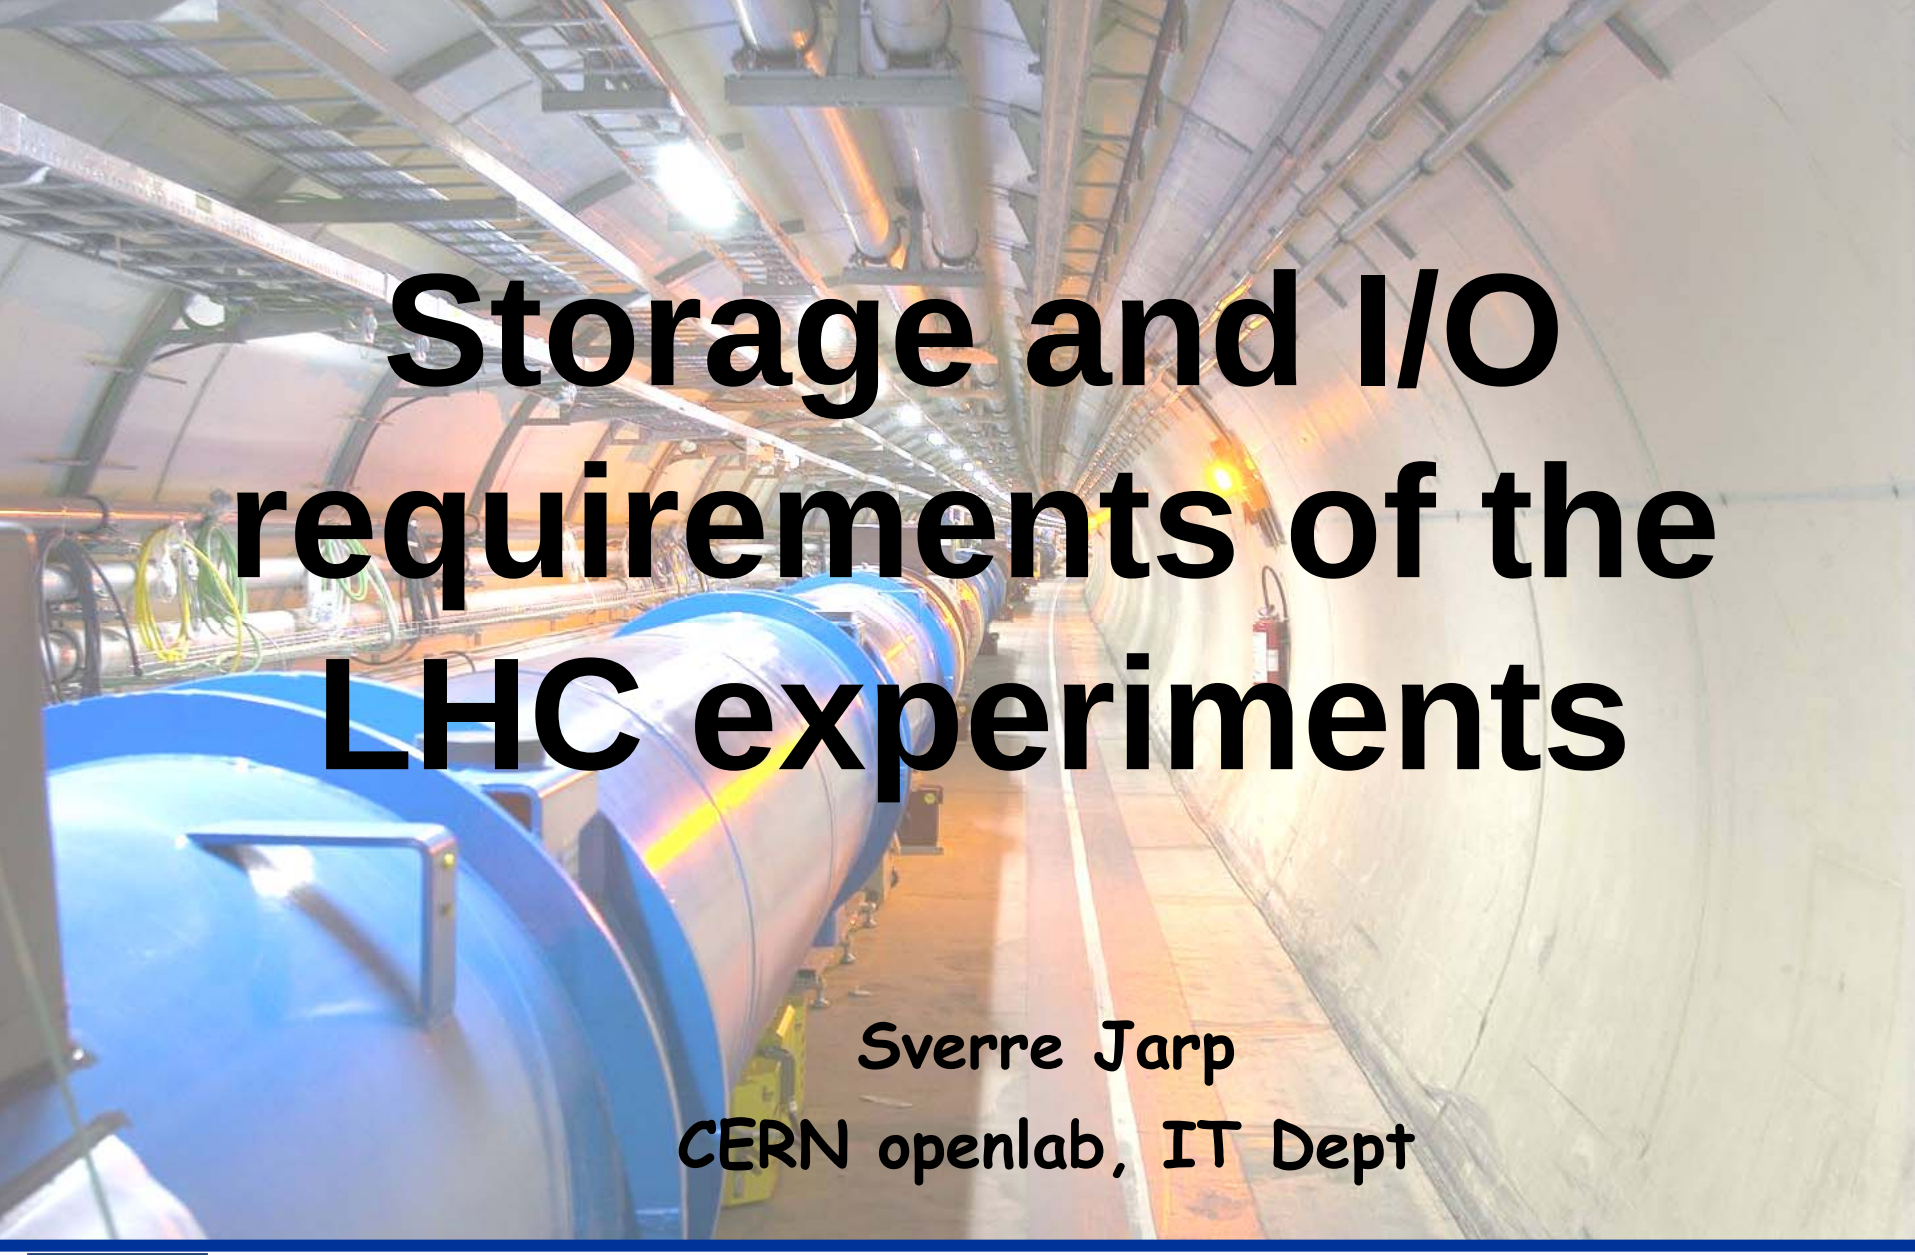

Storage and I/O requirements of the LHC experiments

Sverre Jarp
CERN openlab, IT Dept

"where the Web was born"

Briefly about CERN

What is CERN?

