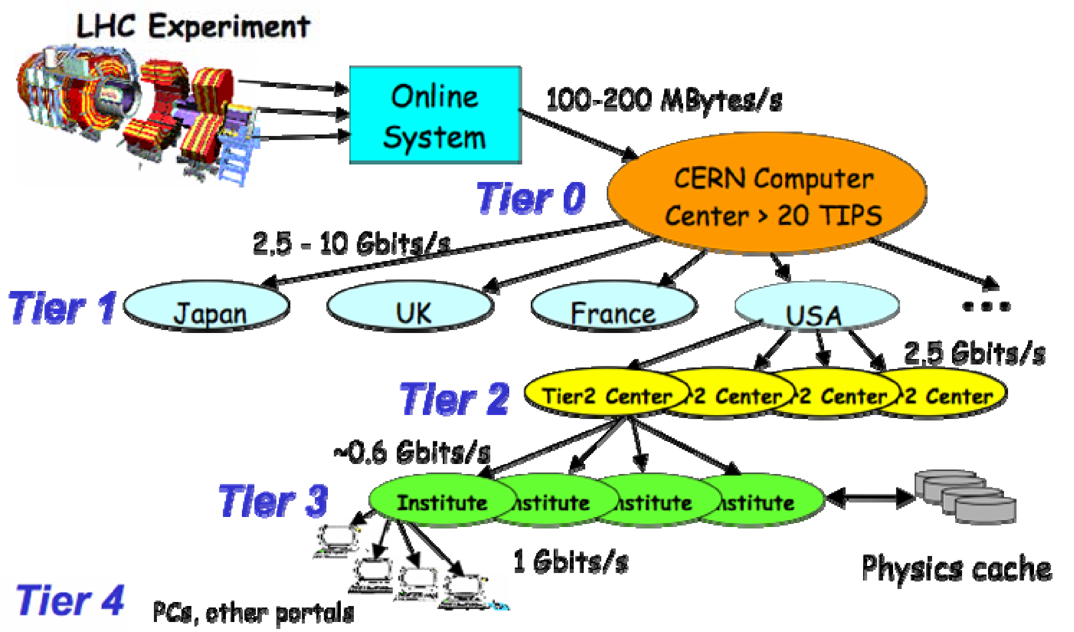

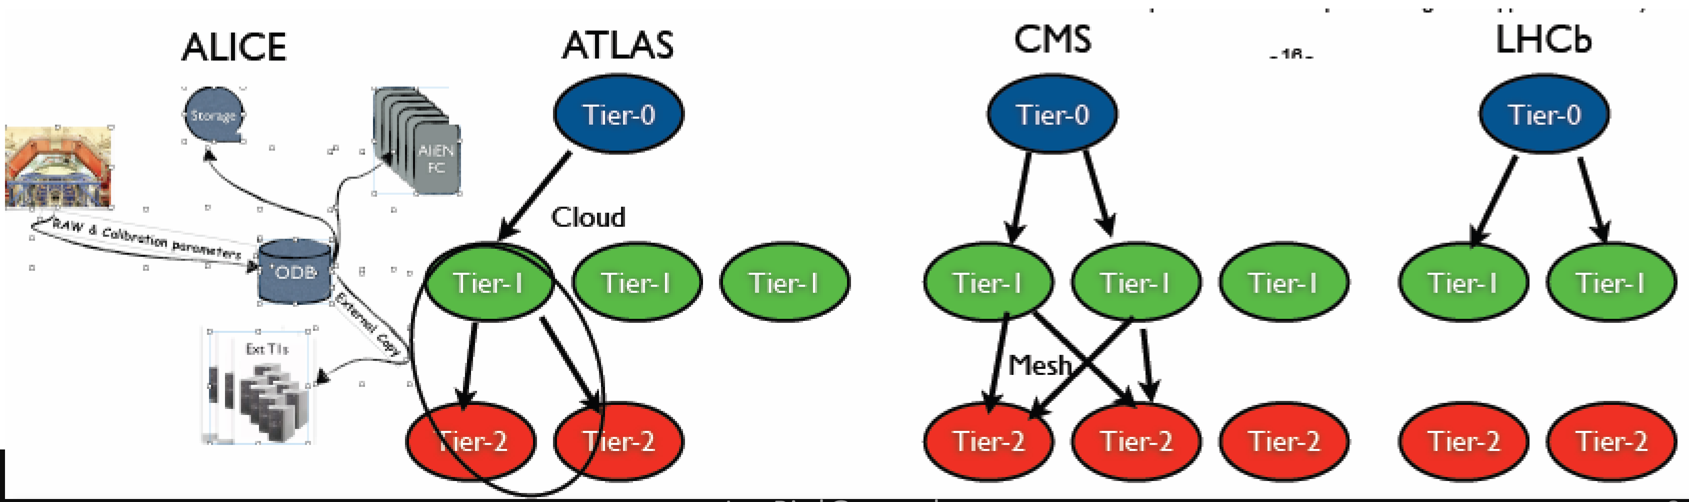

6 months of LHC data

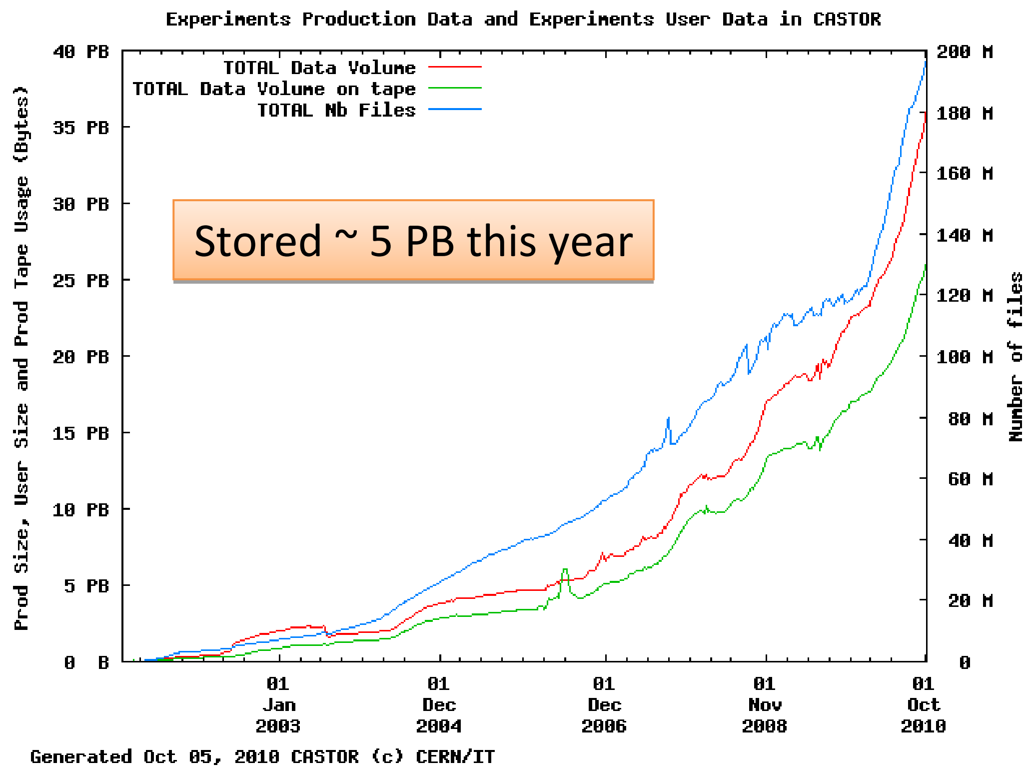

Data written to tape (Gbytes/day)

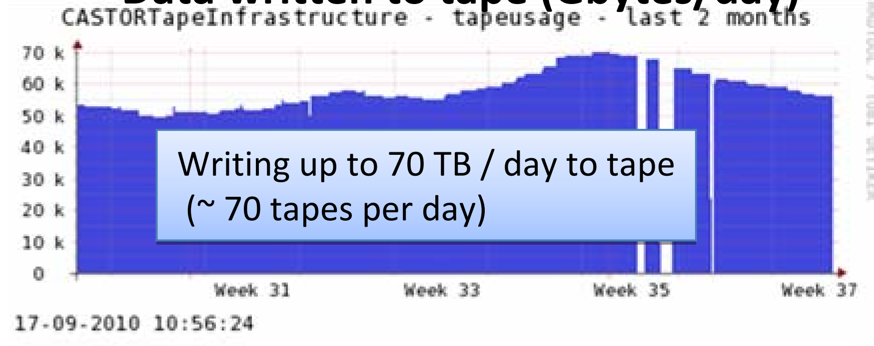

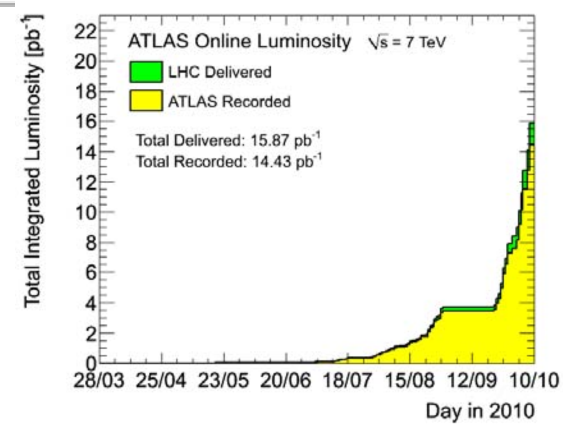

Disk Servers (Gbytes/s)

