where the Web was born

CERN and the LHC Computing Challenge

by

Wolfgang von Rüden

Head, IT Department

HP DutchWorld

12th October 2004

What is CERN?

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

CERN is also:

-2500 staff (physicists, engineers, technicians, …)

- Some 6500 visiting scientists (half of the world's particle physicists)

They come from 500 universities representing 80 nationalities.

What is CERN?

- Physicists smash particles into each other to:
  - identify their components
  - create new particles
  - reveal the nature of the interactions between them
  - recreate the environment present at the origin of our Universe (big bang)

- What for? To answer fundamental questions like:

how did the Universe begin? What is the origin of mass? What is the nature of antimatter?

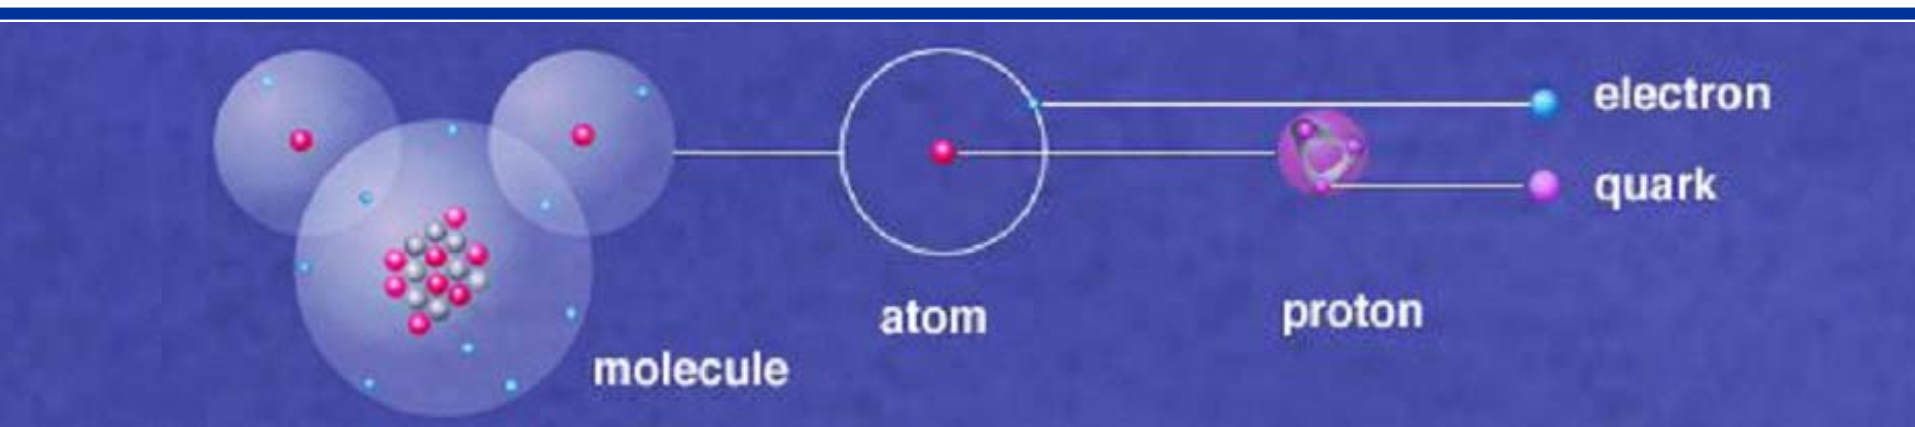

What is CERN?

The special tools for particle physics are:

- ACCELERATORS, huge machines able to speed up 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 re-construct the collisions, to extract the physics data and perform the analysis

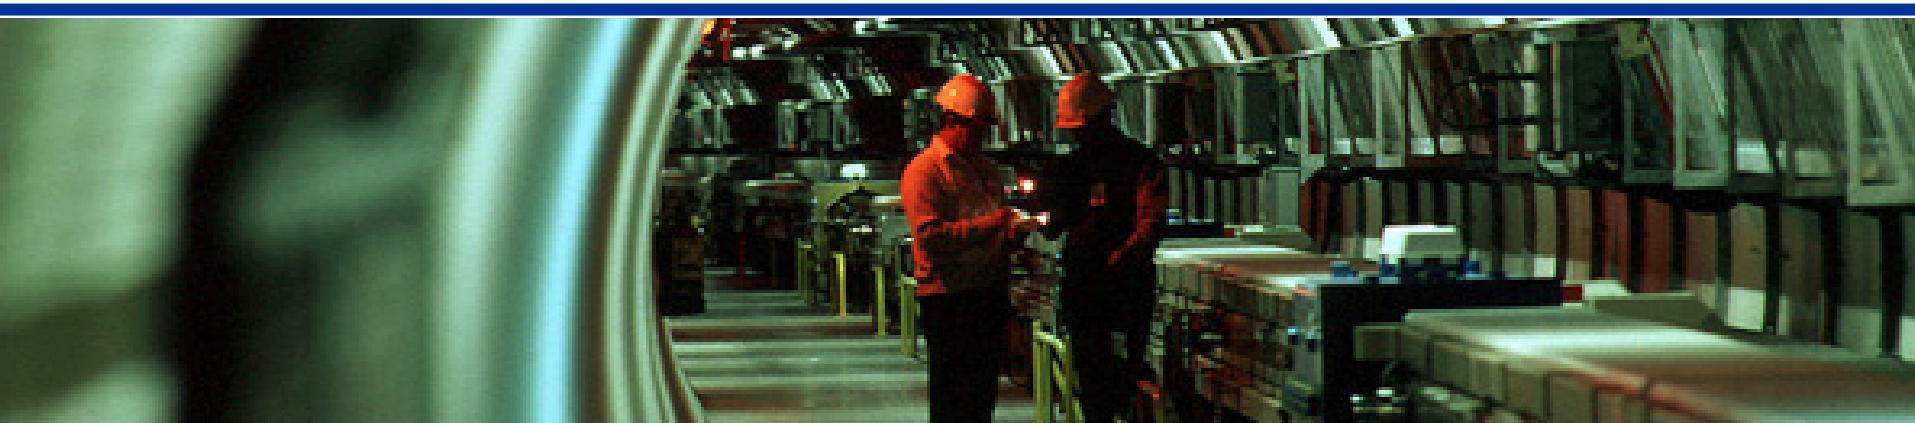

What is CERN?

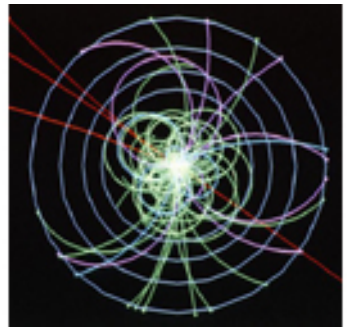

- CERN has made many important discoveries, but our current understanding of the Universe is still incomplete!
- Higher energy collisions are the key to further discoveries of more massive particles ($E=mc^2$)
- One particle predicted by theorists remains elusive: the Higgs boson

What is CERN?

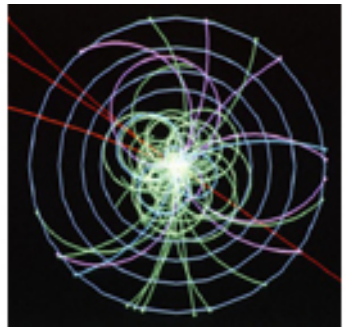

- To answer some of the still open questions, CERN is building a new accelerator, the Large Hadron Collider (LHC)
- The LHC will be the most powerful instrument ever built to investigate elementary particles
- Four very large experiments matching this machine are under construction, ready to make new discoveries in 2007 and beyond
- If the Higgs boson exists, then we will most certainly find it

CERN's accelerator complex

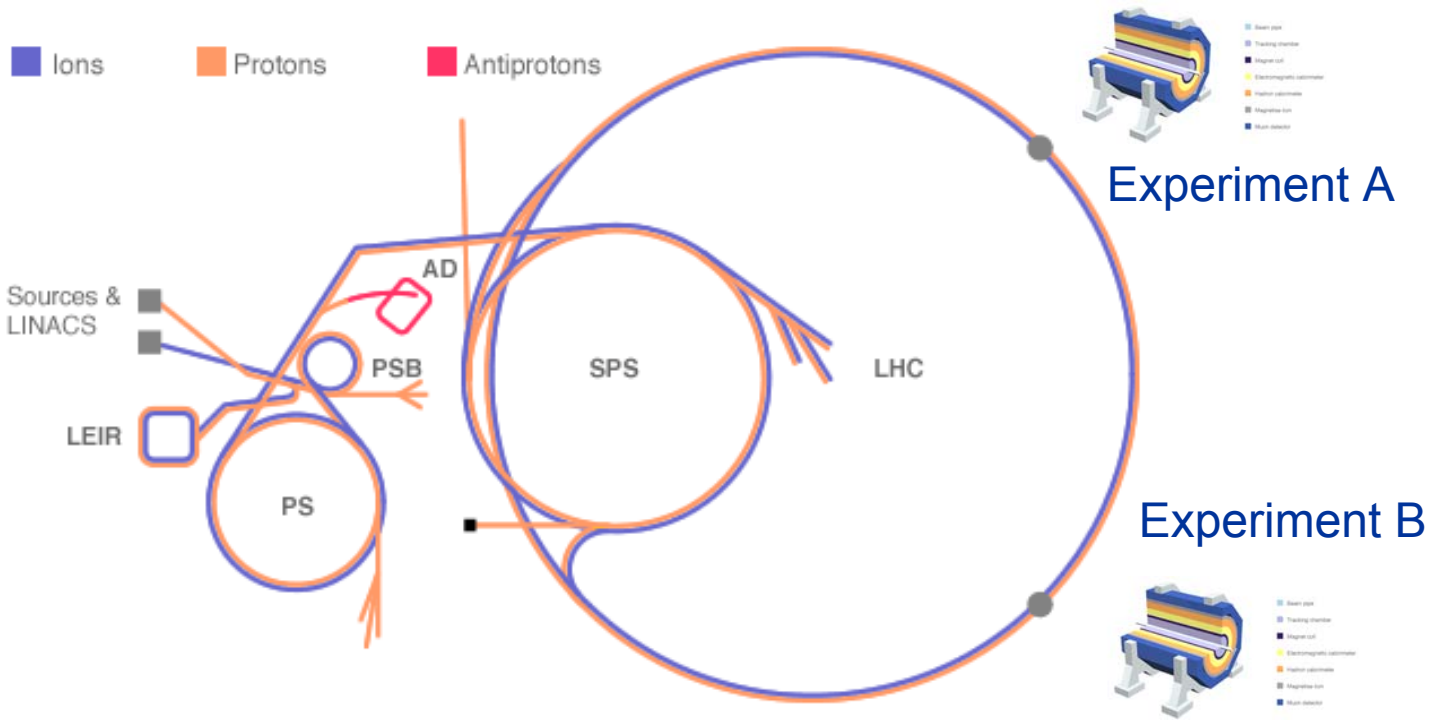

- PS – Proton Synchrotron
- SPS – Super Proton Synchrotron
- LHC – Large Hadron Collider