- CERN is the world's largest particle physics centre
- Particle physics is about:
  - elementary particles, the constituents from which all matter in the Universe is made
  - fundamental forces which hold matter together
- Particles physics requires:
  - special tools to create and study new particles

CERN's tools

The special tools for particle physics are:

- **ACCELERATORS**, huge machines (inside a complex underground structure) - able to accelerate particles to very high energies before colliding them into other particles
- **DETECTORS**, massive instruments which register the particles produced when the accelerated particles collide
- **COMPUTING**, to reconstruct the collisions, to extract the physics data and to perform the analysis

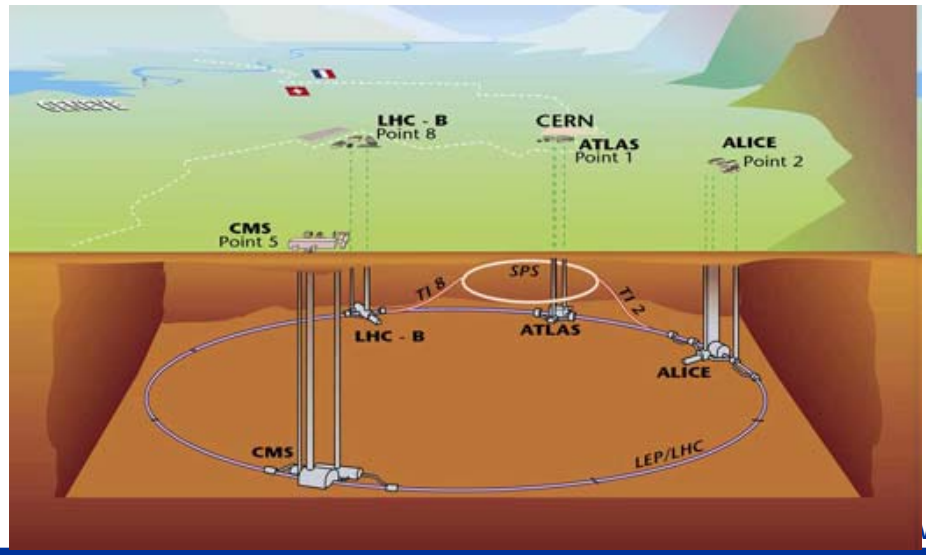

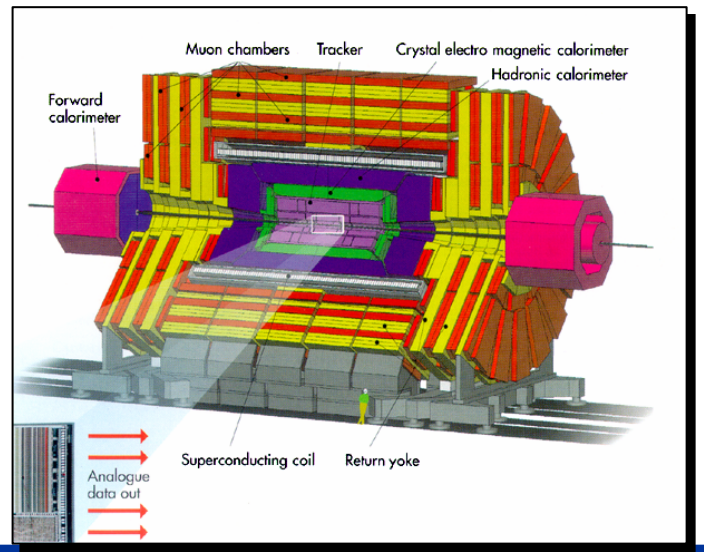

CERN in Numbers

- 2500 Staff
- 6500 Users
- 500 Fellows and Associates
- 80 Nationalities
- 500 Universities
- Budget ~1000 MCHF/year
  (~650 M€/year)

- **20 Member States:**
  Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom.
- **8 Observers:**
  India, Israel, Japan, the Russian Federation, USA, Turkey, the European Commission and UNESCO

What is LHC?

LHC will be switched on in 2007

Four experiments, with detectors as 'big as cathedrals':
ALICE
ATLAS
CMS
LHCb

- It is a particle accelerator that will collide beams of protons at an energy of 14 TeV
- Using the latest super-conducting technologies, it will operate at about – 270°C, just above the absolute zero of temperature
- With its 27 km circumference, the accelerator will be the largest superconducting installation in the world.
- Its two proton beams will interact 40 million times per second

