



The CERN Computer Centre

and the
LHC Computing Grid

Running jobs: 129263.0
Transfer rate: 1.62 GiB/sec



Liceo Scientifico Galileo Galilei – Trieste
Saturday 9th April 2011

Massimo Lamanna
IT Data Storage Services

© 2011 Europa Technologies
US Dept of State Geographer
© 2011 MapLink/Tele Atlas
© 2011 Google

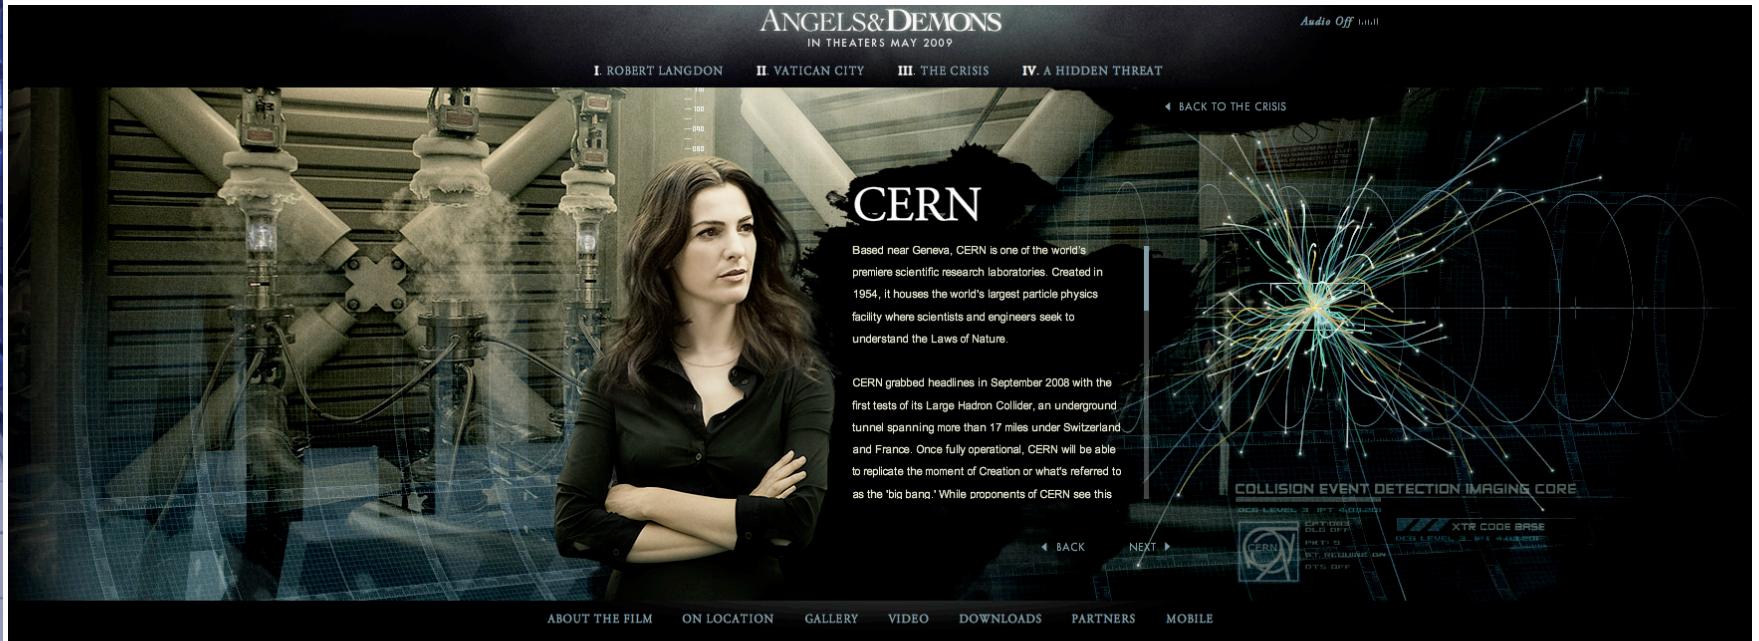
The LHC Computing Grid - April 2011

©2010 Google

Eye alt 14400.82 km

Benventi al CERN!

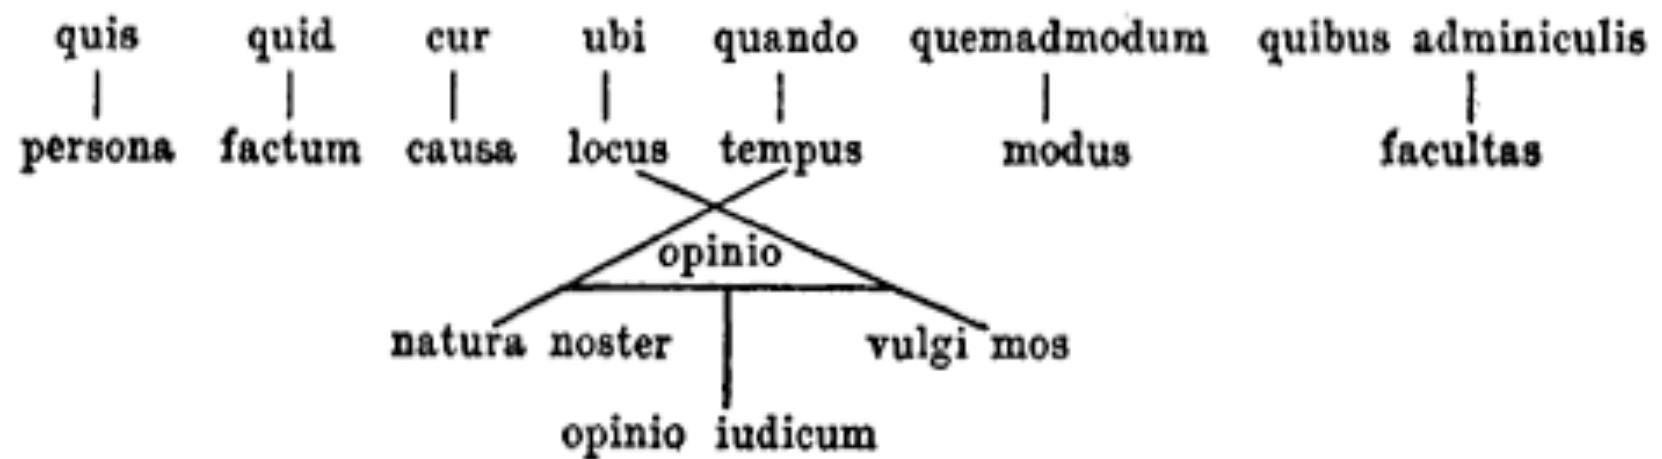
- Liceo Galilei di Trieste



Latino ☺ ... il seminario sarà in italiano con
trasparenze in Inglese ☺

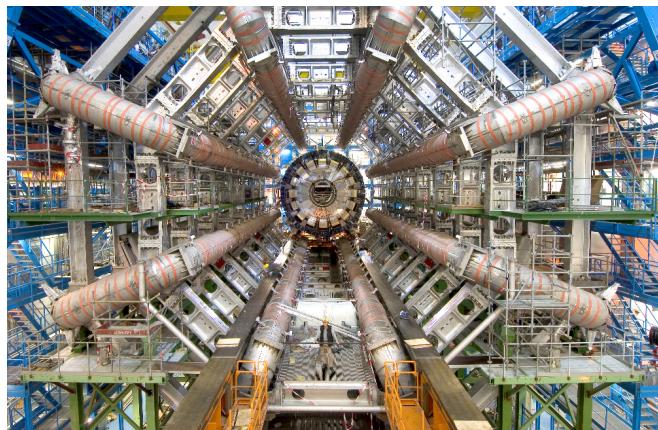


High-Energy Physics Computing



Cicerone et al.

