



Université Mohammed V
Faculté des Sciences
Rabat



Computing challenges of the LHC Big Data in High Energy Physics

Farida Fassi

Faculty of Sciences


Mohammed V University in Rabat, Morocco

Computing Training course
I-COOP+ 2016 project: COOPB20247
3-14, July 2017. IFIC-Valencia, Spain

Why LHC?

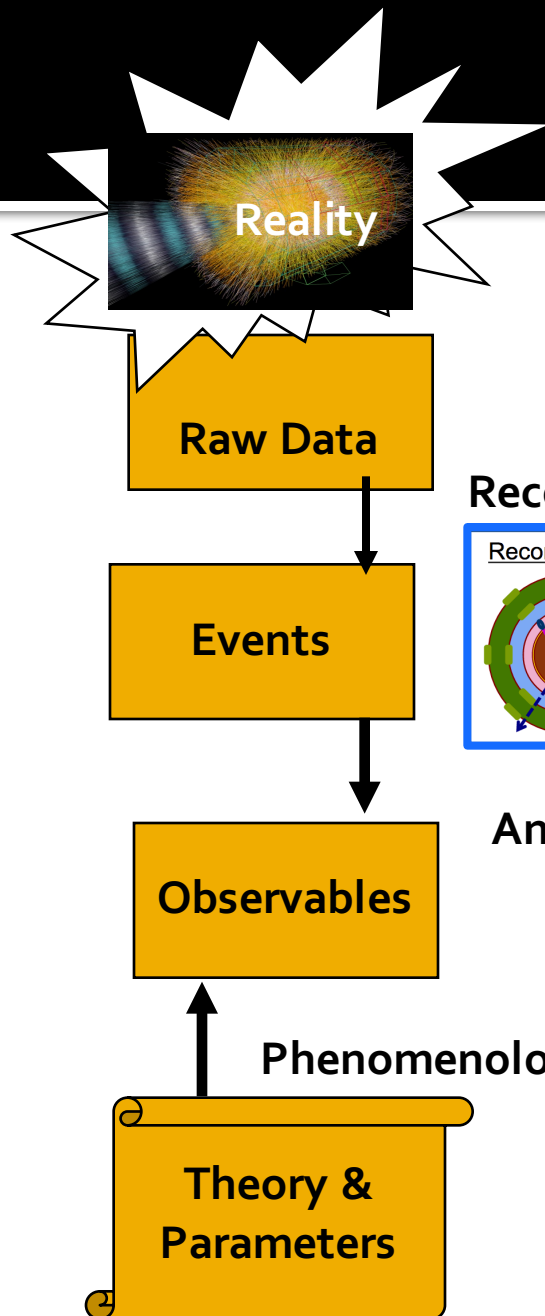
Fundamental questions of Particle Physics...

- ❖ The universe is governed by probabilistic physics
- ❖ One measurement tells us very little
- ❖ **However carefully we set up an experiment**, probabilistic physics decides what we observe to provide an experimental verification of different theories within Particle Physics...
- ❖ *What is the dark matter in the Universe?*
- ❖ *Unification of fundamental forces?*
- ❖ *Understanding space time matter versus antimatter.*
- ❖
- ❖ Astrophysics/cosmological measurements show that **only 5%** of the matter is known.
- ❖ This tiny fraction of matter is well described by the “Standard Model” of Particle Physics.






Our model so far...

We use experiments to inquire about what “reality” (nature) does






Reconstruction




Analysis

We “confront theory with experiment” by comparing what we measured, with what we expected from our hypothesis.



New physics rate at LHC




- ◆ If we want to observe something rare, we may have **to find a few events hidden in vast numbers of other events**
- ◆ Notice the logarithmic scale on the Y-axis: it spans 11 orders of magnitude
- ◆ E.g. you produce 10 Higgs bosons out of 10^{11} billions of collisions
- ◆ The probability increases logarithmically with energy

New physics rate ~ 0.00001 Hz

Event Selection:
1 in 10,000,000,000,000

Tools: The Experiments at the LHC




LHC




CMS

General Purpose,
pp, heavy ions


pp, B-Physics,
CP-Violation
(matter-antimatter
symmetry)



