CERN openlab II

CERN openlab and Intel: Today and Tomorrow

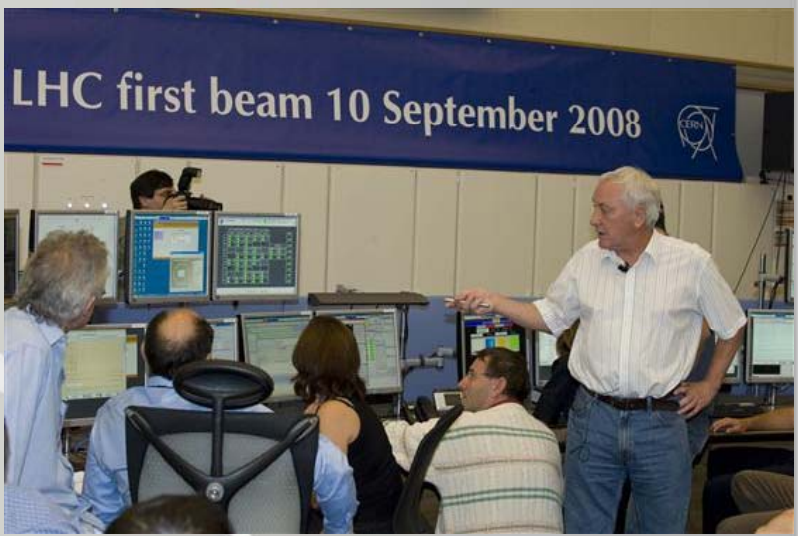

Sverre Jarp
CERN openlab CTO
16 September 2008