LHC experiments (ATLAS)



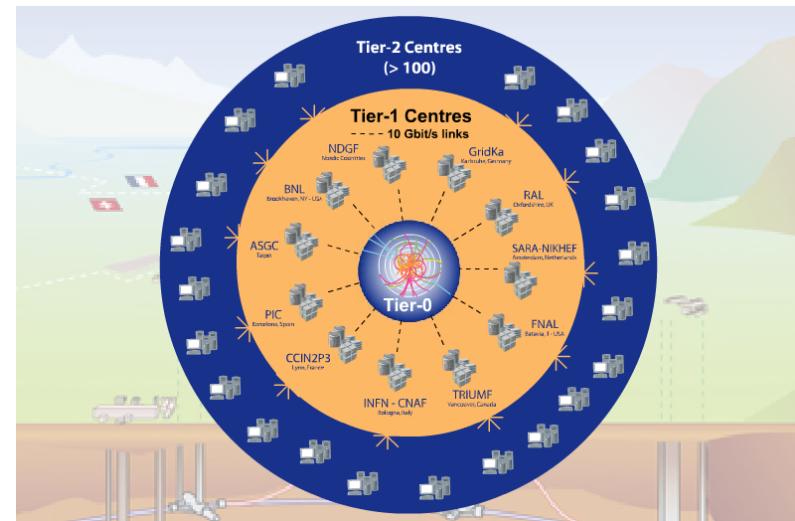
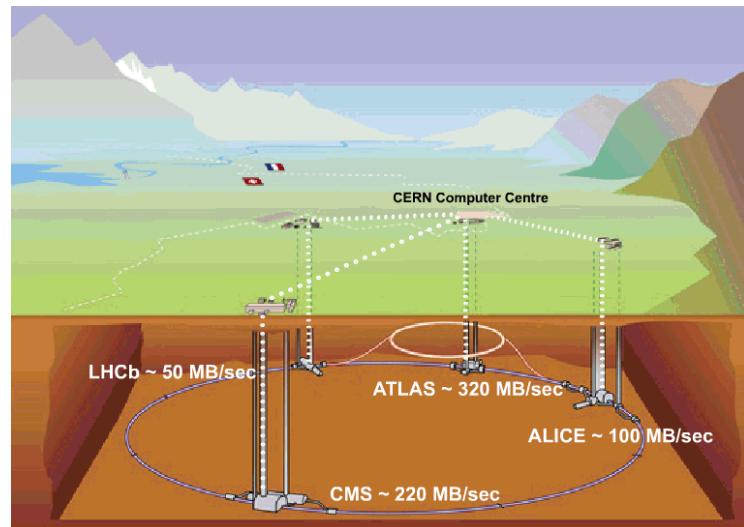
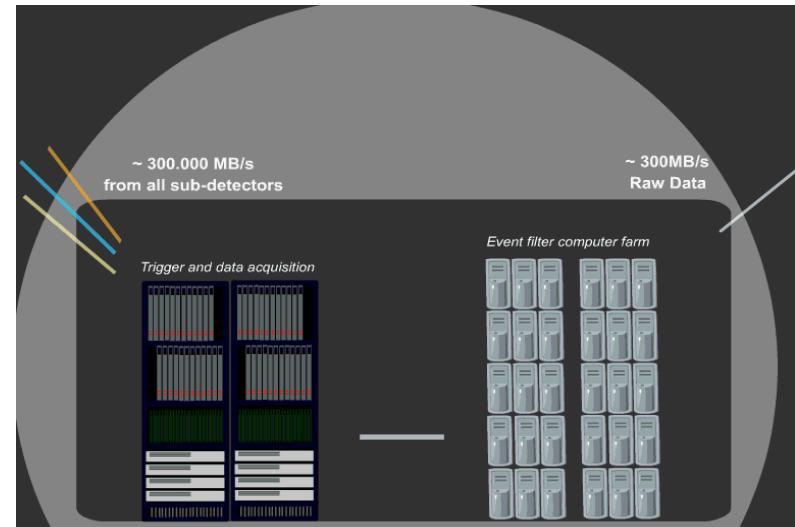
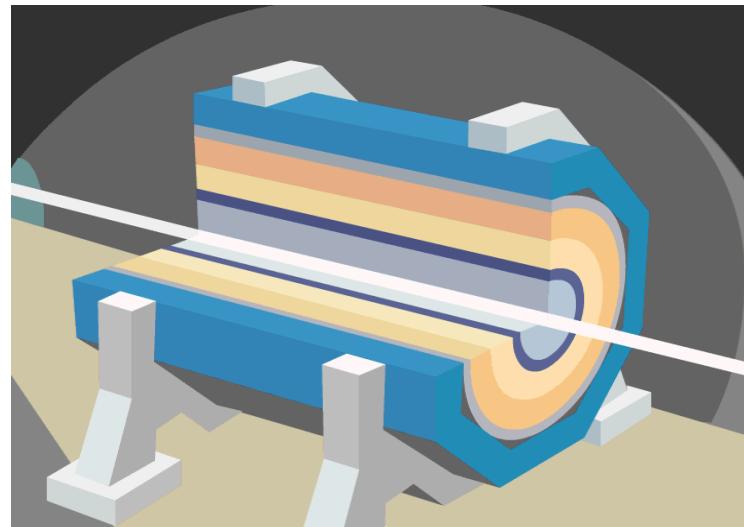
ATLAS: 7,000 tons, 150 million sensors generating data 40 millions times per second i.e. a petabyte/s (1 million GB/s)



ATLAS is around than 3,000 collaborators
From 169 universities from 37 countries
~1000 students!!!