The CERN Site

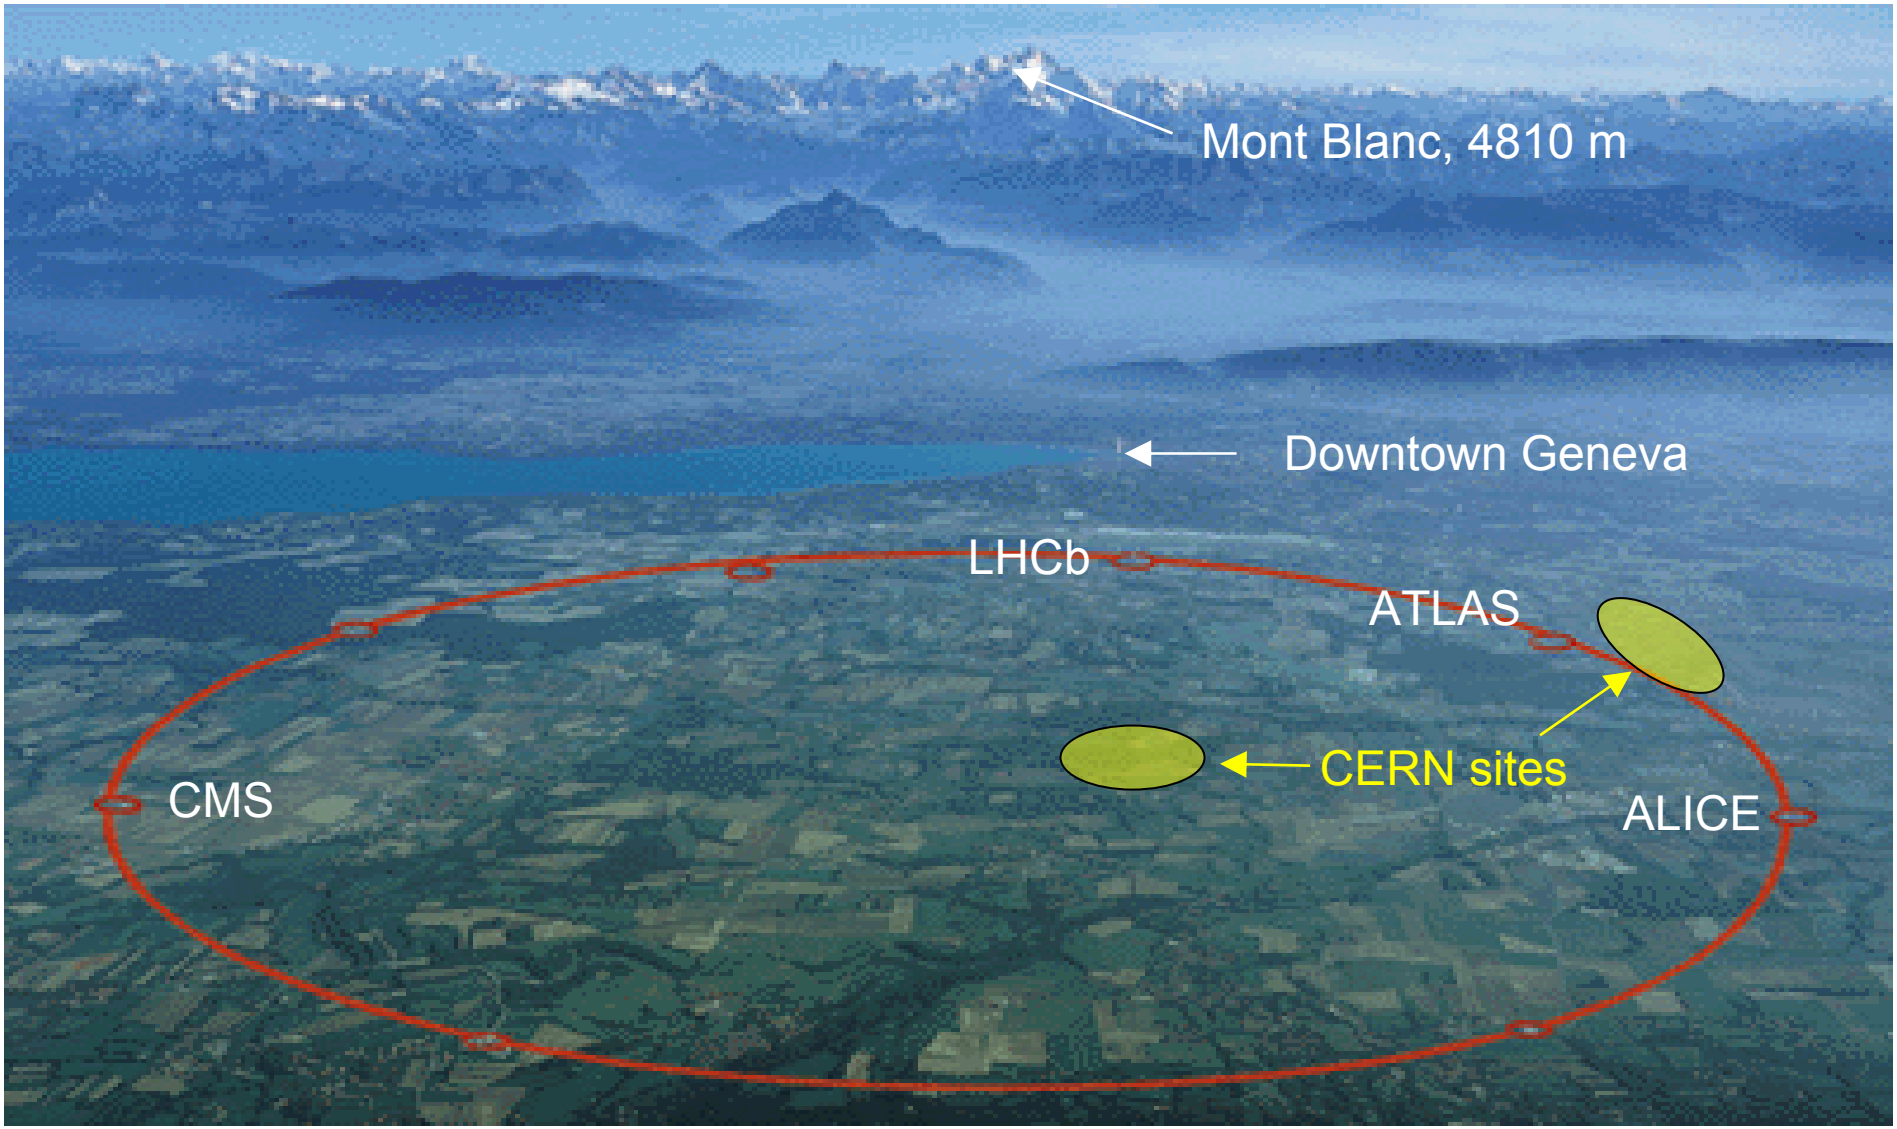

What is LHC?

- LHC 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.

LHC is due to switch on in 2007

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

Typical Experiment Layout

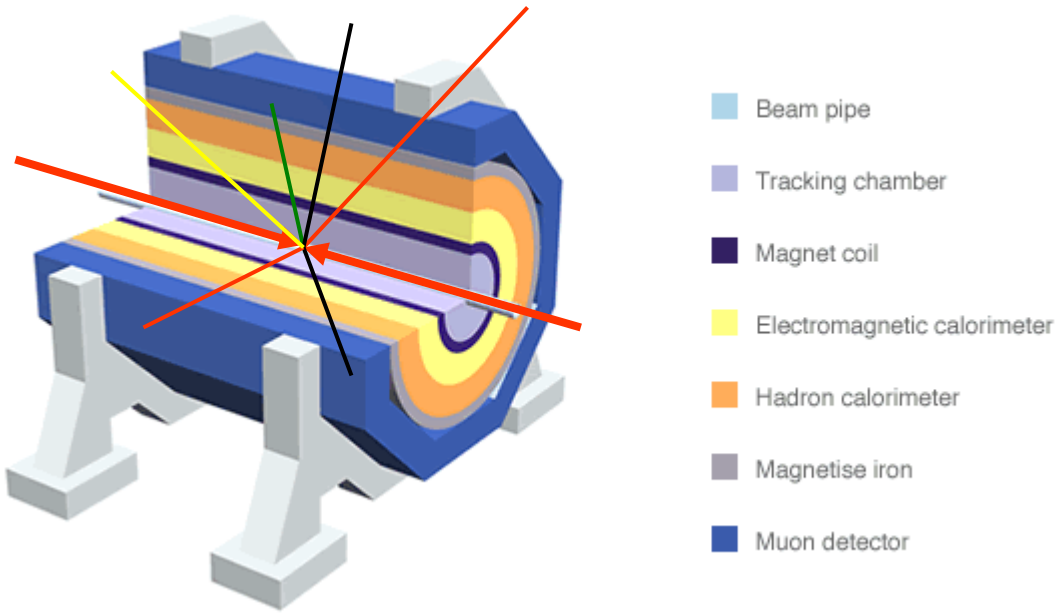

- Complex system of detectors centred around the beam interaction point

Particle Detection Techniques

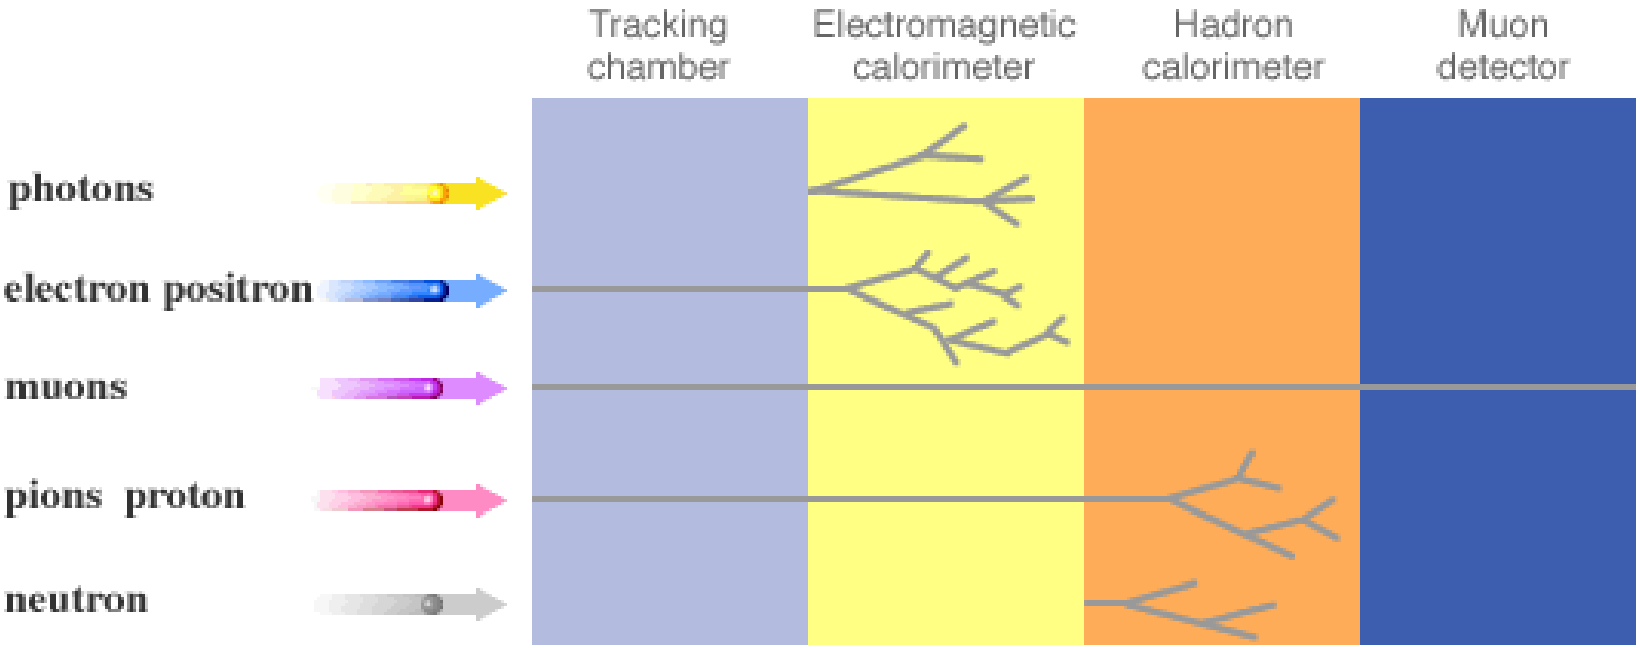

Multiple layers of increasing density to identify particles

ATLAS, one of the four LHC experiments

ATLAS has 150 million measurement channels

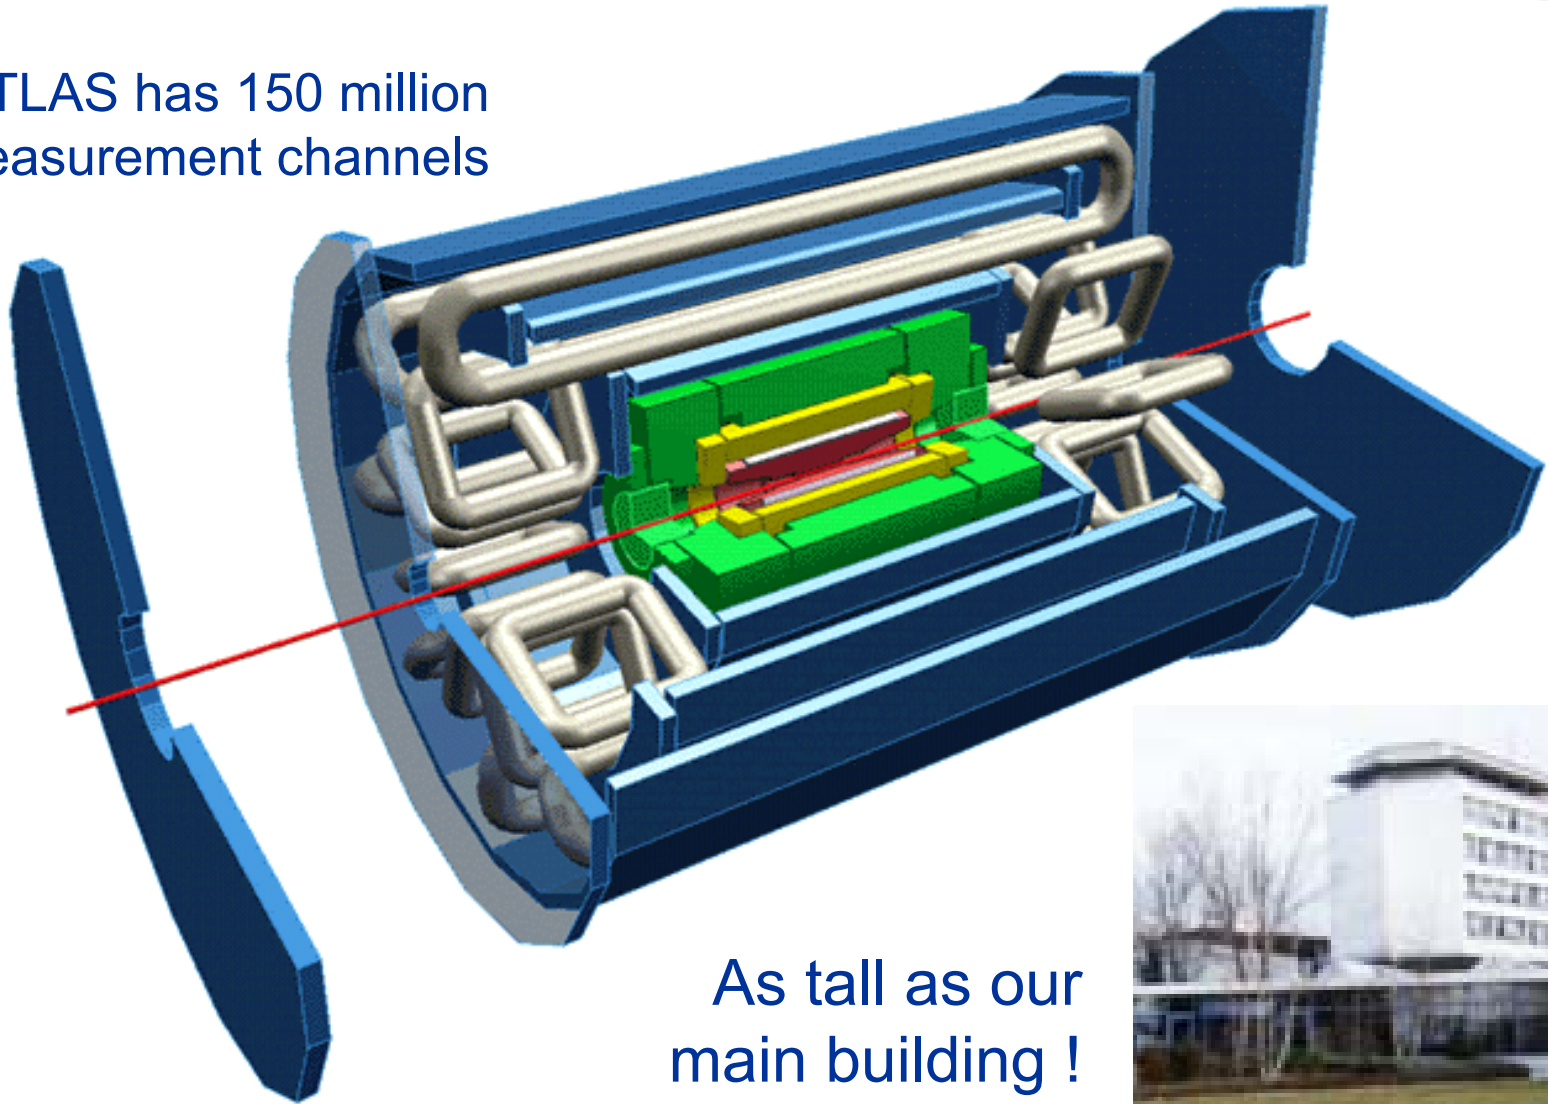

As tall as our main building !