Overview of CERN

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

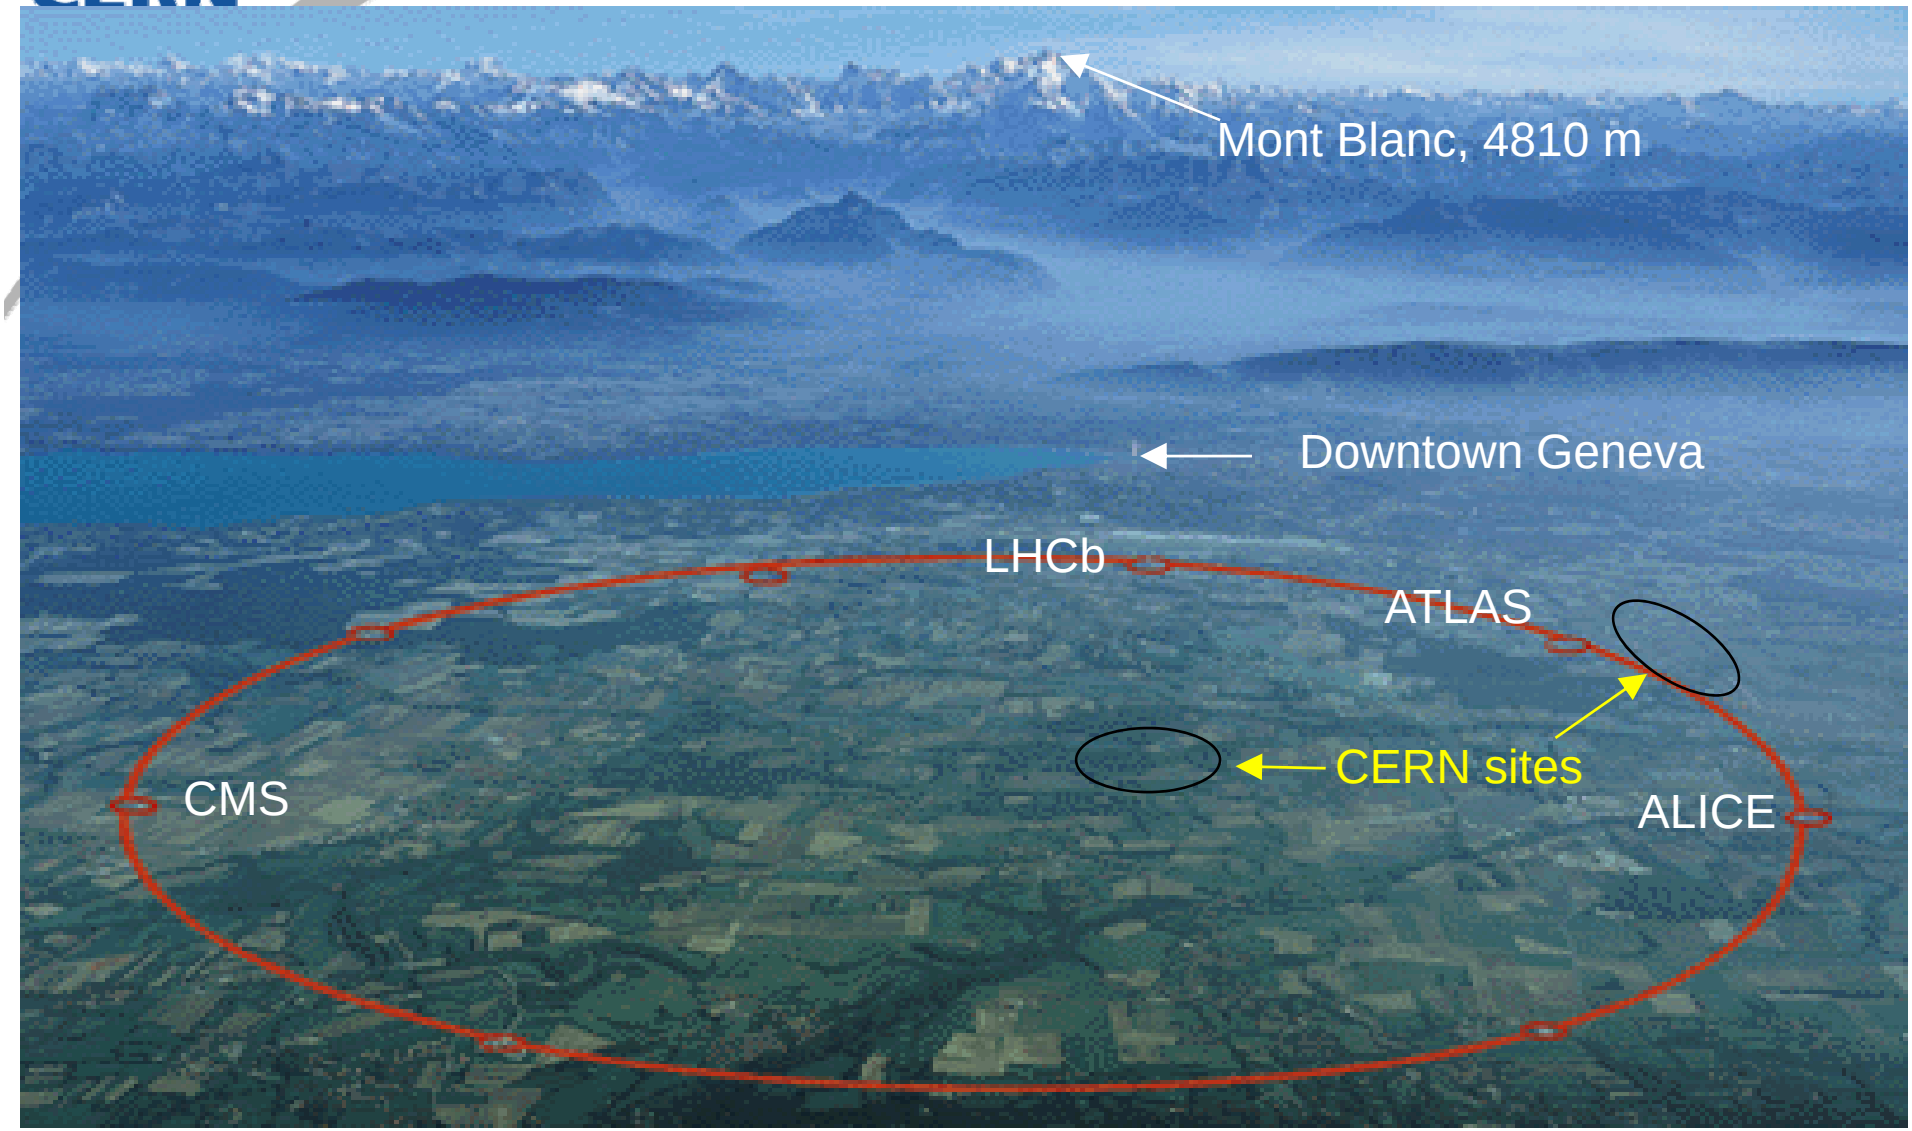

CERN underground