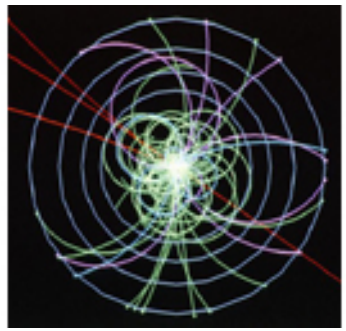

ATLAS construction

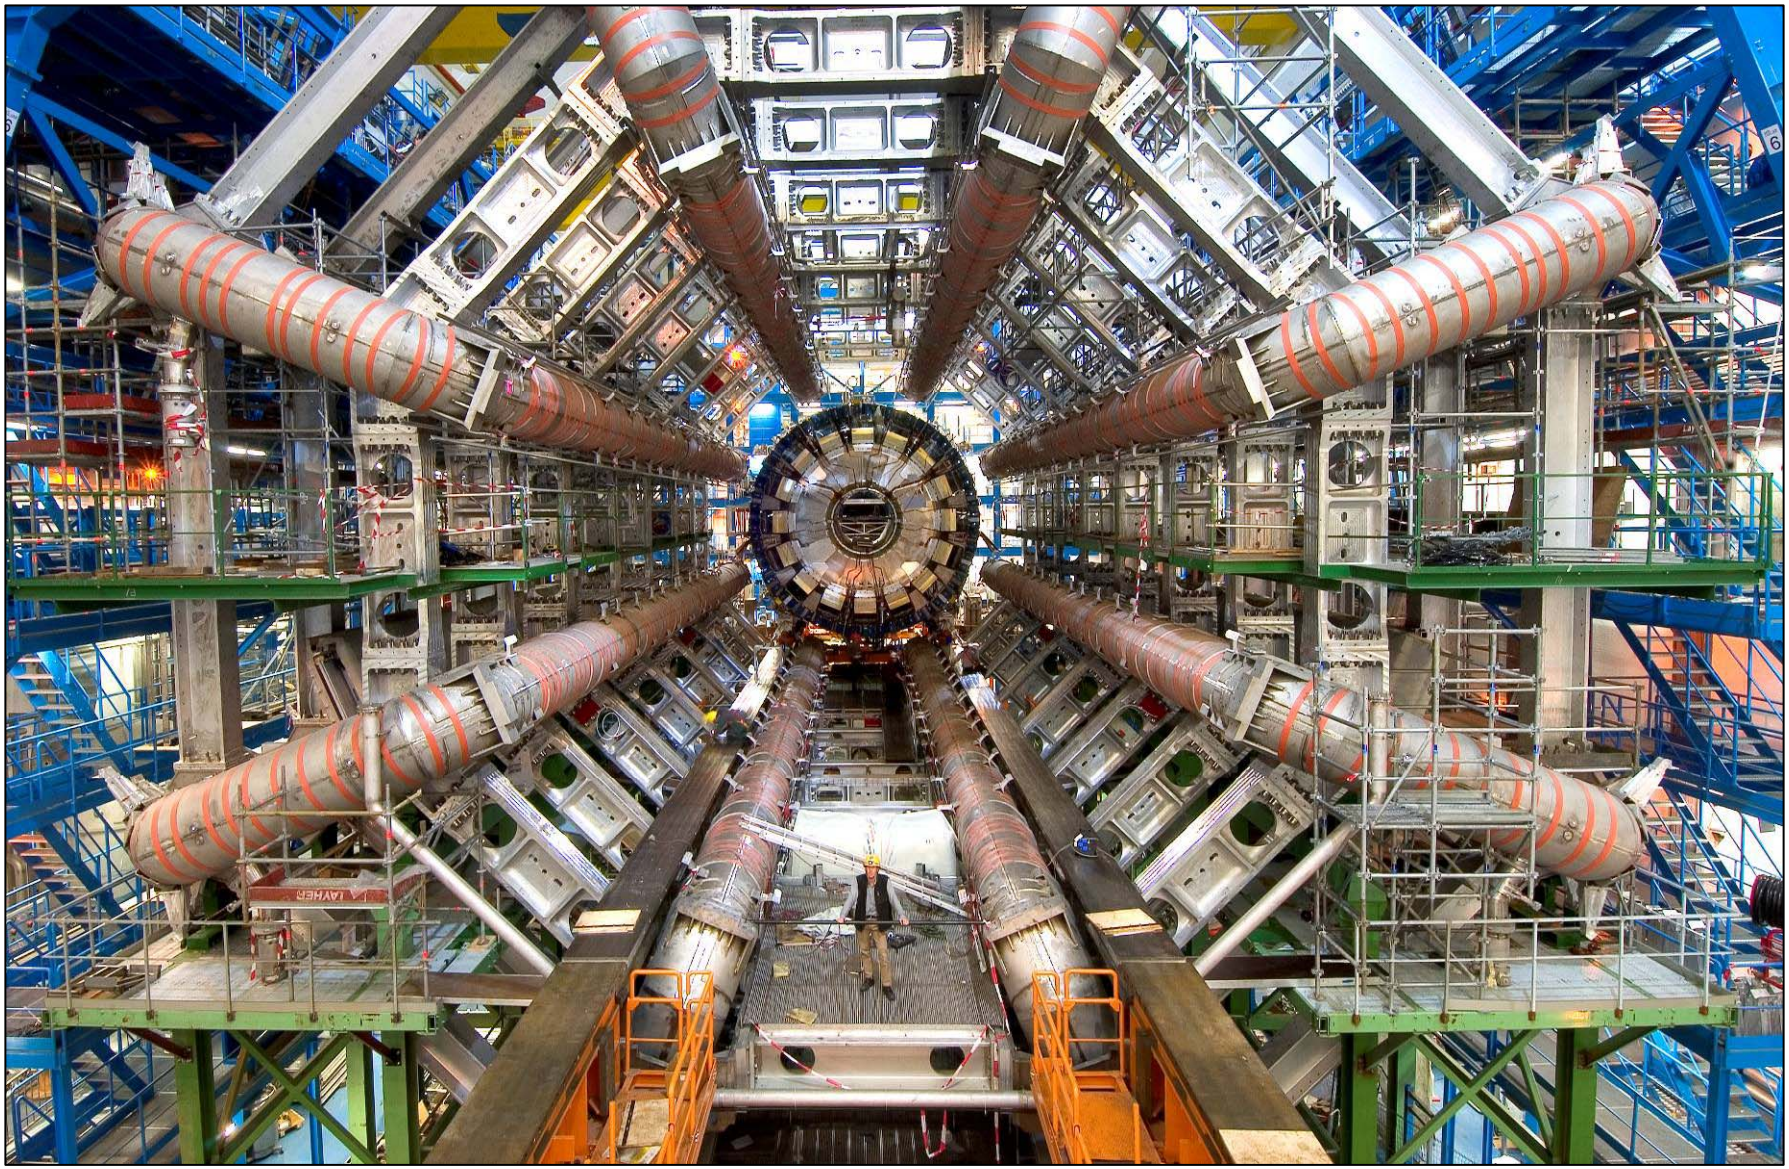

PHYSICS COMPUTING

High Energy Physics Computing Characteristics

- **Independent events (collisions of particles)**
  - trivial (read: pleasant) parallel processing
- **Bulk of the data is read-only**
  - versions rather than updates
- **Meta-data in databases linking to "flat" files**
- **Compute power measured in SPECint (rather than SPECfp)**
  - But good floating-point is important
- **Very large aggregate requirements:**
  - computation, data, input/output
- **Chaotic workload –**
  - research environment - physics extracted by iterative analysis, collaborating groups of physicists

  → Unpredictable → unlimited demand

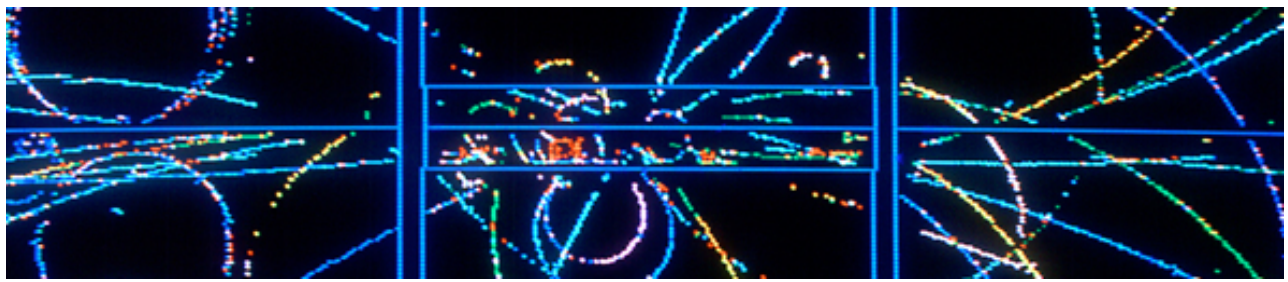

The Computing Environment

- High-throughput computing (based on reliable "commodity" technology)
  - Around 3000 (dual-processor Xeon) PCs with "Scientific Linux"

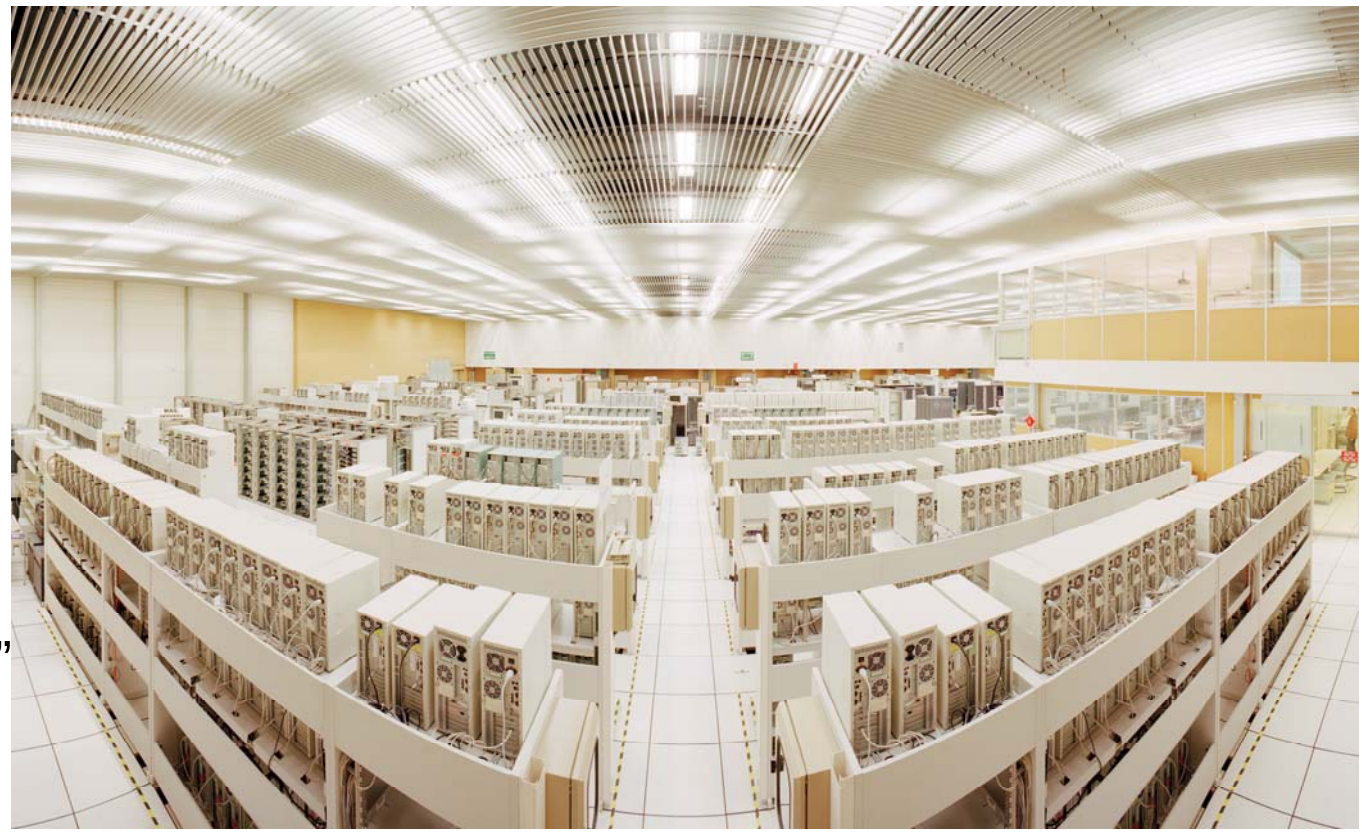

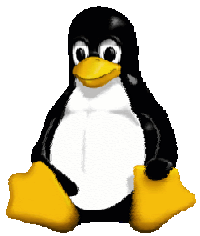

LHC DATA

Online computers filter out a few hundred "good" events per sec. Each event is ~1 MB.

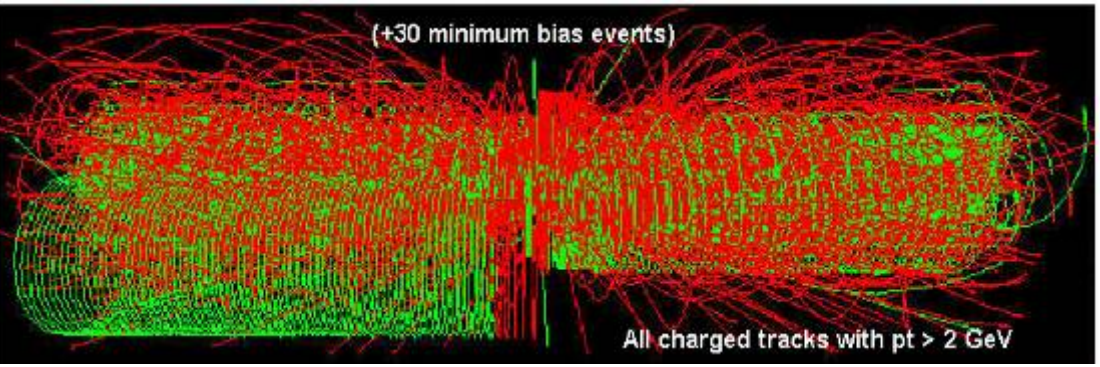

Which are recorded on disk and magnetic tape
at 100-1,000 Megabytes/sec → ~15 Petabytes per year
for all four experiments