LHC-B
Detector



ATLAS

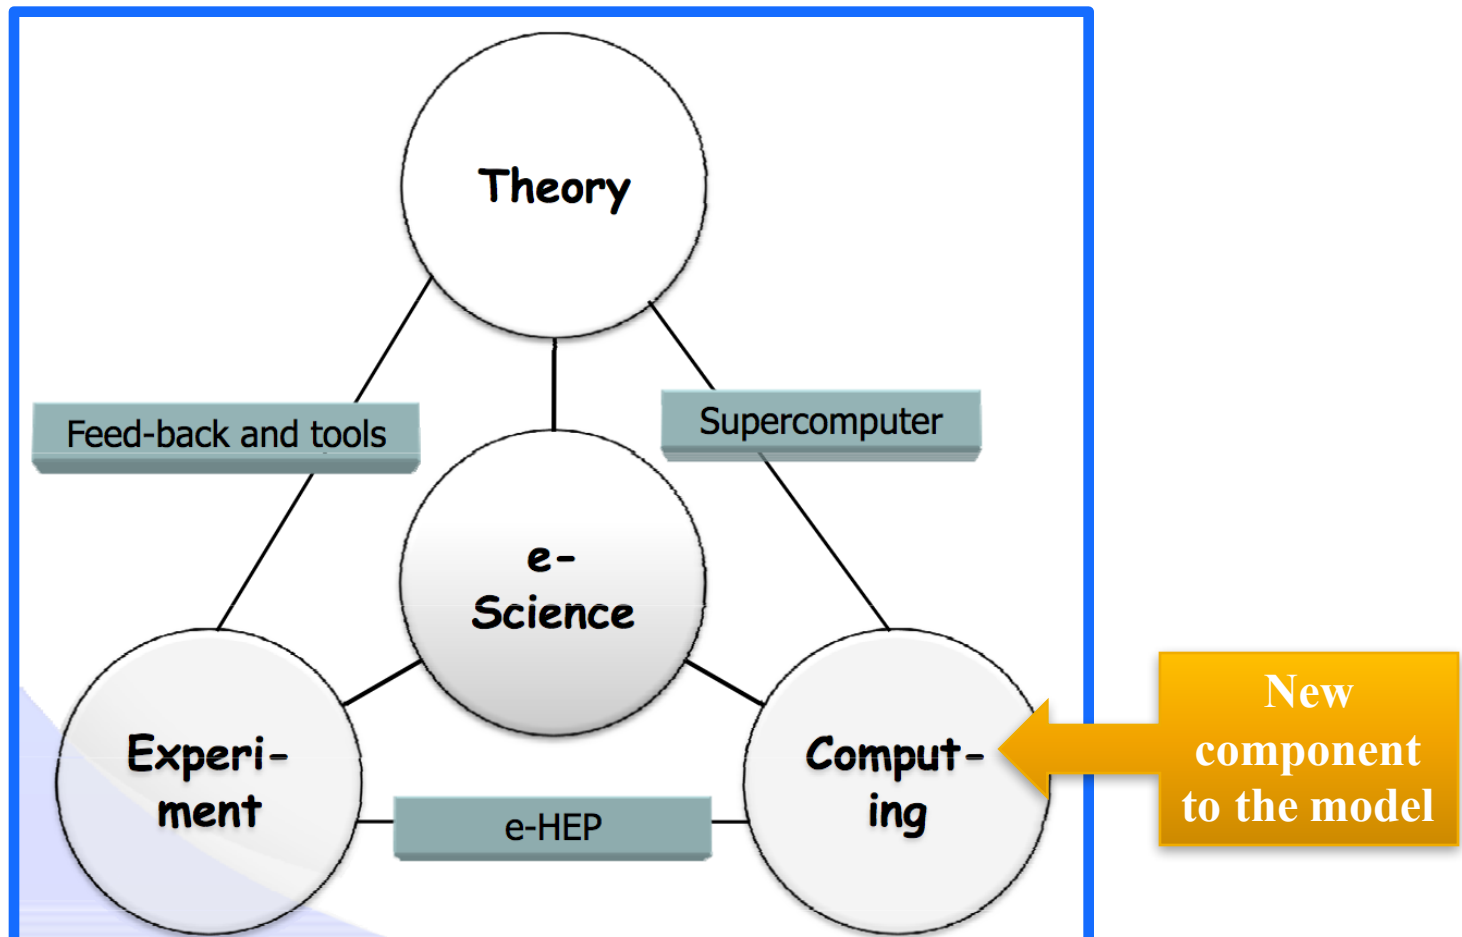


ALICE

Heavy ions, pp
(state of matter of
early universe)

e-Science paradigm concept

- The paradigm of e-Science
 - Experiment-Computing-Theory



e-Science paradigm in High Energy Physics (HEP)

- ❖ Large Hadron Collider (LHC) is the largest and most powerful particle accelerator ever built.


- ❖ Data volumes at the LHC
40 million collisions per second

After filtering, only few hundreds of collisions of interest per second


10^{10} collisions recorded each year

=> more than 25 Petabytes/year of data


- ❖ High Energy Physics (HEP) uses immense data sets that require the computational Grids infrastructure deployed in the framework of the Worldwide LHC Computing Grid (WLCG).




Big Data concept in HEP



- LHC produces tens of trillions of proton-proton collisions every day at the centre of four large detectors
 - ALICE, ATLAS, CMS and LHCb



- ATLAS “sees” bunches of collisions (tens of superimposed events) every 25ns
- That is 40 million/second or about 15 trillion bunch collisions per year
- If all data would be recorded that would lead to 100000 CDs per second

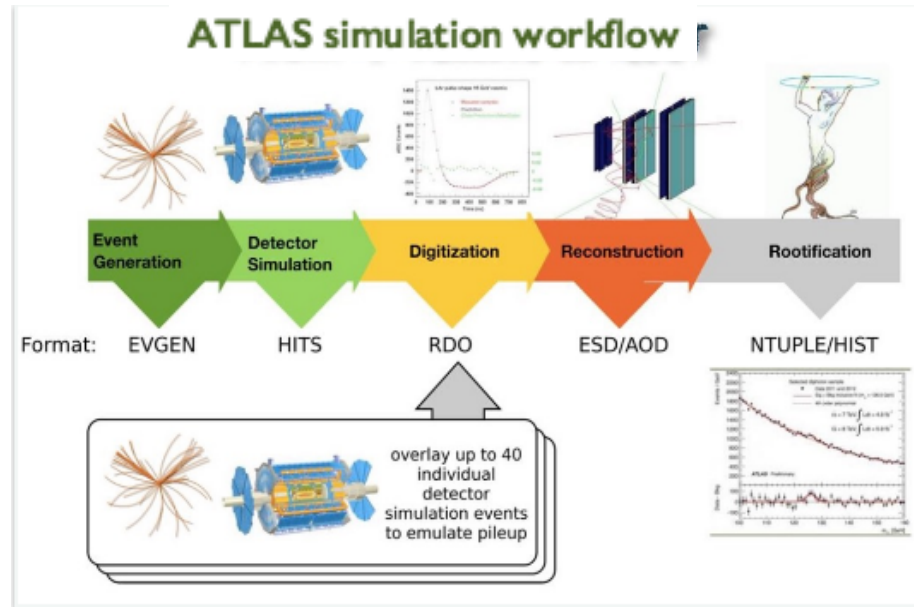


Big Data



Big Data in HEP: Real and Simulated Data

- ❖ The **Raw data** collected from the LHC is only part of the bigger data picture.
- ❖ MonteCarlo Simulation models the evolution of physics processes from collision to digital signals using knowledge from theory and test data.
- ❖ Translate theoretical models into detector observations.
- ❖ Proper treatment of background estimation and sources of systematic errors.
- ❖ **10 billion events simulated by ATLAS to date**
- ❖ Data-driven analysis compares (at statistical level) reconstructed events from real data with those predicted by simulation.