LHC data (simplified)

Per experiment:

- 40 million collisions per second
- After filtering, 100 collisions of interest per second
- A Megabyte of digitised information for each collision = recording rate of 0.1 Gigabytes/sec
- 1 billion collisions recorded = 1 Petabyte/year

Total: ~10.000.000.000.000.000 **= 1% of**

1 Megabyte (1MB)
A digital photo

1 Gigabyte (1GB)
= 1000MB
A DVD movie

1 Terabyte (1TB)
= 1000GB
World annual book production

1 Petabyte (1PB)
= 1000TB
10% of the annual production by LHC experiments

1 Exabyte (1EB)
= 1000 PB
World annual information production

CMS

LHCb

ATLAS

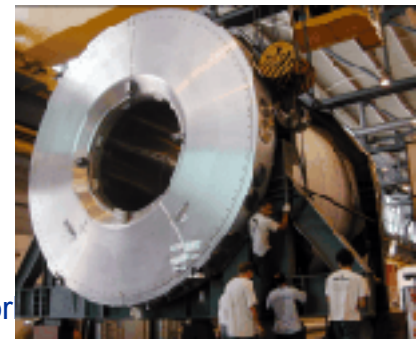

ALICE

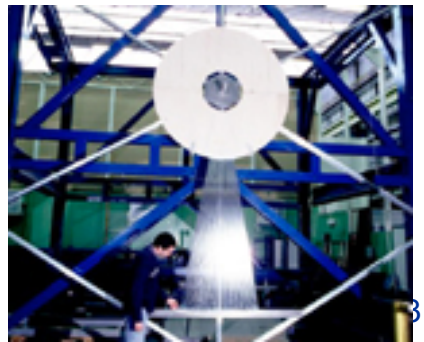

LHC data (simplified)

LHC data correspond to about 20 million CDs each year

Where will the experiments store all of these data?

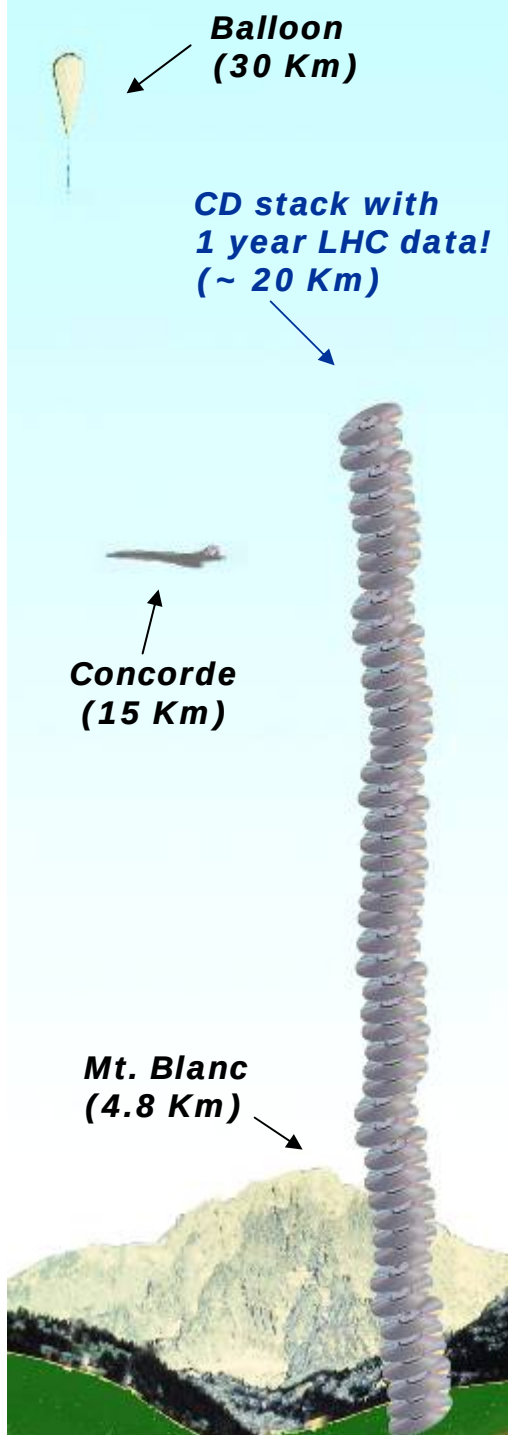

LHC data processing

LHC data analysis requires a computing power equivalent to ~ 70,000 of today's fastest PC processors

Where will the experiments find such a computing power?

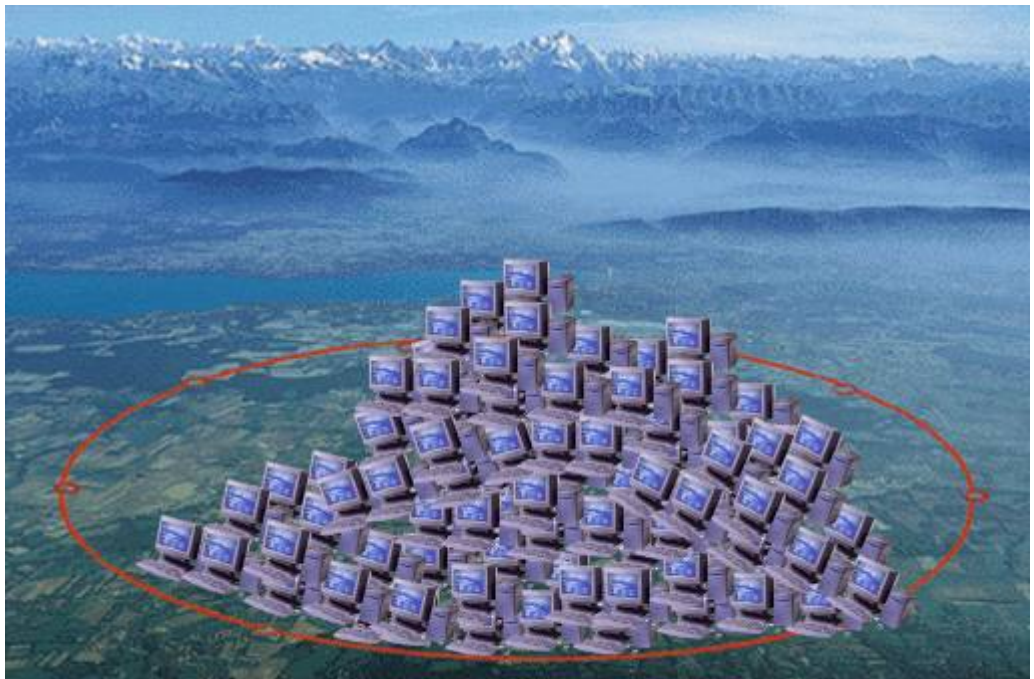

Data Handling and Computation for Physics Analysis

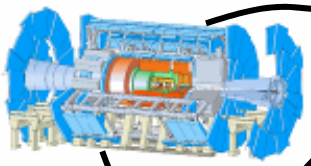

detector

selection & reconstruction

reconstruction

event summary data

processed data

raw data

event reprocessing

batch physics analysis

analysis

analysis objects
(extracted by physics topic)

event simulation

simulation

interactive physics analysis

les.robertson@cern.ch

High Throughput Computing

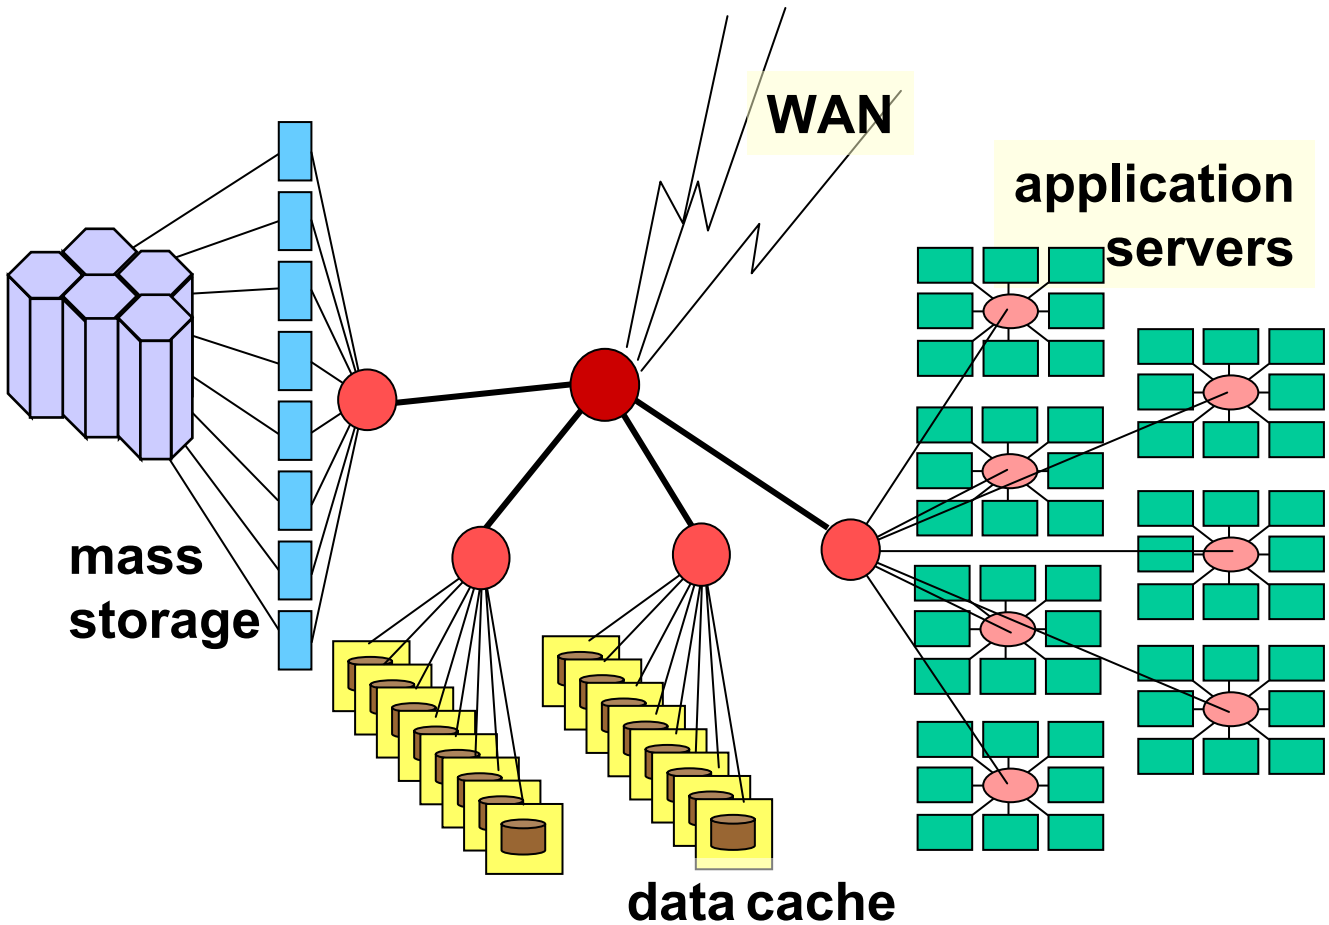

simple, flexible architecture

- easy to integrate mass market components
- easy evolution to new technologies

High Throughput Computing

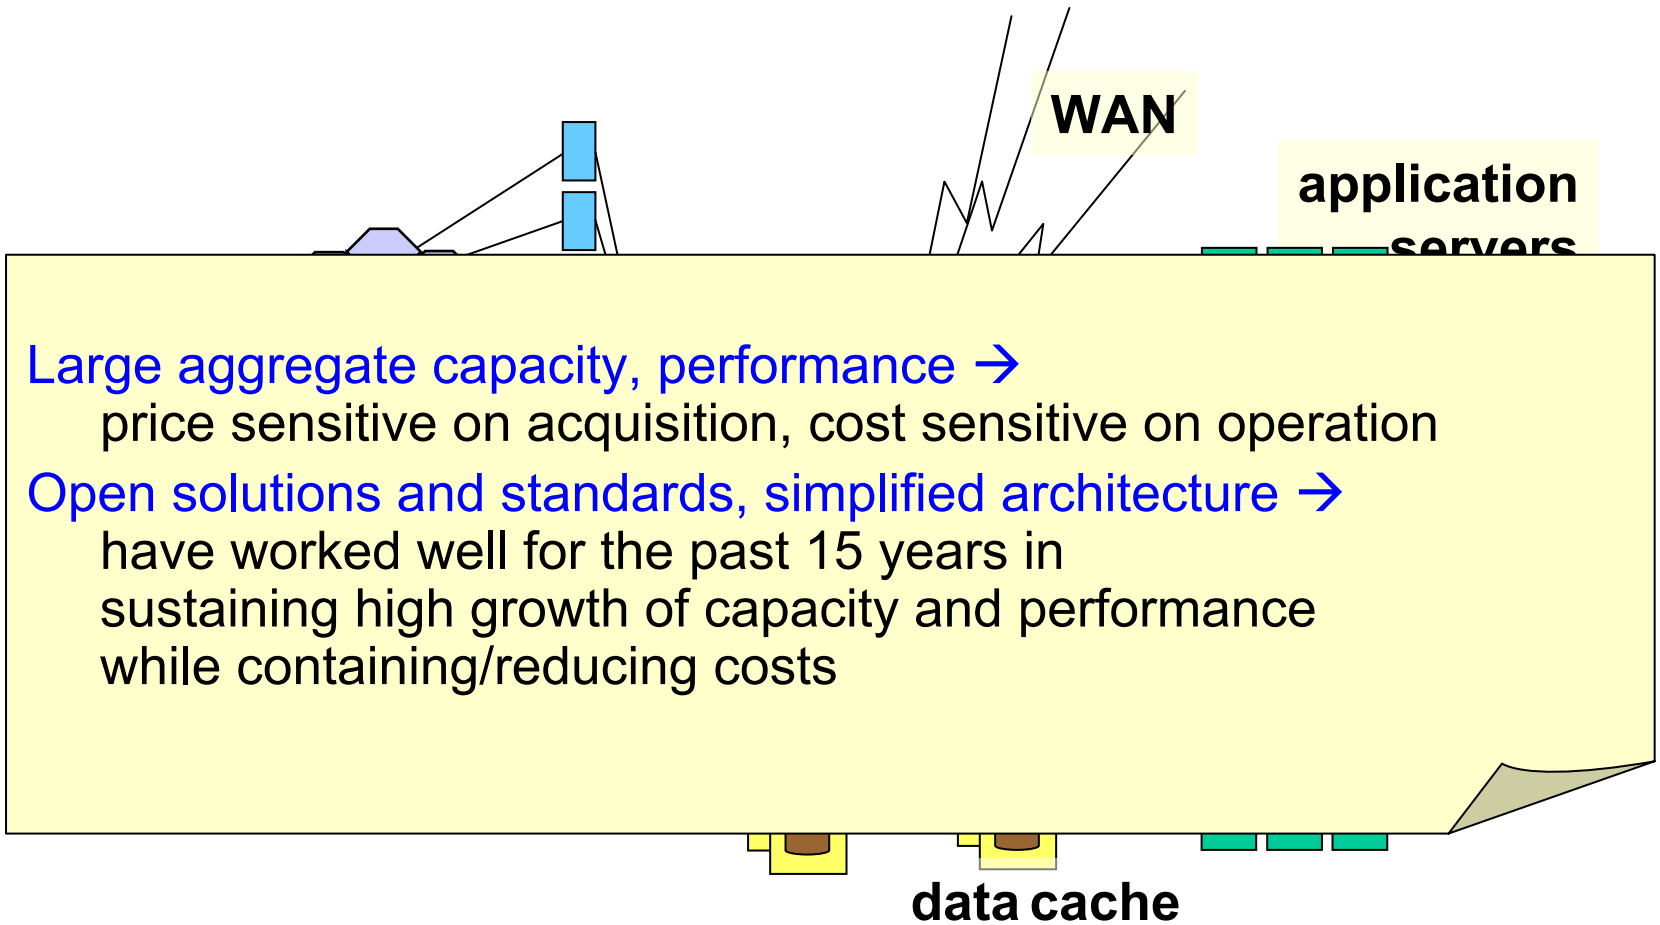

simple, flexible architecture

- easy to integrate mass market components
- easy evolution to new technologies

Computing at CERN today

- High-throughput computing based on reliable “commodity” technology
- About 2000 dual processor PCs
- More than 3 Petabyte of data on disk (10%) and tapes (90%)

Nowhere near enough!