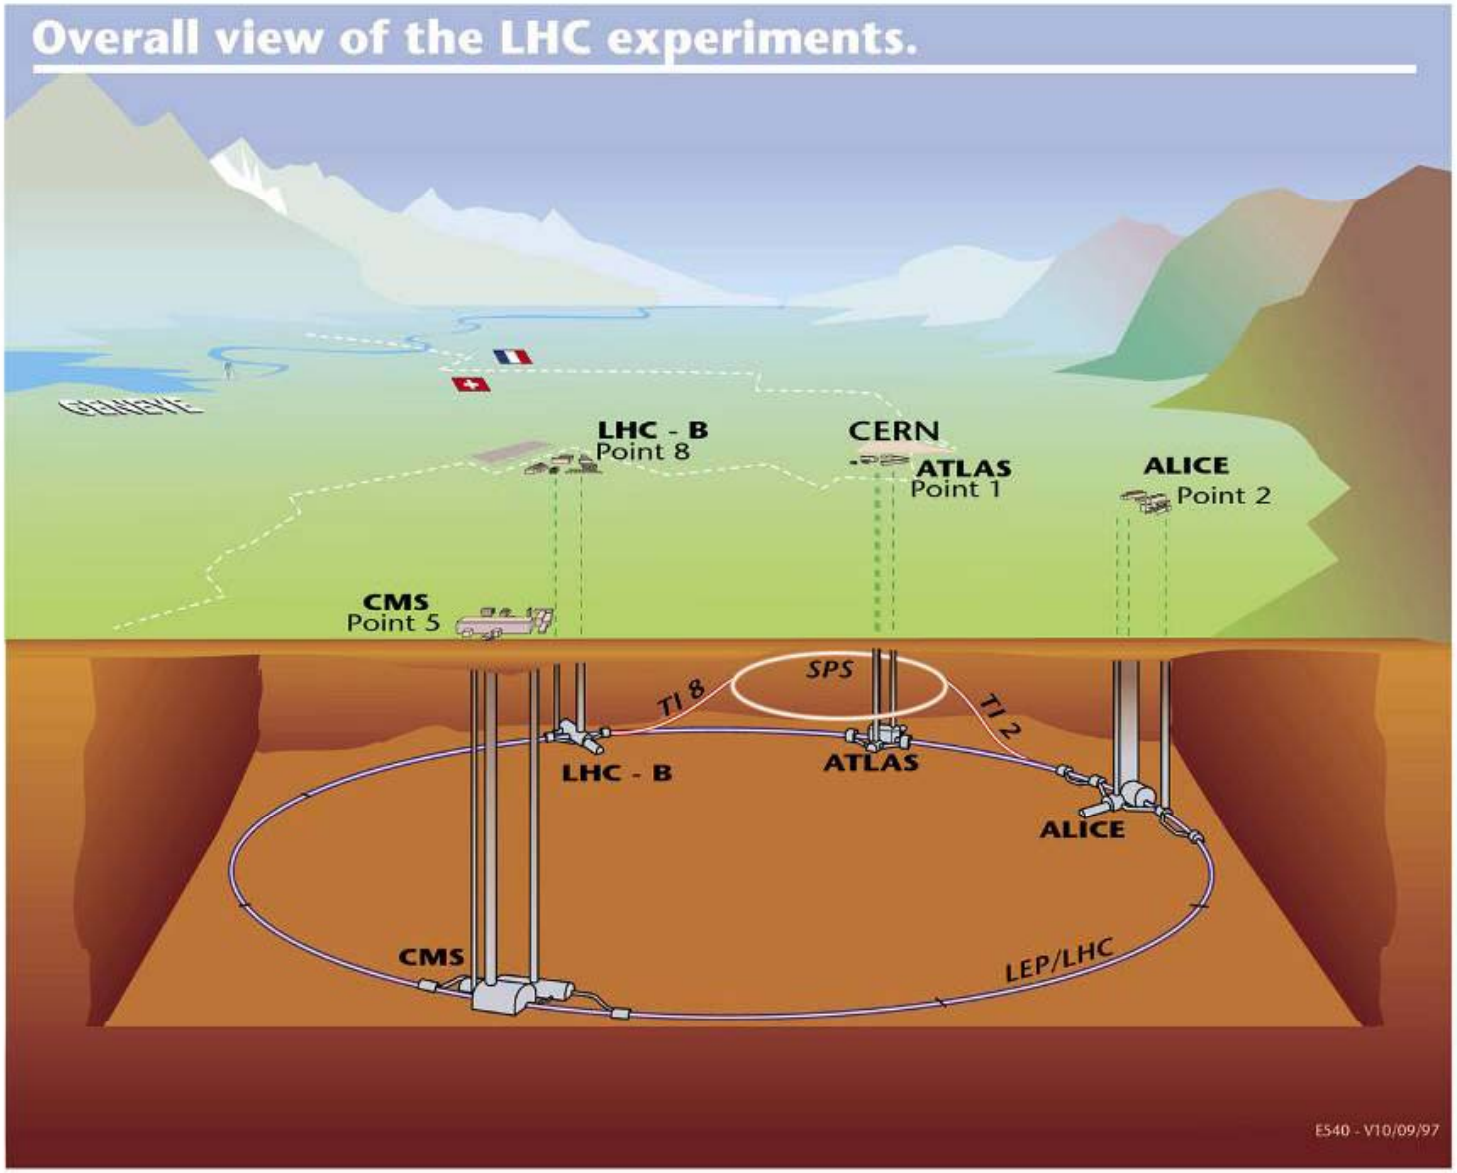

What is LHC?

- The Large Hadron Collider can collide beams of protons at an energy of 14 TeV (inaugurated on 10 September!)
- Using the latest super-conducting technologies, it operates 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