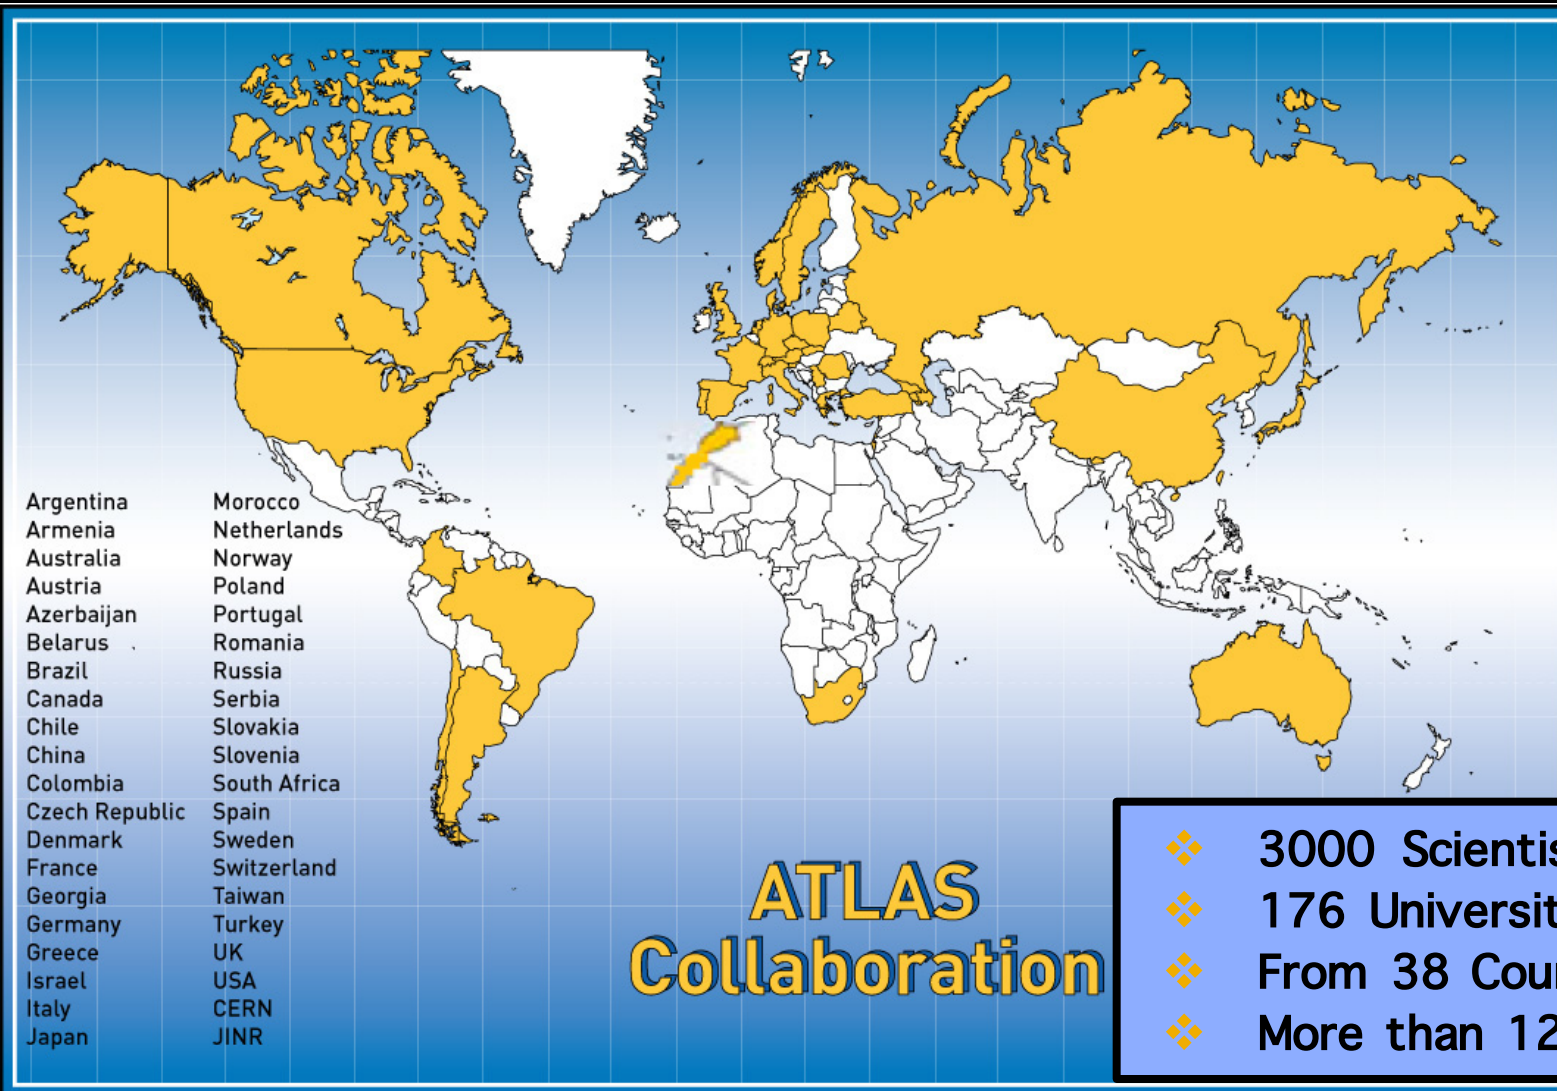
Big Data in HEP: Real and Simulated Data

- **Raw data** rate from LHC detectors
- **About one Petabyte per second**
- **This would cost about 1 trillion euros** for storage

Simulation in HEP


- The fundamental physics processes are only measurable after non-invertible transformations from:
 - Physics (e.g. quarks/Jets 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**
- **Need much more simulation than data to make the uncertainty contribution from simulation statistics negligible.**

HEP community more and more global



- ❖ 3000 Scientists
- ❖ 176 Universities and Labs
- ❖ From 38 Countries,
- ❖ More than 1200 students

LHC Big Data challenges...

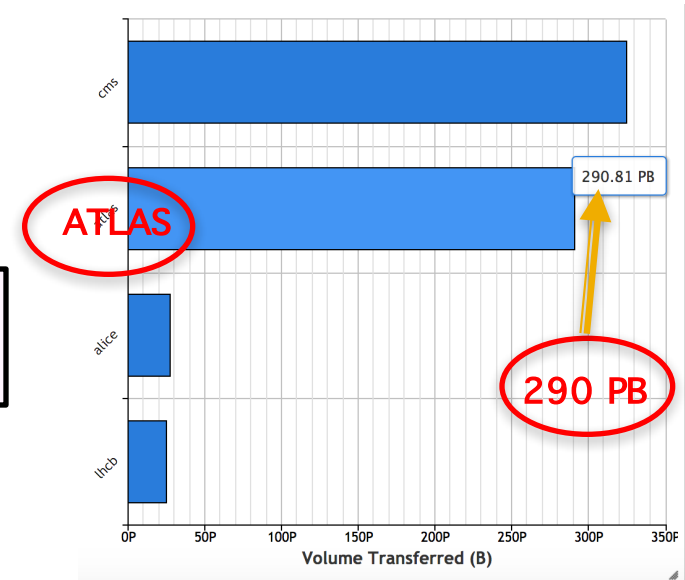


❖ The LHC has already delivered billions of recorded collision events.

- ❖ Over **290 PB** of data transferred
- ❖ **>100 PB** more storage needed for data replication, simulation and Analysis derivation.

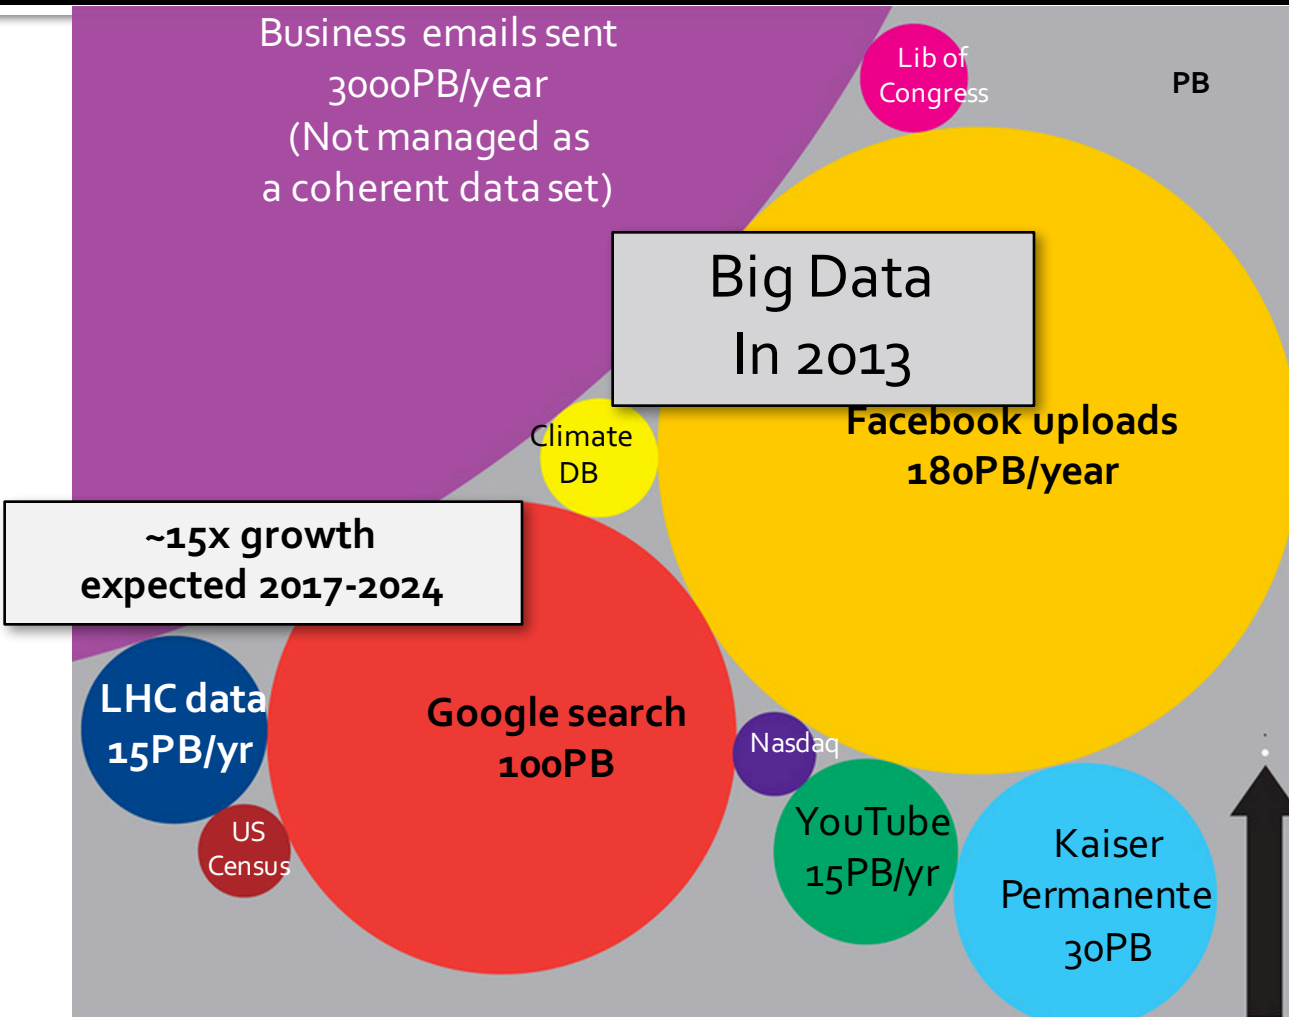
❖ Enormous challenge for the experiments for data collection, storage and processing