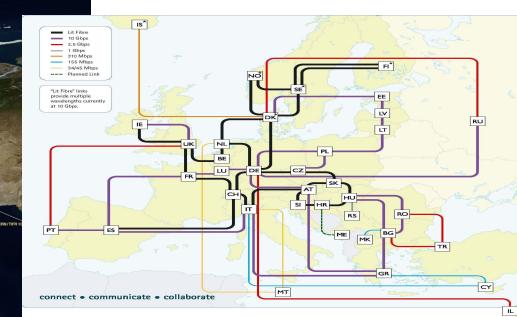
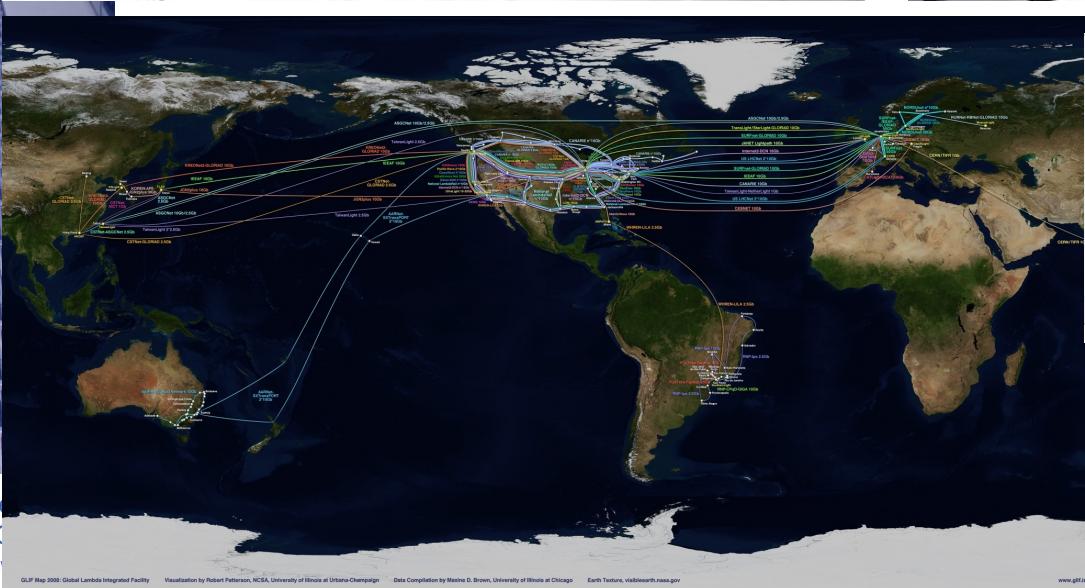
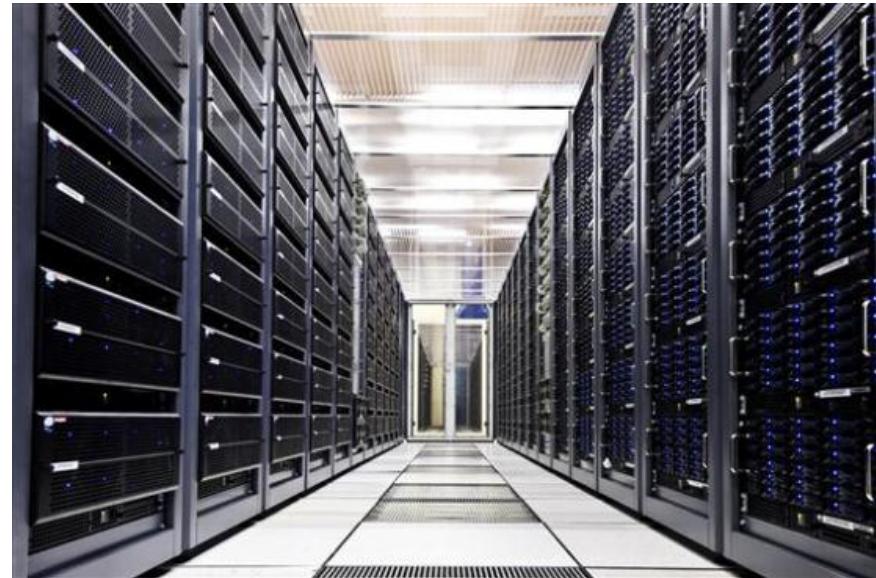
The role of the CERN Computer Centre





CERN Computer Centre: Storage, Distribution and Processing (Reconstruction and Analysis)

CERN IT Department



Géant: the pan-european Research and Education Network

CERN IT Department
CH-1211 Geneva 15
Switzerland
www.cern.ch/it

GLIF Map 2008

Visualization by Robert Patterson, NCSA, University of Illinois at Urbana-Champaign

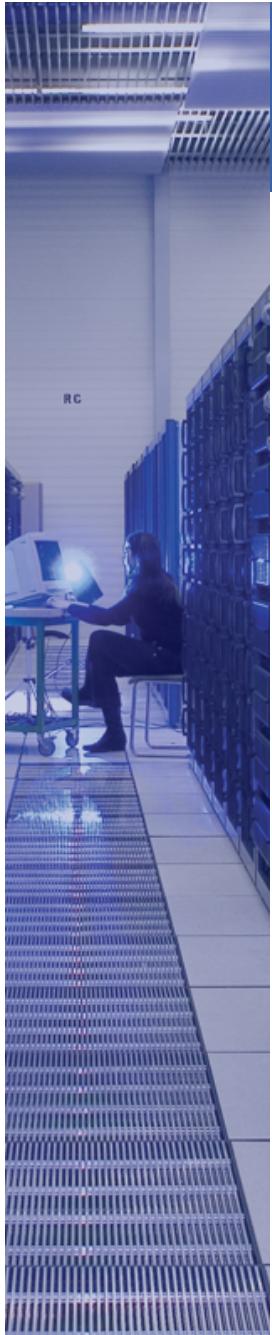
Data Compilation by Maxine D. Brown, University of Illinois at Chicago

Earth Texture, visibleearth.nasa.gov

www.géant.eu

LHCOPN: dedicated links with major computer centres worldwide

The LHC Computing Grid - April 2011



The LHC Data Challenge

- The accelerator will run for 20 years
- Experiments *are* producing about **15 Million Gigabytes** of data each year (about 20 million CDs!)
- LHC data analysis requires a computing power equivalent to **~100,000 of today's fastest PC processors**
- Requires many cooperating computer centres, as CERN can *only* provide **~20% of the capacity**



A challenge for physics...

... and a challenge for technology research and industry as well

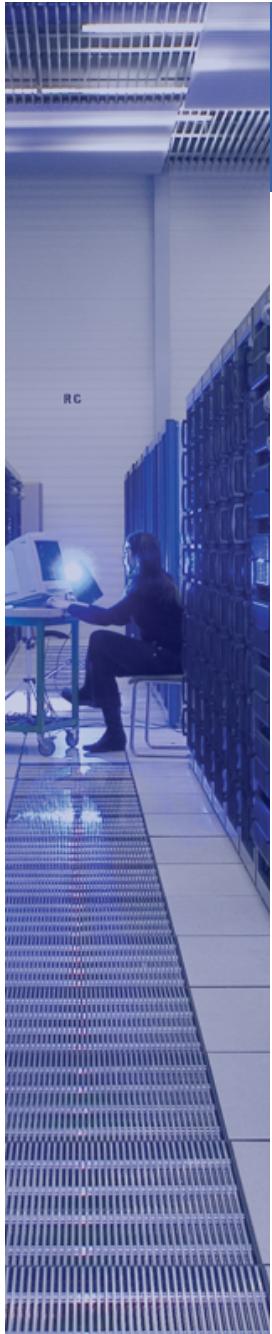


“Historical” example...

1990s:

The web was invented at CERN!
The machine used by Tim Berners-Lee in 1990 to develop and run the first WWW server, multi-media browser and web editor.