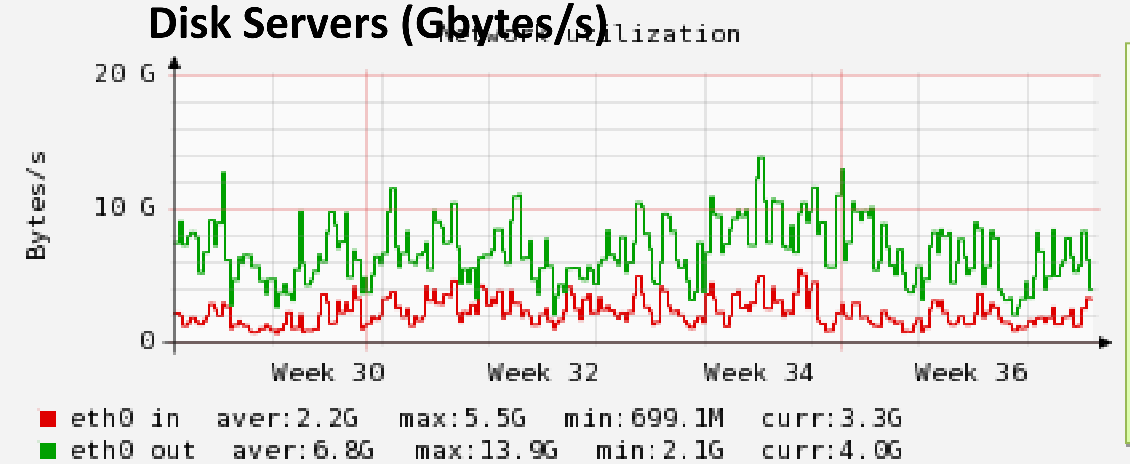

Tier 0 storage:

- Accepts data at average of 2.2 GB/s; peaks > 6 GB/s
- Serves data at average of 7 GB/s; peaks > 14 GB/s
- **CERN Tier 0 moves ~ 1 PB data per day**

WLCG Usage

- Use remains consistently high
  - 1 M jobs/day; 100k CPU-days/day

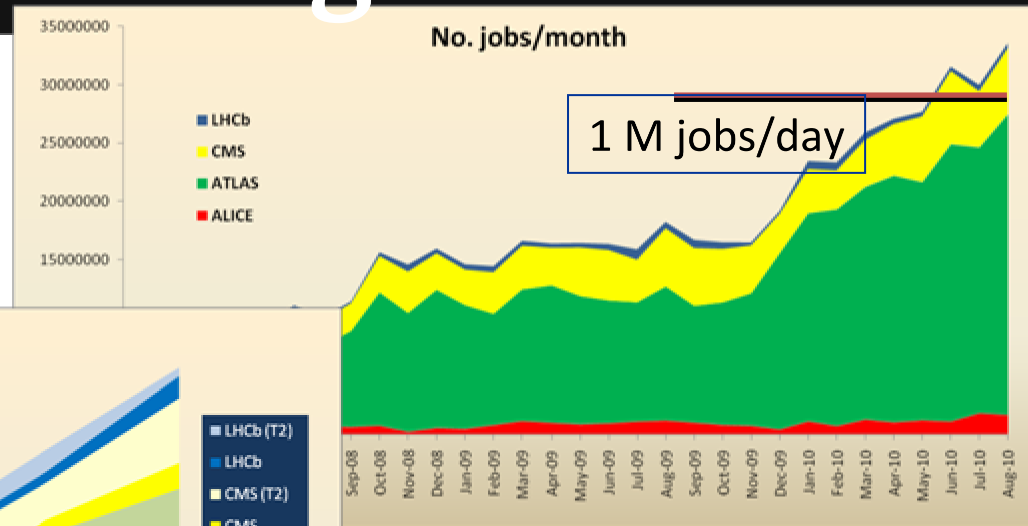

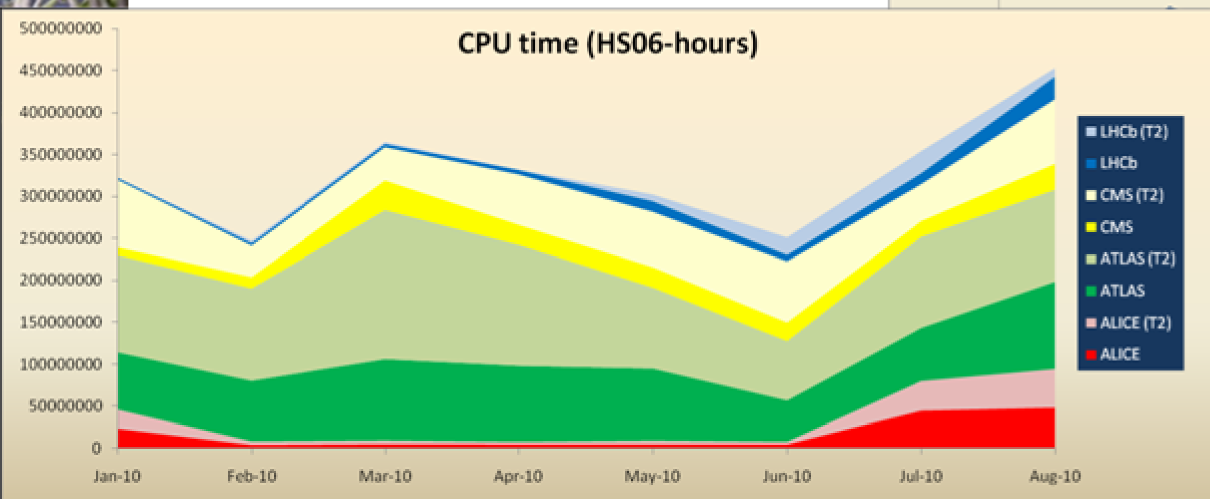

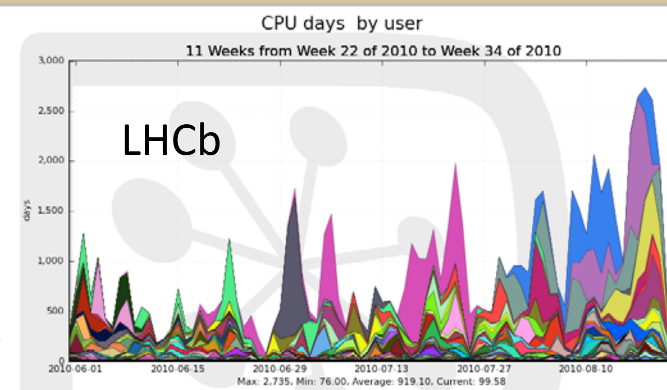

As well as LHC data, large simulation productions ongoing

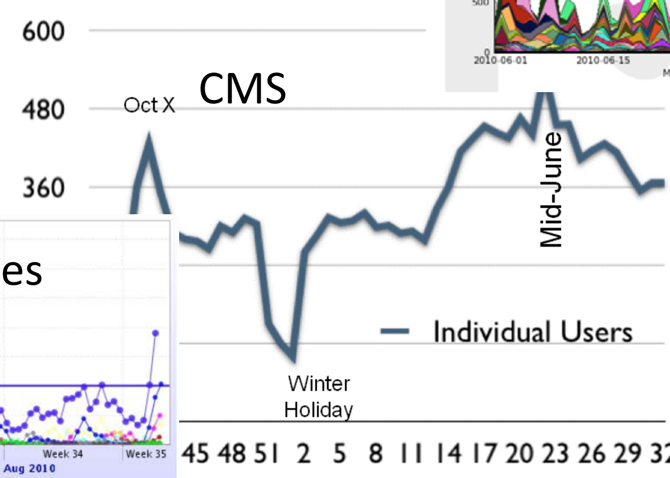

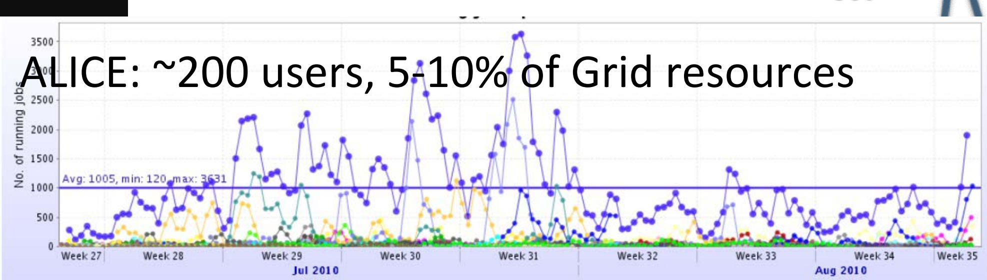

Large numbers of analysis users

CMS ~500, ATLAS ~1000, LHCb/ALICE ~200

CPU – July

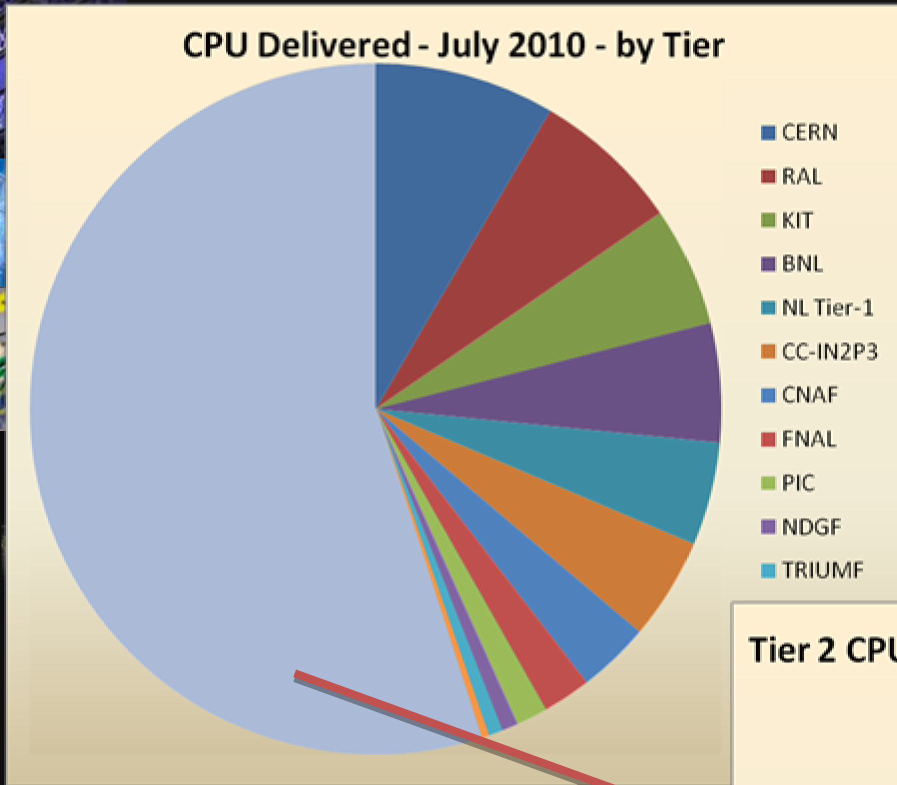

- Significant use of Tier 2s for analysis
  - frequently-expressed concern that too much analysis would be done at CERN is not reflected