- ❖ **What is this data?**
- ❖ read-out of $\mathcal{O}(100\text{M})$ detector channels
- ❖ 150 Million sensors deliver data
 - ❖ 40 Million times per second
- ❖ Raw data rate from LHC detector: **1PB/s**
- ❖ This translates to Petabytes of data recorded world-wide (Grid)




❖ Grid Computing is a critical tool to address the Big Data processing challenge and produce timely physics results...then success of LHC scientific program!!!

Comparison with other Big Data Applications




>800 PB LHC data moved across World in 2016

ATLAS data set




Current ATLAS data set, all data products: 250+ PB

How to cope with LHC Big Data challenges!!!



The Worldwide LHC Computing Grid (WLCG): Solution to LHC Big Data challenge...


- **WLCG** is a global distributed computing Infrastructure, based on the Grid technologies.
- WLCG provides seamless access to computing power and data storage capacity distributed over the globe.
- Computer centres worldwide arranged in a Tier structure.



Tier	Sites	Role	Example
0	1	Central Facility for data processing	CERN
1	12	Regional computing centres with high quality of service, 24/7 operations, large storage and compute resources	RAL
2	140	Computing centres within a regional cloud used primarily for data analysis and simulation	Edinburgh (ECDF)

**ATLAS Tiered
Computing mode**

Tier-0 at CERN



- ❖ The main goal is to make use of the resources available and integrated into a single infrastructure accessible by all LHC, no matter where they are

From Big Data to Physics Discovery

- *Higgs boson is a major scientific discovery acknowledged by the 2013 Nobel Prize in physics*

Global Effort → Global Success

Results today only possible due to extraordinary performance of accelerators – experiments – Grid computing

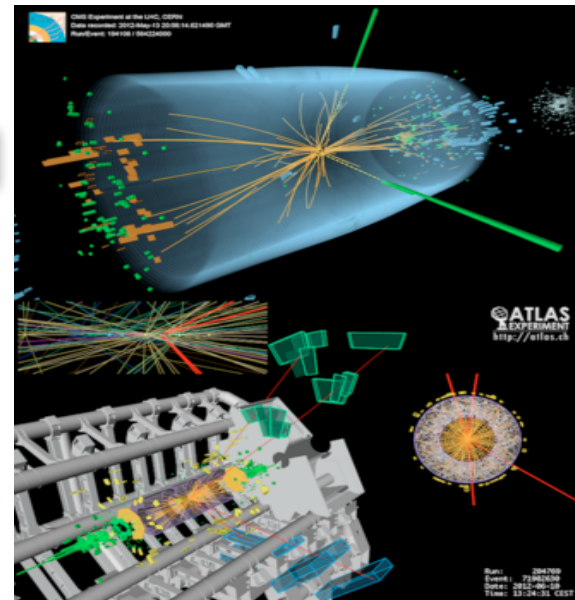
Observation of a new particle consistent with a Higgs Boson (but which one...?)

Historic Milestone but only the beginning

Global Implications for the future



R-D Heuer



Grid computing enables the rapid delivery of physics results


LHC Run 2 performance...

- LHC Run 2 performance is above expectations
- all factors driving computing have increased above expected levels

	2016 exp.	2016	2017 exp.	2018 exp.
Peak Luminosity ($10^{34} \text{ cm}^{-2}\text{s}^{-1}$)	1.0	1.5	1.7-1.9	1.7-1.9
Integr. Luminosity (fb^{-1})	25	40	~45	~45

- For 2016, the available resources were sufficient!!!
 - **More tapes at CERN have been needed**
- Analysis for 2017, 2018

- Expectations are increased, the requirements are ~20% above previous estimates
- Unprecedented peak instantaneous luminosity > 40% beyond LHC design
- **Data accumulation ~60% beyond 25 fb^{-1} goal for 2016**



Future LHC Computing Needs!



CPU:


- x60 from 2016

- CPU needs (per event) will grow with track multiplicity (pileup) and energy.
- Storage needs are proportional to accumulated luminosity.
- Grid resources are limited by funding and fully utilized.

Data:

- RAW 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB



LHC RAW Data Volume



HL-LHC computing challenges



- **HL-LHC data analysis**
 - To extract physics results requires to handle/analyze a lot more data!
- **Tests started with new technologies**
 - **“Big Data” technology (new toolkits and systems to support analysis of datasets in industry)**
 - **Cloud Computing and High Performance Computing (HPC)**
- **Educates our community to use industry-based technologies**
- **Use tools developed in larger communities reaching out side of our field**