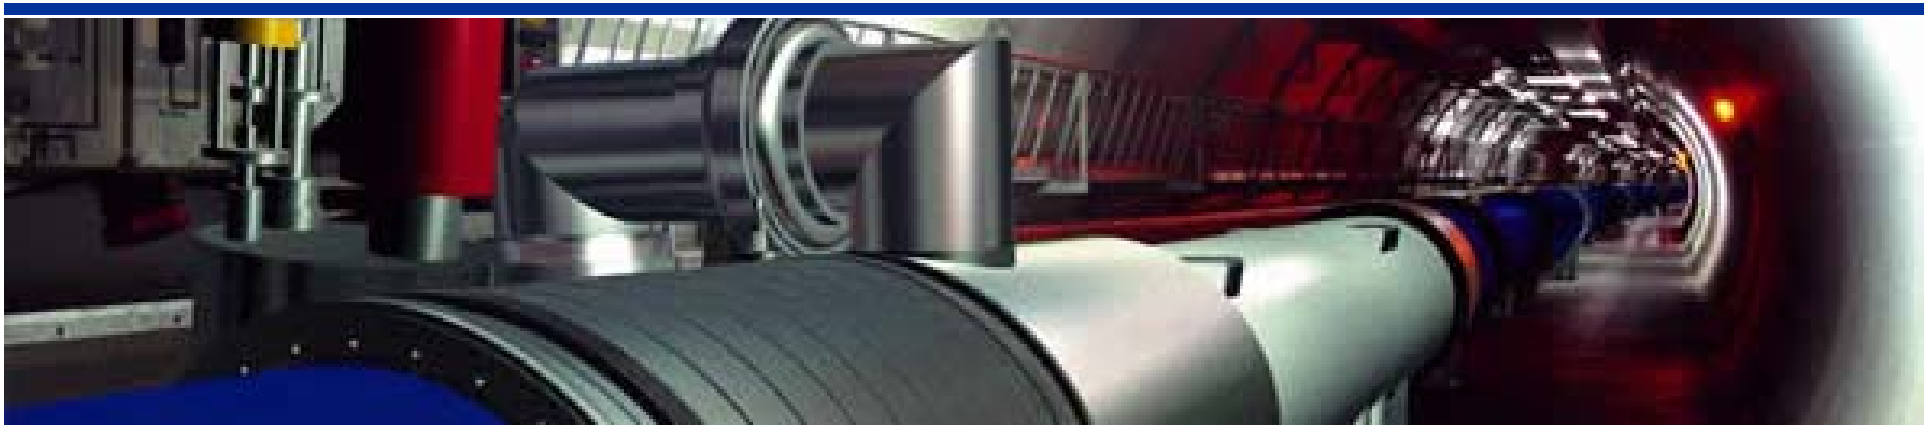

ATLAS

- General purpose LHC detector – 7000 tons

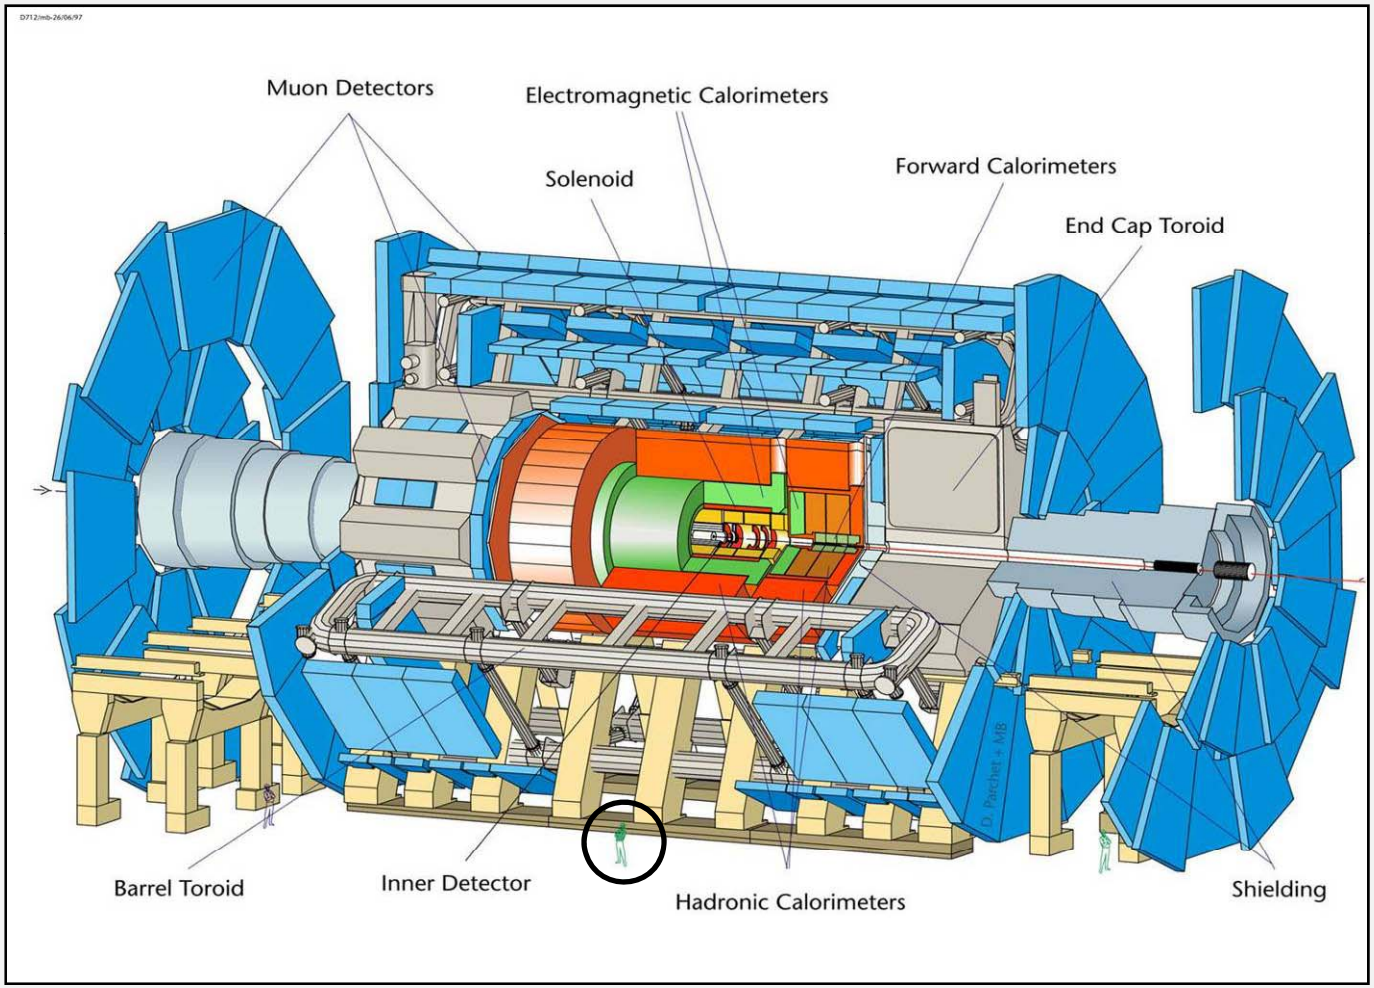

Particle Physics - 101

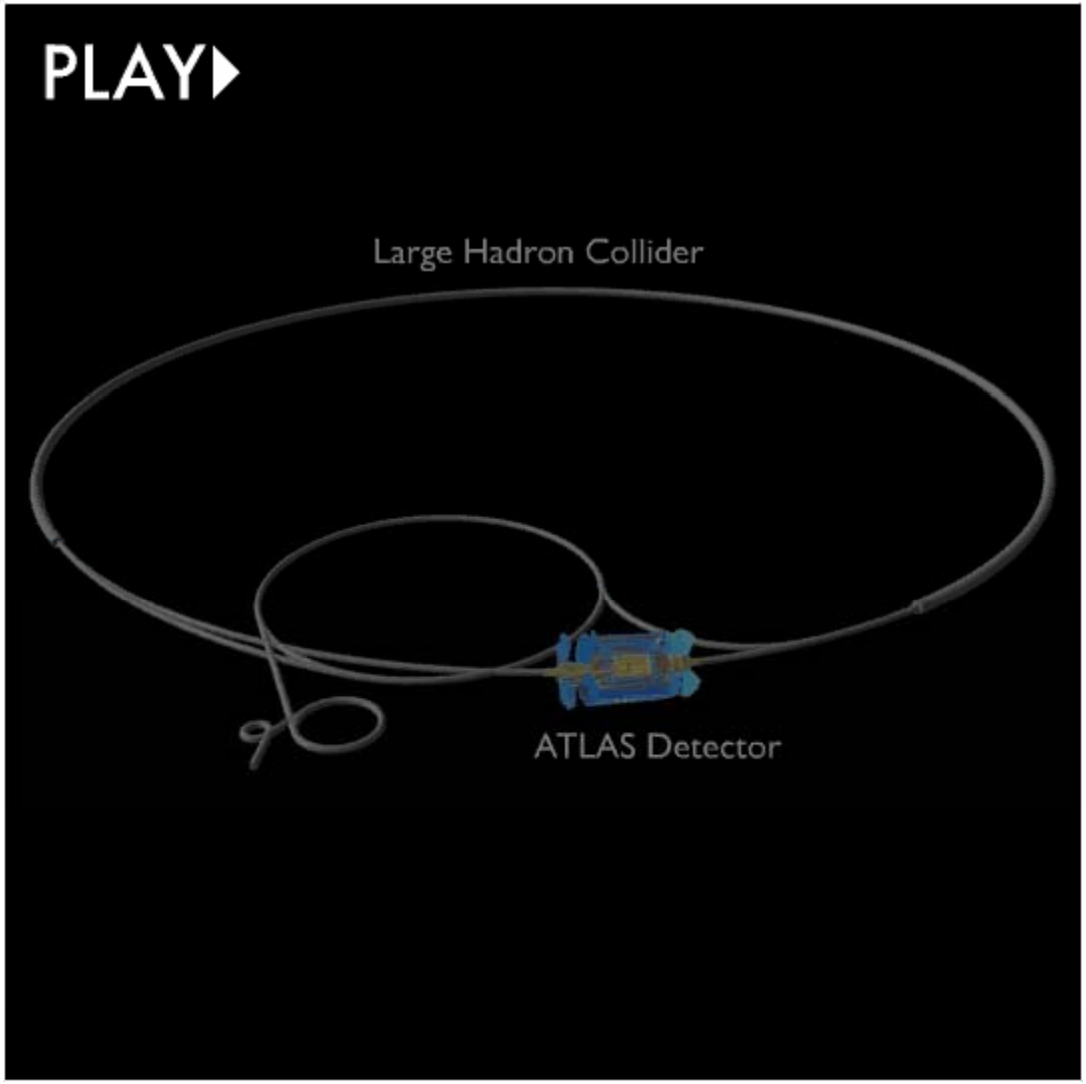

CMS

LHCb