Building up LHC computing at CERN

Networking by 2007:
~ 70 Gb/s to external centres
~ 100 Gb/s general networking

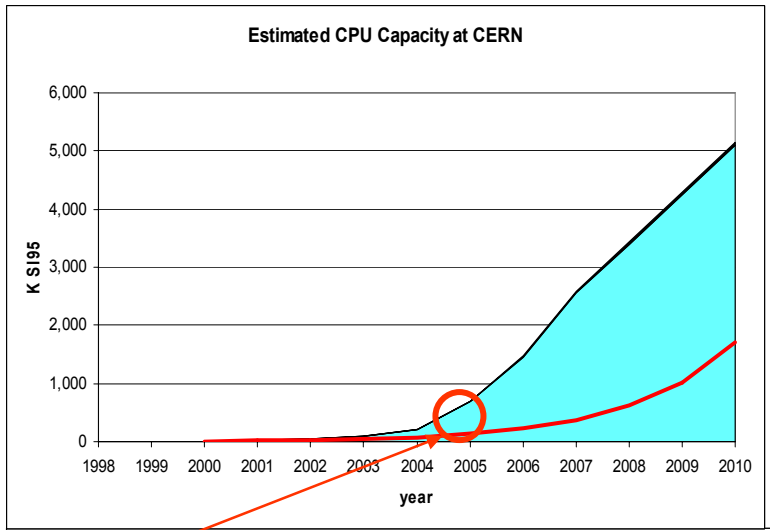

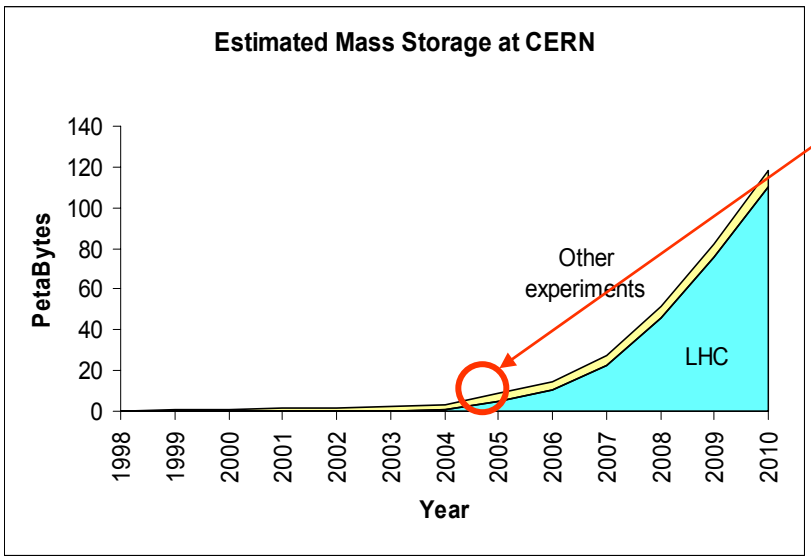

today

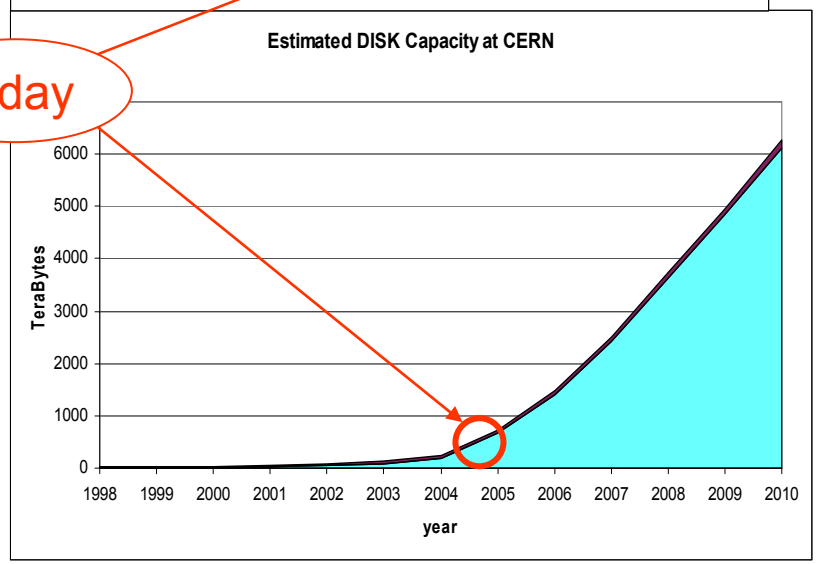

Preparing for 7,000 boxes in 2008

2.5 MW power

Preparing for 7,000 boxes in 2008

Preparing for 7,000 boxes in 2008

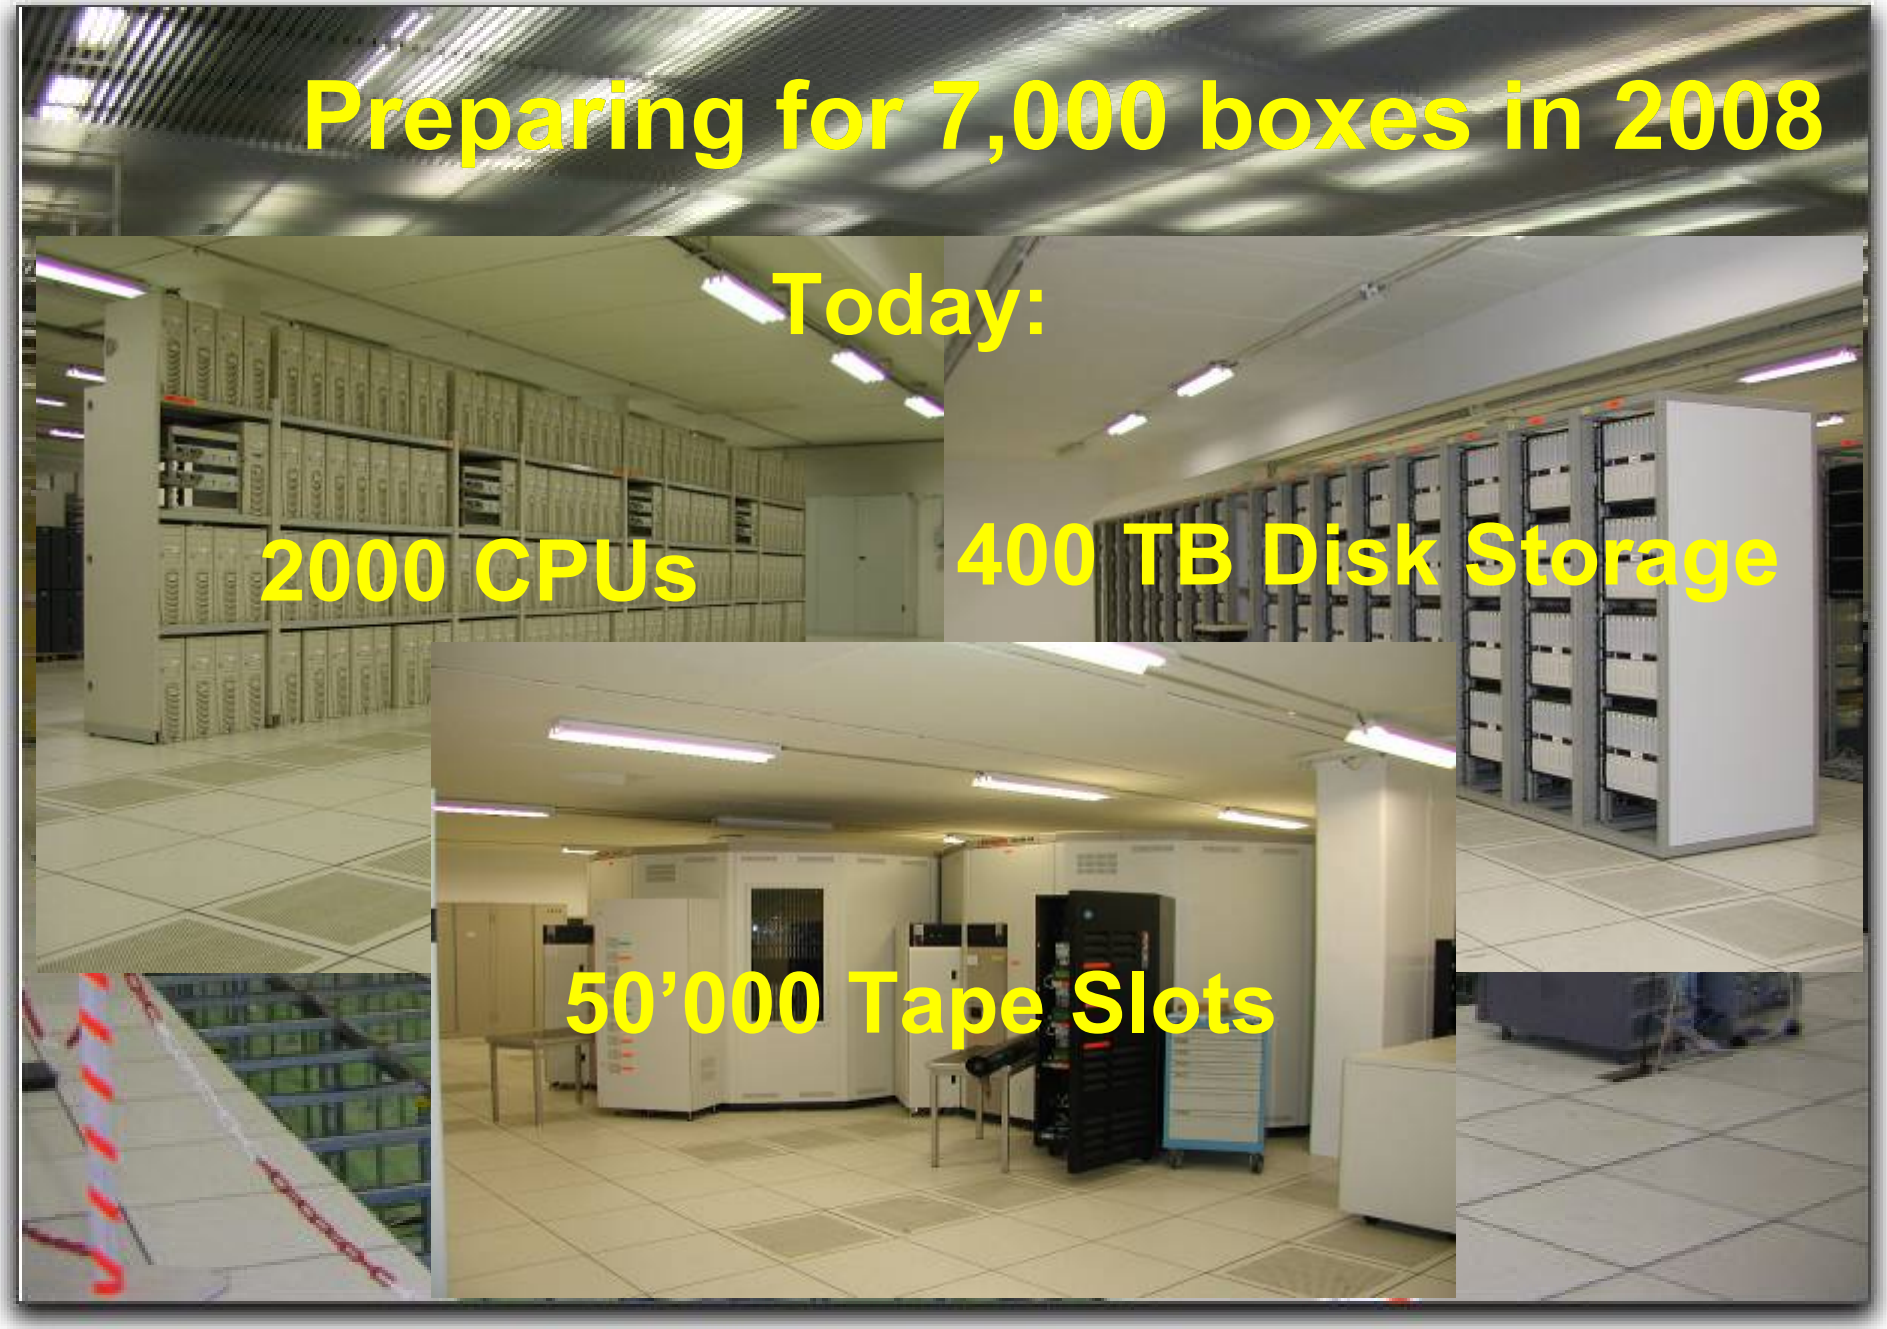

Computing for LHC

Europe:
~270 institutes
~4500 users

Elsewhere:
~200 institutes
~1600 users

- **Problem:** even with computer centre upgrade, CERN can only provide a fraction of the necessary resources

- **Solution:** computing centres, which were isolated in the past, will now be connected, **uniting the computing resources of particle physicists in the world using GRID technologies!**

