CERN openlab II

An Overview of CERN's Approach to Energy Efficient Computing

Sverre Jarp
CERN openlab CTO
25 June 2008

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
    - Accelerators
    - Particle detectors
    - Powerful computers

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.

The CERN Site

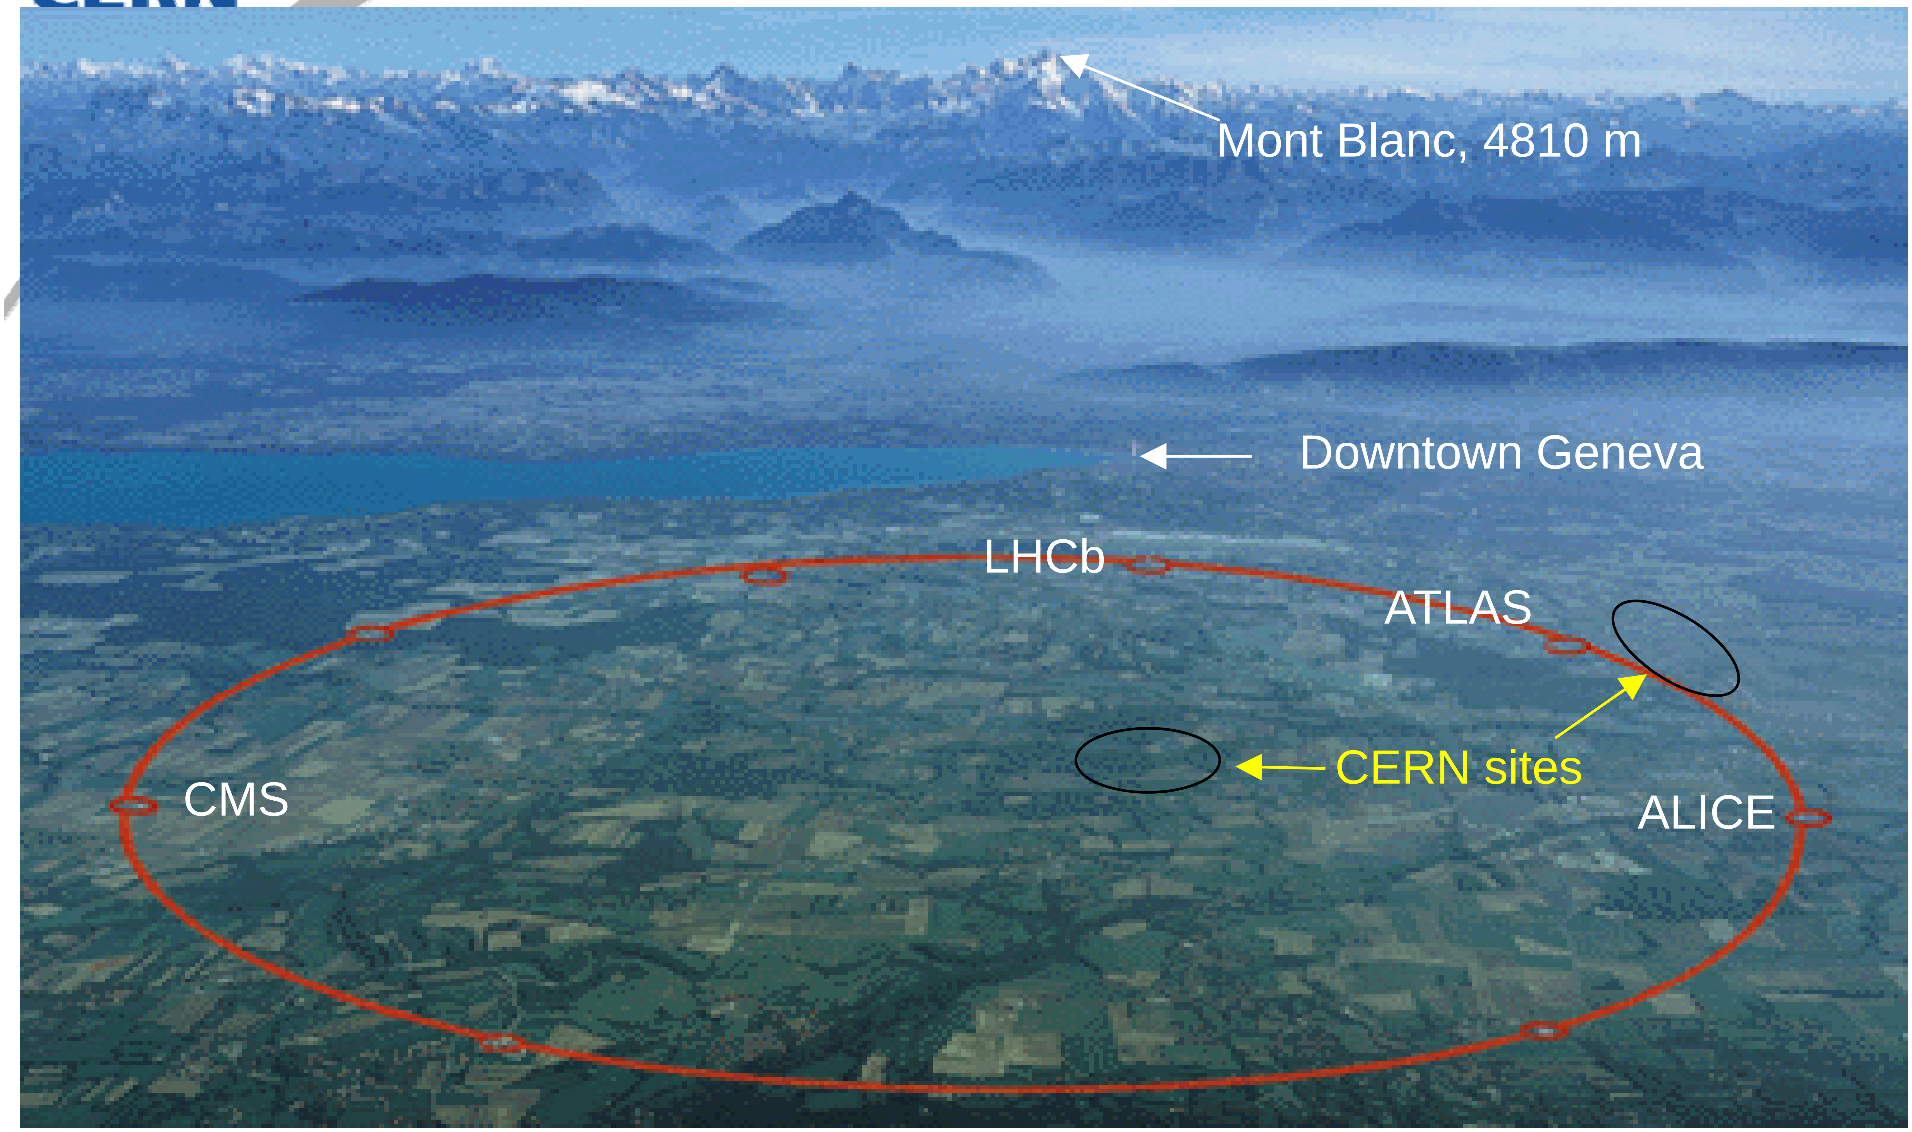

What is LHC?