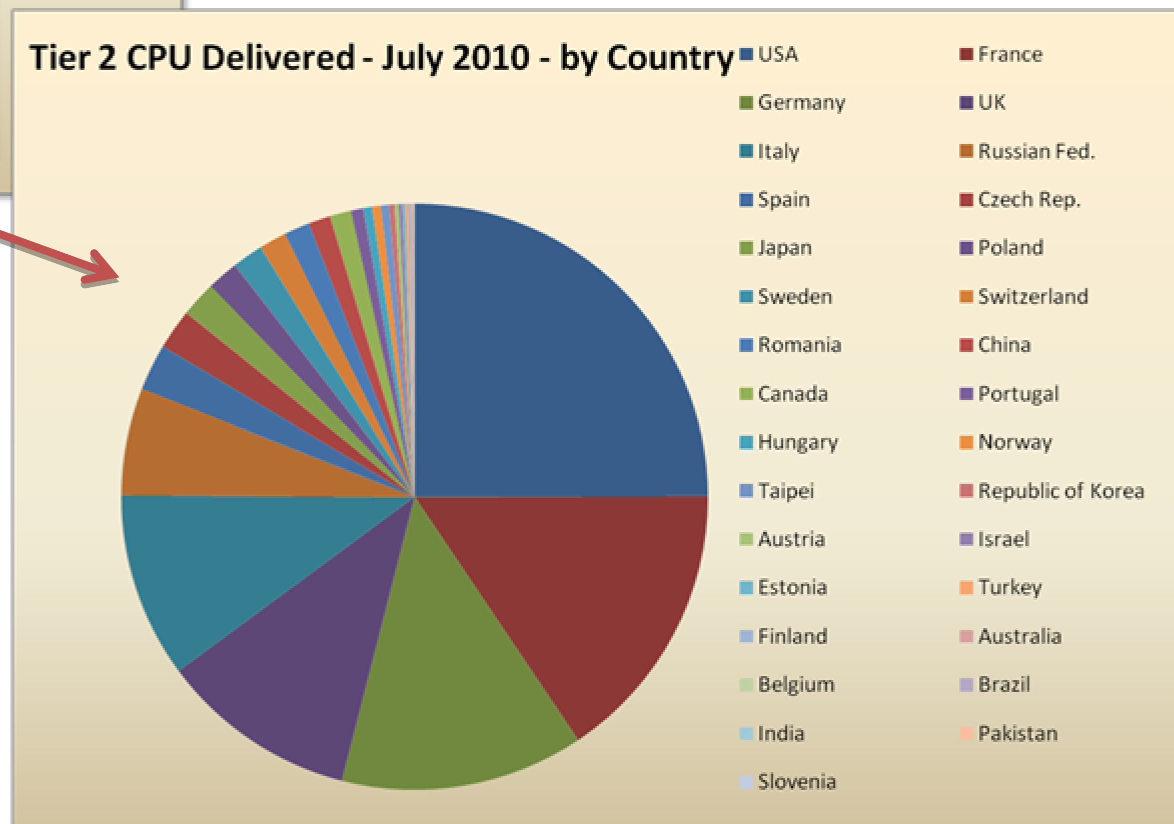

- Tier 0 capacity underused in general
  - But this is expected to change as luminosity increases

ICHEP Paris 2010

Data transfer

- Data transfer capability today able to manage much higher bandwidths than expected/feared/planned

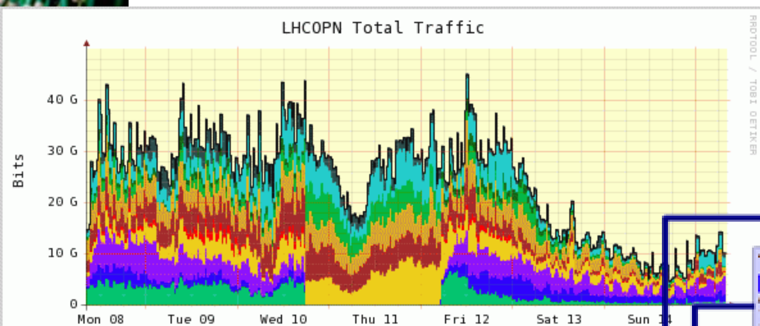

Fibre cut during STEP'09:
Redundancy meant no interruption

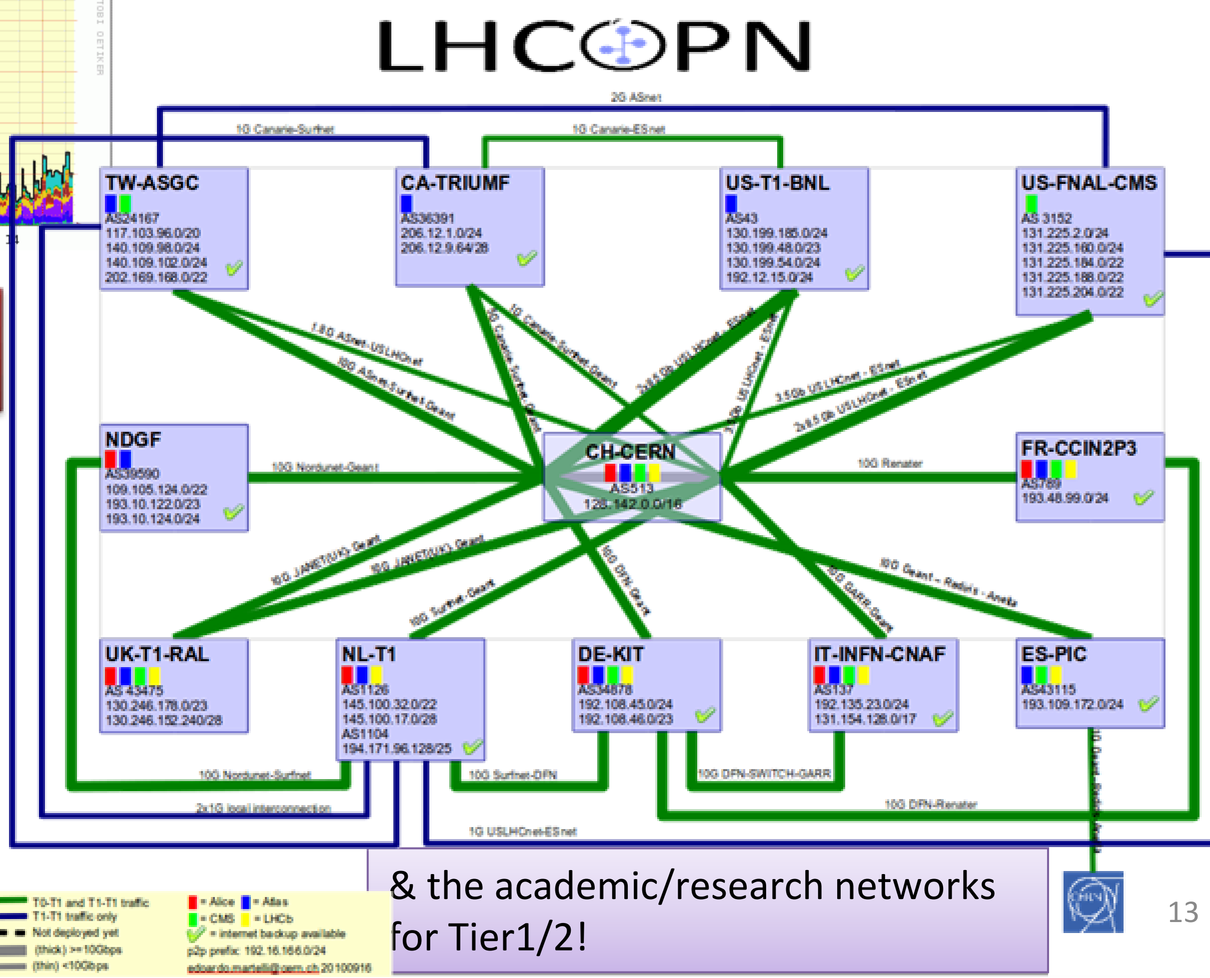

Data transfer:

•SW: gridftp, FTS (interacts with endpoints, recovery), experiment layer

•HW: light paths, routing, coupling to storage

•Operational: monitoring

& the academic/research networks for Tier1/2!