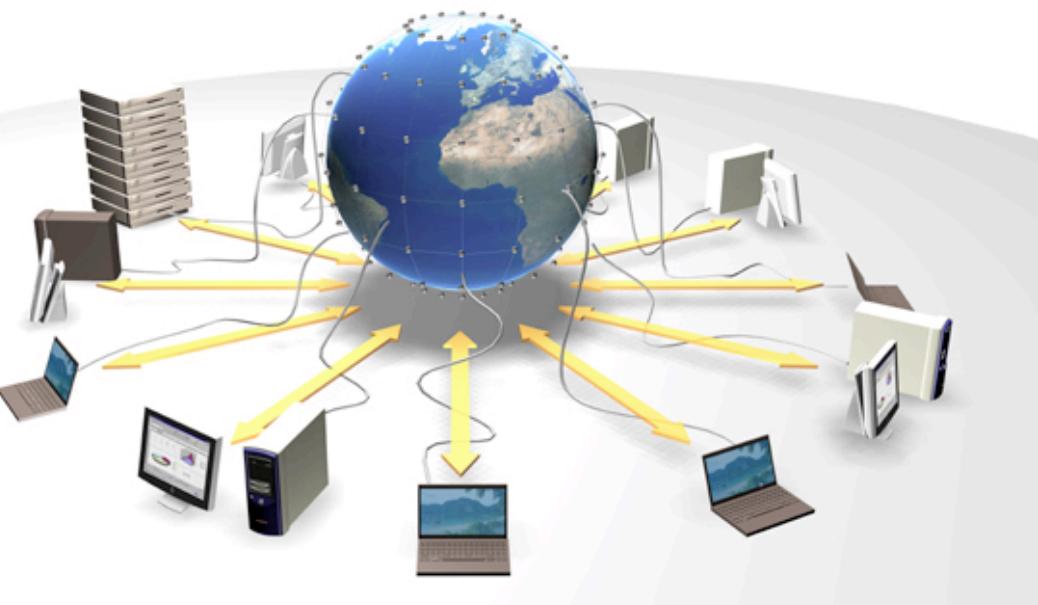


A new solution: the Grid

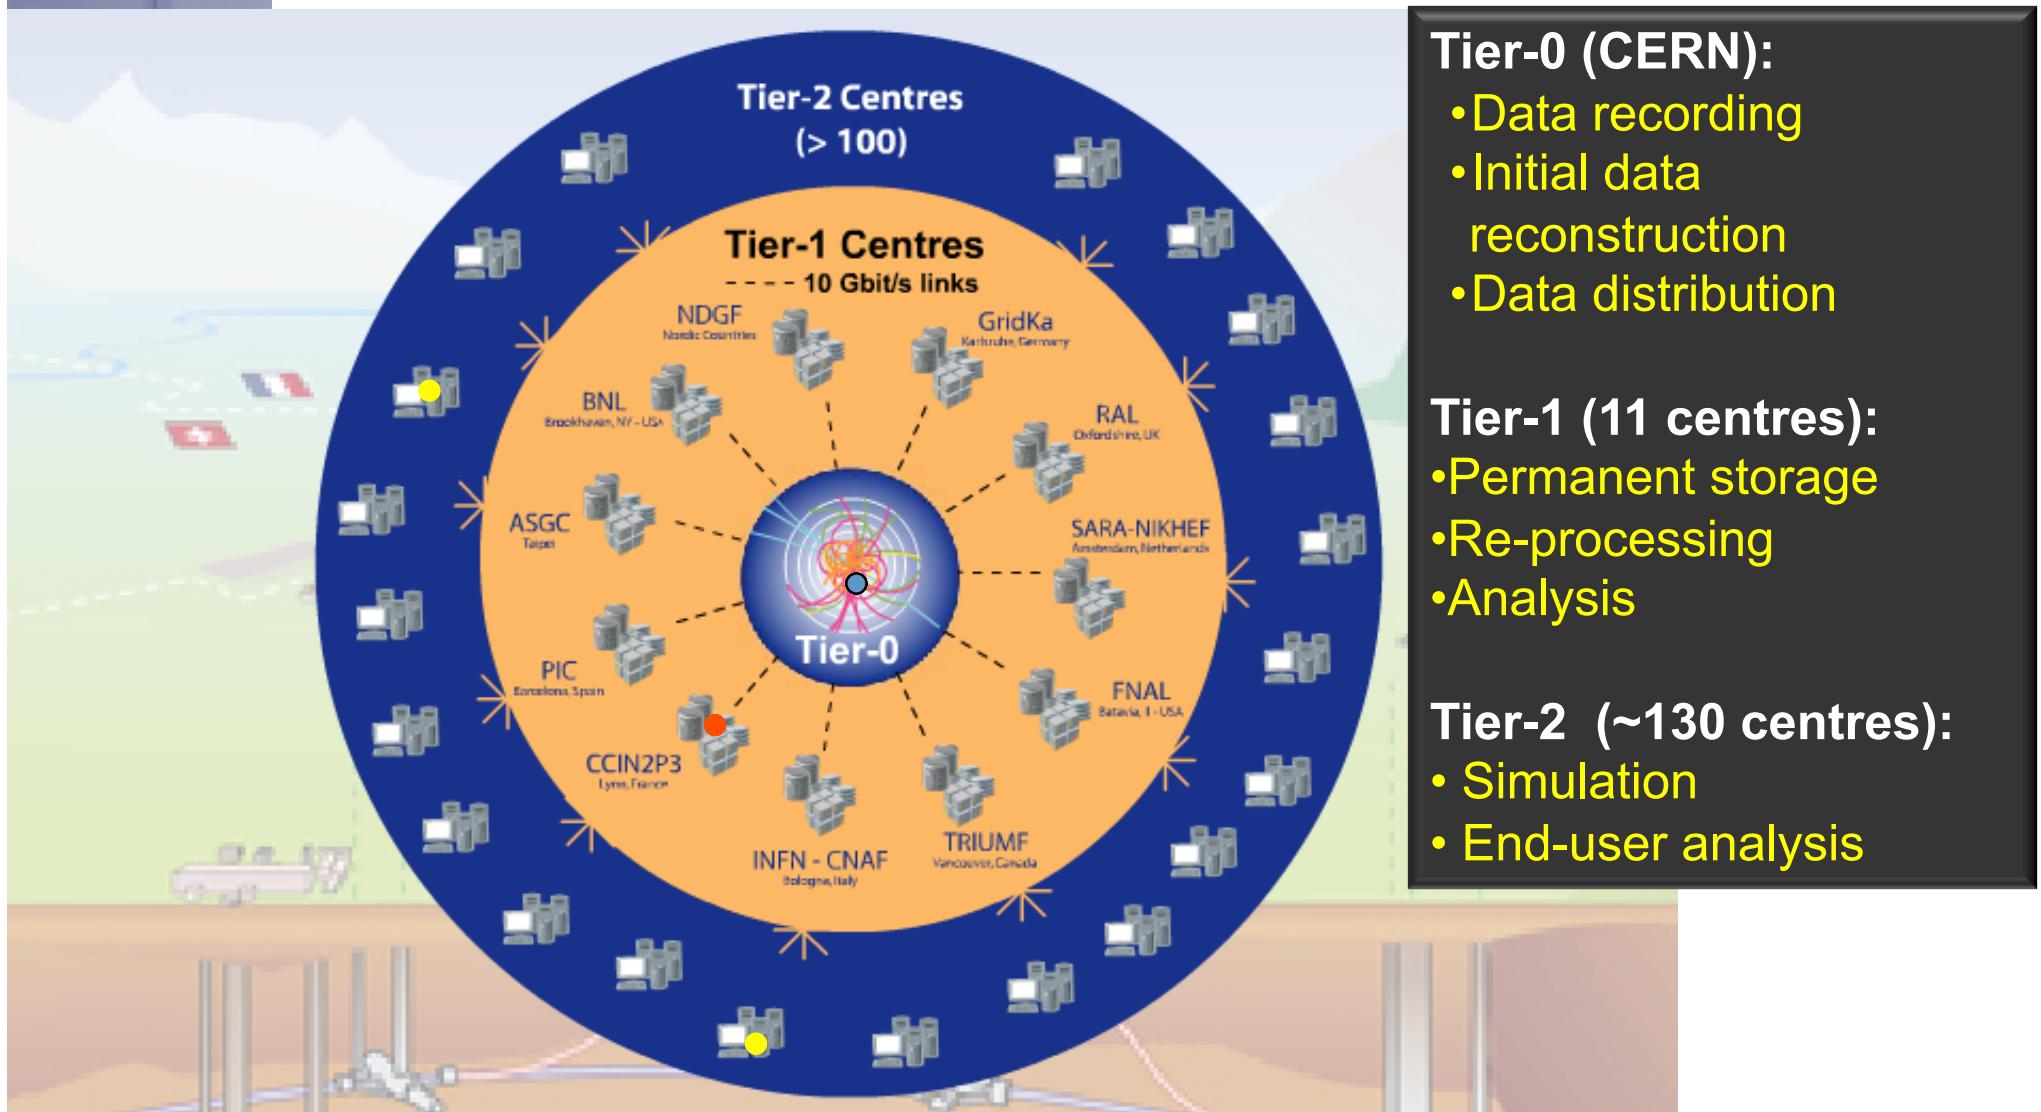
- Use the Grid to unite computing resources of particle physics institutes around the world