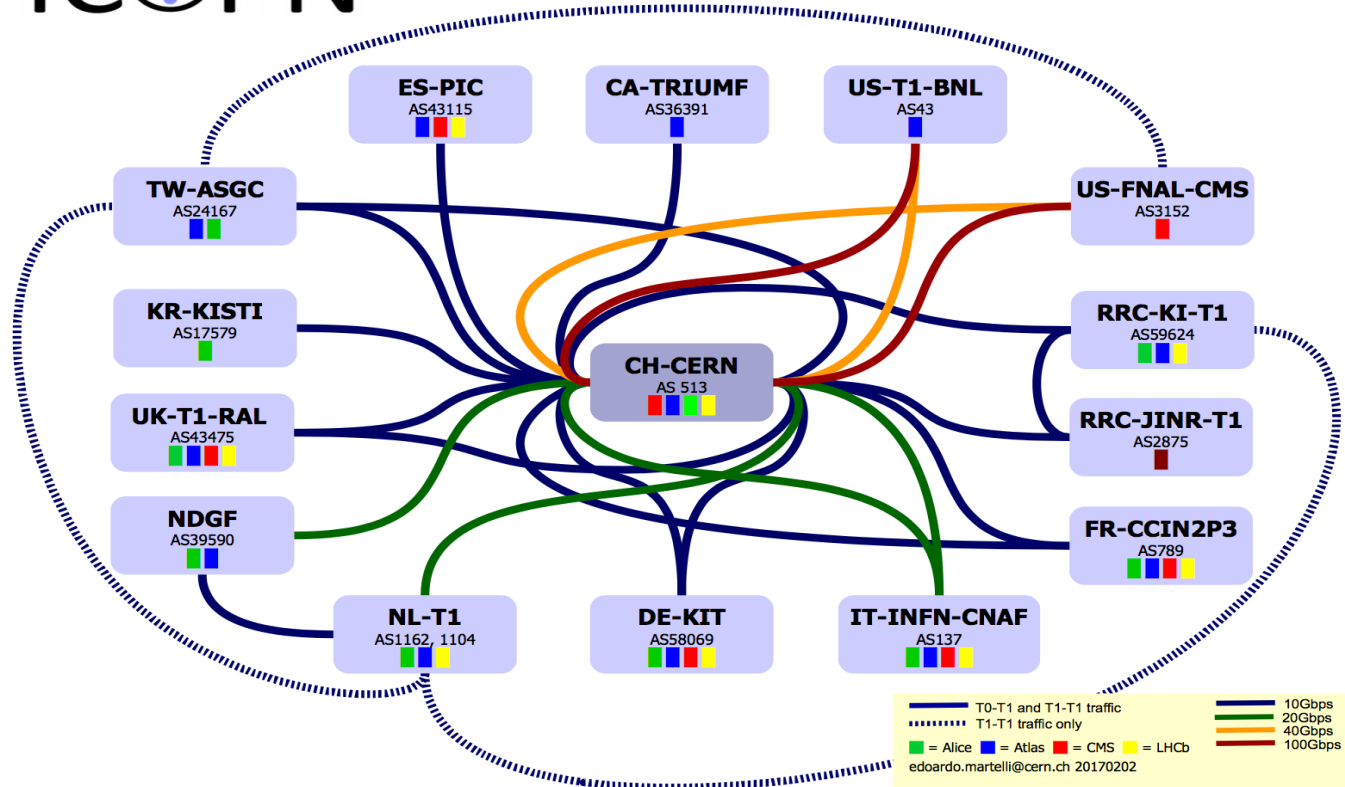


Advanced Networks for HEP...

- **LHC Run 1 brought us a centennial discovery: the Higgs Boson**
- **LHC Run 2 will bring us (at least) greater knowledge, and perhaps greater discoveries: Physics beyond the Standard Model.**
- **Advanced networks will continue to be a key to the discoveries in HEP**
- **Technology evolution might fulfill the short term needs!**
- **A new paradigm of global circuit based networks will need to emerge during LHC Run2**
 - **New approaches and a new class of global networked systems to handle Exabyte scale data are needed (building on LHCONE)**
 - **Worldwide deployment of such systems by 2023 will be essential for the High Luminosity LHC HL-LHC**

LHCOPN - Large Hadron Collider Optical Private Network

LHCOPN



LHCOPN evolution

- The LHCOPN is kept as the main network to exchange data among Tier0 and Tier1s
- **Links being upgraded to multiple 10Gbps and 100Gbps (gigabyte per second)**

3rd 100G Geneva-Budapest link



T-Systems – GEANT#1 – GEANT#2

Technical challenges:

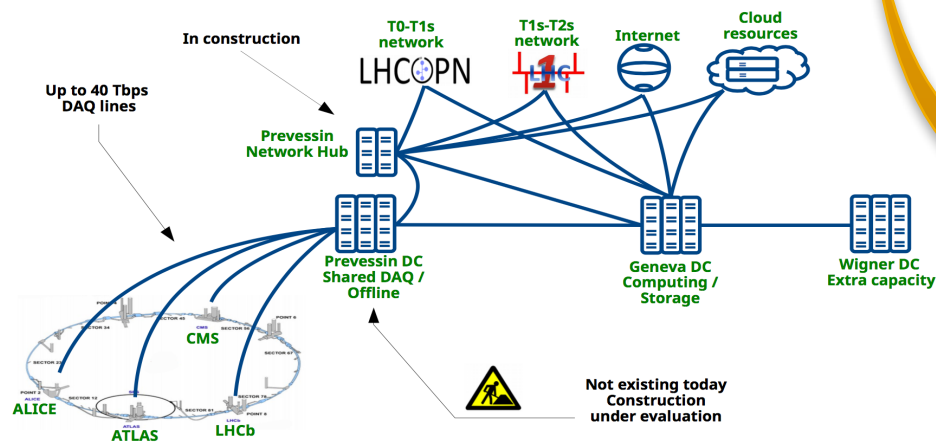
❖ **Implications on networks needed**

❖ Optimization of the physics output vs cost

❖ Software, algorithms, computing models, distributed infrastructure

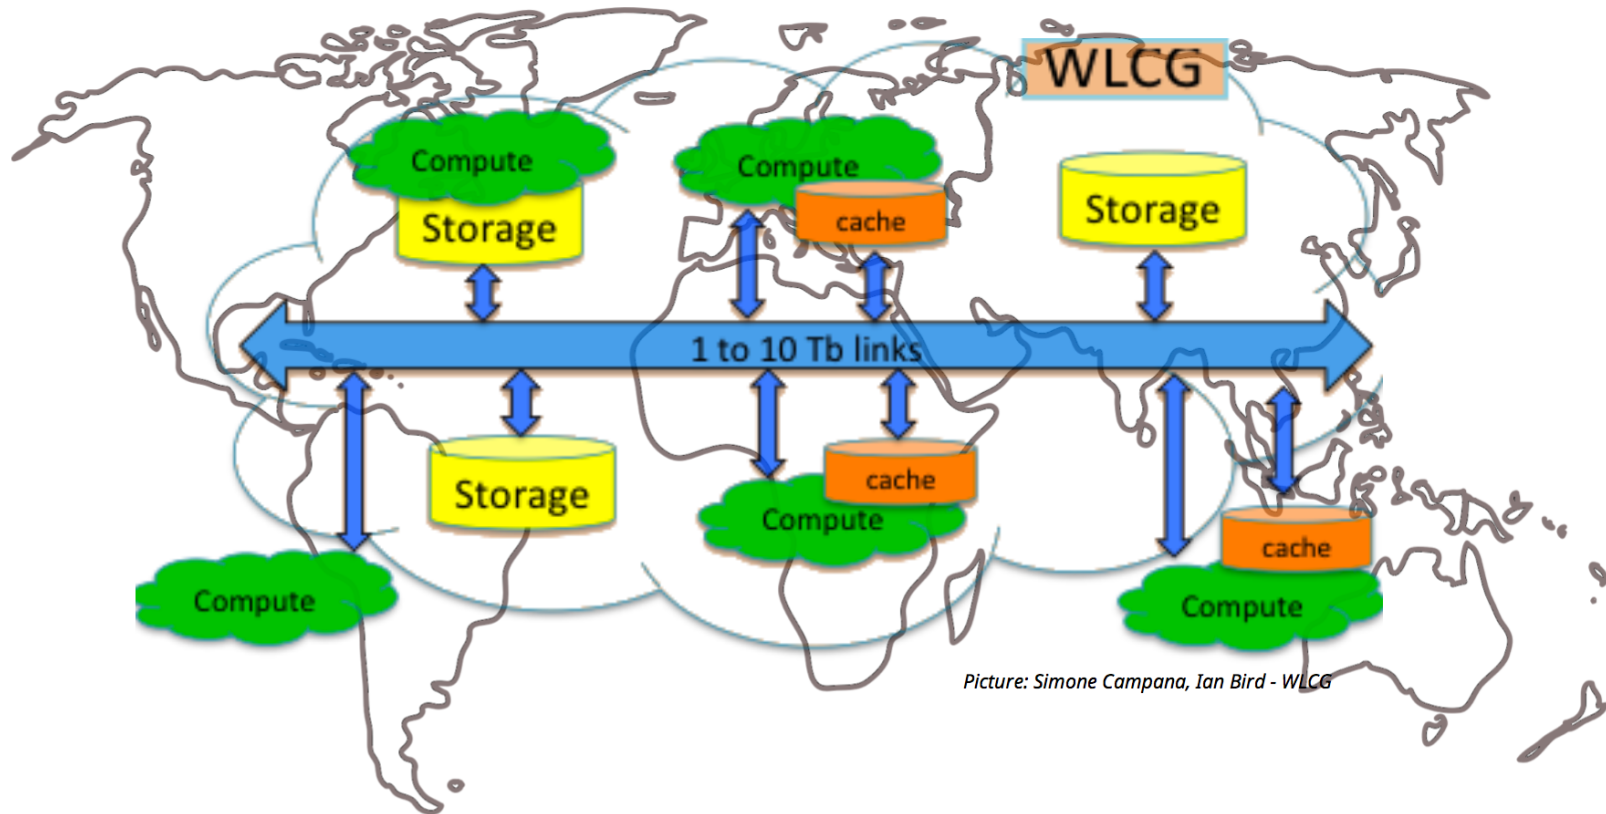


IT Information Technology Department



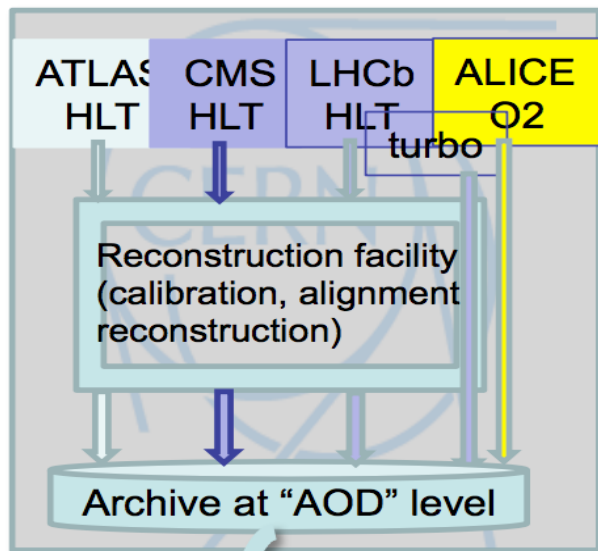
3x100 GB/s Links
CERN ↔ Wigner Data Centre
(Budapest)
as extension of CERN T0 in
production

Possible change of computing model for Run3



More performance Bandwidth dedicated to HEP data transfers is critical for the new model

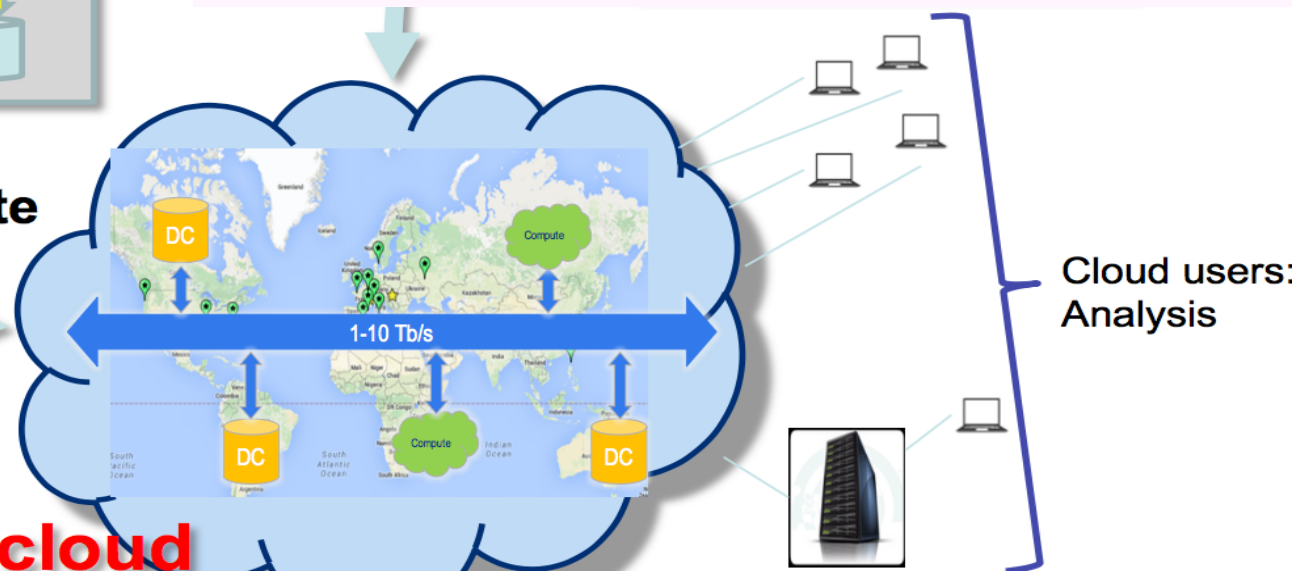
Possible change of computing model for Run3 and beyond