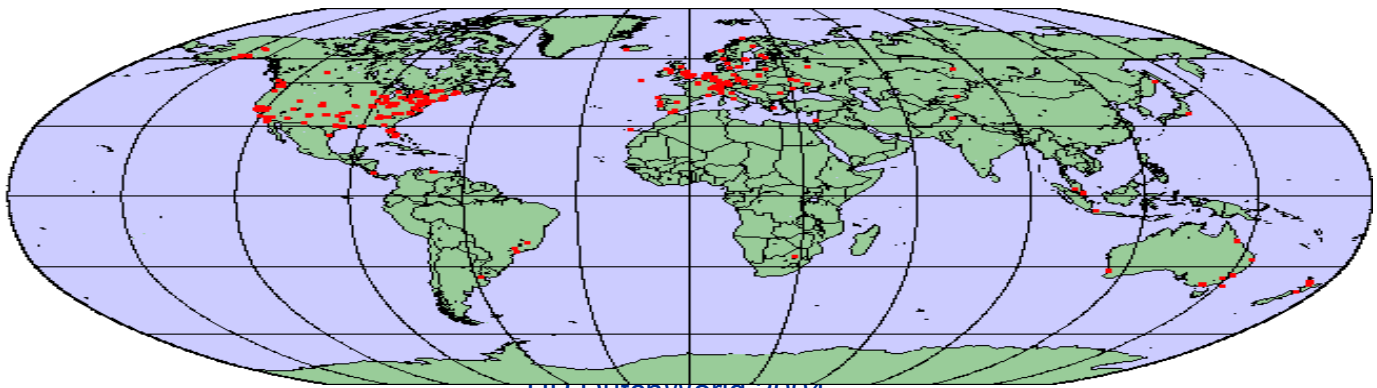

Grid @ CERN

- **LHC Computing Grid (LCG) – the flagship project**
- **Enabling Grids for E-Science in Europe (EGEE)**
  - **Has started in April 2004 with 70 partners and 32M€ EU funding**
  - **Will provide the next generation middleware**
  - **Will run a 24/7 Grid service together with LCG**
- **CERN openlab for DataGrid applications**
  - **Funded by CERN and Industry**
  - **Main project: opencluster**
  - **New project: openlab security (under preparation)**

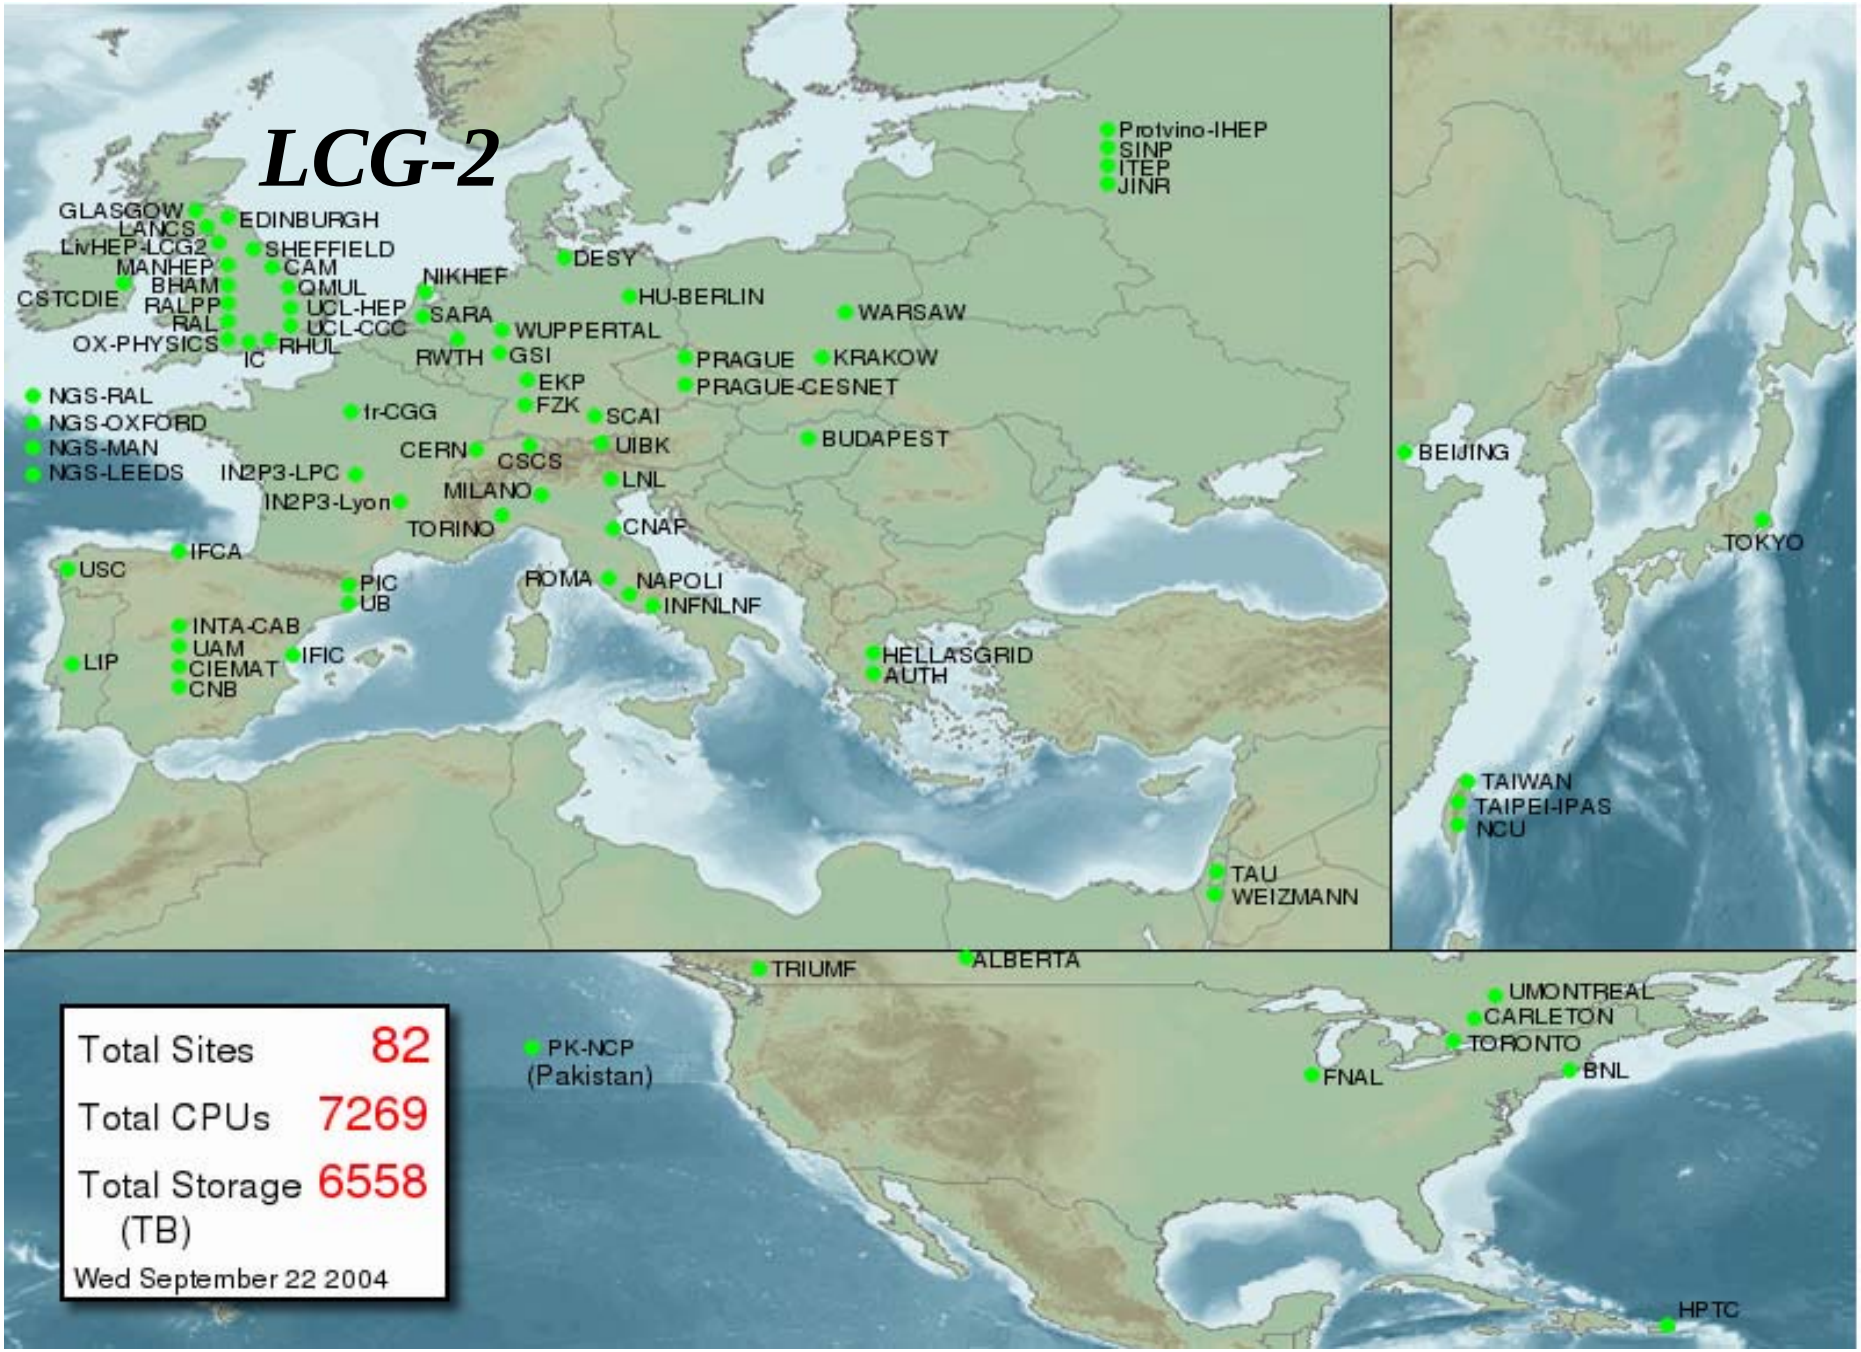
LCG-2
Protvino-IHEP
SINP
ITEP
JINR
GLASGOW
EDINBURGH
LANCS
LivHEP-LCG2
SHEFFIELD
MANHEP
CAM
BHAM
QMUL
CSTCDIE
RALPP
UCL-HEP
RAL
UCL-CCG
OX-PHYSICS
RHUL
IC
NIKHEF
SARA
RWTH
DESY
WUPPERTAL
GSI
HU-BERLIN
WARSAW
PRAGUE
KRAKOW
PRAGUE-CESNET
EKP
FZK
SCAI
UIBK
BUDAPEST
NGS-RAL
NGS-OXFORD
NGS-MAN
NGS-LEEDS
fr-CGG
CERN
CSCS
IN2P3-LPC
IN2P3-Lyon
MILANO
LNL
TORINO
CNAF
IFCA
USC
PIC
UB
ROMA
NAPOLI
INFNLNF
INTA-CAB
UAM
CIEMAT
CNB
LIP
IFIC
HELLASGRID
AUTH
BEIJING
TOKYO
TAIWAN
TAIPEI-IPAS
NCU
TAU
WEIZMANN
TRIUMF
ALBERTA
UMONTREAL
CARLETON
TORONTO
BNL
FNAL
PK-NCP
(Pakistan)
HPTC
Total Sites 82
Total CPUs 7269
Total Storage (TB) 6558
Wed September 22 2004