The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations

The **Grid** is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe



WLCG Tiers Organization



Tier-0 (CERN):

- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (11 centres):

- Permanent storage
- Re-processing
- Analysis

Tier-2 (~130 centres):

- Simulation
- End-user analysis

Computer Centre By Numbers

23 Mar 2011 Wed 21:30:26

Service information

full name: Computer Centre By Numbers**short name:** CCBYNUM**group:** IT-CF-FPP**site:** CERN**email:** imre.szebenyi@cern.ch**manager:** Imre Szebenyi 

Service availability ([more](#))

availability: 

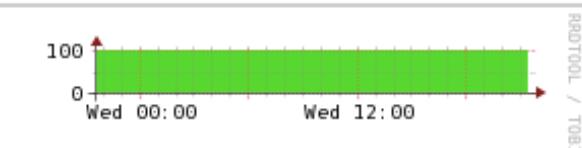
percentage: 100%

status: **available**last update: 21:20:34, 23 Mar 2011
(10 minutes ago)

expires after: 1440 minutes

 [rss feed with status changes](#)

availability in the last 24 hours ([more](#)):



Additional service information ([more](#))

Number of processors: 15,678

Number of cores: 64,198

Memory capacity (TiB): 158

Memory modules: 55,990

Raw HDD capacity (TiB): 63,254

Number of HDD's: 64,074

Number of systems: 11,730

Number of RAID controllers: 3,742

Number of enclosures: 1,417

SPEC CPU2006: 482,431

Number of racks: 827

Number of virtual machines: 1,624

Number of Fibre channel ports: 828

Number of 1G ports: 16,936

Number of 10G ports: 558

Current power consumption (kW): 2,506

Current power consumption (kVA): 2,640

- 24x7 operator and system admin support

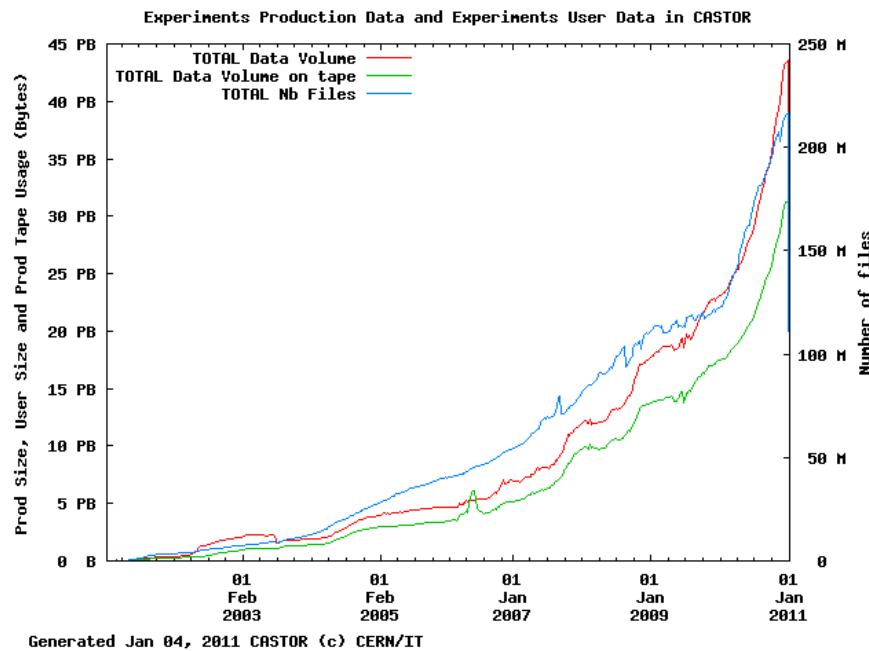
- Management and Automation framework for large scale Linux clusters

- Hardware installation & retirement

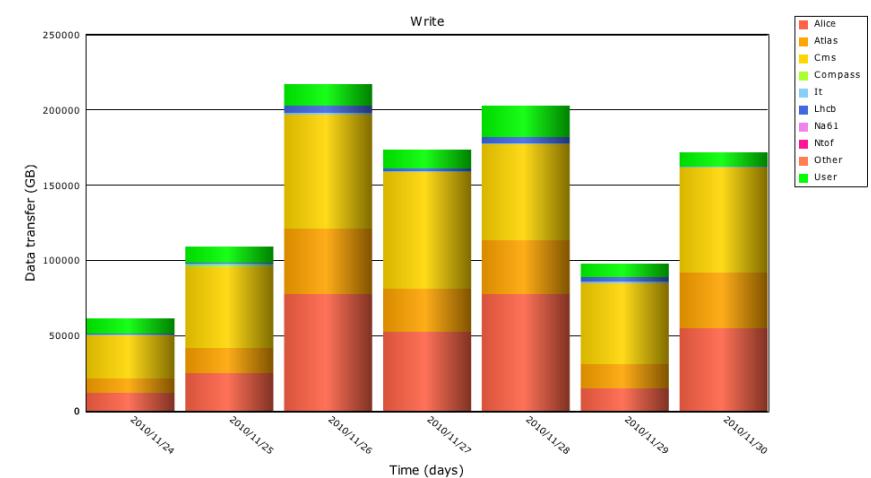
- ~7,000 hardware movements/year; ~1000 disk failures/year