Data transfers

Final readiness test (STEP'09)

Preparation for LHC startup

LHC physics data

2009: STEP09 + preparation for data

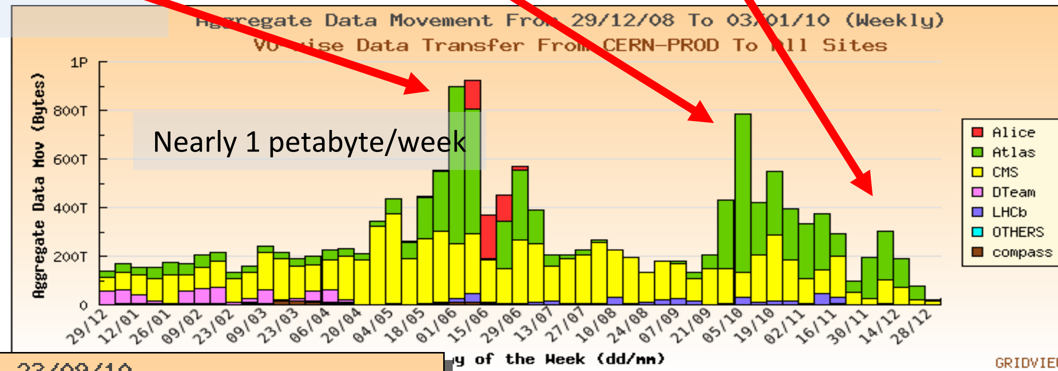

Nearly 1 petabyte/week

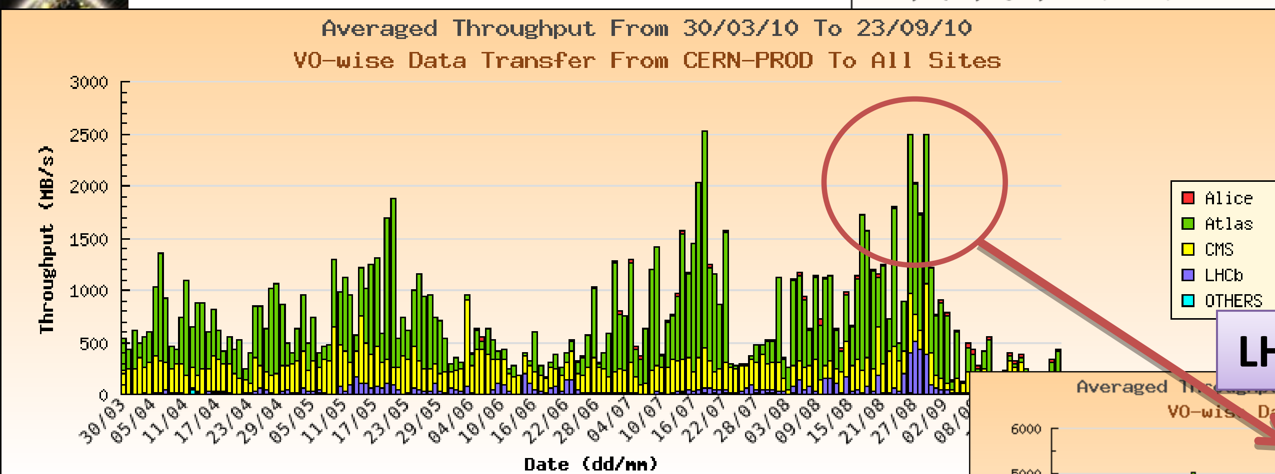

LHC running: April – Sept 2010

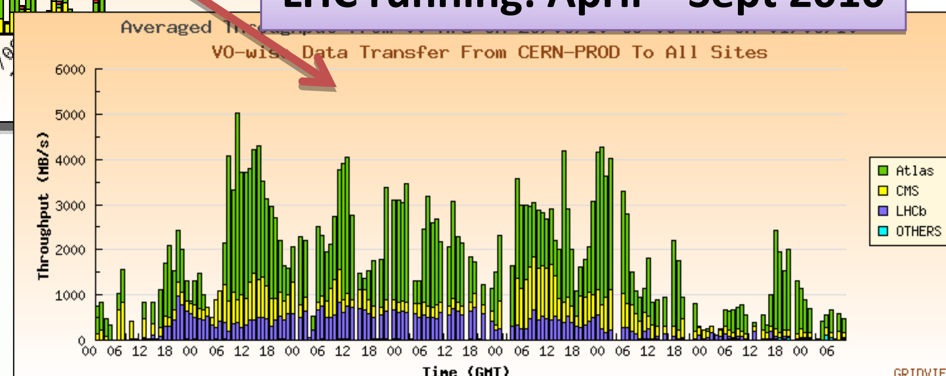

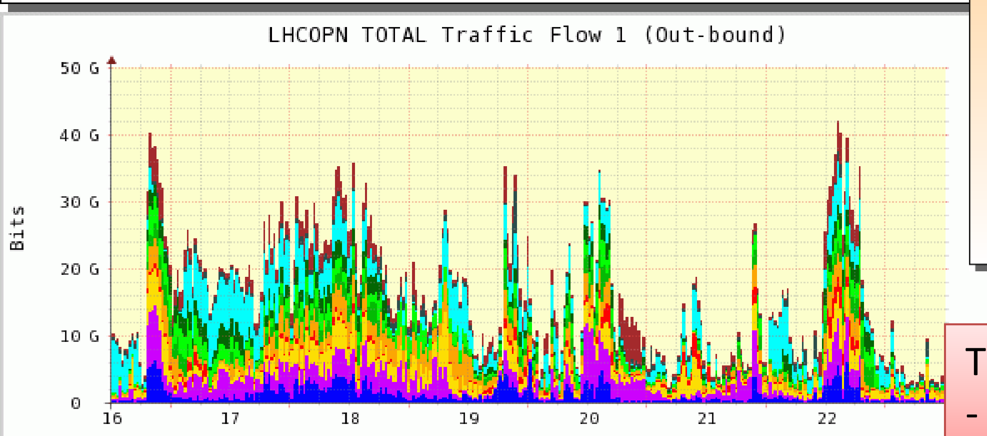

Traffic on OPN up to 70 Gb/s!

- ATLAS early reprocessing campaigns

Data transfers: total

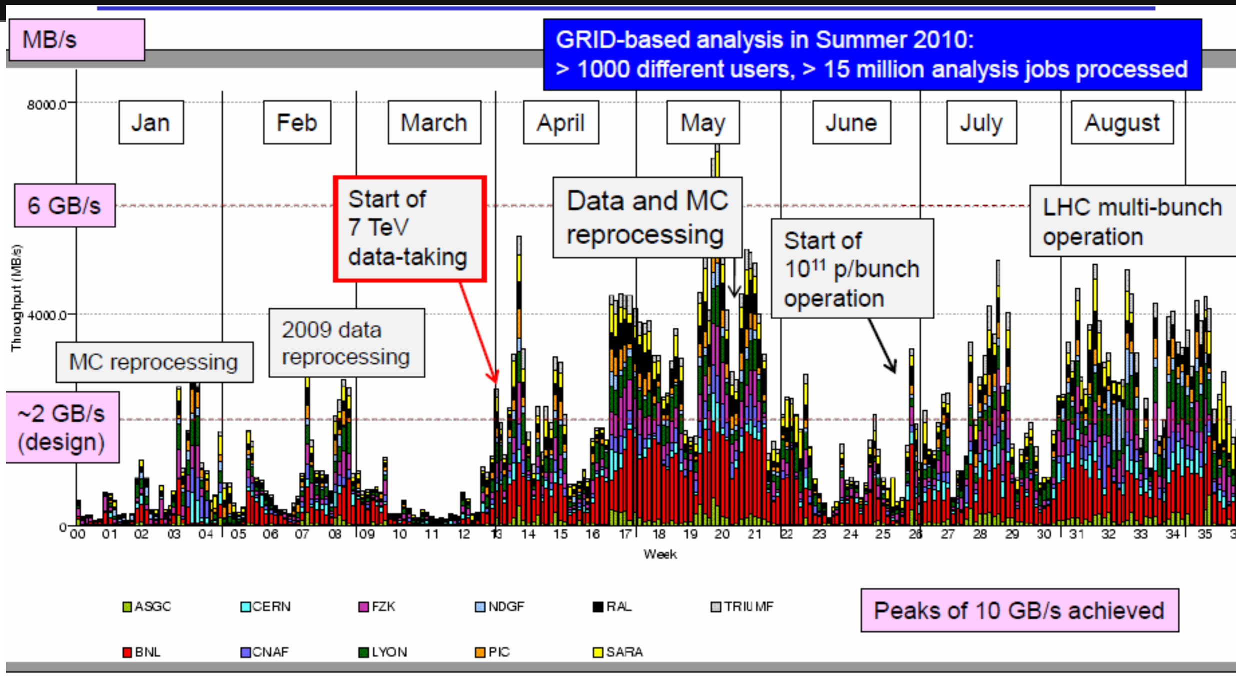

ATLAS Throughput

For all experiments: early data has been available for analysis within hours of data taking