- The Large Hadron Collider will collide beams of protons at an energy of 14 TeV (<u>in the summer of 2008</u>)
- Using the latest super-conducting technologies, it will operate at about – 271°C, just above the temperature of absolute zero.
- With its 27 km circumference, the accelerator is the largest superconducting installation in the world.

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

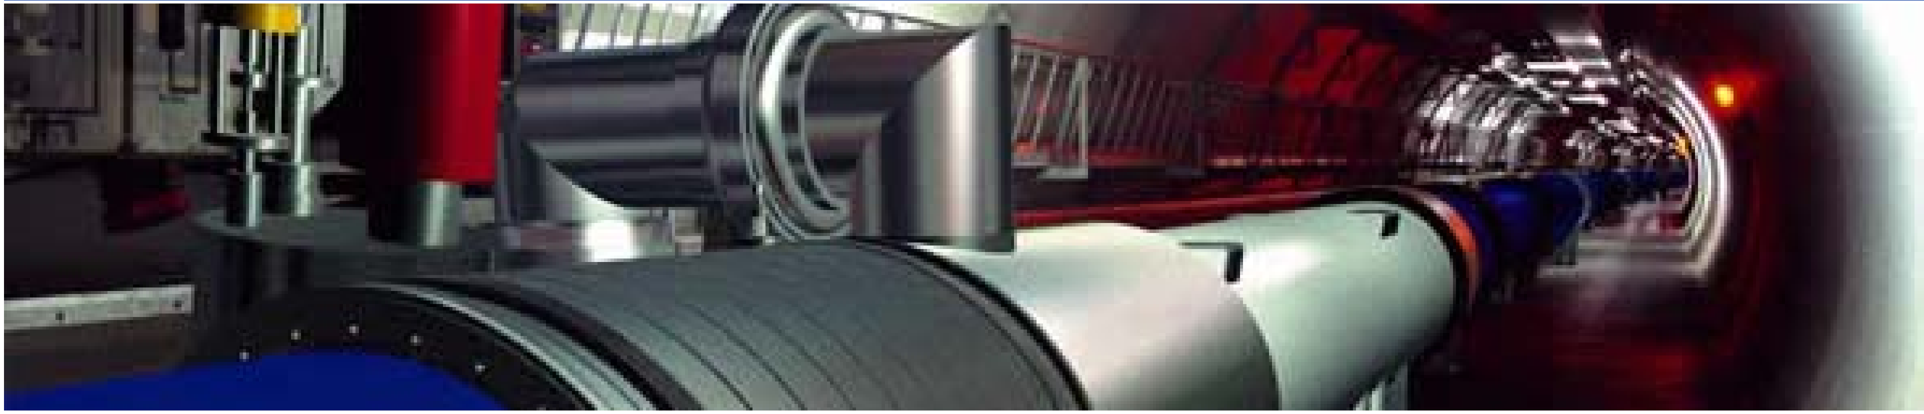

Compact Muon Solenoid
(CMS)