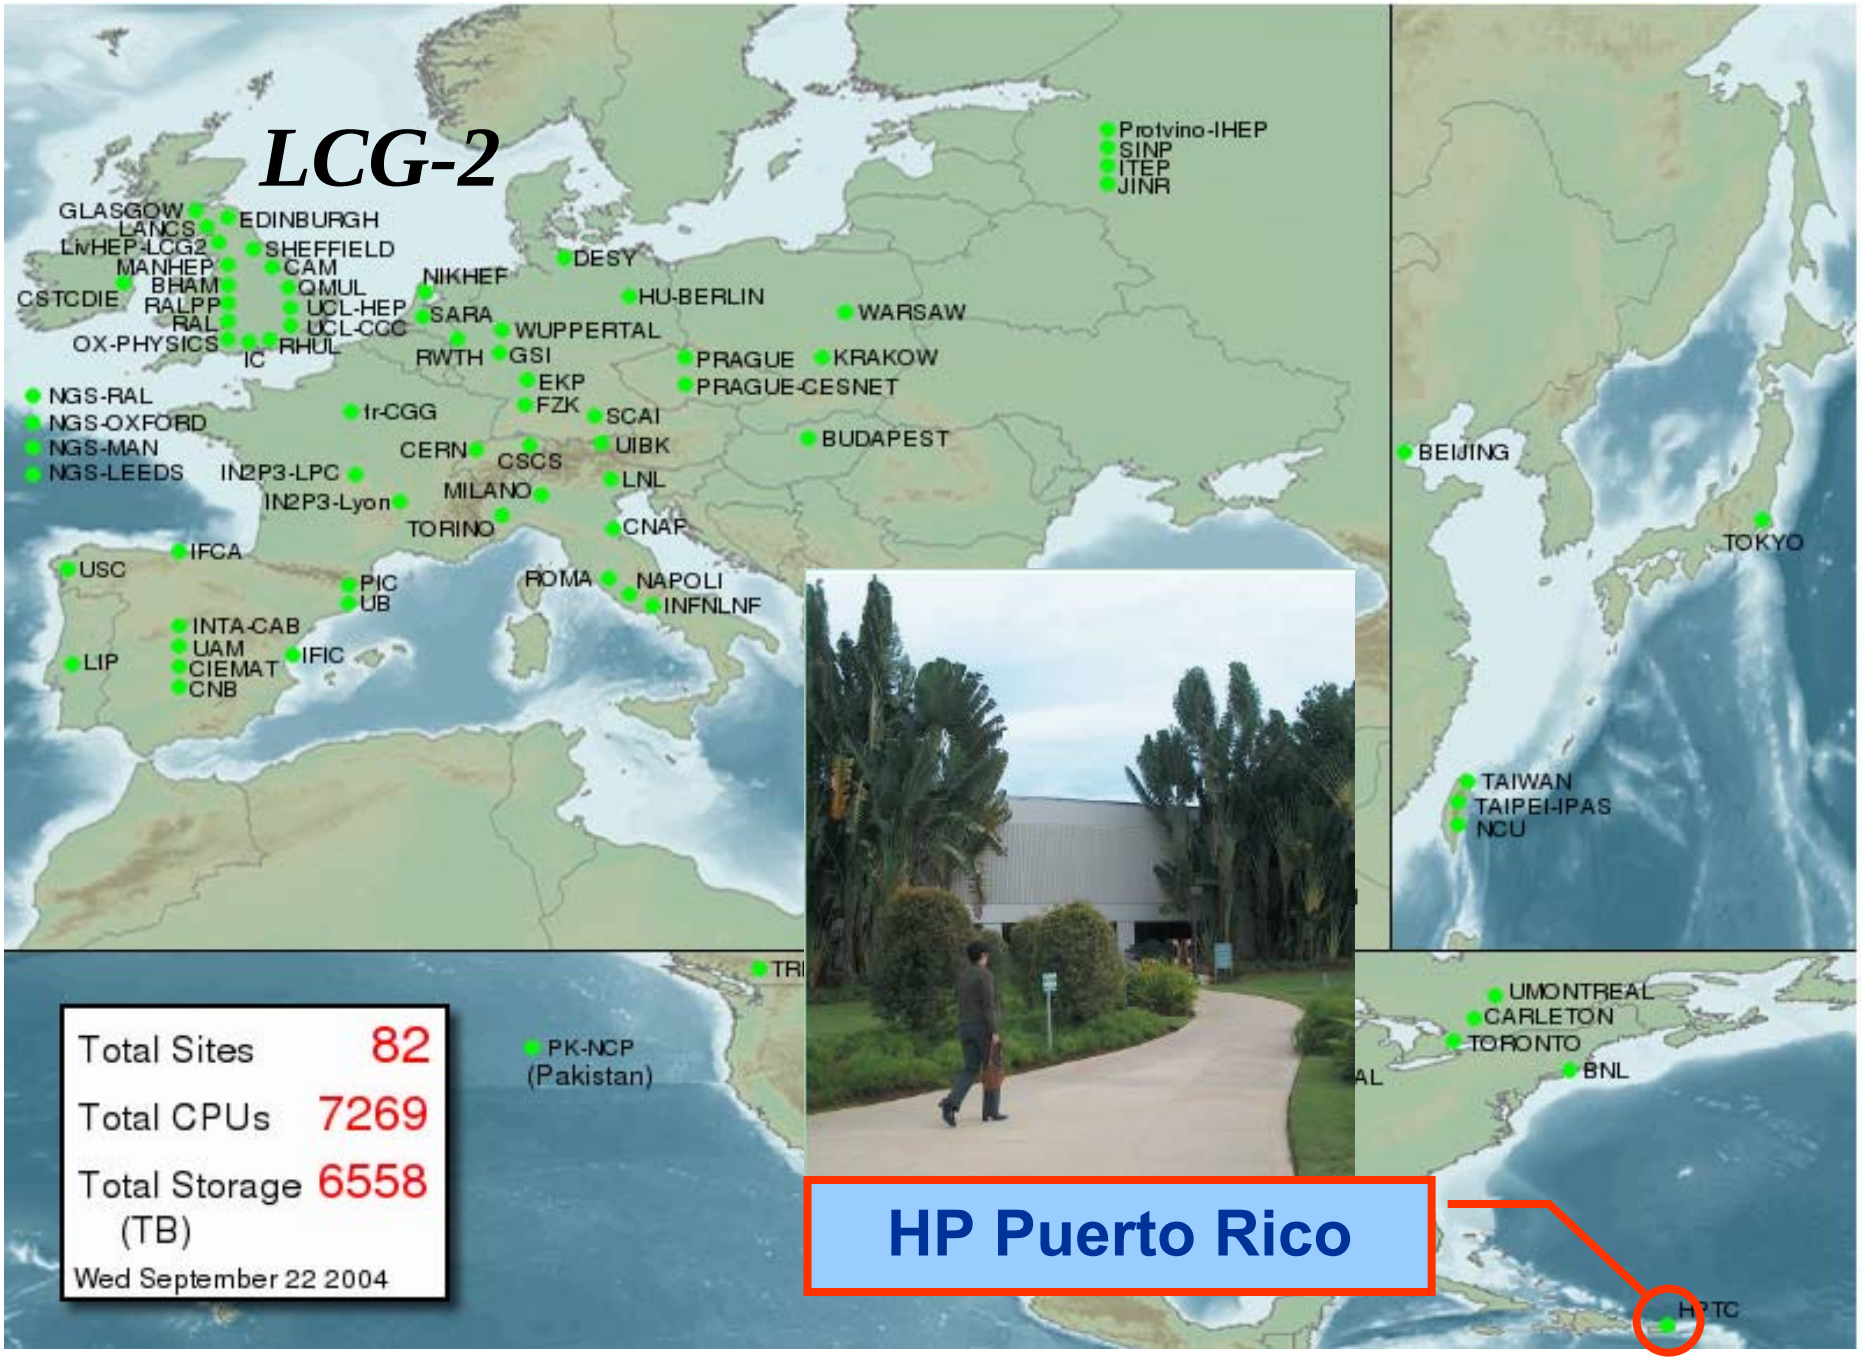
LCG-2
Protvino-IHEP
SINP
ITEP
JINR
GLASGOW
EDINBURGH
LANCS
LivHEP-LCG2
SHEFFIELD
MANHEP
CAM
BHAM
QMUL
CSTCDIE
RALPP
UCL-HEP
RAL
UCL-CCG
OX-PHYSICS
RHUL
IC
NIKHEF
SARA
DESY
HU-BERLIN
WARSAW
WUPPERTAL
RWTH
GSI
PRAGUE
KRAKOW
EKP
PRAGUE-CESNET
FZK
SCAI
NGS-RAL
fr-CGG
NGS-OXFORD
NGS-MAN
NGS-LEEDS
CERN
CSCS
UIBK
BUDAPEST
IN2P3-LPC
LNL
MILANO
IN2P3-Lyon
TORINO
CNAF
BEIJING
TOKYO
IFCA
USC
PIC
UB
ROMA
NAPOLI
INFNLNF
INTA-CAB
UAM
CIEMAT
CNB
IFIC
LIP
TAIWAN
TAIPEI-IPAS
NCU
UMONTREAL
CARLETON
TORONTO
BNL
PK-NCP
(Pakistan)
Total Sites 82
Total CPUs 7269
Total Storage (TB) 6558
Wed September 22 2004
HP Puerto Rico

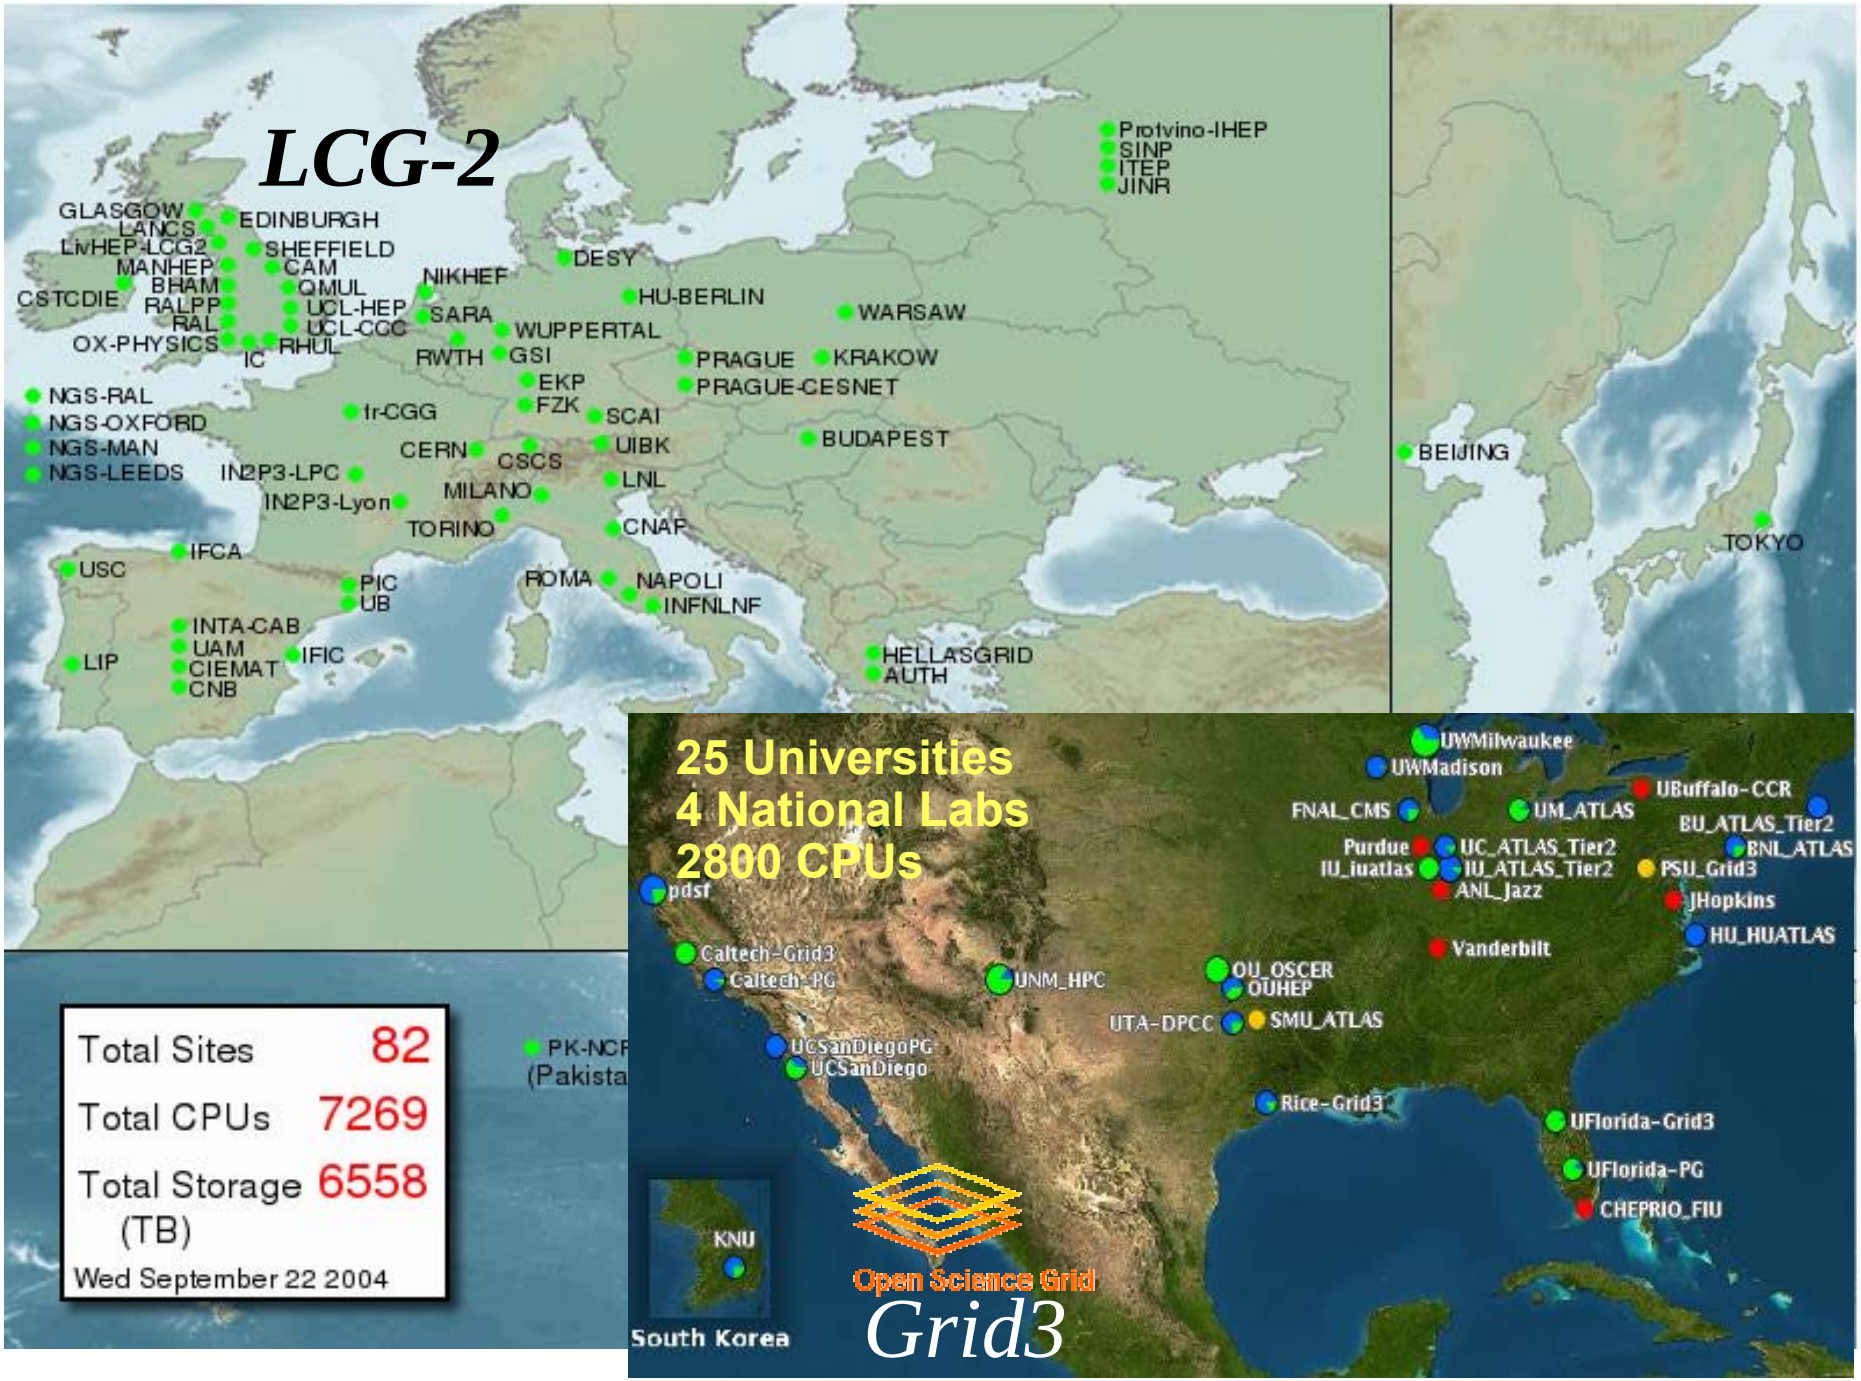
LCG-2
Protvino-IHEP
SINP
ITEP
JINR
GLASGOW
EDINBURGH
LANCS
LivHEP-LCG2
SHEFFIELD
MANHEP
CAM
BHAM
QMUL
CSTCDIE
RALPP
UCL-HEP
RAL
UCL-CCC
OX-PHYSICS
RHUL
IC
NIKHEF
SARA
RWTH
DESY
HU-BERLIN
WUPPERTAL
GSI
WARSAW
PRAGUE
KRAKOW
EKP
PRAGUE-CESNET
FZK
SCAI
NGS-RAL
fr-CGG
NGS-OXFORD
NGS-MAN
NGS-LEEDS
CERN
CSCS
UIBK
BUDAPEST
IN2P3-LPC
MILANO
LNL
IN2P3-Lyon
TORINO
CNAF
BEIJING
TOKYO
IFCA
USC
PIC
UB
ROMA
NAPOLI
INFNLNF
INTA-CAB
UAM
CIEMAT
CNB
IFIC
LIP
HELLASGRID
AUTH
PK-NCP
(Pakista
Total Sites 82
Total CPUs 7269
Total Storage (TB) 6558
Wed September 22 2004
25 Universities
4 National Labs
2800 CPUs
UWMilwaukee
UWMadison
UBuffalo-CCR
FNAL_CMS
UM_ATLAS
BU_ATLAS_Tier2
Purdue
UC_ATLAS_Tier2
BNL_ATLAS
IU_iuatlas
IU_ATLAS_Tier2
PSU_Grid3
ANL_Jazz
pdsf
JHopkins
HU_HUATLAS
Vanderbilt
Caltech-Grid3
Caltech-PG
UNM_HPC
OU_OSCER
OUHEP
UTA-DPCC
SMU_ATLAS
UCSanDiegoPG
UCSanDiego
Rice-Grid3
UFlorida-Grid3
UFlorida-PG
CHEPREO_FIU
KNU
South Korea
Open Science Grid
Grid3