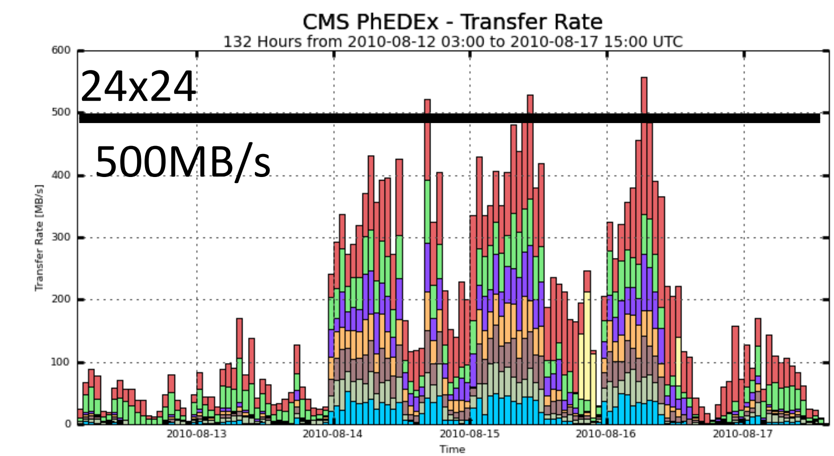

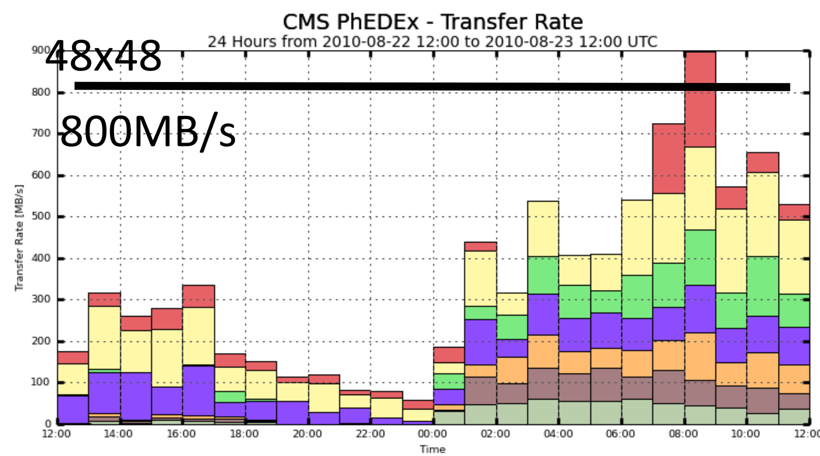

– CMS saw effect of going from 24x24 to 48x48 bunches: same trigger

Service quality: defined in MoU

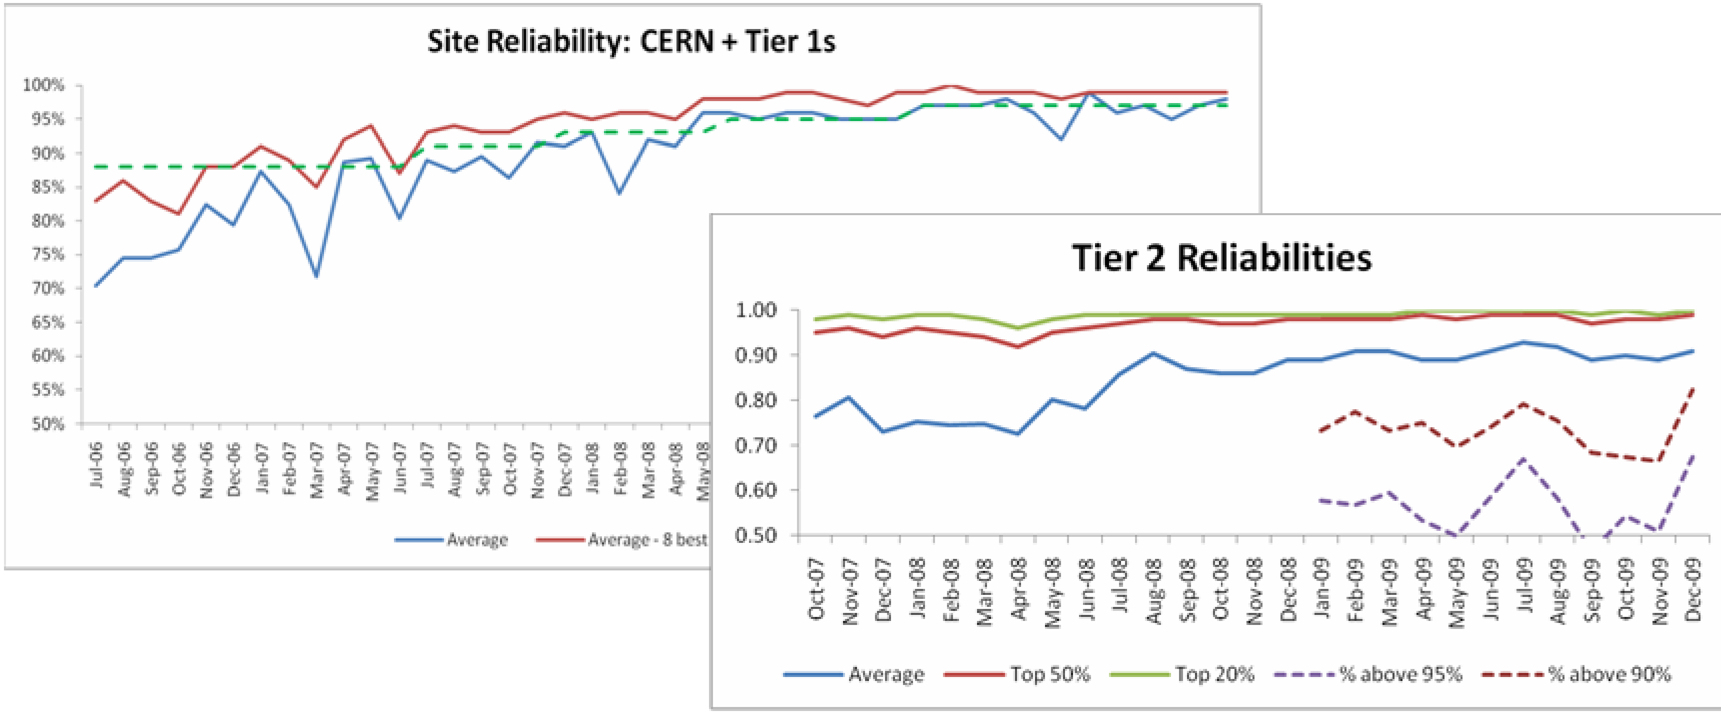

- MoU defines key performance and support metrics for Tier 1 and Tier 2 sites
  - Reliabilities are an approximation for some of these
  - Also metrics on response times, resources, etc.
- The MoU has been an important tool in bringing services to an acceptable level

From testing to data:

Independent Experiment Data Challenges

e.g. DC04 (ALICE, CMS, LHCb)/DC2 (ATLAS) in 2004 saw first full chain of computing models on grids

Service Challenges proposed in 2004

To demonstrate service aspects:

- -Data transfers for weeks on end
- -Data management
- -Scaling of job workloads
- -Security incidents ("fire drills")
- -Interoperability
- -Support processes

2004

2005

SC1 Basic transfer rates

SC2 Basic transfer rates

2006

SC3 Sustained rates, data management, service reliability

SC4 Nominal LHC rates, disk→ tape tests, all Tier 1s, some Tier 2s

2007

2008

CCRC'08 Readiness challenge, all experiments, ~full computing models

2009

STEP'09 Scale challenge, all experiments, full computing models, tape recall + analysis

2010

- Focus on real and continuous production use of the service over several years (simulations since 2003, cosmic ray data, etc.)
- Data and Service challenges to exercise all aspects of the service – not just for data transfers, but workloads, support structures etc.