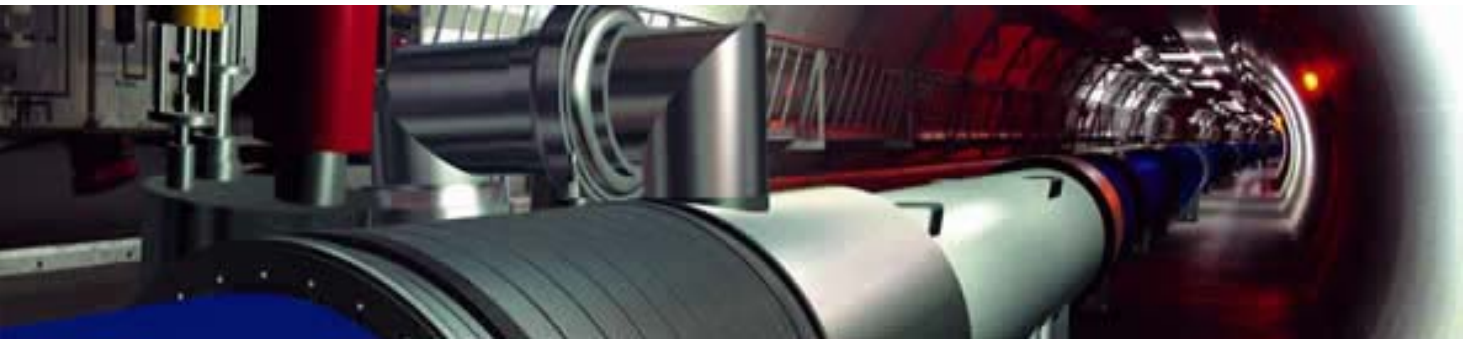

LHC data handling

- LHC data corresponds to about 20 million CDs each year
- Permanent storage → magnetic tape
- Transient storage → NAS servers
  - 20% of total

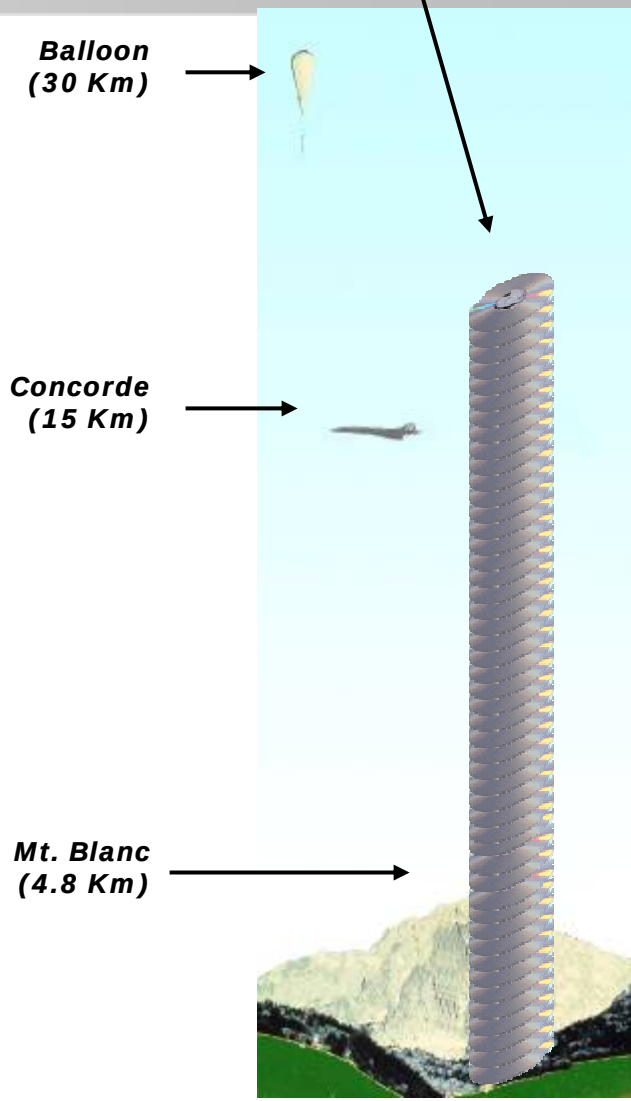

Data Flow for Physics Analysis

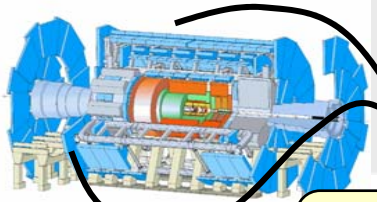

detector

event filter (selection & reconstruction)

reconstruction

event summary data

10 %

processed data

100 %

raw data

event reprocessing

batch physics analysis

analysis

1 %

analysis objects
(extracted by physics topic)

event simulation

simulation

interactive physics analysis

Data Handling in detail
(at Tier0 – the centre)

Multiple Data Activities

1. data transfer from the ‘DAQ’ buffer to the “Tier0” buffer
2. data transfers from the T0 buffer to the reconstruction farm and derived data back to the T0 buffer
3 data migration from the T0 buffer to the tape system
4 data export from the T0 buffer to the sister Labs

Each of the flows has its own characteristics and is different for each of the 4 experiments

multi-dimensional ‘impedance’ matching problem

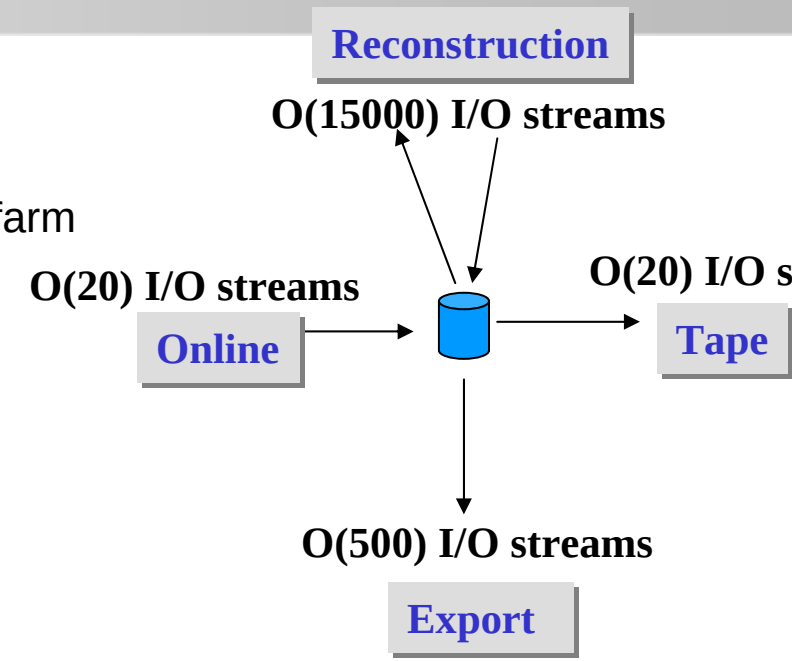

	Online [MB/s]	Tape [MB/s]	Reconstr. [MB/s]	Export [MB/s]	Total [MB/s]
ALICE HI	1250	1250	300	300	~ 3000 HI
ATLAS	320	440	540	780	~ 2100
CMS	225	270	270	315	~ 1100
LHCb	60	40	35	35	~170