Distribute

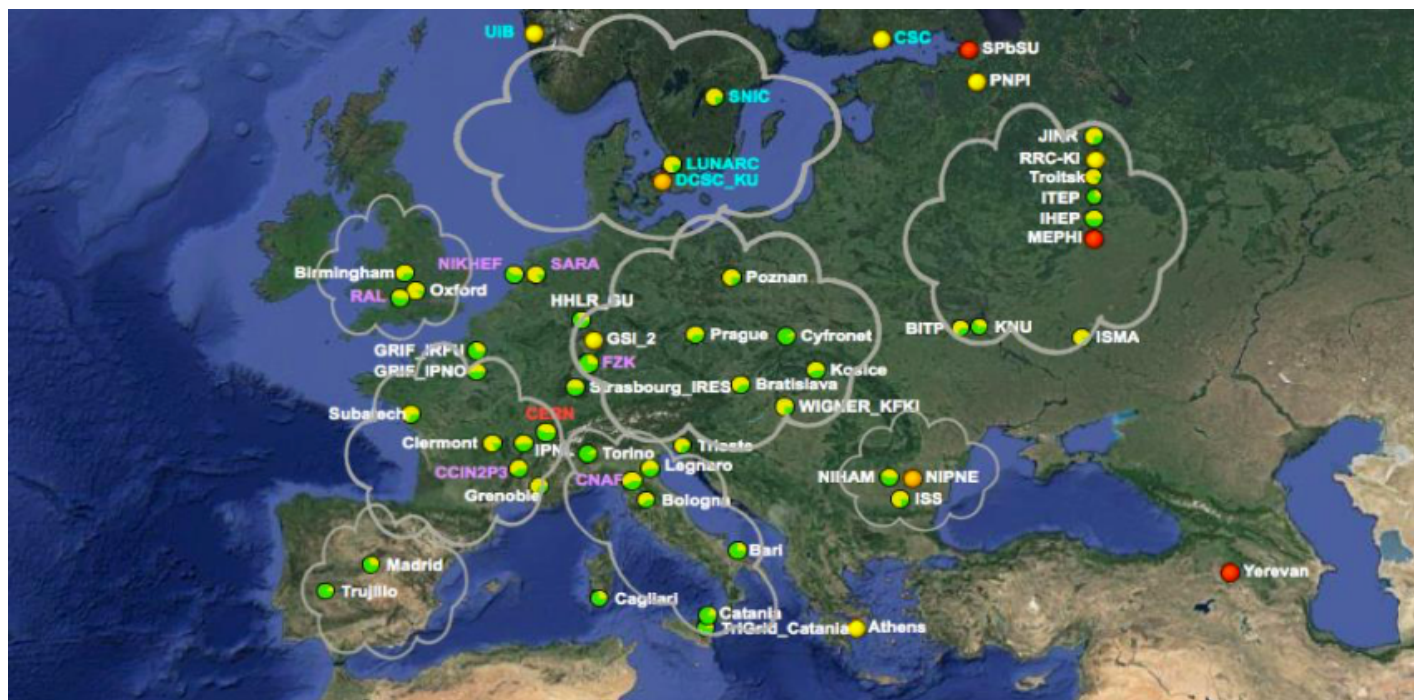
LHC Data cloud
Storage and computing

- **Need to be able to deploy data to a diverse set of resources**
 - Clouds, WLCG sites, HPC centers, etc
- **Will need to a combination of real time delivery and advanced data caching**
- **In order to replicate samples of hundreds TB in hours we will need the systems optimized end-to-end and a very high capacity network in between.**



ALICE cloud model: Reducing Complexity


- ❖ Virtually joining together the sites based on proximity (latency) and network capacity into Regional Data Clouds
- ❖ Each cloud/region provides reliable data management and sufficient processing capability
- ❖ Dealing with handful of clouds/regions instead of the individual sites



Resources external to WLCG: Amazon, Google, HCP, etc

- Cloud Resources

- private (e.g. institute clusters)
- commercial (Amazon, Google,



z.B. Open Stack



Beispiel: Nutzung der ATLAS & CMS HLT-Farmen während des LS 1

left R&D Phase since long

→ „Grid of Clouds“ is a reality

High Performance Computers were designed for massively parallel applications (different from HEP use case) but we can benefit from single core job slots

- Large HPCs use a variety of architecture

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


Titan:

The largest supercomputer available for scientific applications



Titan System (Cray XK7)			
Peak Performance	27.1 PF 18,688 compute nodes	24.5 PF GPU	2.6 PF CPU
System memory	710 TB total memory		
Interconnect	Gemini High Speed Interconnect		3D Torus
Storage	Lustre Filesystem		32 PB
Archive	High-Performance Storage System (HPSS)		29 PB
I/O Nodes	512 Service and I/O nodes		

LHC collaborations have members with access to these machines and to many others!!!



ATLAS@home: Volunteer Computing

- **People volunteering their PC's spare CPU cycles for science**
 - **most commonly used software is BOINC**
- **ATLAS MC simulation jobs inside a CernVM**
- **Jobs are taken from ATLAS job management system and submitted to BOINC server through ARC CE**
- **Fully integrated with PanDA**
 - **Tasks are assigned to BOINC_MCORE queue**
 - **Jobs and data are hosted on BOINC**
 - **server, volunteers do not talk to Panda/Grid services**
 - **aCT/ARC CE provides the bridge to PanDA**
- **Currently provides 1-2% of simulation resources Grid-wide**