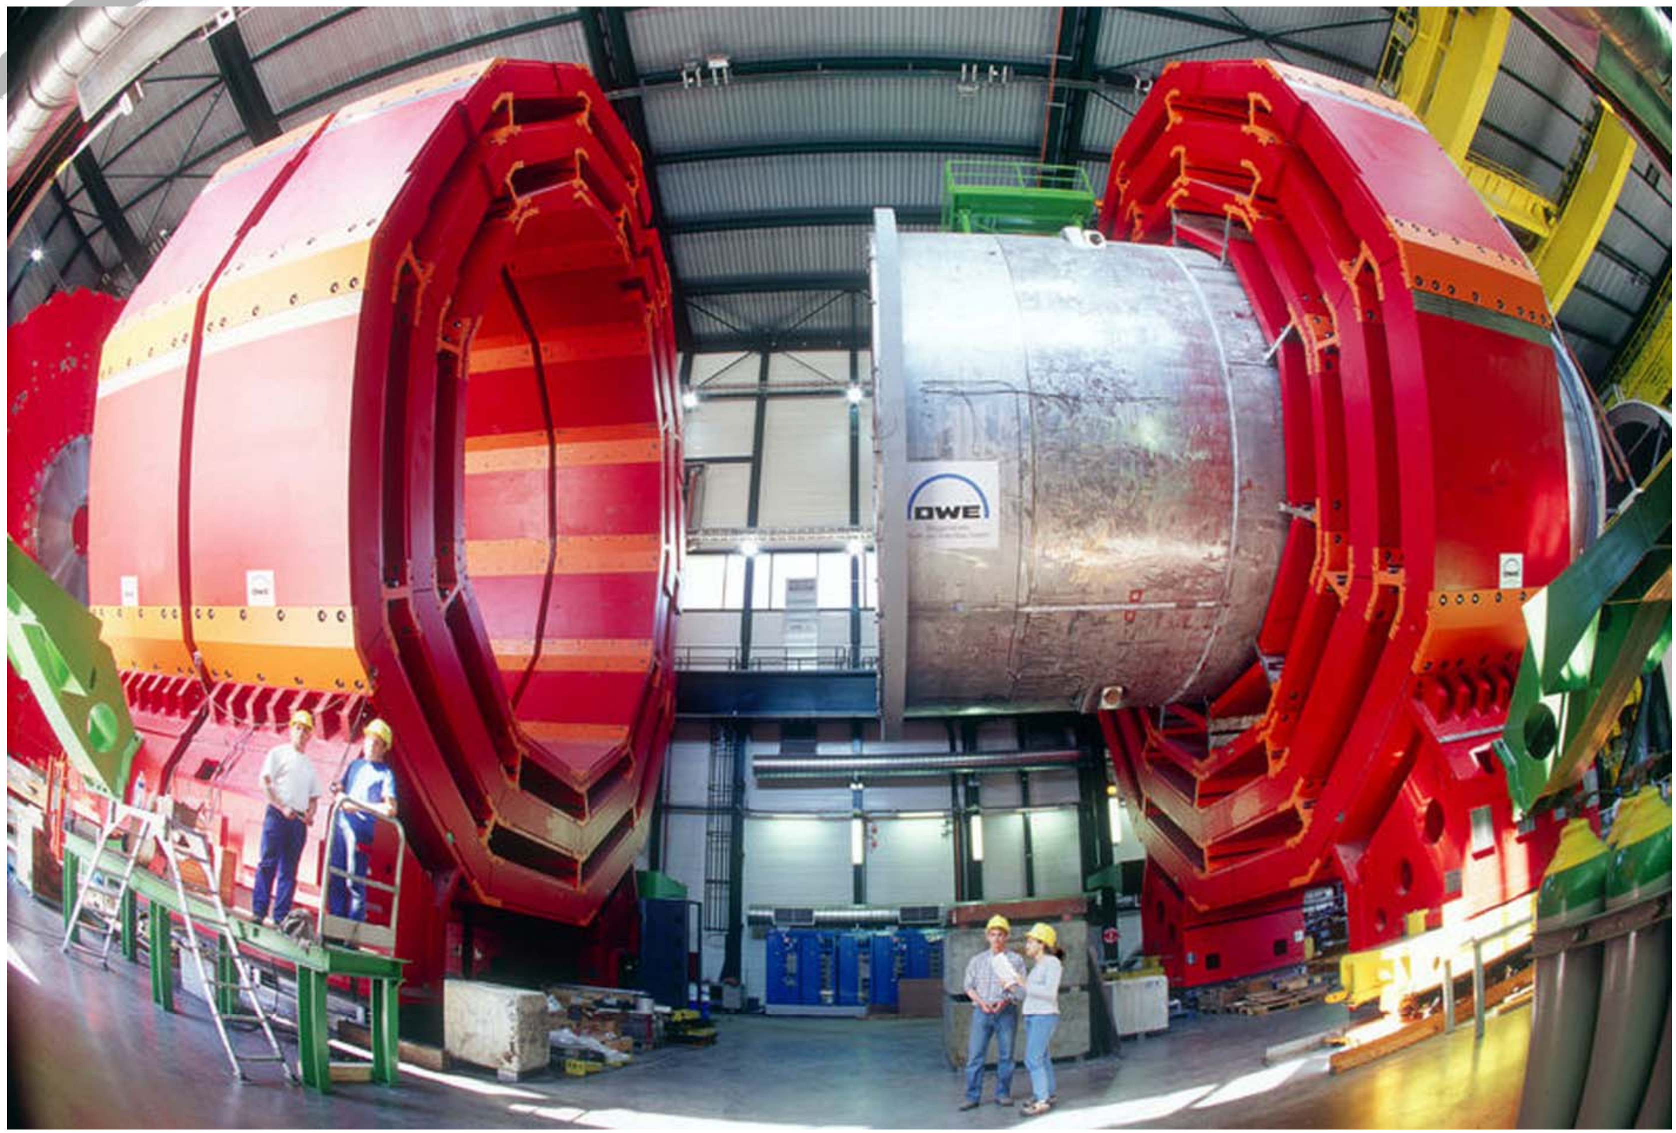

LHC data (simplified)

Per experiment:

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

CMS

LHCb

ATLAS

ALICE

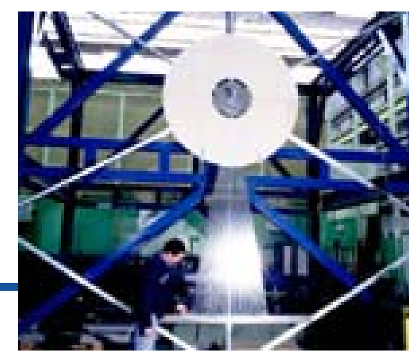

Computing at CERN today

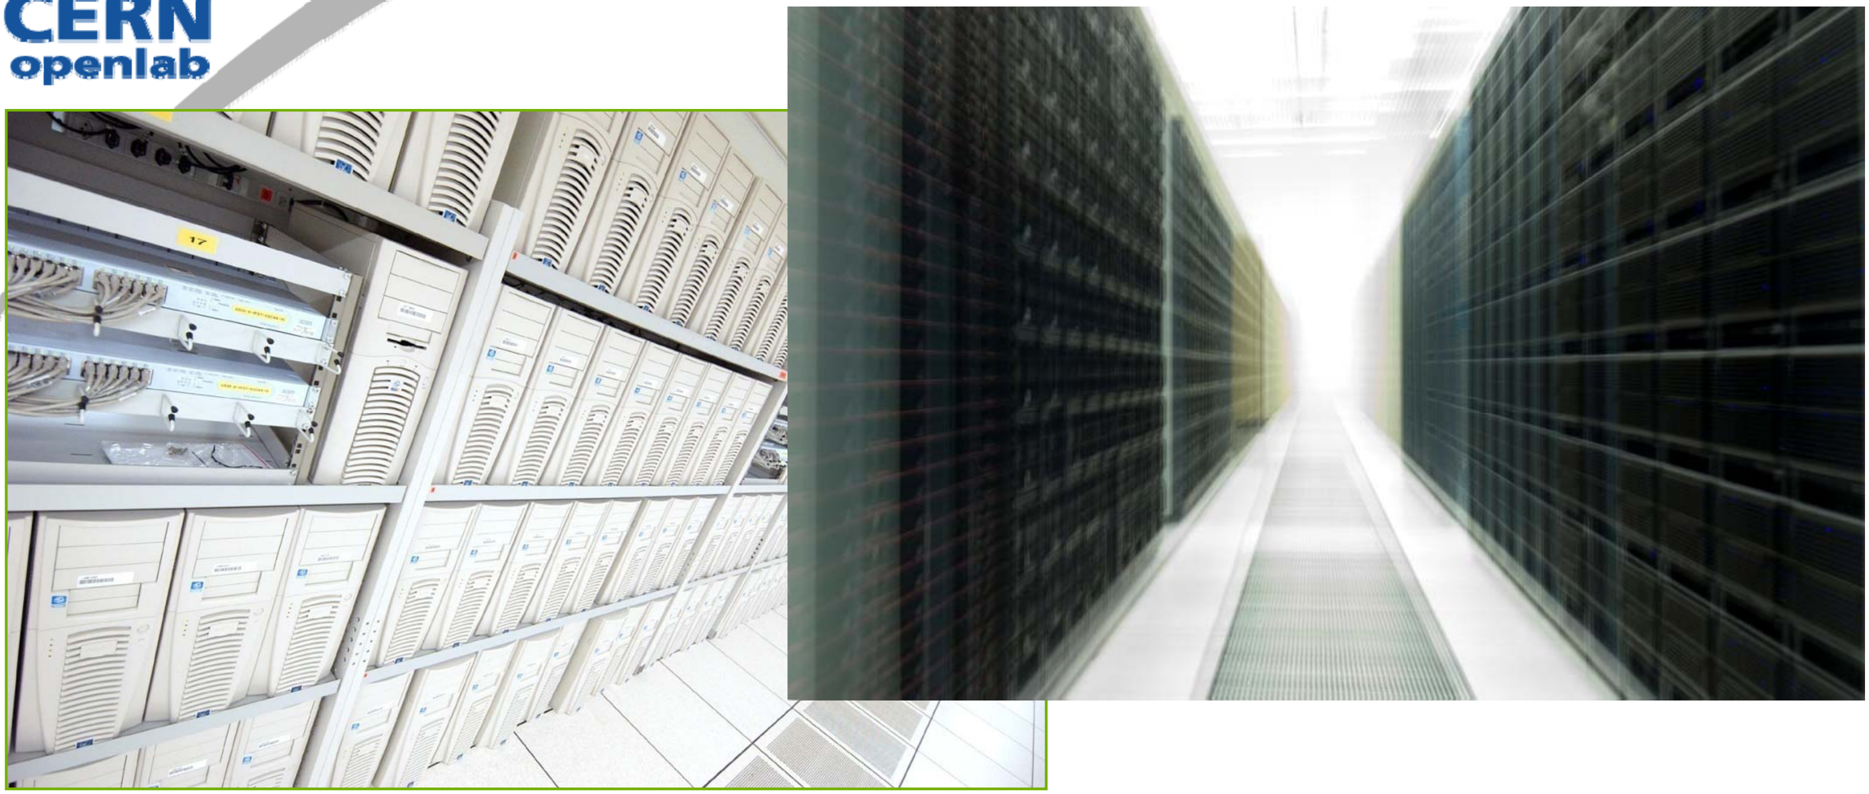

- High-throughput computing based on reliable "commodity" technology
- About 3000 dual-socket multi-core PC servers running Linux
- More than 5 Petabytes of data on tape; 20% cached on disk