LCG-2
NORDUGRID
Grid Solution for Wide Area Computing and Data Handling
30 sites
3200 cpus
25 Universities
4 National Labs
2800 CPUs
Total Sites 82
Total CPUs 7269
Total Storage (TB) 6558
Wed September 22 2004
Open Science Grid
Grid3
South Korea


LCG-2
NORDUGRID
Grid Solution for Wide Area Computing and Data Handling
FINLAND
SWEDEN
NORWAY
RUSSIA
ESTONIA
DENMARK
Trondheim
Umeå
Helsinki
St.Petersburg
Gjøvik
Bergen
Oslo
Uppsala
Tallinn
Tartu
Stockholm
Linköping
Göteborg
Ålborg
Lund
Odense
Copenhagen
30 sites
3200 cpus
-IHEP
BEIJING
TOKYO
GLASGOW
LANCS
EDINBURGH
LivHEP-LCG2
SHEFFIELD
MANHEP
CAM
BHAM
QMUL
CSTCDIE
RALPP
UCL-HEP
RAL
UCL-CCC
OX-PHYSICS
RHUL
IC
RWTH
NIKHE
SARA
NGS-RAL
NGS-OXFORD
NGS-MAN
NGS-LEEDS
fr-CGG
CERN
IN2P3-LPC
IN2P3-Lyon
MIL
TORINO
IFCA
USC
LIP
INT
UA
CIE
CN
Soon to come:
HP Labs Palo Alto and Bristol
HP supported sites in Singapore and China
Total Sites 82
Total CPUs 7269
Total Storage (TB) 6558
Wed September 22 2004
PK-NCP
(Pakista
Caltech-PG
UNM_HPC
OU_OSCER
OUHEP
UTA-DPCC
SMU_ATLAS
UCSanDiegoPG
UCSanDiego
Rice-Grid3
UFlorida-Grid3
UFlorida-PG
CHEPREO_FIU
KNU
South Korea
Open Science Grid
Grid3
Buffalo-CCR
BU_ATLAS_Tier2
BNL_ATLAS
SU_Grid3
JHopkins
HU_HUATLAS

CERN

openlab for DataGrid applications

In partnership with and sponsored by

IBM

intel

hp invent

ORACLE

enterasys

Networks that Know

CERN openlab

- IT Department's main R&D focus
- Framework for collaboration with industry
- Evaluation, integration, validation
  - of cutting-edge technologies that can serve LCG
- Initially a 3-year lifetime
  - As of 1.1.2003
  - Later: Annual prolongations
- Slogan: "You make it, we break it".

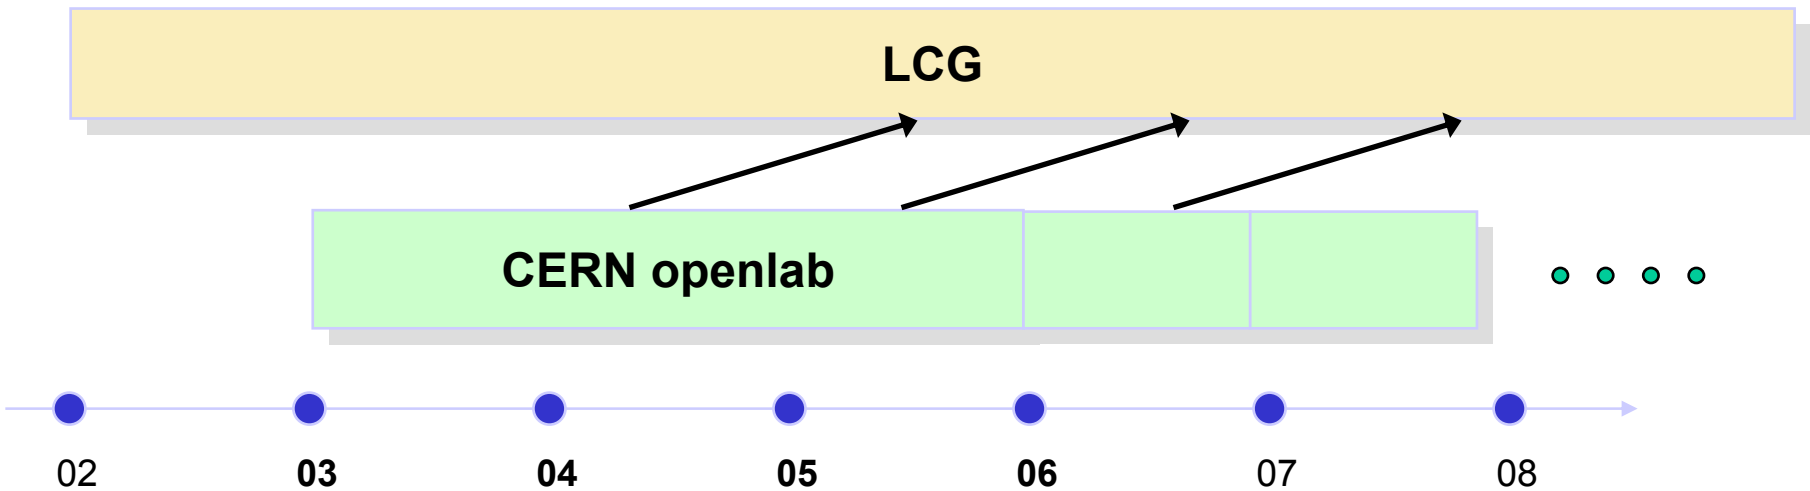

openlab participation

- **Five Partners (contributing ≥ 1.5 M€ over 3 years)**
  - Enterasys:
    - 10 GbE core routers
  - HP:
    - Integrity servers (103 * 2-ways, 2 * 4-ways)
    - Two post-doc positions
  - IBM:
    - Storage Tank file system (SAN FS), currently with 28 TB
  - Intel:
    - Large number of 64-bit Itanium processors & 10 Gbps NICs
    - 64-bit Nocona system w/PCI-Express
  - Oracle:
    - 10g Database software w/add-ons
    - Two post-doc positions
- **One contributor (contributing ≥ 170 k€ for 1 year)**
  - Voltaire
    - 96-way Infiniband switch and necessary HCAs

High Throughput Cluster Prototype openlab/LCG

- Experience with likely ingredients in LCG:
  - -- 64-bit programming
  - -- next generation I/O (10 Gb Ethernet, Infiniband, etc.)
- High performance cluster used for evaluations, and for data challenges with experiments
- Flexible configuration – components moved in and out of production environment
- Co-funded by industry and CERN

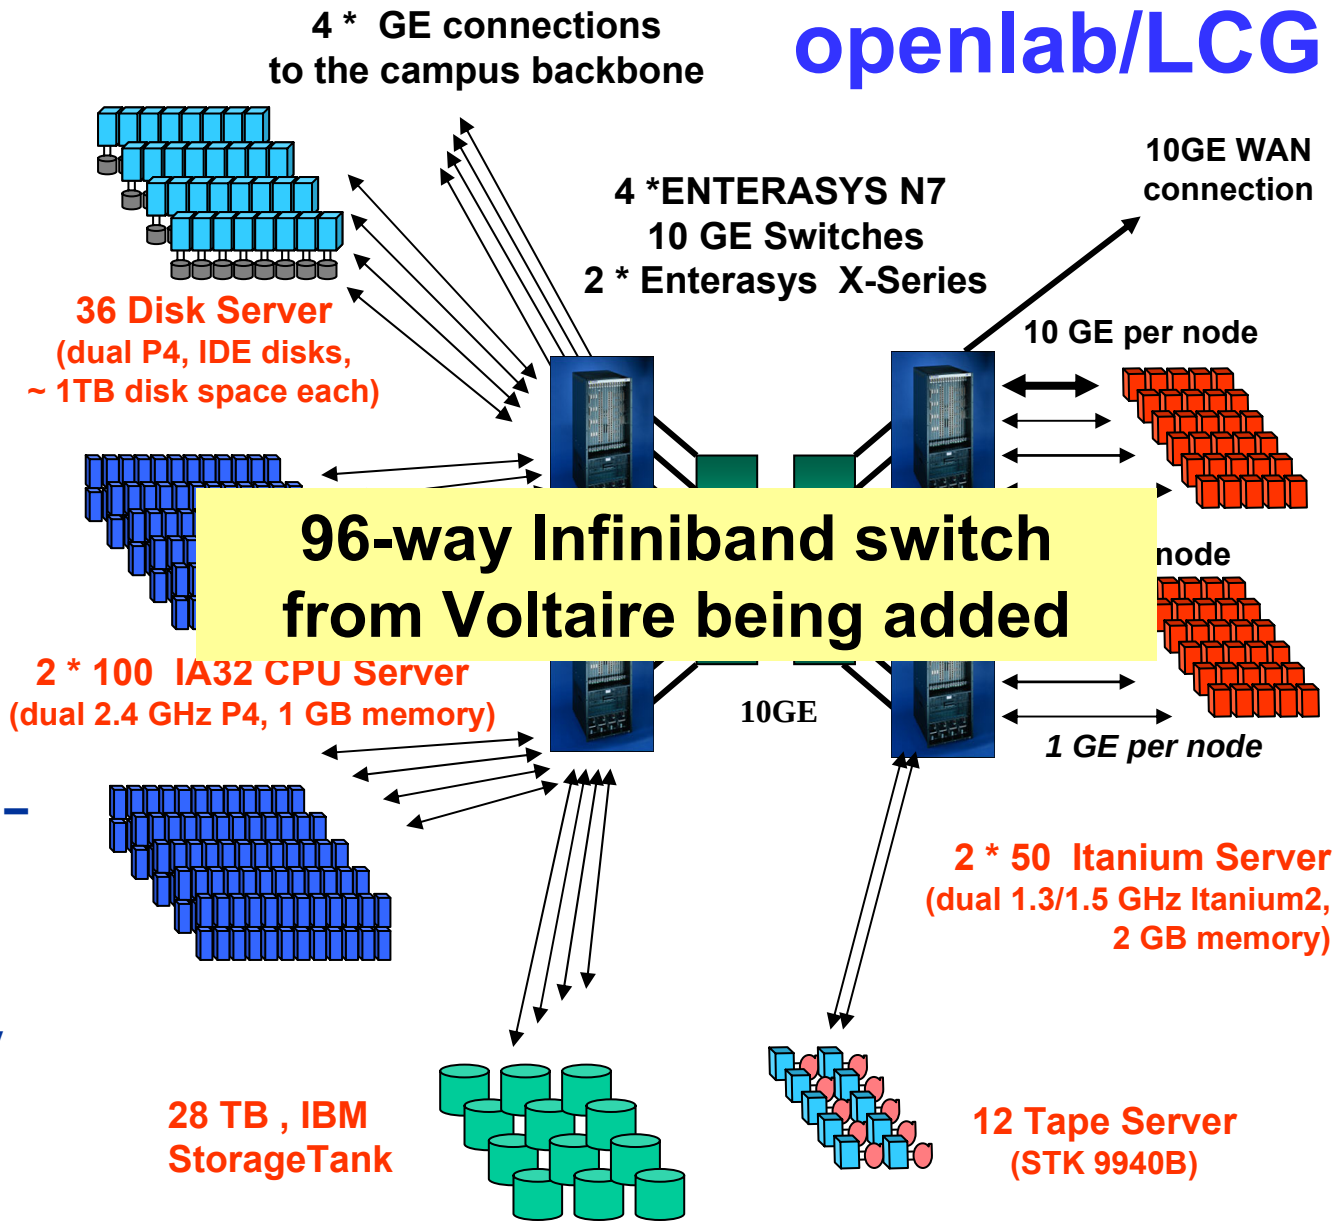

HP/Intel's opencluster CPUs

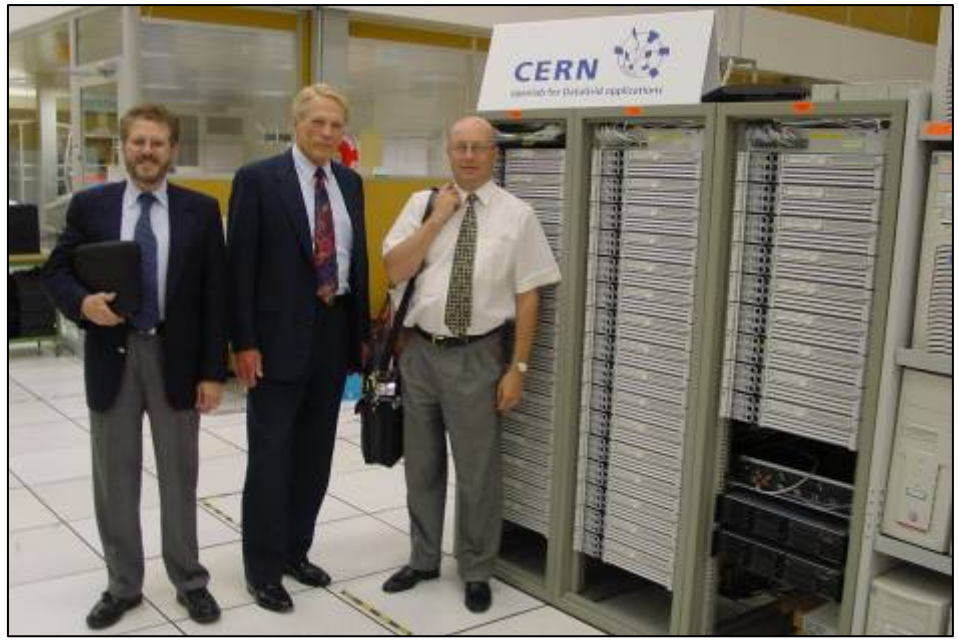

Use of Itanium systems

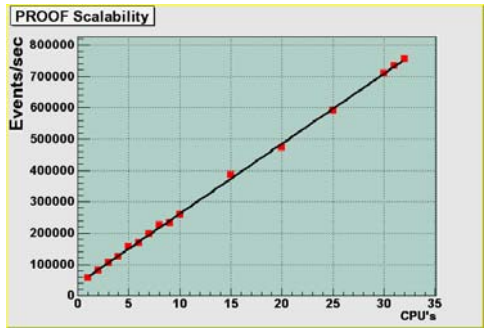

- **Why Itanium?**
  - Choice made in already in 2002, only solution available
  - Pure 64-bit approach forces “complete conversion” to new mode
    - Ported programs: ROOT, CLHEP, GEANT4, ALIROOT, LCG2, etc.
  - HP Itanium servers have excellent stability and I/O capabilities
  - We use standard “Scientific Linux CERN 3” (RedHat compatible)
    - Intel and GNU compilers
  - Very good performance monitoring tools, for both application and system performance
  - SPECint performance is adequate (~1300 SPECint)
  - Eagerly awaiting dual-core “Montecito” processors next year
- **When to switch to Itanium ?**
  - Price/performance break-even expected for mid 2007 (cost argument)
  - Whenever the addressing range requires it (≥ 4GB memory)
  - Get ready in time, certifying applications is a lot of work!