General Data & Control Flow

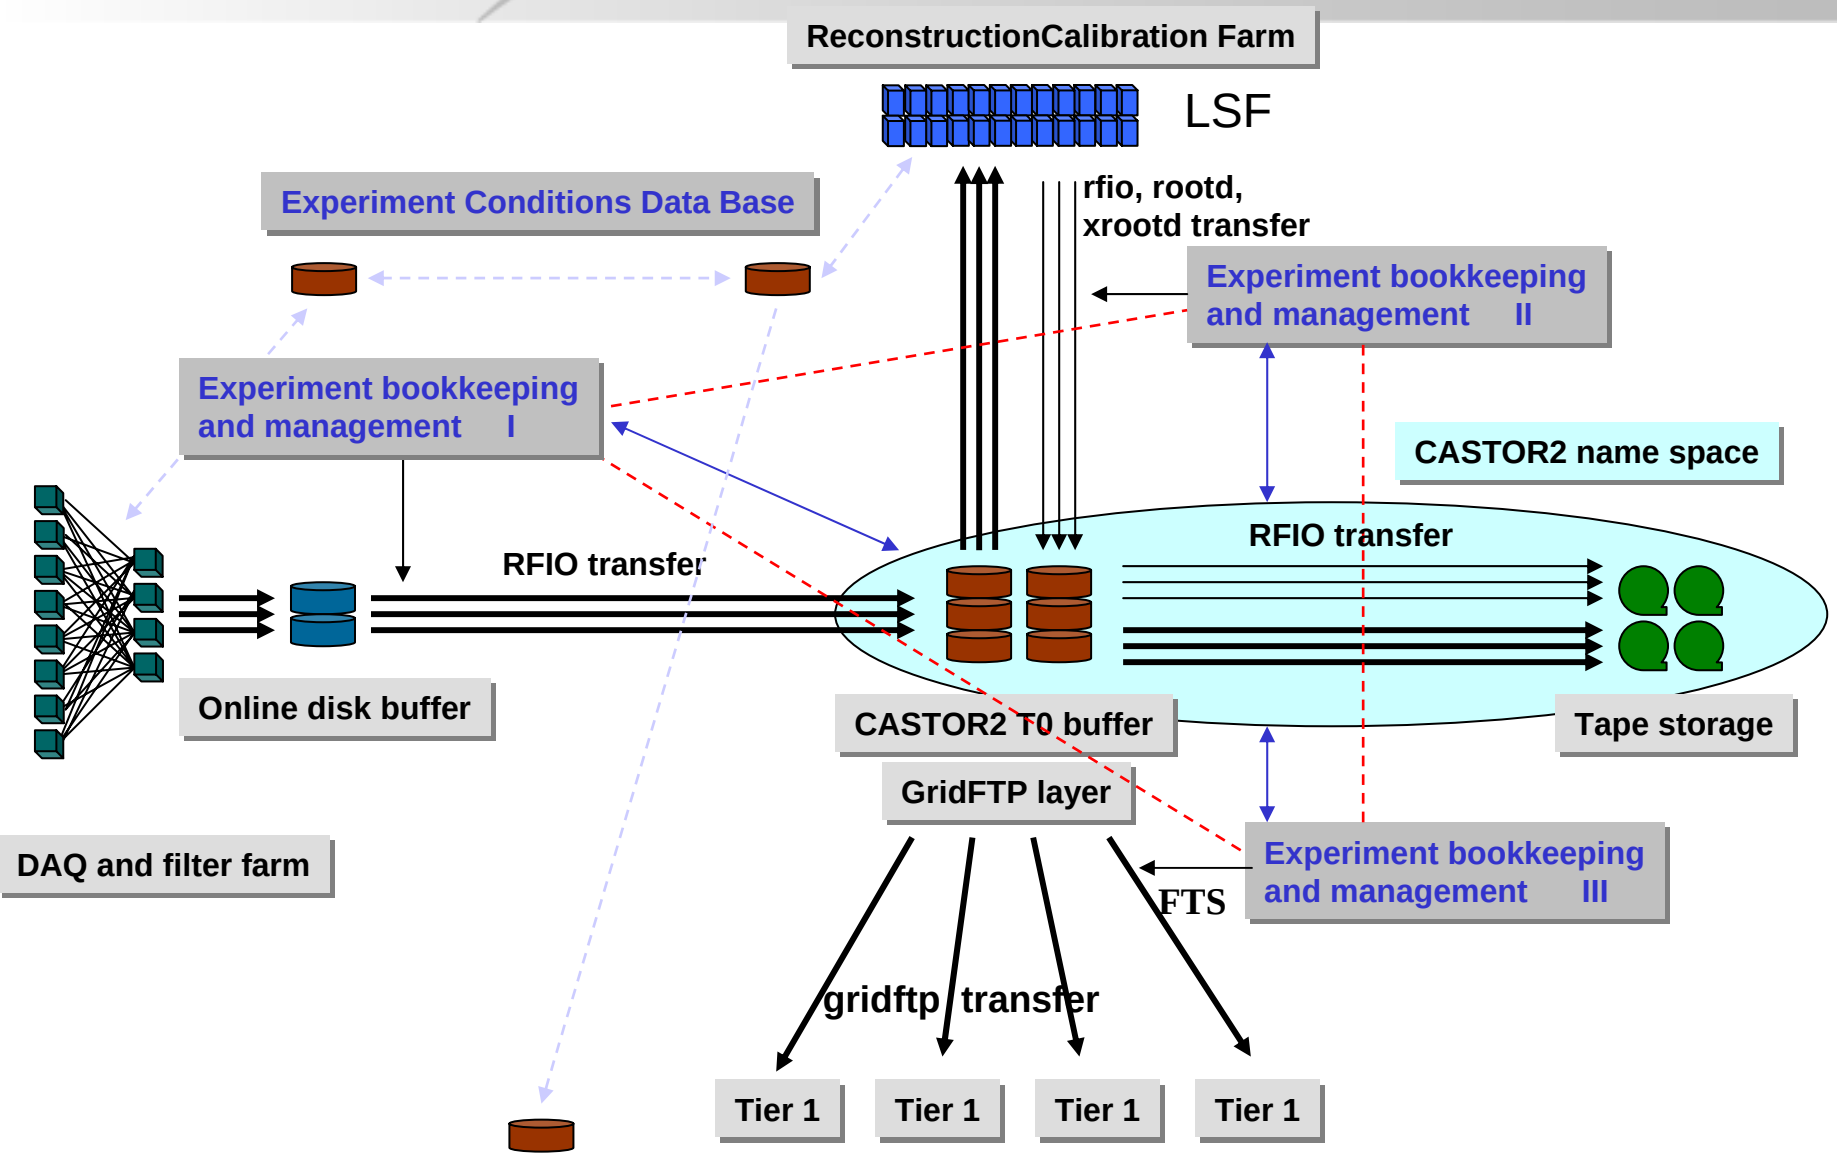

Network Topology (Q1-06)