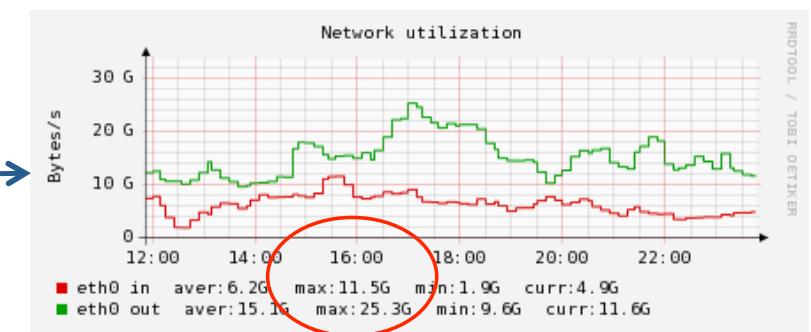
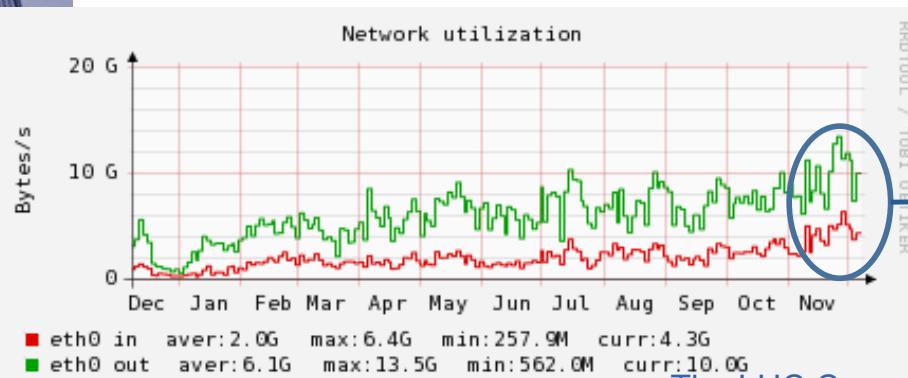
2010 data taking as seen in the CERN CC



Stored ~ 15 PB in 2010 with peaks at 220 TB/day during PbPb

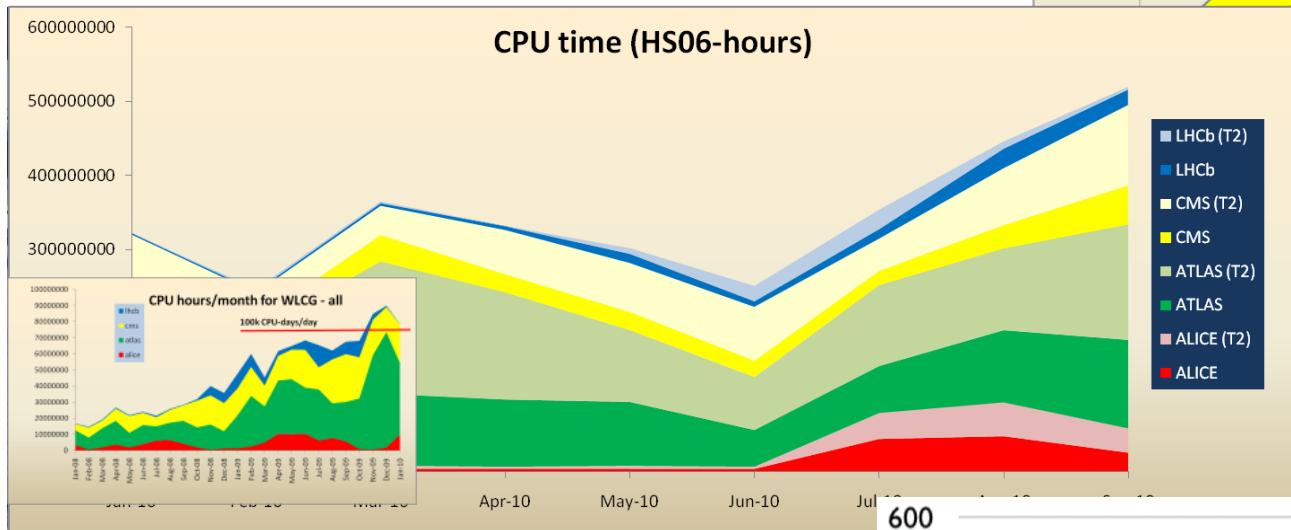


Tier-0 Bandwidth
Average in: 2 GB/s with peaks at 11.5 GB/s
Average out: 6 GB/s with peaks at 25 GB/s

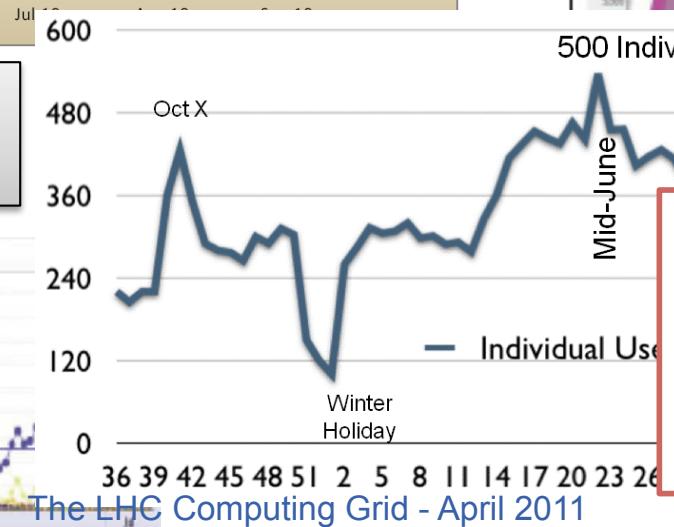
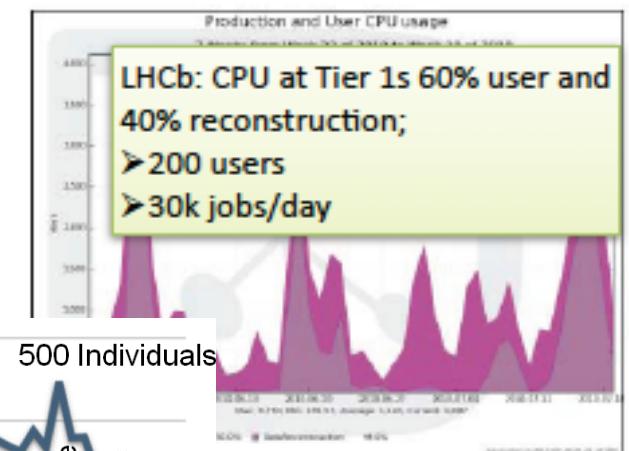
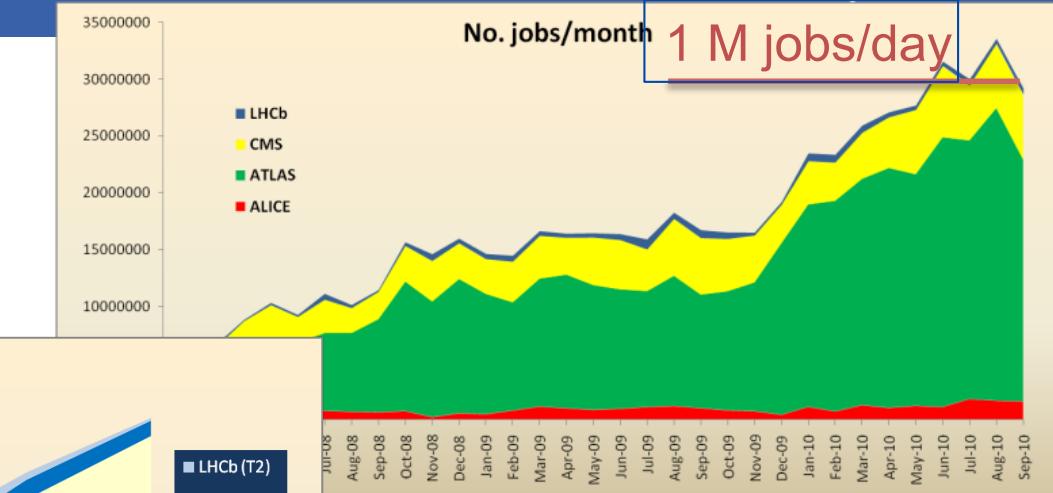