Rediscovering the Standard Model
From basics to W/Z to multi-jets to the Top

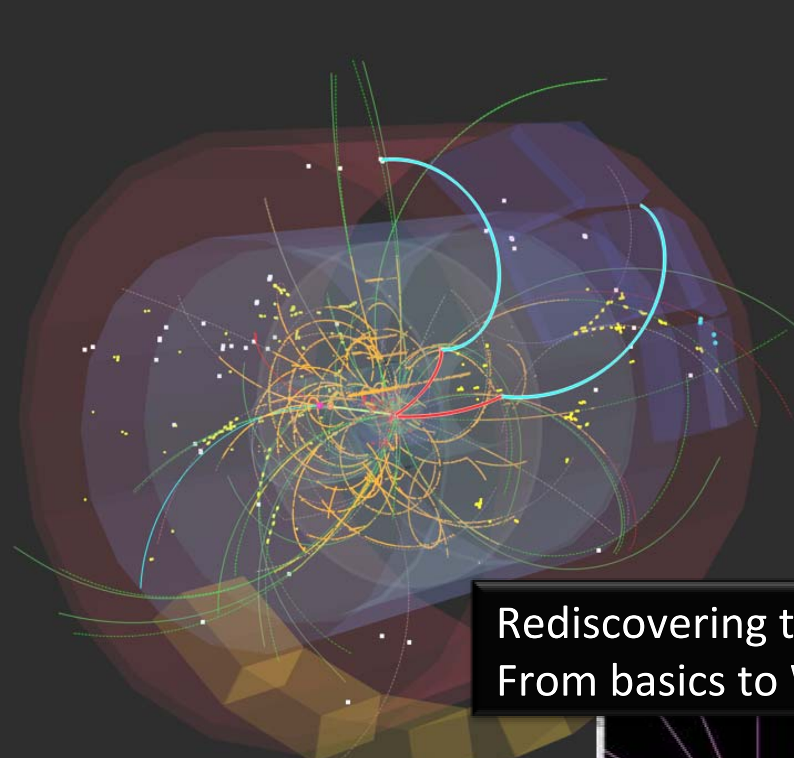

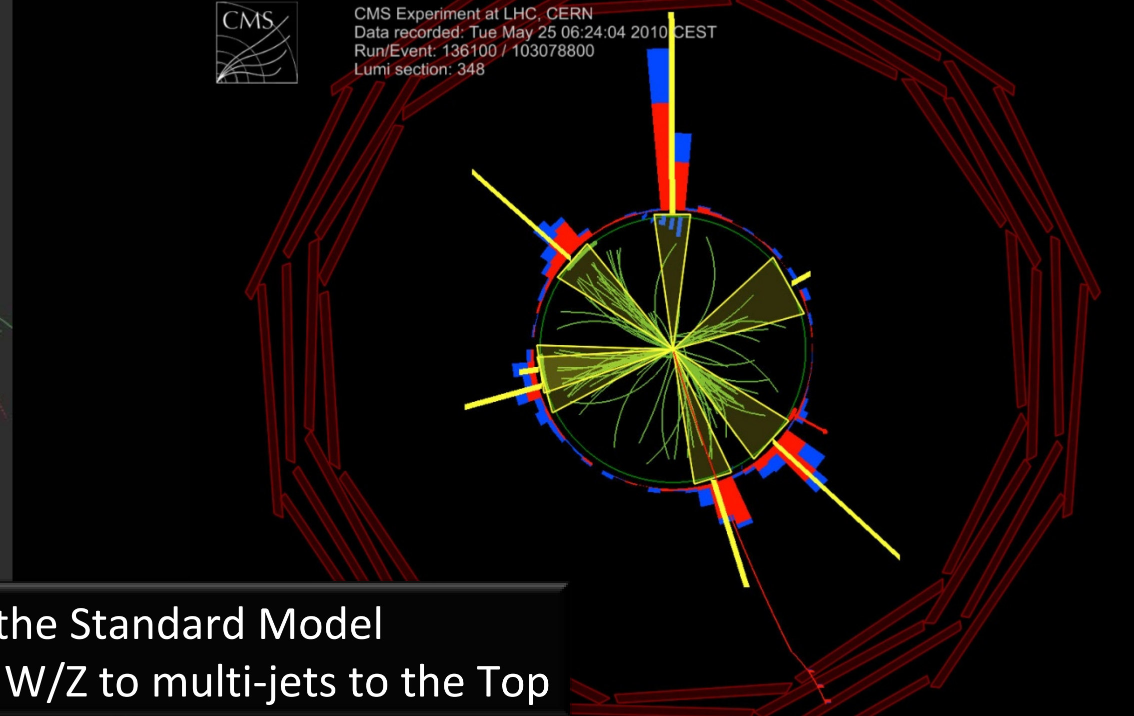

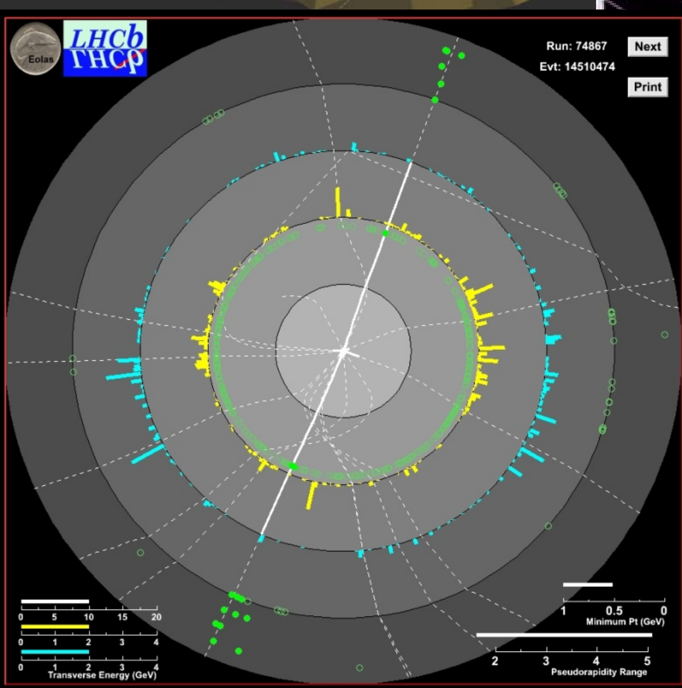

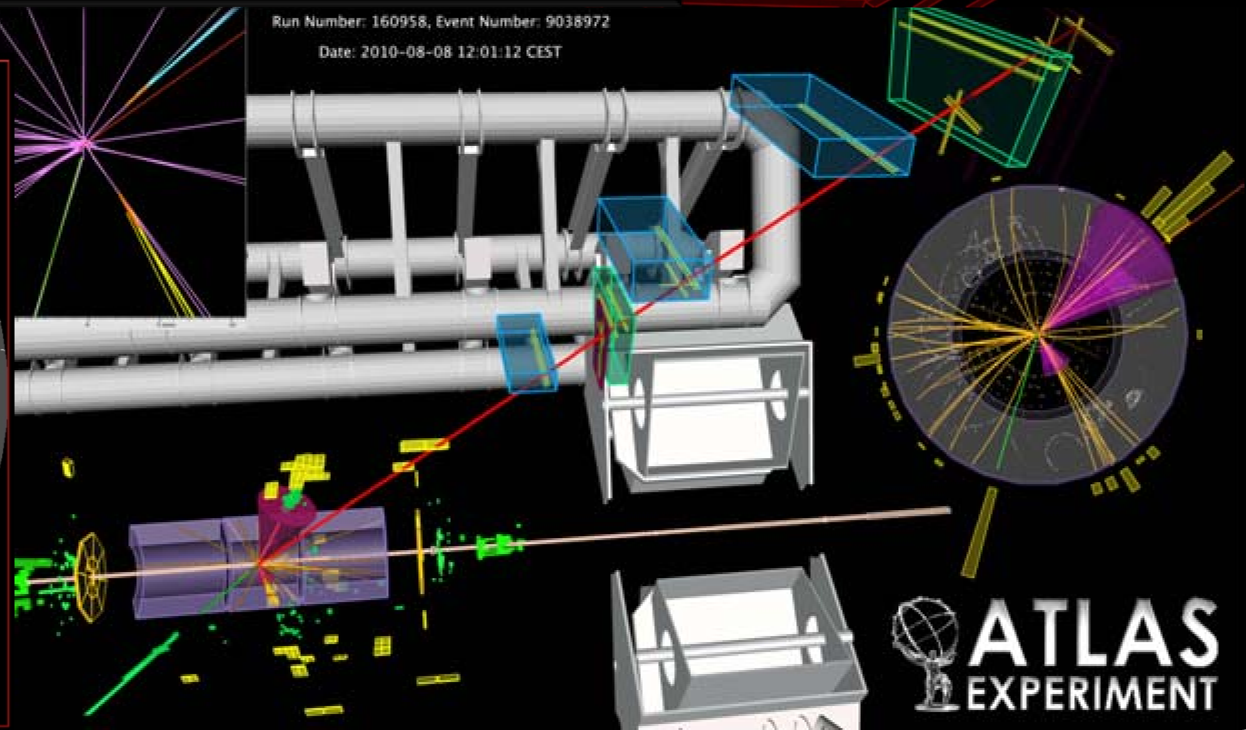

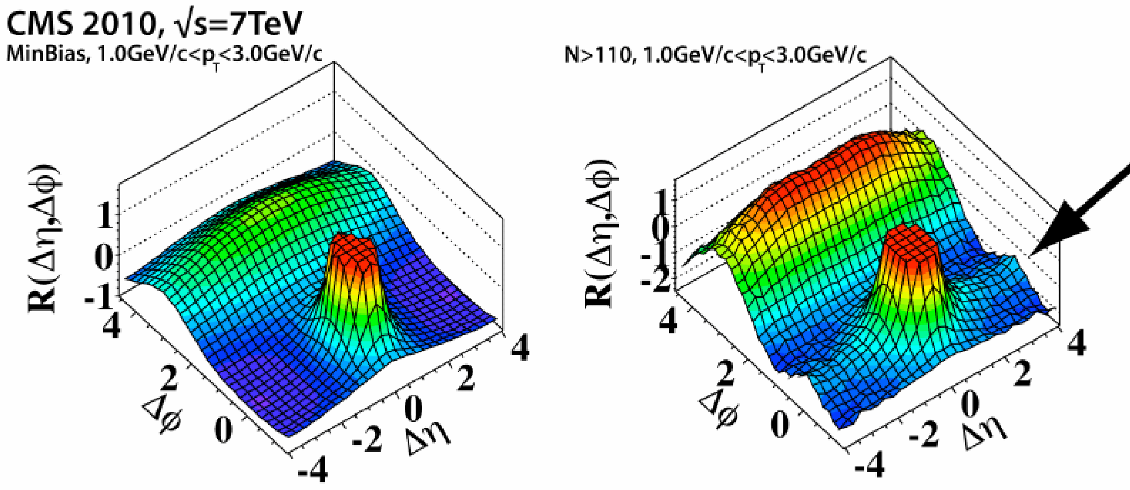

Figure 2: The variation of R with Δη and Δφ, for proton-proton collisions in CMS. Left: for minimum bias collisions; Right: for collisions that produced at least 110 charged particles.

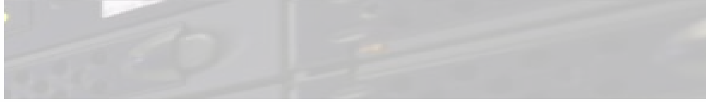

ATLAS sets world's best limits on q*

7 September 2010

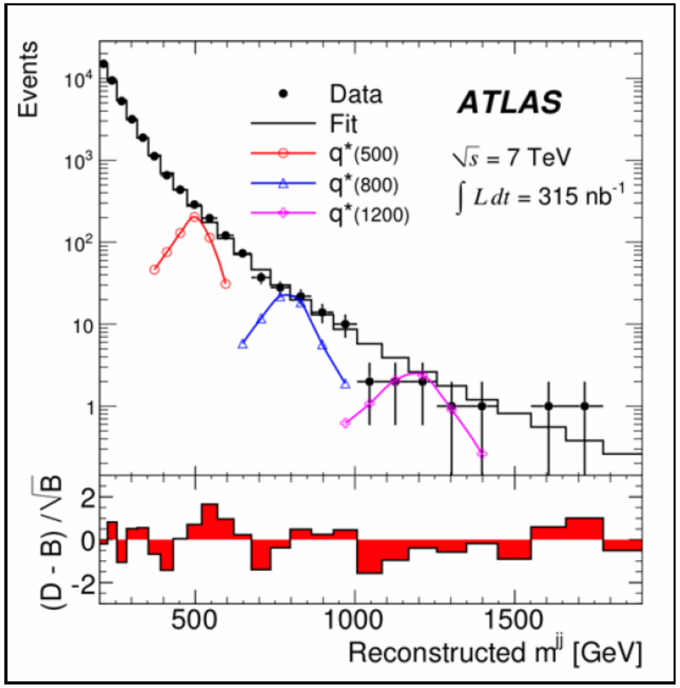

Reconstructed dijet mass (filled points) fitted to a smooth background distribution, with predicted q* signal indicated at three different masses.