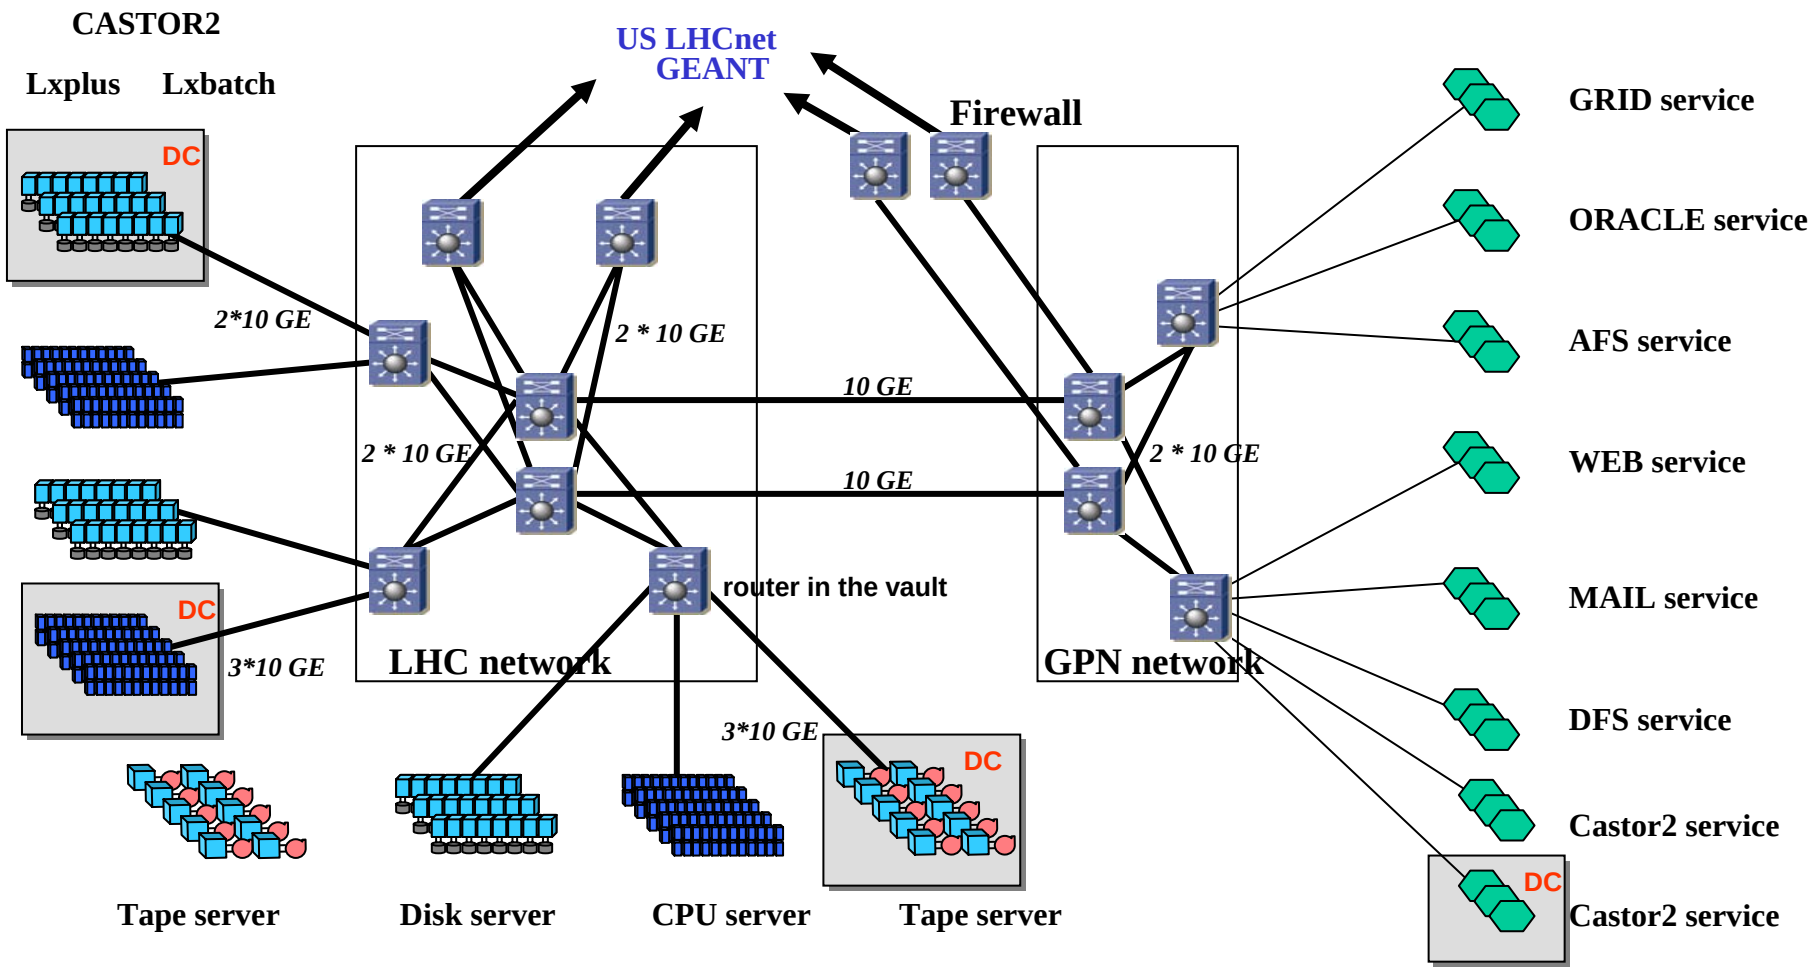

- Dedicated resources for the Data Recording Challenges (CPU, disk, tape, Castor2)
- 'logical' separation between DC setup and production systems

LHC Computing Grid

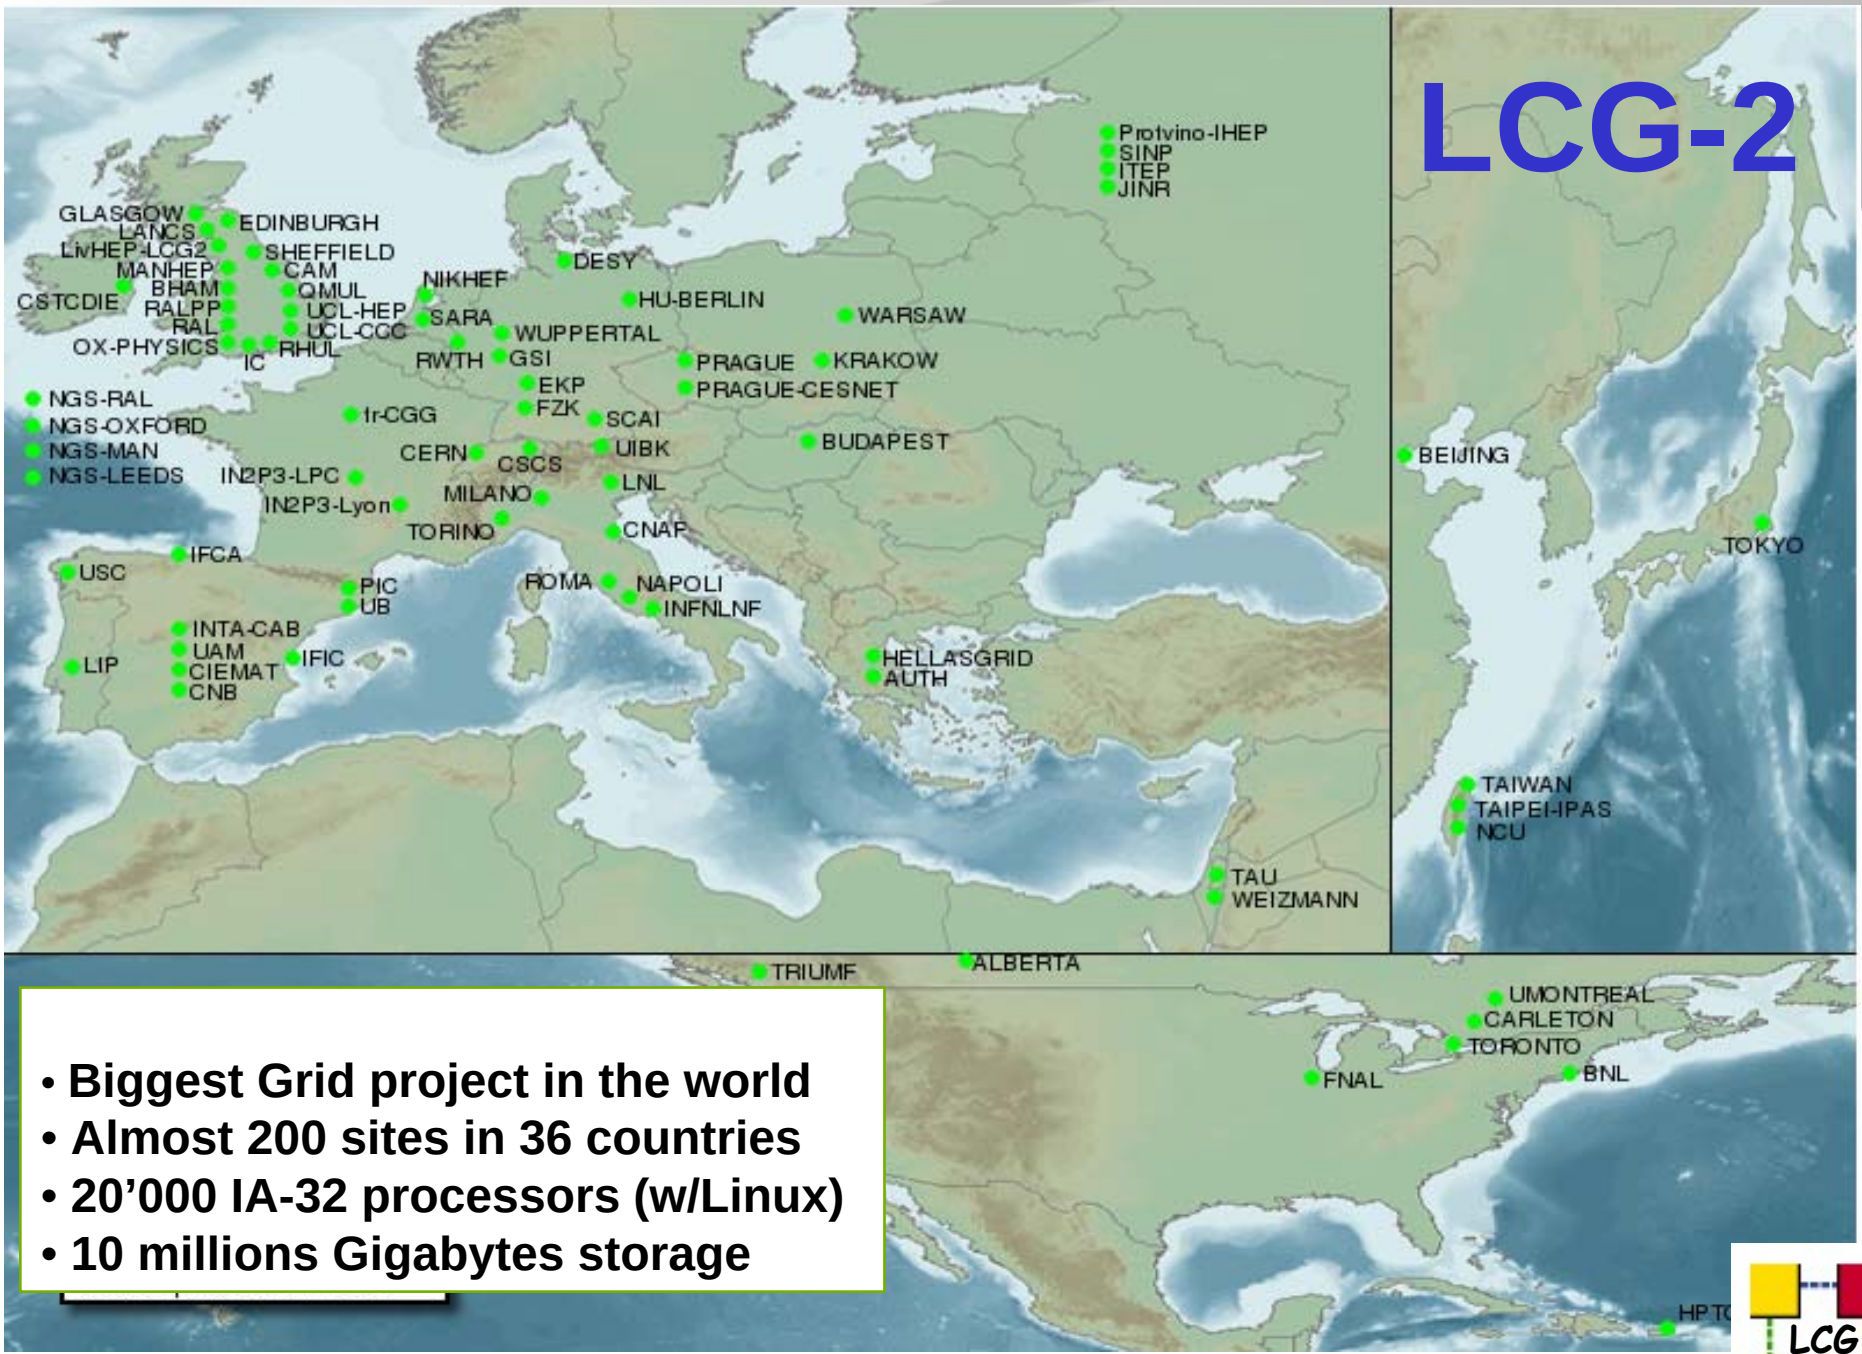
LCG-2
Protvino-IHEP
SINP
ITEP
JINR
GLASGOW
EDINBURGH
LANCS
LivHEP-LCG2
SHEFFIELD
MANHEP
CAM
BHAM
QMUL
CSTCDIE
RALPP
UCL-HEP
RAL
UCL-CCG
OX-PHYSICS
RHUL
IC
NIKHEF
SARA
RWTH
DESY
HU-BERLIN
WUPPERTAL
GSI
WARSAW
PRAGUE
KRAKOW
PRAGUE-CESNET
EKP
FZK
SCAI
UIBK
BUDAPEST
NGS-RAL
NGS-OXFORD
NGS-MAN
NGS-LEEDS
fr-CGG
CERN
CSCS
IN2P3-LPC
IN2P3-Lyon
MILANO
LNL
TORINO
CNAF
IFCA
USC
PIC
UB
ROMA
NAPOLI
INFNLNF
INTA-CAB
UAM
CIEMAT
CNB
IFIC
LIP
HELLASGRID
AUTH
BEIJING
TOKYO
TAIWAN
TAIPEI-IPAS
NCU
TAU
WEIZMANN
TRIUMF
ALBERTA
UMONTREAL
CARLETON
TORONTO
FNAL
BNL
• Biggest Grid project in the world
• Almost 200 sites in 36 countries
• 20’000 IA-32 processors (w/Linux)
• 10 millions Gigabytes storage