WLCG Usage

- Use remains consistently high
 - 1 M jobs/day; >>100k CPU-days/day
 - Actually much more inside pilot jobs

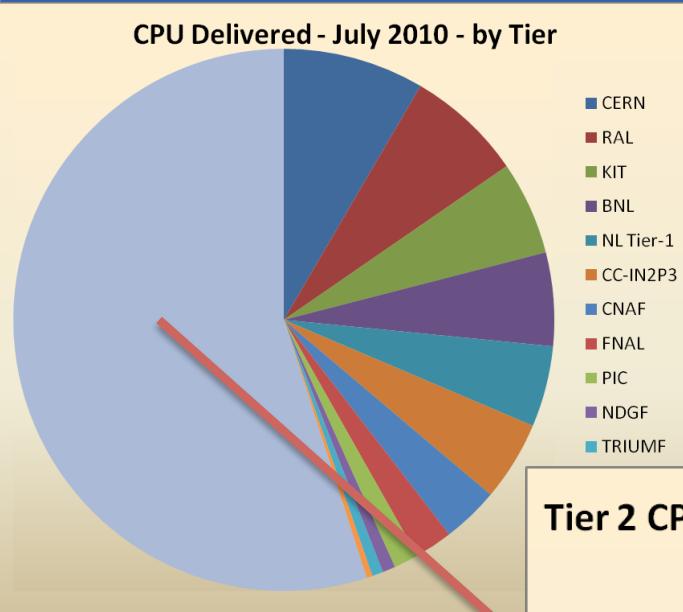
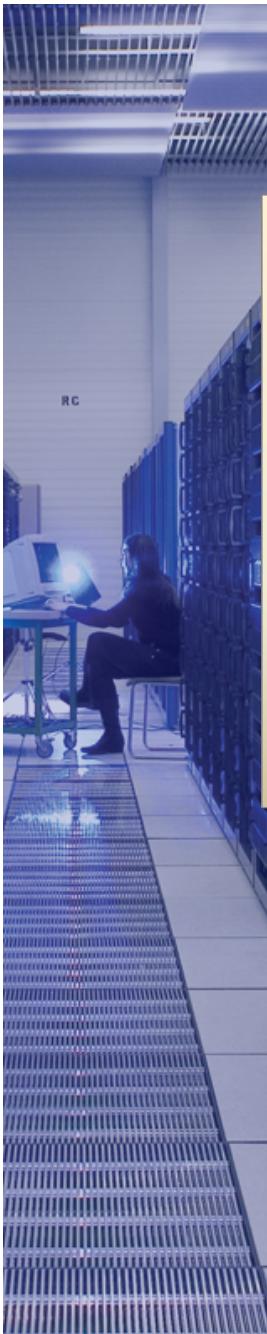


As well as LHC data, large simulation productions ongoing



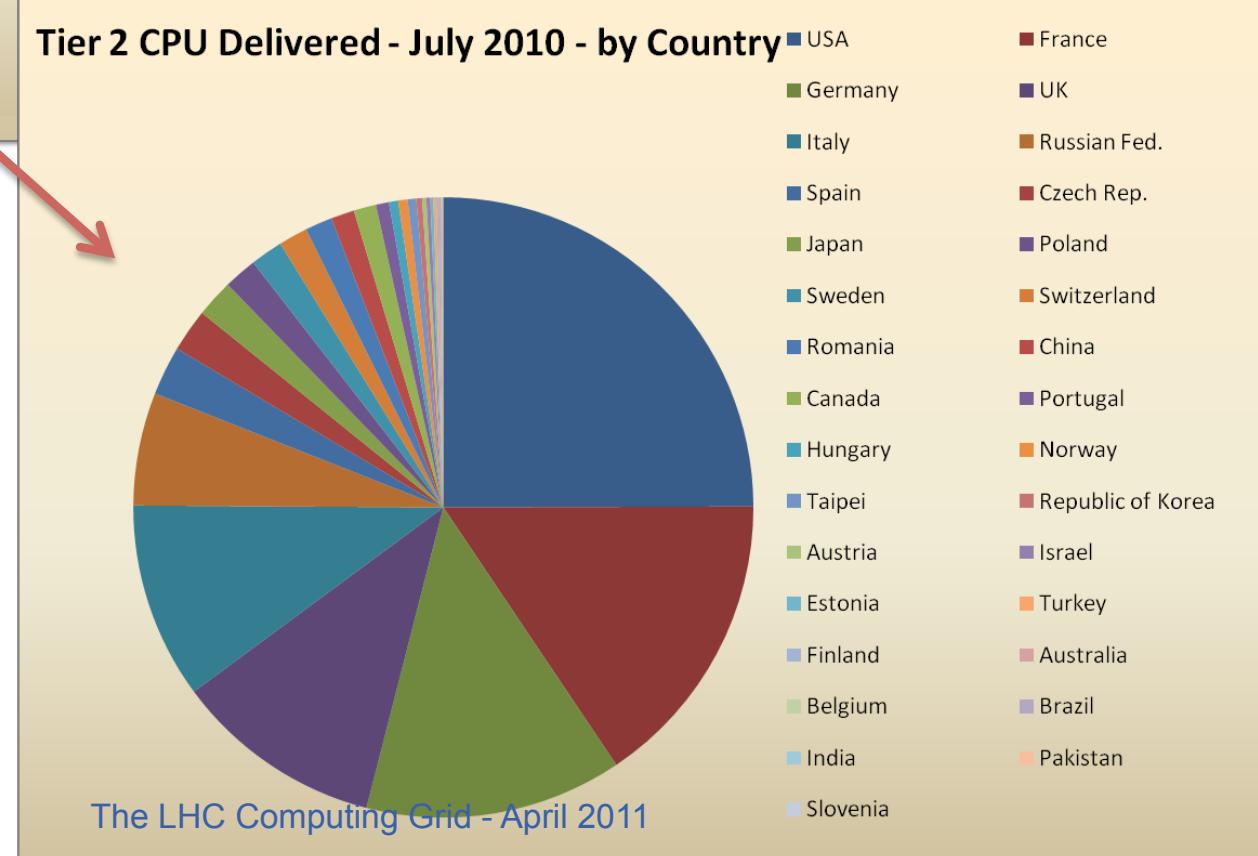
- Large numbers of analysis users
 - CMS ~800,
 - ATLAS ~1000,
 - LHCb/ALICE ~250