Expected Power Evolution

Demand will grow continuously through the LHC era

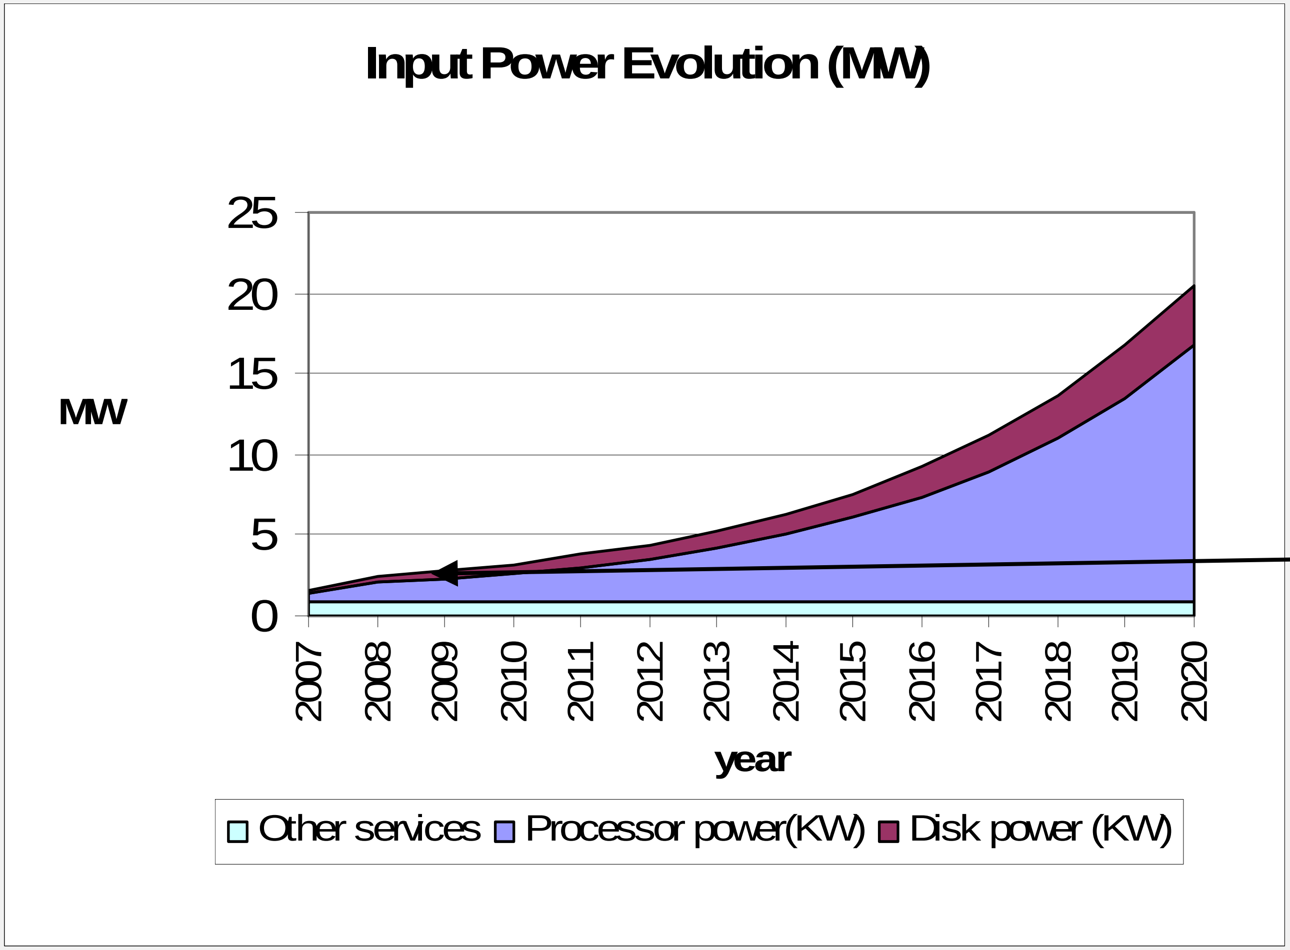

Power limit of present Computer Centre!

How have we reacted?

- CERN and the High Energy Community have reacted quickly to cope with the issues of very high computing demands inside a power-constrained environment:
  - Established a Worldwide Computing Grid
  - → Made the current Computer Centre energy-efficient
  - Started planning new Computer Centres

LHC Computing Grid

- Largest Grid service in the world !