Wide Area Network

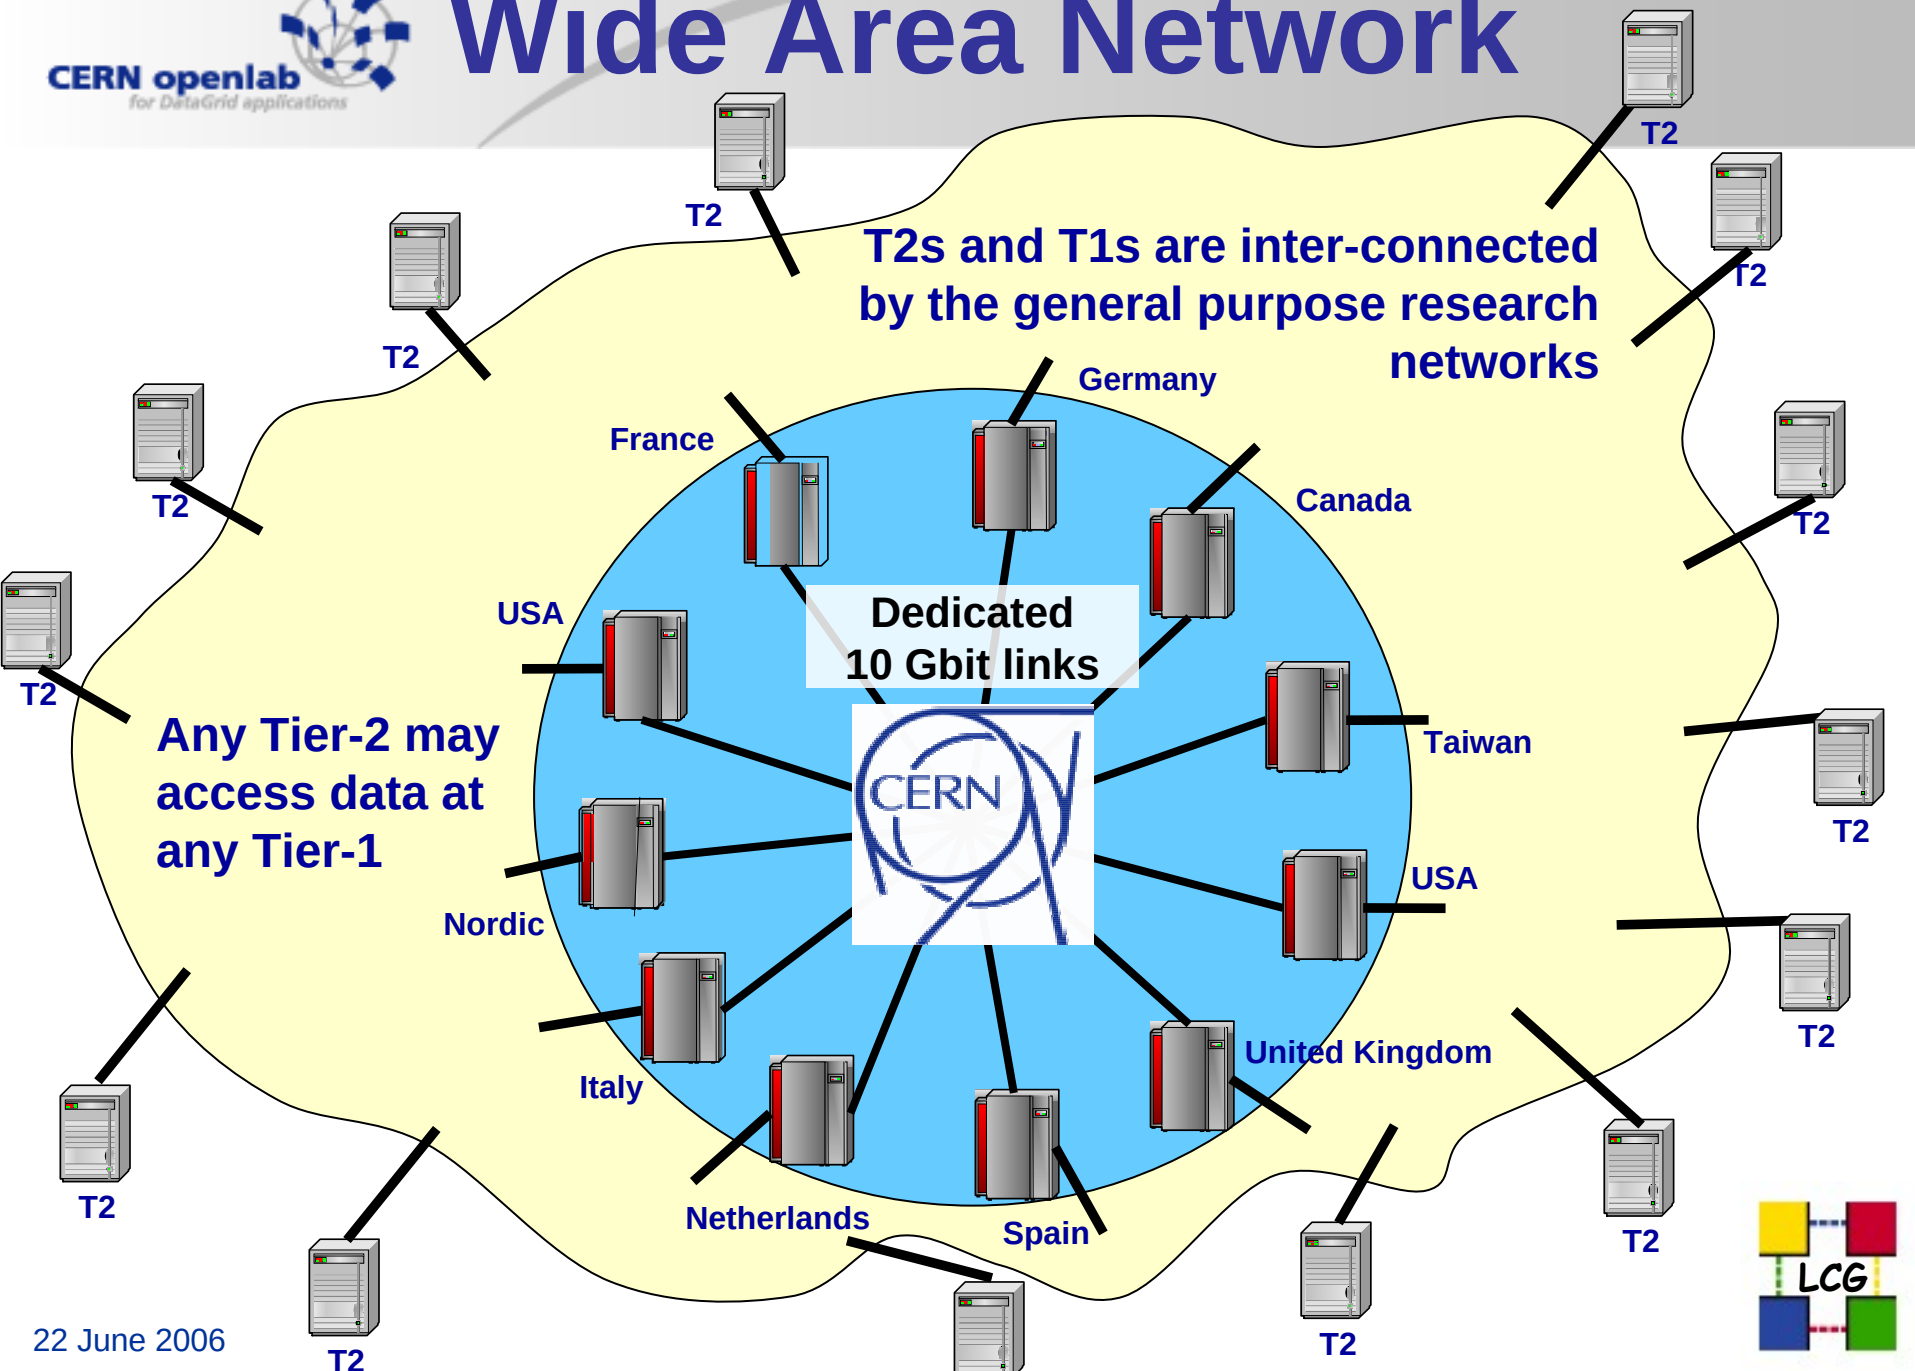

CERN openlab

www.cern.ch/openlab

PARTNERS

CONTRIBUTORS

CERN openlab

- CERN-IT department's main R&D focus
- Framework for collaboration with industry
- Evaluation, integration, validation
  - of cutting-edge technologies that can serve the LHC Computing Grid (LCG)
- Sequence of 3-year agreements
  - 2003 – 2005: the "opencluster" project
  - 2006 – 2008: openlab Phase II with new projects

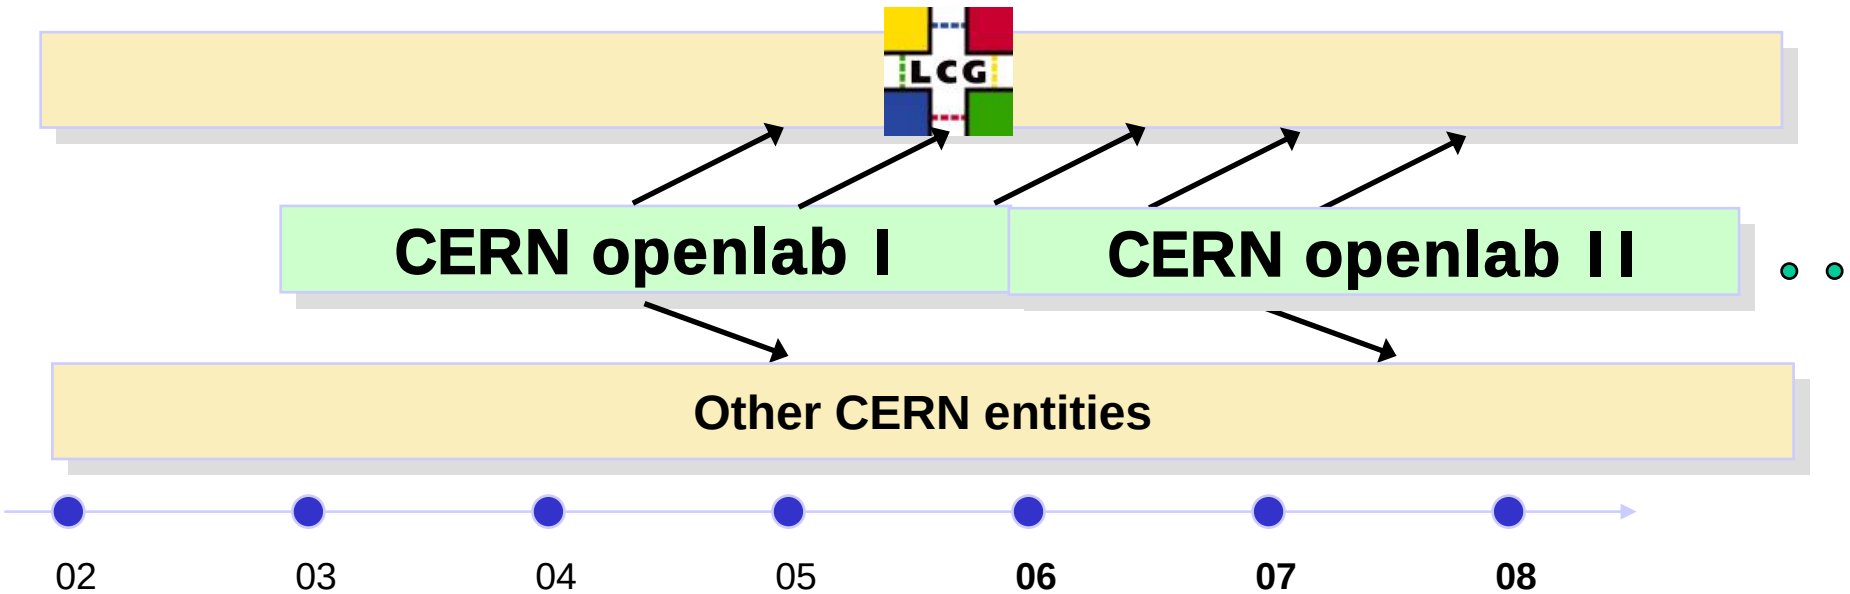

Computational Fluid Dynamics

- **Based on Itanium cluster with Infiniband switches from Voltaire**
- **CFD:**
  - A numerical analysis of fluid flow, heat transfer and associated phenomena in physical systems
  - Always limited by available computing resources
  - Reduces design and engineering costs by avoiding prototype studies
  - Calculation improved by almost an order of magnitude