ATLAS@Home as a lightweight site solution

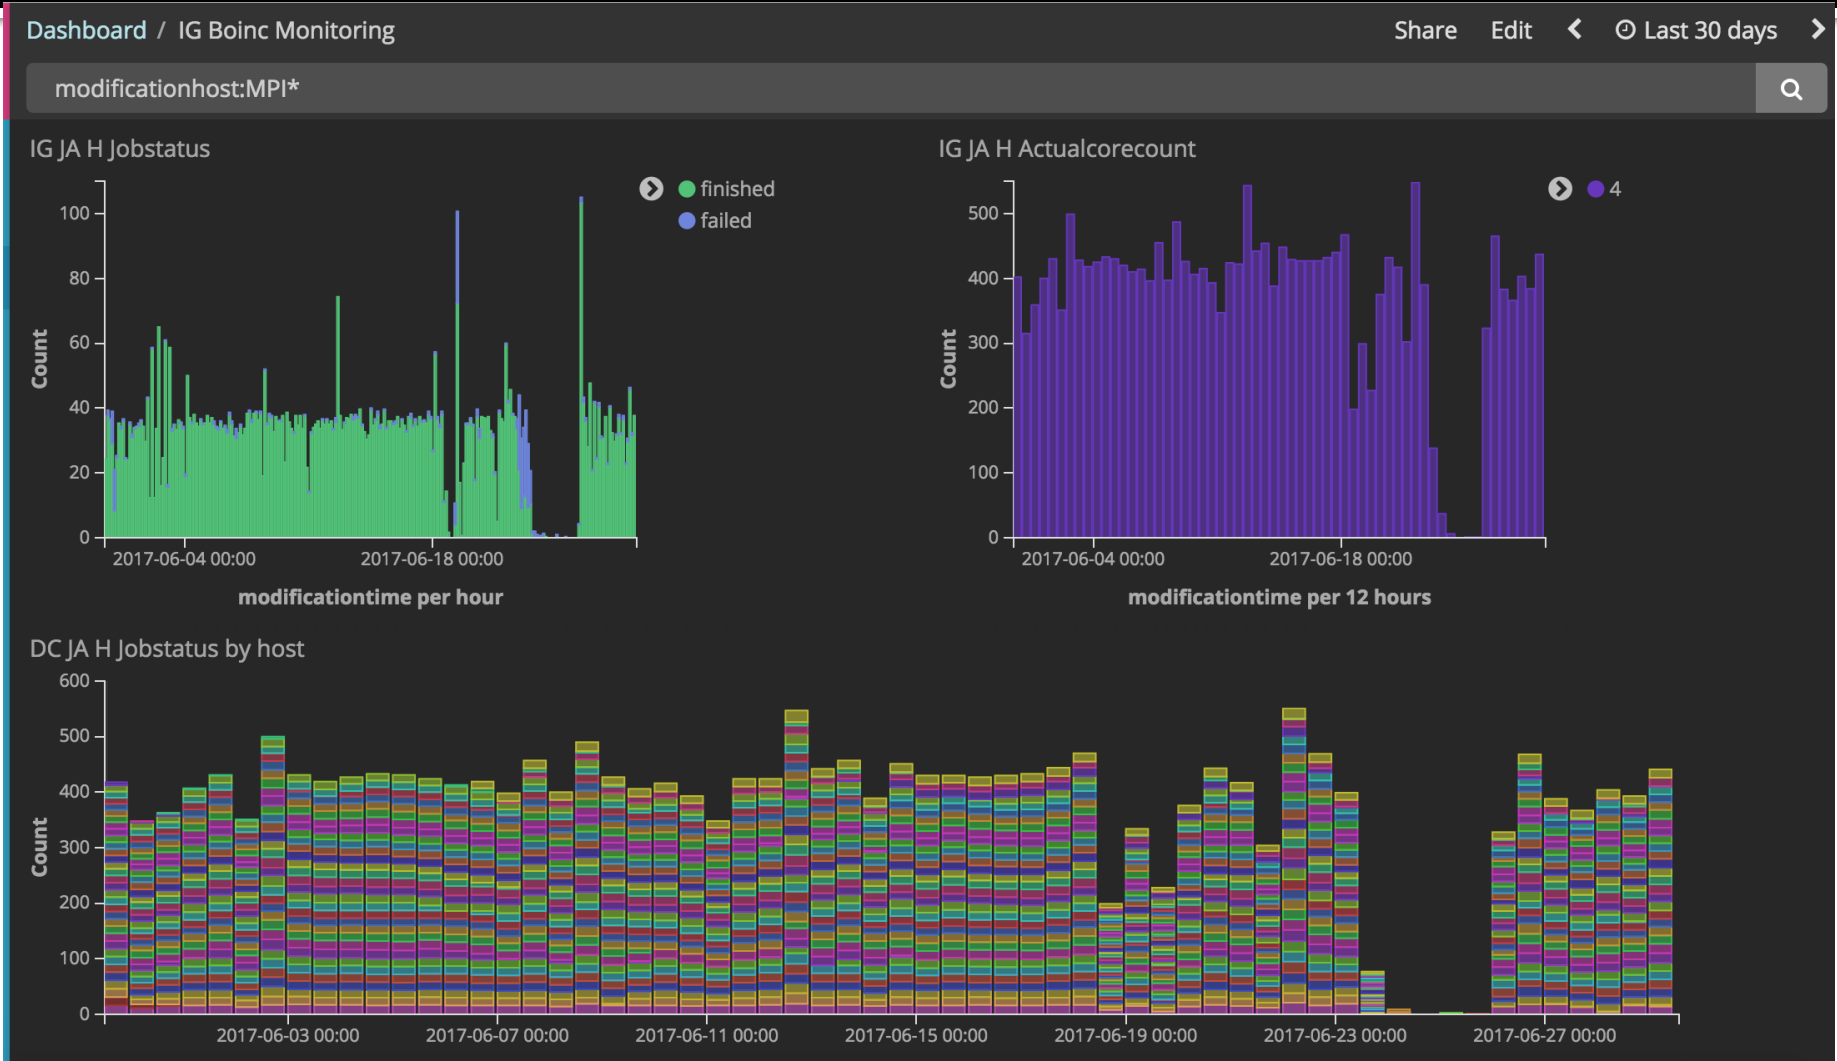
- **How to contribute to ATLA@home:**
 - **Create account on <http://atlasathome.cern.ch>**
 - **Install BOINC client and VirtualBox**
 - **Connect to ATLAS project**
 - **(Multicore) jobs start, using as many cores as available**

ATLAS@Home

- is a very easy way for sites with a small amount of spare resources to contribute to ATLAS Office PCs, old machines, under used machines, new machines awaiting commissioning,
- Just install 2 pieces of software on each node
- No CE, SE, etc services required

As a testbed we may test this at FSR!!!

Kibana dashboard prototype



Cloud-enabling Technologies in HEP

■ CernVM:

- Virtual machine based on Scientific Linux (maintained by CERN)
- Very lightweight, can be directly deployed on various cloud sites

■ HTCondor:

- Free and open-source batch system commonly used in HEP
- Excellent with integrating dynamic worker nodes

- **xRootD** and data federations for remote access

■ Cloud manager (e.g. **ROCED** [KIT]):

- Cloud scheduler that supports multiple cloud APIs
- (OpenStack, Amazon EC2 and other commercial providers)
- Easily extendable thanks to modular design
- Parses HTCondor ClassAds and boots VMs on cloud sites
- depending on the number of queued jobs

■ CernVM-FS:

- On-demand HTTP based file system (Caching via HTTP Proxy)
- Many big experiments use it to deploy software to WLCG compute centres
- works excellently also on cloud sites



European Science Cloud

Helix Nebula Science Cloud



■ Helix Nebula Science Cloud

- European hybrid cloud platform that will support high-performance, data- intensive scientific computing
- for end-users from many research communities:
 - **High-energy physics, astronomy, life sciences, etc**



- Sponsored by 10 of Europe's leading public research organizations and co-funded by the European Commission (H2020).
- **Procurers: CERN, CNRS, DESY, EMBL-EBI, ESRF, IFAE, INFN, KIT, SURFSara, STFC**

Conclusions

- ❖ **Grid Computing has helped deliver physics rapidly at the first LHC Runs...**
 - ❖ **The decade of the Grid and beginning of Big Data challenges**
- ❖ **Big Data era: Entering a phase of computing evolution**
- ❖ **Challenges for computing:**
 - ❖ **scale & complexity**
 - ❖ **will continue to increase dramatically**
- ❖ **New computing models and more efficient software have to be developed**
 - ❖ **Additional resources are needed**
 - ❖ **cloud computing, High-Performance computing, etc**
- ❖ **The distributed computing model allows us to incorporate clouds and supercomputing centers and to use them efficiently now and for LHC Run3 and beyond**

Backup

Phenomenology

- A good theory contains very few numbers
- But it can predict a large number of reactions
- Getting those predictions from the theory is called “phenomenology”

10.4. W and Z decays

The partial decay width for gauge bosons to decay into massless fermions $f_1 \bar{f}_2$ is

$$\Gamma(W^+ \rightarrow e^+ \nu_e) = \frac{G_F M_W^3}{6\sqrt{2}\pi} \approx 226.5 \pm 0.3 \text{ MeV} , \quad (10.41a)$$

$$\Gamma(W^+ \rightarrow u_i \bar{d}_j) = \frac{CG_F M_W^3}{6\sqrt{2}\pi} |V_{ij}|^2 \approx (707 \pm 1) |V_{ij}|^2 \text{ MeV} , \quad (10.41b)$$

$$\Gamma(Z \rightarrow \psi_i \bar{\psi}_i) = \frac{CG_F M_Z^3}{6\sqrt{2}\pi} [g_V^{i2} + g_A^{i2}] \quad (10.41c)$$

$$\approx \begin{cases} 300.3 \pm 0.2 \text{ MeV } (u\bar{u}), & 167.24 \pm 0.08 \text{ MeV } (\nu\bar{\nu}), \\ 383.1 \pm 0.2 \text{ MeV } (d\bar{d}), & 84.01 \pm 0.05 \text{ MeV } (e^+e^-), \\ 375.9 \pm 0.1 \text{ MeV } (b\bar{b}). \end{cases}$$

From Particle
Data Book

Our modified theory predicts a different rate for $Z \rightarrow \mu\mu$

- This gives us a way to prove or disprove it!