• **Almost 150 sites in 35 countries**

• **100’00 IA processor cores (w/Linux)**

• **Tens of petabytes of storage**

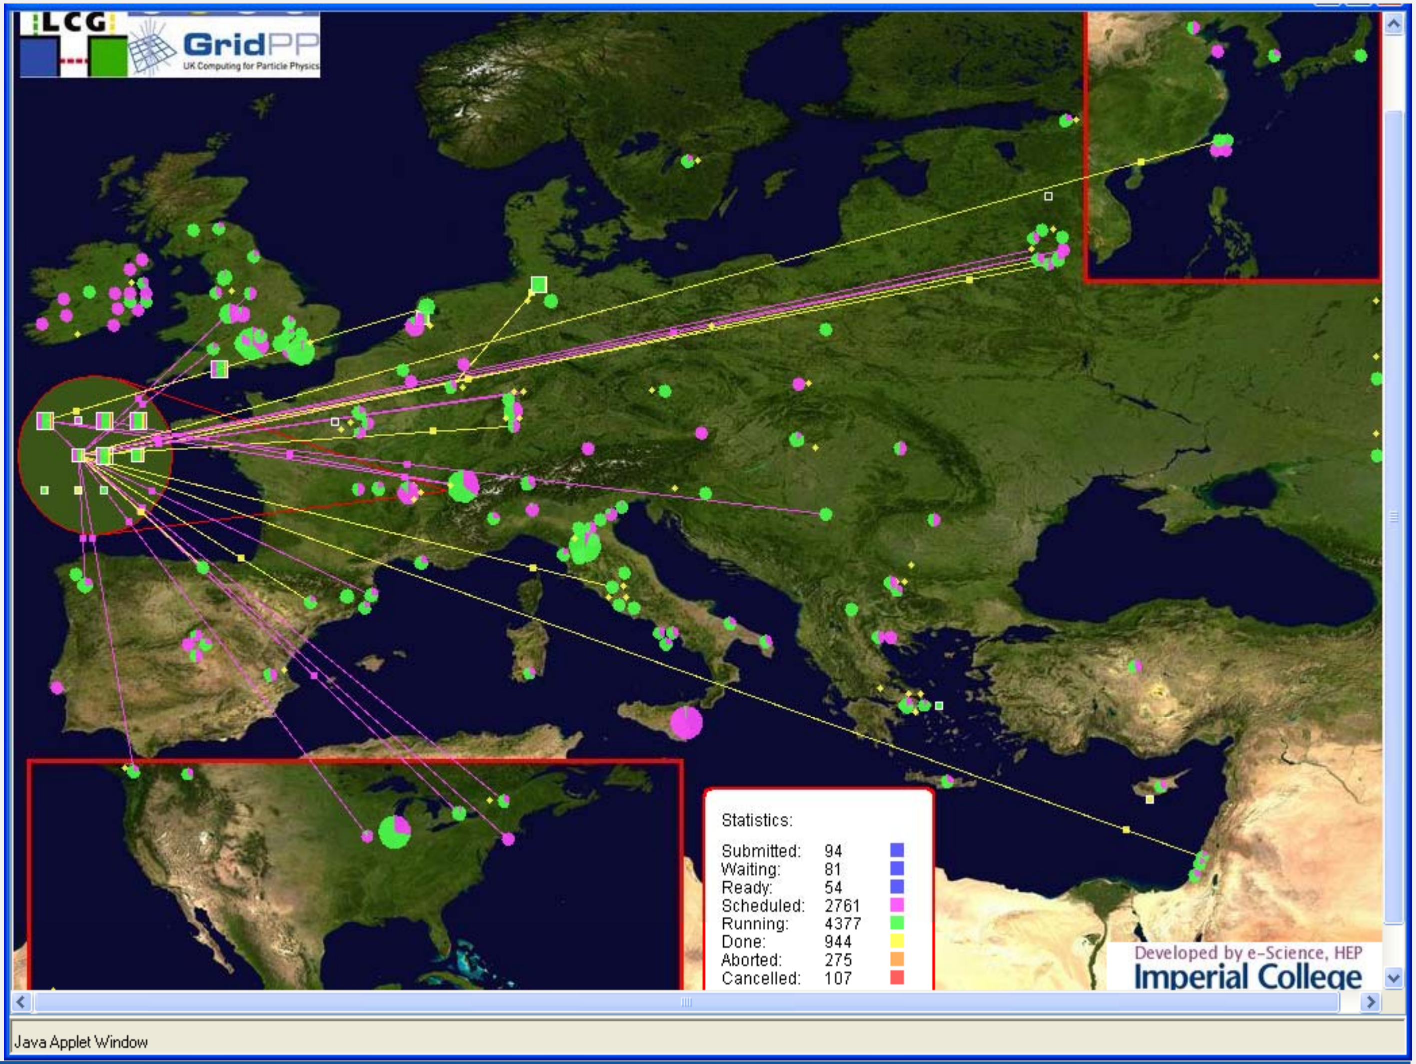

CERN's Power Saving Strategies (1)

- Make power efficiency part of the tendering process:
  - Use representative workloads
    - Including idle time
  - Measure power in the primary AC circuit
  - Compare performance/watt for different servers
  - Include power and cooling costs when purchasing new systems

CERN's Power Saving Strategies (2)

- New server configurations:
  - Consider multi-core processors, whenever possible
    - But watch out for additional memory or I/O requirements

  - For throughput computing, consider processors with below-peak frequency
  - Select high-quality, efficient power supplies
    - A high power factor at both the high and the low range
  - Consider blade servers
    - Typically, they come with more efficient power supplies and fans

CERN's Power Saving Strategies (3)

- Optimize the Data Centre for power efficiency
  - Reduce mixing of hot and cold air
    - Align racks to create hot and cold aisles
  - Consider sealing cold aisles
  - Consider extracting heat directly from hot aisles

CERN's Power Saving Strategies (4)

- Evaluate software technologies
  - Use virtualization when server workload is low
  - Improve performance of the applications
    - Better compilers (for example, the Intel C++ compiler)
  - Use software that can put idle processors to sleep
  - Evaluate multi-threading
    - Be aware that software complexity may increase
  - Evaluate grid computing
    - Spread the total workload across multiple facilities

Conclusions

- Power and cooling issues are now part of daily life in computing!
- Probably the most important items:
  - Understand how power affects the total cost of ownership (TCO)
  - Acquire appropriate server technology
    - Recently, CERN acquired 2.33 GHz quad-core Harpertown Blade & 1U servers
  - Make Computer Centres cooling-efficient
  - Exploit relevant software and virtualization options
- Foster an "energy efficiency" culture in your company!

BACKUP

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 scales with SPECint (not SPECfp)
  - But good floating-point (30% of total) 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

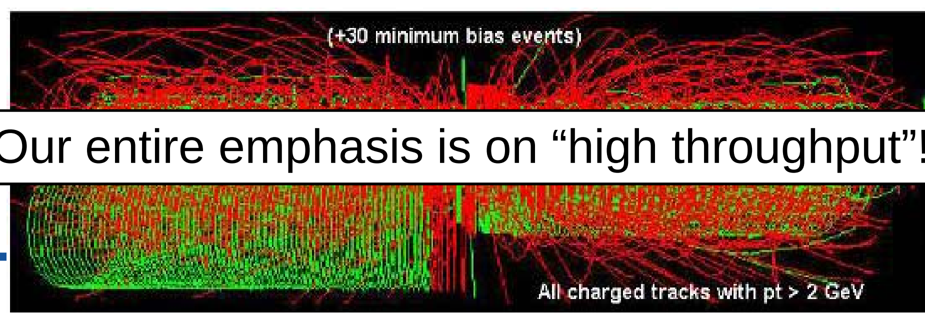

Our entire emphasis is on “high throughput”!