Use of Itanium systems

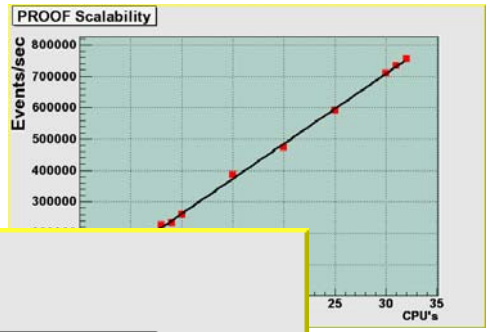

- W
- W

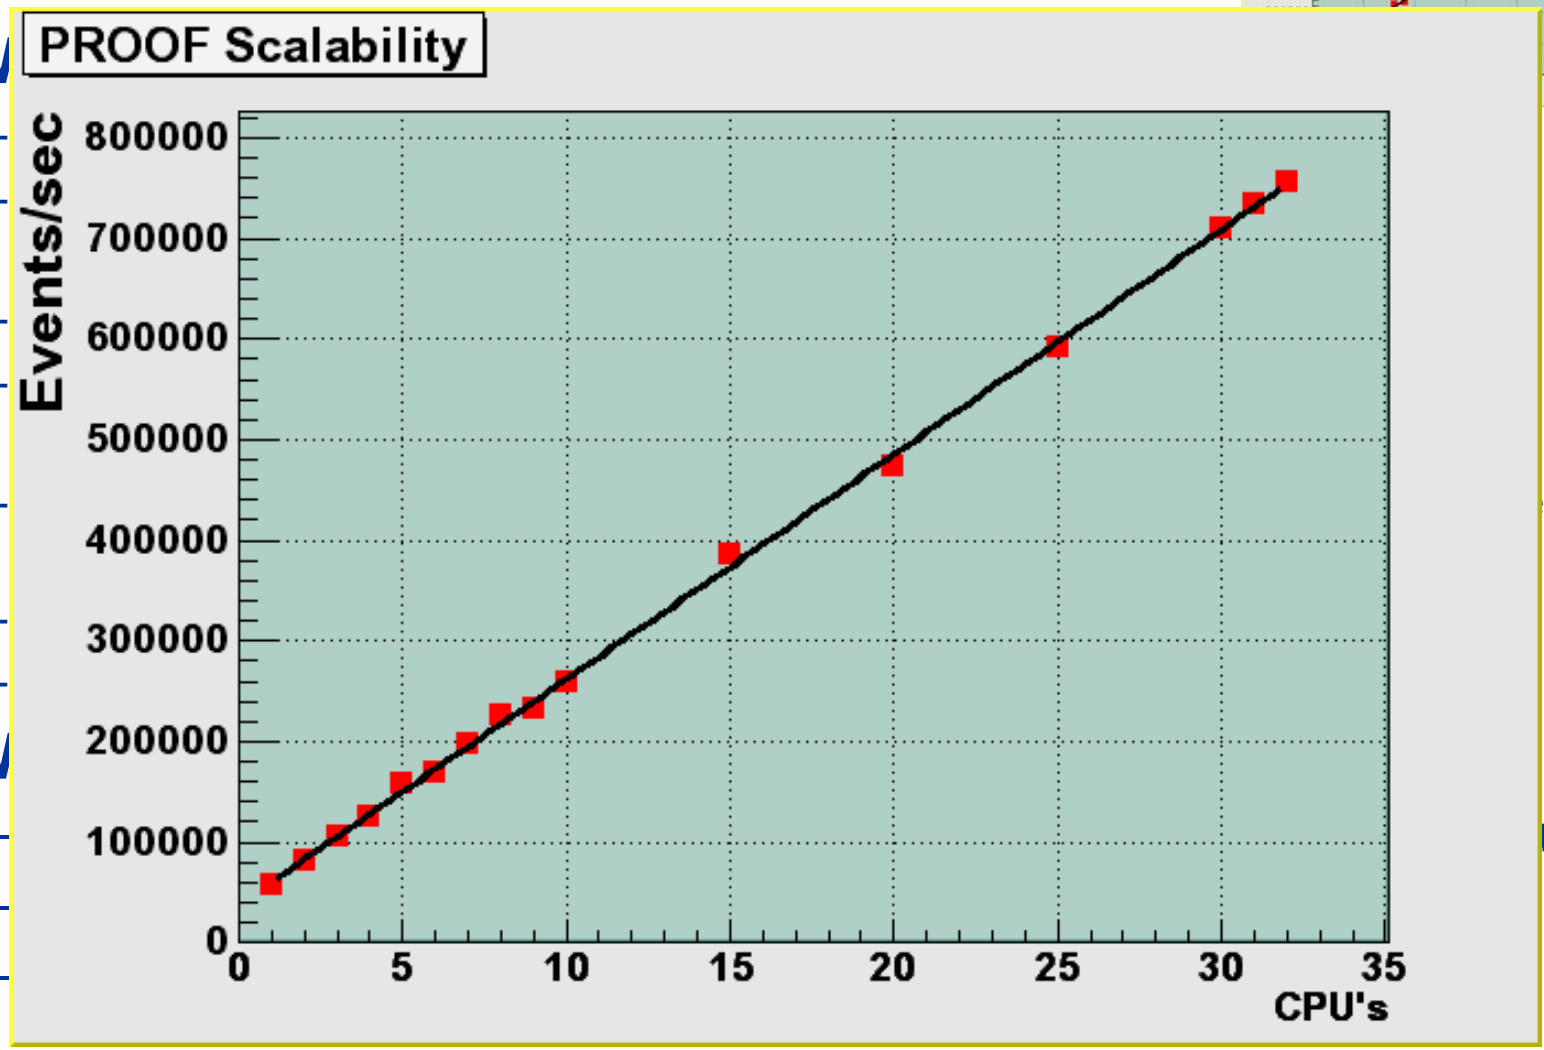

Gridification, a success story

- **Starting point: The software chosen for LCG had been developed only with IA32 (and specific Red Hat versions) in mind**
- **Two openlab members worked for many months to complete the porting of LCG-2 software to Itanium**
  - Result: All major components now work on Itanium/Linux:
    - Worker Nodes, Compute Elements, Storage Elements, User Interface, etc.
  - Code, available via Web-site, transferred to HP sites (initially Puerto Rico and Bristol)
  - Changes given back to software maintenance teams
- **Porting experience summarized in white paper**

A good step forward to a heterogeneous Grid !

10 Gbps WAN tests
(between CERN and California Institute of Technology)

- **Initial breakthrough during Telecom-2003**
  - with IPv4 (single/multiple) streams: 5.44 Gbps
    - Linux, Itanium-2 (RX 2600), Intel 10Gbps NIC
  - Also IPv6 (single/multiple) streams
- **In June 2004**
  - Again IPv4, and single stream (Datatag/Openlab):
  - 6.55 Gbps with Linux, Itanium-2 (RX4640), S2IO NIC
- **In September 2004:**
  - Same conditions as before:
  - 7.29 Gbps

But SuNET with a much longer lightpath has just grabbed the record, even if they only reach 4.3 Gbps. **We will be back!**

What do our partners say ?

- **"Through this collaboration with the CERN DataGrid, HP's researchers and engineers will be put to the test to truly push the envelope in developing advanced Grid computing technologies." Jim Duley, Director for Technology programs, HP University Relations**
- **"This is the perfect environment for us to enhance our Storage Tank Technology to meet the demanding requirements of large scale Grid computing systems." Jai Menon, IBM Fellow and co-director of IBM's Storage Systems Institute.**
- **"CERN's DataGrid project is an ideal application for Intel's most powerful processor yet, the Itanium. The awesome computer power required will find a formidable engine in the Itanium." Steve Chase, Director, Business and communication Solutions group of Intel.**
- **"The aggregate data throughput for LHC will exceed one terabit per second. Enterasys is confident that its 10-Gigabit Ethernet Technology will enable CERN to unlock the full potential of its DataGrid." John Roese, CTO of Enterasys Networks.**
- **"Leading-edge Grid technologies developed at CERN will be road-tested as part of its LHC project. As these technologies the come into the commercial mainstream, both we and our customers will benefit even further." Sergio Giacoletto, Executive VP, Oracle Systems Europe, Middle East and Africa.**

Change – what does it mean ?

- **Stan Williams, HP labs Palo Alto, during the *Computing in High Energy Physics Conference*, Interlaken, two weeks ago:**

  "During the 50 years of CERN's existence, computing performance has improved by 8 orders (100 million) of magnitude".

  "Today's computers are roughly a factor of one billion less efficient at doing their job than the laws of fundamental physics state that they could be."

- **Whatever we have seen so far, this is just the beginning. The development continues to accelerate.**
- **Change will be a constant fact of life, we better learn to manage it properly**
- **The highest cost factor is still the personnel cost**

During the last 20 years …

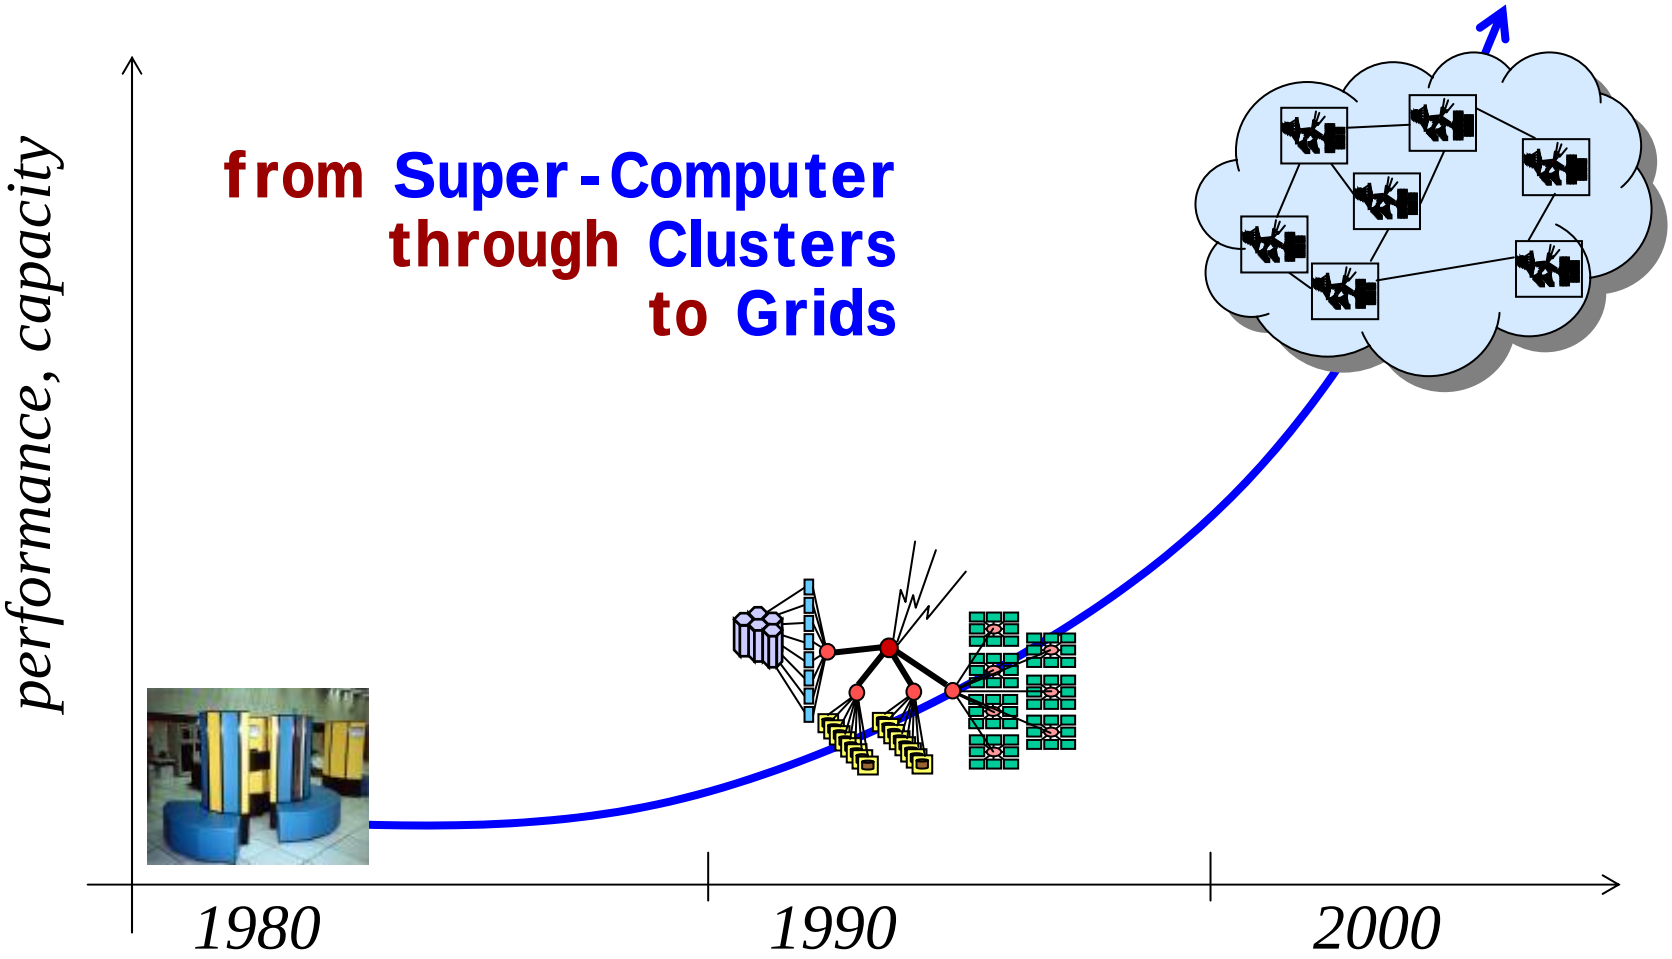

Conclusions

- High Energy physics, and LHC in particular, has enormous computing needs
- The amount of data (15PB/year) that LHC will generate is far beyond that of any previous project
- Science is still a driving force for advancing computing techniques, but science can't be successful without industry
- The CERN openlab is a good example of successful partnership.
- Grids are becoming a reality; it is time now to get ready for them
- Computing and communication have developed faster than any other technology over the past 50 years, by at least a factor of one billion !!
- There is no end in sight! Physics limits will allow for another factor of 100 Million during the next 20 years
- Change will happen with or without you, so you better get ready for it
- We expect at least four major changes during the LHC lifetime

Thank you for your attention