    - From, for instance, one month to less than four days
  - Model dimensions increased from 0.5 to 3 M cells
- **Very important contribution to all the LHC experiments**
  - and others

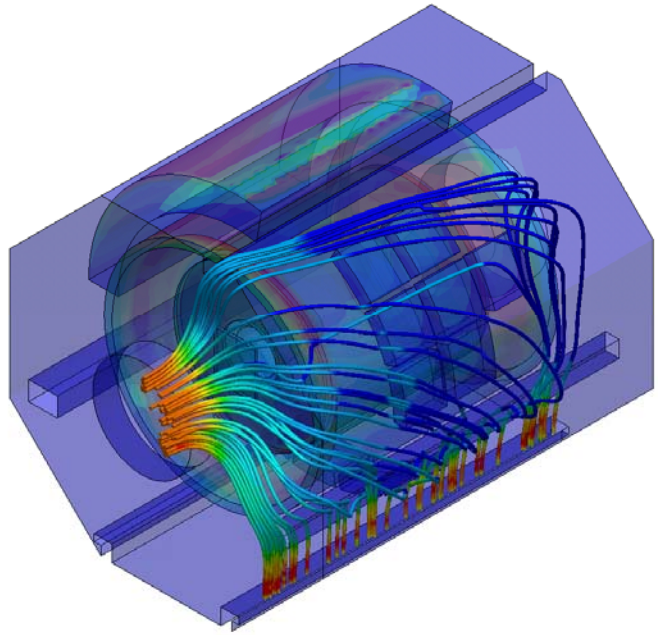

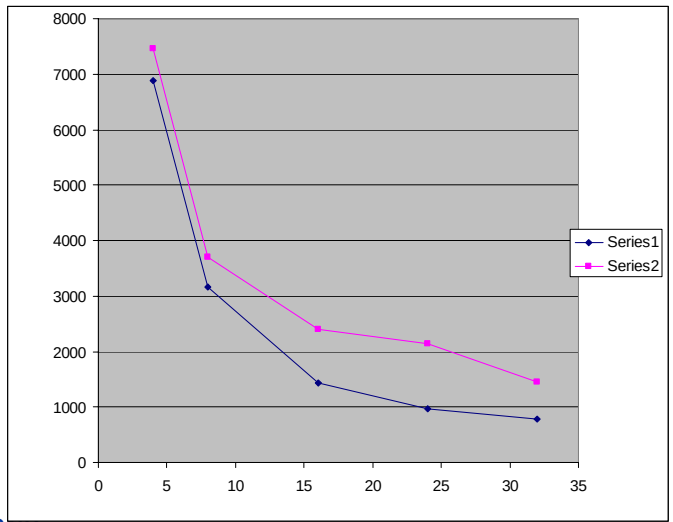

Conclusions

- **CERN is busily preparing for the arrival of LHC data in one year's time!**
  - New and exciting technologies will be used to cope with the data
    - 10 Gb networking
    - Terabyte disk and tape technology
    - 64-bit processors with multicore and virtualization capabilities
  - Our Grid offers seamless integration, all around the globe
    - Together with our partners (EU, industrial partners, other Physics Labs, other sciences) we expect to continue to come up with interesting proofs-of-concept and technological spin-off !
- **High Throughput Computing is "on the move" !**

LHC Computing