Data Handling and Computation for Physics Analysis

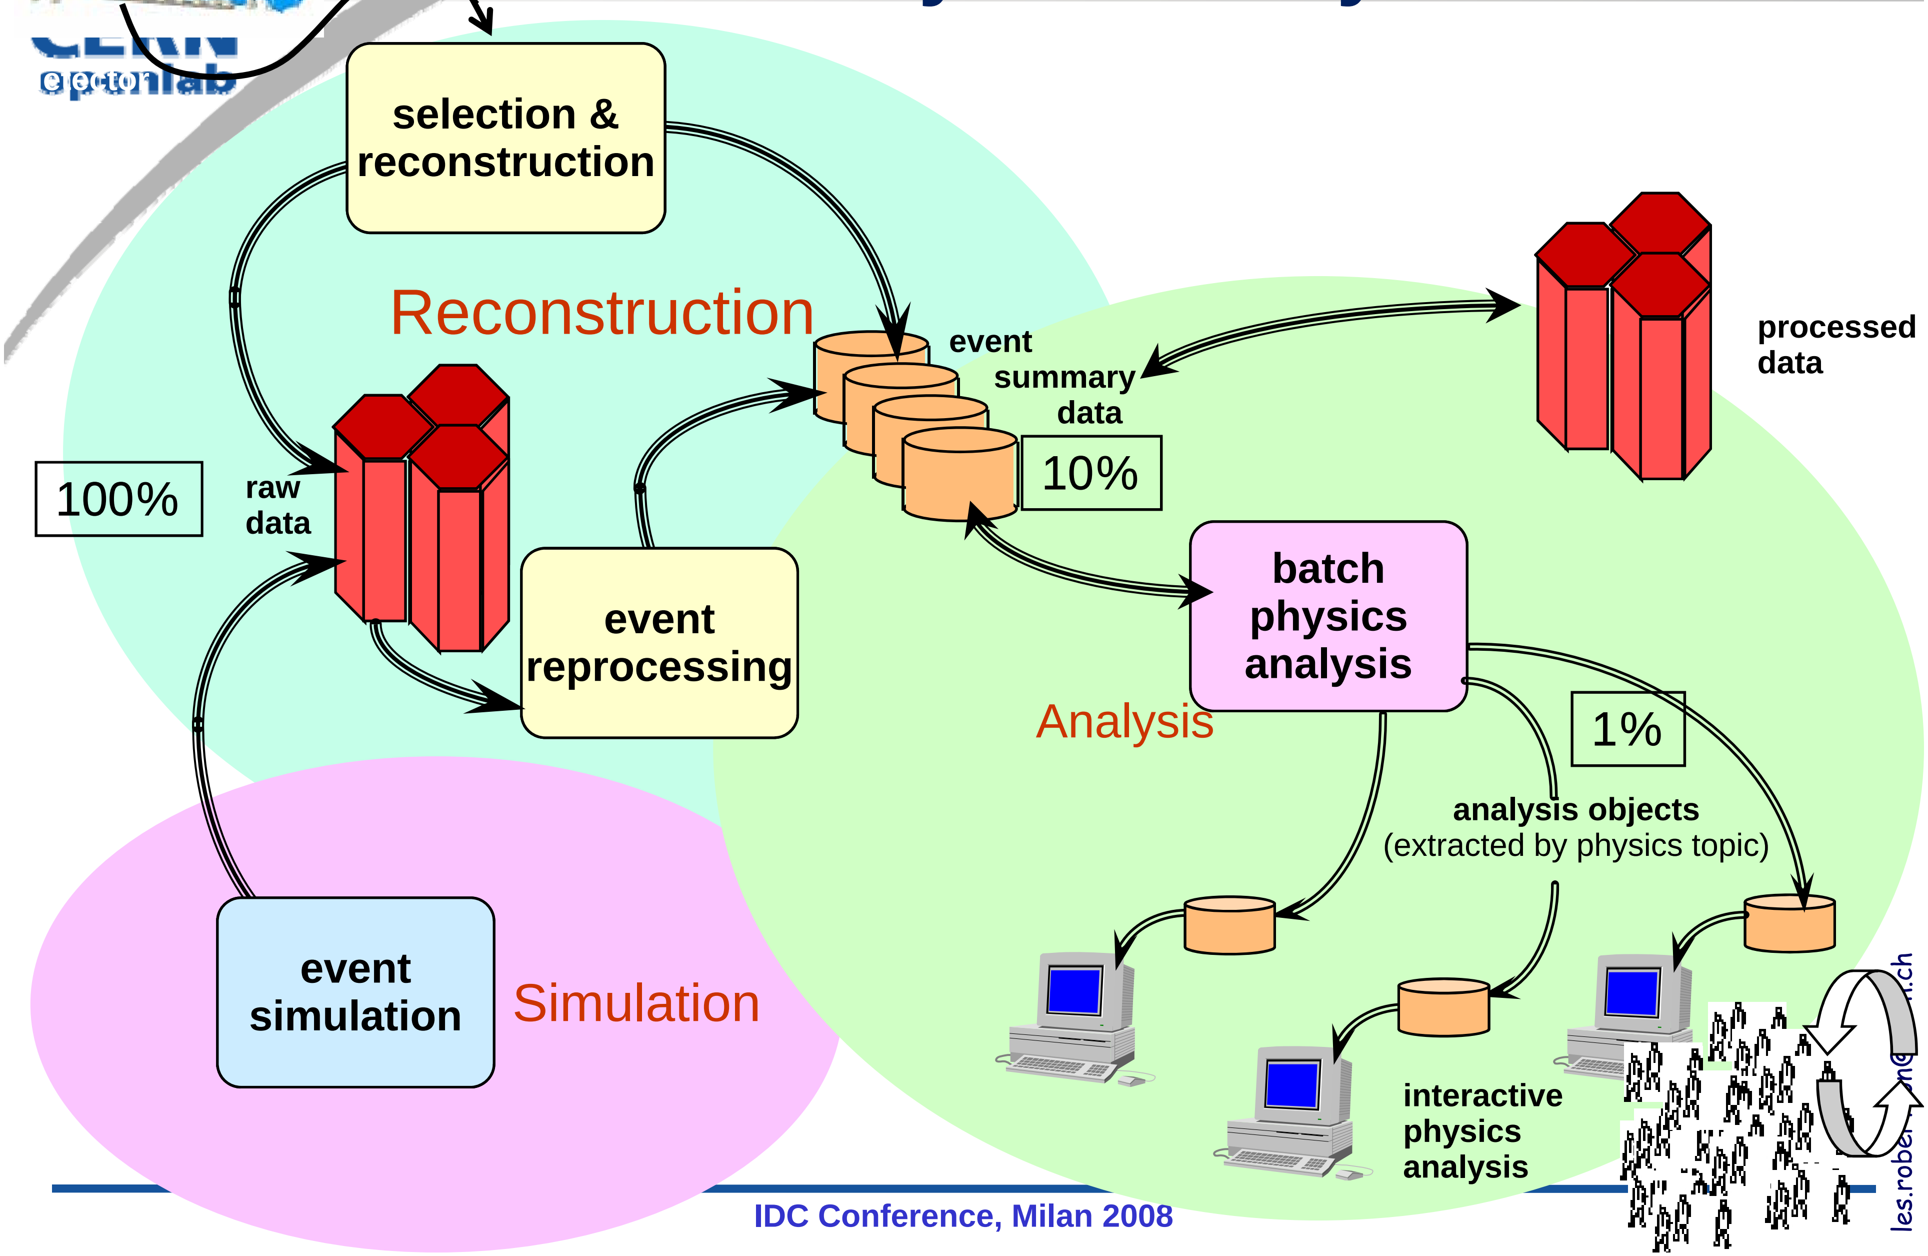

SHIFT architecture

(Scalable Heterogeneous