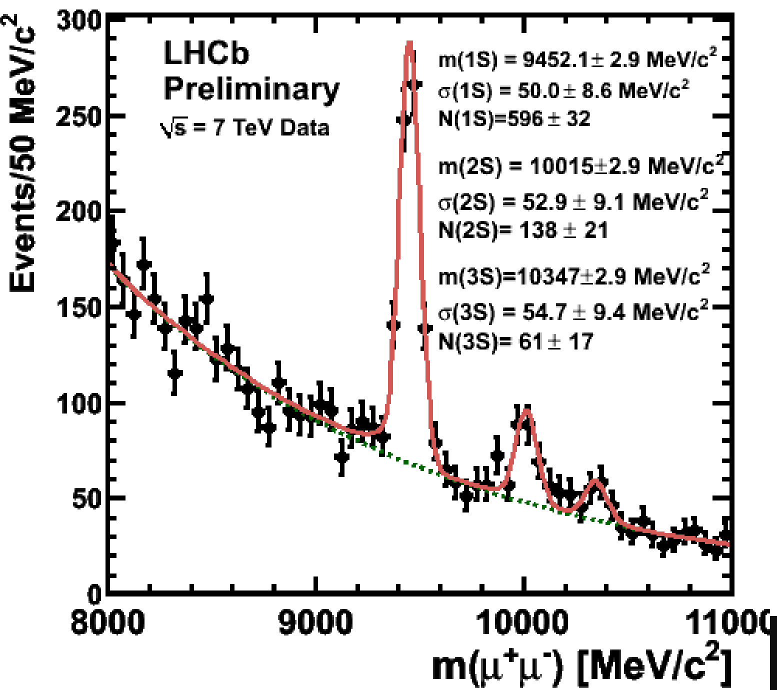

- Progress == real physics output in a very short time ...
- "rediscovery" of Standard Model
- Starting to see improvements on state of the art, and hints ...

Some observations

- ✓ Experiments have truly distributed models
- ✓ Network traffic is close to what was planned – and the network is extremely reliable
- ✓ Significant numbers of people (hundreds) successfully doing analysis – at Tier 2s
- ✓ Physics output in a very short time - unprecedented
- ✓ Today resources are plentiful, and not yet full; This will surely change ...
- ✓ Needs a lot of support and interactions with sites – heavy but supportable

Other observations

- Availability of grid sites is hard to maintain...
  - 1 power "event"/year at each of 12 Tier0/1 sites is 1/month
  - DB issues are common – use is out of the norm?
  - Still takes considerable effort to manage these problems
- Problems are (surprisingly?) generally not middleware related ...
- Actual use cases today are far simpler than the grid middleware attempted to provide for
- Advent of "pilot jobs" changes the need for brokering
- Hardware is not reliable, no matter if it is commodity or not; RAID controllers are a spof
  - We have 100 PB disk worldwide – something is always failing
- We must learn how to make a reliable system from unreliable hardware
- Applications must (really!) realise that:
  - the network is reliable,
  - resources can appear and disappear,
  - data may not be where you thought it was
    - even at 0.1% this is a problem if you rely on it at these scales!

Providing reliable data management is still an outstanding problem

What is WLCG today?

An infrastructure demonstrated to be able to support LHC data processing and analysis: consisting of (in order of importance!)

- Above all a collaboration:
  - Single point of coordination, communication, requirements synthesis, ...
  - Vehicle to coordinate resources and funding
  - Extremely useful in organising with technology and service providers
  - In EC-speak WLCG is a highly structured community – via the MoU
- A service:
  - WLCG provides a common operations coordination and management – on top of EGEE/EGI, OSG, and others
  - Security coordination – operational and policy development
  - World-wide trust federation of CA's and VO's
  - The Policy framework was indispensable in actually deploying WLCG across the world
- An implementation of a distributed computing infrastructure
  - Today with grid technology and higher level (WLCG and experiment-specific) middleware
  - Tomorrow ...

Must be able to evolve the technical implementation (grid→ XX) without breaking the collaboration and service

Evolution and sustainability

- Need to adapt to changing technologies; e.g.:
  - Major re-think of storage and data access
  - Use of many-core CPUs (and other processor types?)
  - Filesystems, etc.
  - Virtualisation as a solution for job management
    - Brings us in line with industrial technology
    - Integration with public and commercial clouds
- Network infrastructure
  - This is the most reliable service we have
  - Invest in networks and make full use of the distributed system
- Grid Middleware
  - Complexity of today's middleware compared to the actual use cases
  - Evolve by using more "standard" technologies: e.g. Message Brokers, Monitoring systems are first steps

Areas for evolution – 1

- End to end usable and transparent networks and data movement:
  - for the 100 Gigabit era;
- Data issues:
  - Data management and access
  - How to make reliable systems from commodity (or expensive!) hardware
  - Fault tolerance
  - Data preservation and open access
- Global AAI:
  - SSO
  - Evolution/replacement/hiding of today's X509
  - Use existing ID federations?
  - Integrate with commercial/opensource software?

Evolution – 2

- Sustenance and evolution of the current grid middleware services for maintainability and effectiveness (e-infrastructure):
  - Evolve towards sustainable and externally provided/supported services for fabric layer
  - Improvement in efficiencies for use, support and maintenance
- Technology evolution from research to end-to-end production deployment in research/scientific codes:
  - Multi-core
  - Virtualized environments.
  - Commercial clouds.
  - GPUs
- Green computing
  - Innovations in end-to-end application, middleware and fabric design, technology and process that increase energy efficiency
  - Use of remote data centres

Evolution of Data Management

- 1st workshop held in June
  - Recognition that network as a very reliable resource can optimize the use of the storage and CPU resources
    - The strict hierarchical MONARC model is no longer necessary
  - Simplification of use of tape and the interfaces
  - Use disk resources more as a cache
  - Recognize that not all data has to be local at a site for a job to run – allow remote access (or fetch to a local cache)
    - Often faster to fetch a file from a remote site than from local tape
- Data management software will evolve
  - A number of short term prototypes have been proposed
  - Simplify the interfaces where possible; hide details from end-users
- Experiment models will evolve
  - To accept that information in a distributed system cannot be fully up-to-date; use remote access to data and caching mechanisms to improve overall robustness
- Timescale: 2013 LHC run

Some observations

- Small subset of data that is distributed is actually used
- Don’t know a priori which data type will be used most (but this will change in time)
- Data is popular for a short time only