Cooperating effort



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

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





Impact of the LHC Computing Grid in Europe

CERN
IT
Department

eGEE
Enabling Grids
for E-sciencE

Archeology
Astronomy
Astrophysics
Civil Protection
Comp. Chemistry
Earth Sciences
Finance
Fusion
Geophysics
High Energy Physics
Life Sciences
Multimedia
Material Sciences
...

Scheduled = 21539
Running = 25374

- LCG has been the driving force for the European multi-science Grid EGEE (Enabling Grids for E-sciencE)
- EGEE was a global effort, and the largest Grid infrastructure worldwide
- Co-funded by the European Commission (Cost: ~170 M€ over 6 years, funded by EU ~100M€)
- Now moving to a permanent European Infrastructure: EGI

>300 sites
48 countries
>200,000 cores
>20 PetaBytes
>10,000 users
>150 VOs
>150,000 jobs/day

The LHC Computing Grid - April 2011



GridPP
UK Computing for Particle Physics

21:13:50 UTC

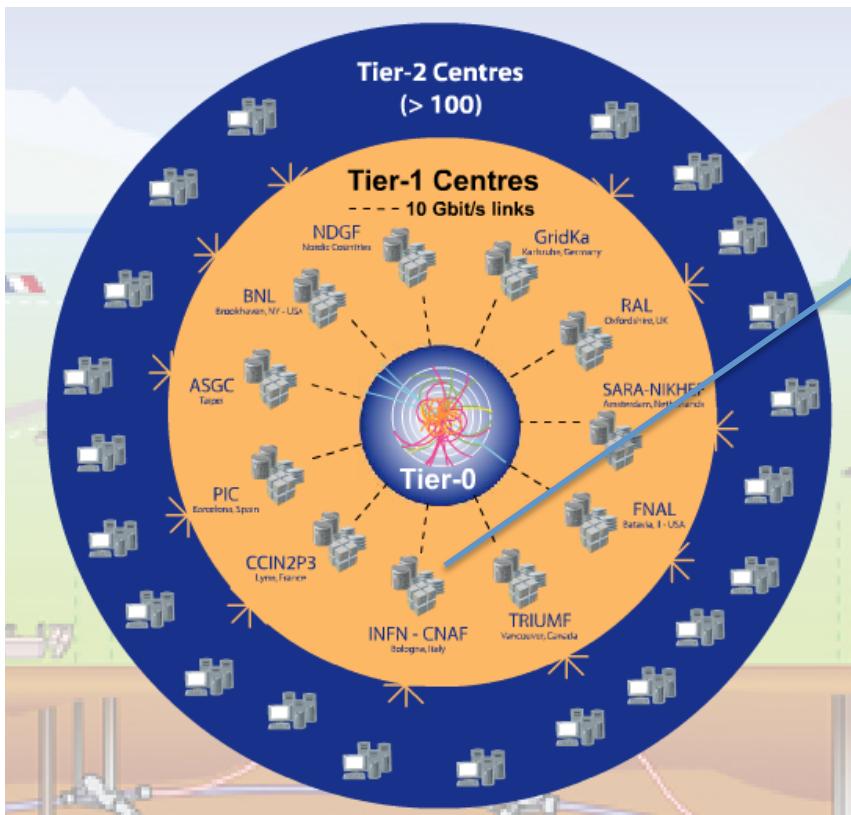


More info:

INFN (Istituto Nazionale Fisica Nucleare):

<http://www.infn.it>

IGI (Italian Grid Initiative): <http://www.italiangrid.org/>



The LHC Computing Grid - April 2011





Health-e-Child

The central diagram illustrates the **GRID** (Global Research and Information Data) architecture, which integrates various medical informatics components:

- Similarity Search**: A network graph showing nodes connected by lines, with a callout explaining "Intra slice topology".
- Temporal Modeling**: A network graph showing nodes connected by lines, with a callout explaining "Each visit's internal dependencies are represented as edges between node of the same layer (intra slice topology)" and "Temporal dependencies between visits are represented as edges between nodes belonging to different layers (inter slice topology)".
- Visual Data Mining**: A heatmap visualization with a legend for "Visual aids: • Coloring • Size • Labels • Hierarchy" and a callout stating "Hierarchical hierarchy is reflected automatically".
- PA - Cerebrum vs. Cerebellum**: A brain scan comparison between cerebrum (blue) and cerebellum (purple).
- Genetics Profiling**: A brain scan with a color scale from red to green.
- Treatment Response**: Two X-ray images of a knee joint labeled "Example: treatment response" and "Follow-up [1 year]".
- Inferring Outcome**: A flowchart titled "Temporal reasoning: Inferring outcome" showing data flow between various nodes.
- Biomechanical Models Tumor Growth Modelling**: A brain scan with a highlighted orange area and a timeline showing "March", "September", and "September Simulation". Below it, a caption reads: "O Clatz, PV Bonelius, H Delingette, M Sermesant, SK Warfield, G Melamud, N Ayache. Brain Tumor Growth Simulation. IEEE TMI 2005".

Personalised Simulation: A screenshot of a software interface titled "Personalised Simulation: Results" showing a 3D model of a heart with a green shaded region labeled "Dyskinetic area".

Semantic Browsing: A 3D grid-based visualization showing nodes connected by lines, with a callout stating "Hierarchical hierarchy is reflected automatically".

Biomechanical Model: A brain scan with a highlighted orange area and a timeline showing "March", "September", and "September Simulation". Below it, a caption reads: "O Clatz, PV Bonelius, H Delingette, M Sermesant, SK Warfield, G Melamud, N Ayache. Brain Tumor Growth Simulation. IEEE TMI 2005".

Similarity Search: A screenshot of a software interface titled "Similarity Search" showing a network graph of nodes.

Temporal Modelling: A screenshot of a software interface titled "...& Temporal Modeling" showing a network graph of nodes.

Visual Data Mining: A screenshot of a software interface titled "Visual Data Mining" showing a heatmap visualization.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it



For more information about the Grid:



www.cern.ch/lcg



www.eu-egi.org/



www.eu-egee.org



Grid Café
www.gridcafe.org



www.gridcafe.org

Thank you for your kind attention!