Integrated Facility)

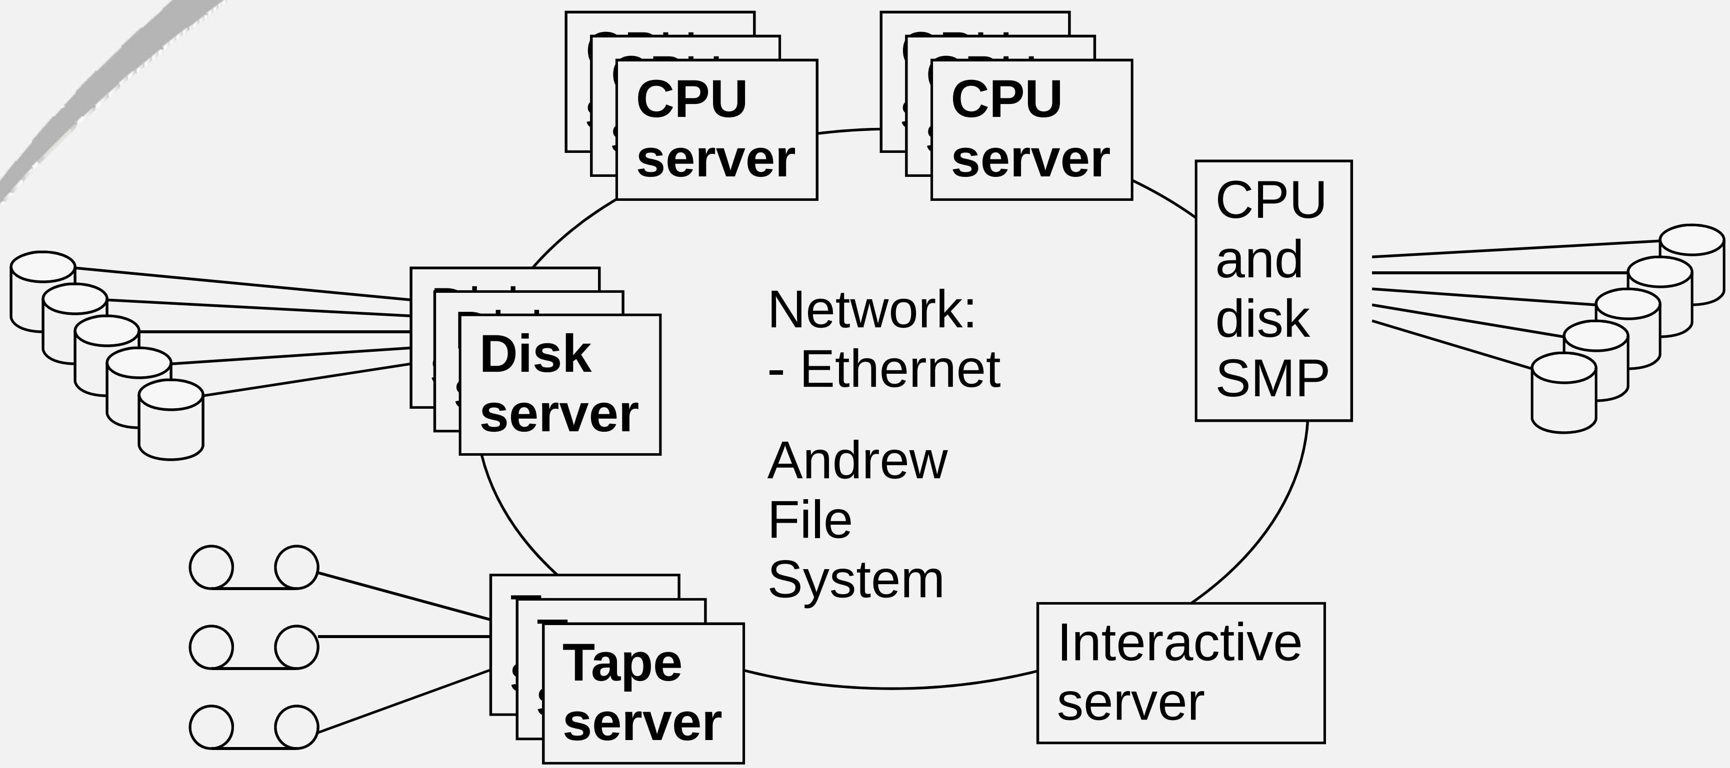

In 2001 SHIFT won the 21st Century Achievement Award issued by Computerworld