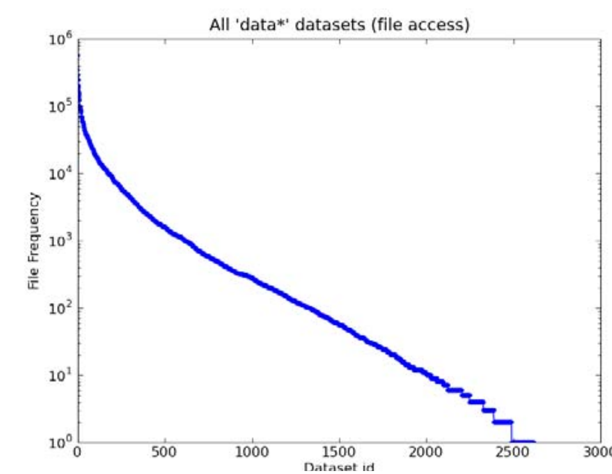

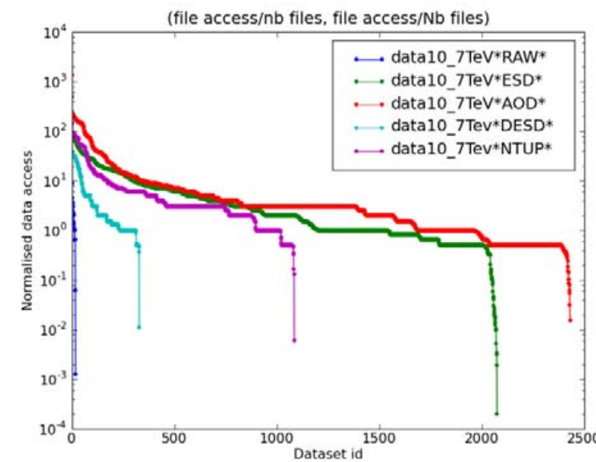

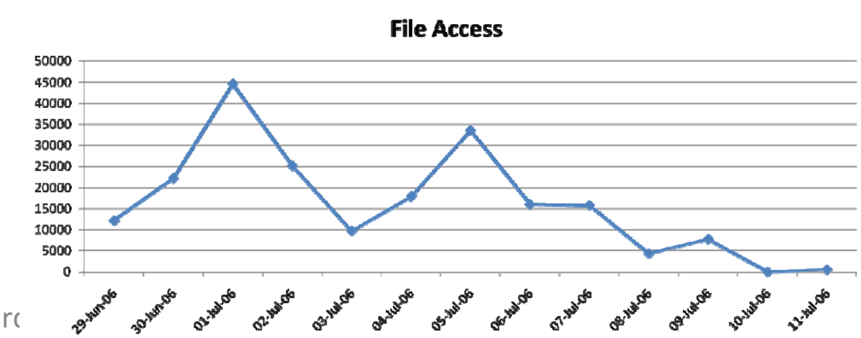

Data placement model

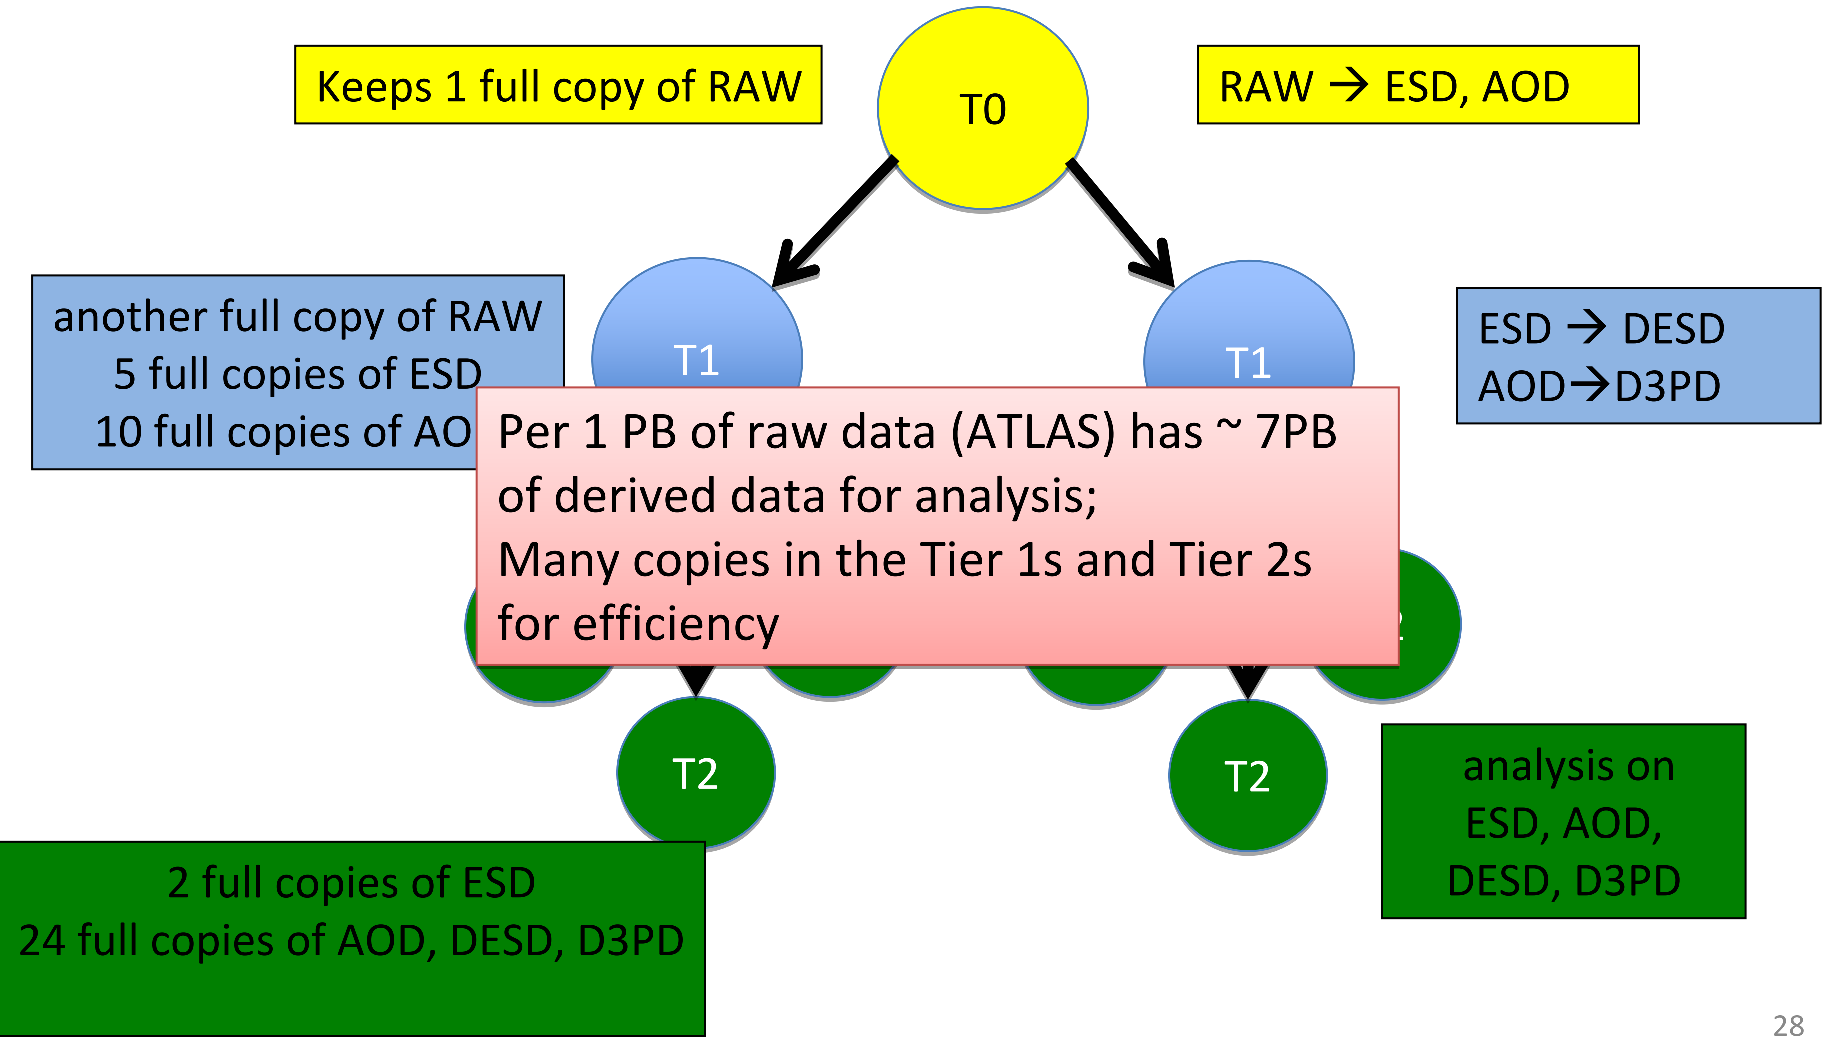

Evolution of data placement

- Move towards caching of data rather than strict planned placement
- Download the data when required
  - Selects popular datasets automatically
  - When datasets no longer used will be replaced in the caches
- Data sources can be any (Tier 0, 1, 2)
- Can still do some level of intelligent pre-placement
- Understanding a distributed system built on unreliable and asynchronous components means
  - Accepting that catalogues may be not fully updated
  - Data may not be where you thought it was
  - Thus must allow remote access to data (either by caching on demand and/or by remote file access)

ICHEP PARIS 2010

Network Requirements

Part of the requirements are already well covered by the OPN.

For controlled (re-) processing:

- Data Distribution from Tier-0 to Tier-1s
  - Initial data from the detector and from first pass reconstruction

- Data Distribution from Tier-1 to all other Tier-1's
  - After re-processing of the initial data in the Tier-1's

- Data Distribution from Tier-1s to some Tier-2s
  - After re-processing to distribute derived data

For uncontrolled data analysis:

- Data Distribution from all Tier-1s to all Tier-2s
  - For further derived data for/from analysis

- Data Distribution from any Tier-2 to any other Tier-2
  - For further derived data for/from analysis

To allow for a full caching model additional services are needed.

Tier-2 Analysis Bandwidth Requirements

- Based on **CPU** capacity
  - A typical Tier-2 site with 1000 cores, a typical rate of 25 Hz for AOD analysis, …

1 Gb/s

- Based on **cache turnover** after re-processing
  - A typical 1 week turnover of a typical 400 TB cache, …

5 Gb/s

- Based on analysis efficiency and **user expectations**
  - A typical 1 day latency for a 25 TB analysis sample, …..

3 Gb/s

Tier-2 Connectivity Categories

- Minimal
  - Small Tier-2s, well suited for end-use analysis

1 Gb/s

- Nominal
  - Nominal sized Tier-2s , big analysis samples can be updated regularly

5 Gb/s

- Leadership
  - Large Analysis Centers, supporting many users, frequent cache turnovers

10 Gb/s

Meant is shared, best effort connectivity,
not guaranteed bandwidth between each of the sites

So what do we need from networks?

- Bandwidth, Reliability, Connectivity
  - Not forgetting that we have collaborators on 6 continents
    - Including outer edges of Europe
  - Have set up wg to express these requirements in conjunction with network communities
- But we also need a service:
  - Monitoring
    - Is largely missing today – we have a hard time to understand if there is a network problem
  - Operational support
    - Is a complex problem with many organisations involved. Who owns problem? How can we (user) track progress? Etc. etc.

Sharing?

- Today we have enough bandwidth (on the OPN and many T1-T2 and T2-T2 links)
- Have FTS which allows us to share bandwidth between experiments (if they all use it!)
  - But don't really use that feature today
  - Network today is not the choke point ...
- Will we need to do bandwidth sharing? Or over-provision?
  - And at which level is that sharing to be done?
  - This sounds like a complexity which we should try and avoid

Conclusions

- Distributed computing for LHC is a reality and enables physics output in a very short time
- Experience with real data and real users suggests areas for improvement –
  - The infrastructure of WLCG can support evolution of the technology
- Networks will be a key for that evolution
  - We must fully embrace the distributed model