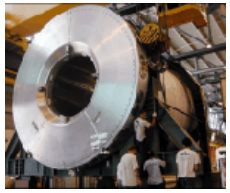

ATLAS

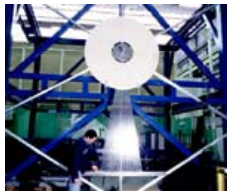

ALICE

Data management and computing

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

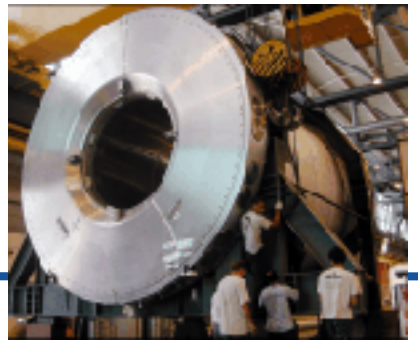

ALICE

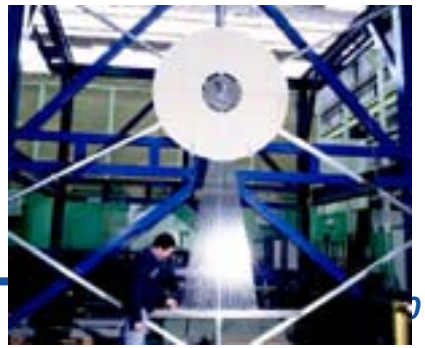

Computing at CERN today

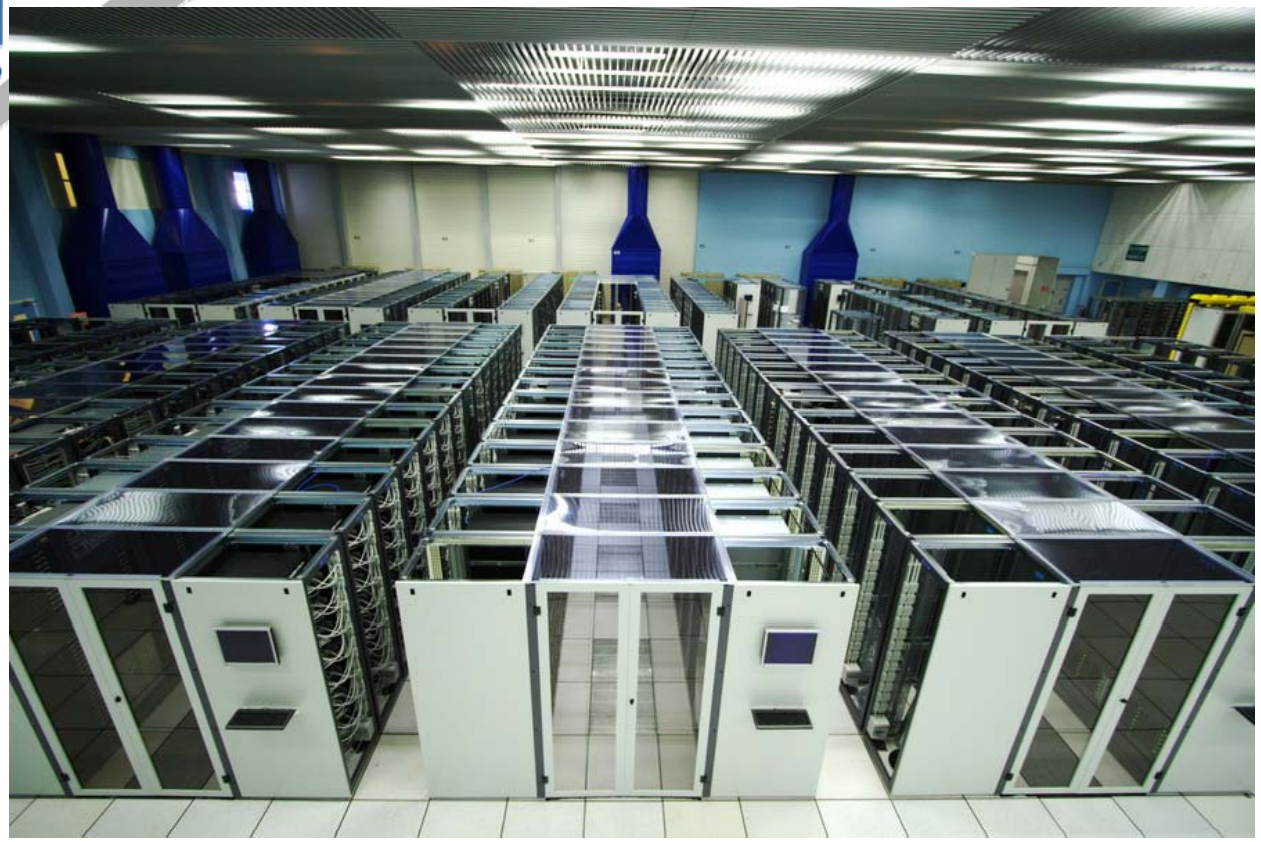

Nowhere near enough!

- High-throughput computing based on reliable “commodity” technology
- About 3500 multi-socket multi-core PC servers running Linux
- More than 10 Petabytes of data on tape; 30% cached on disk

Expected Power Evolution

Demand will grow continuously through the LHC era

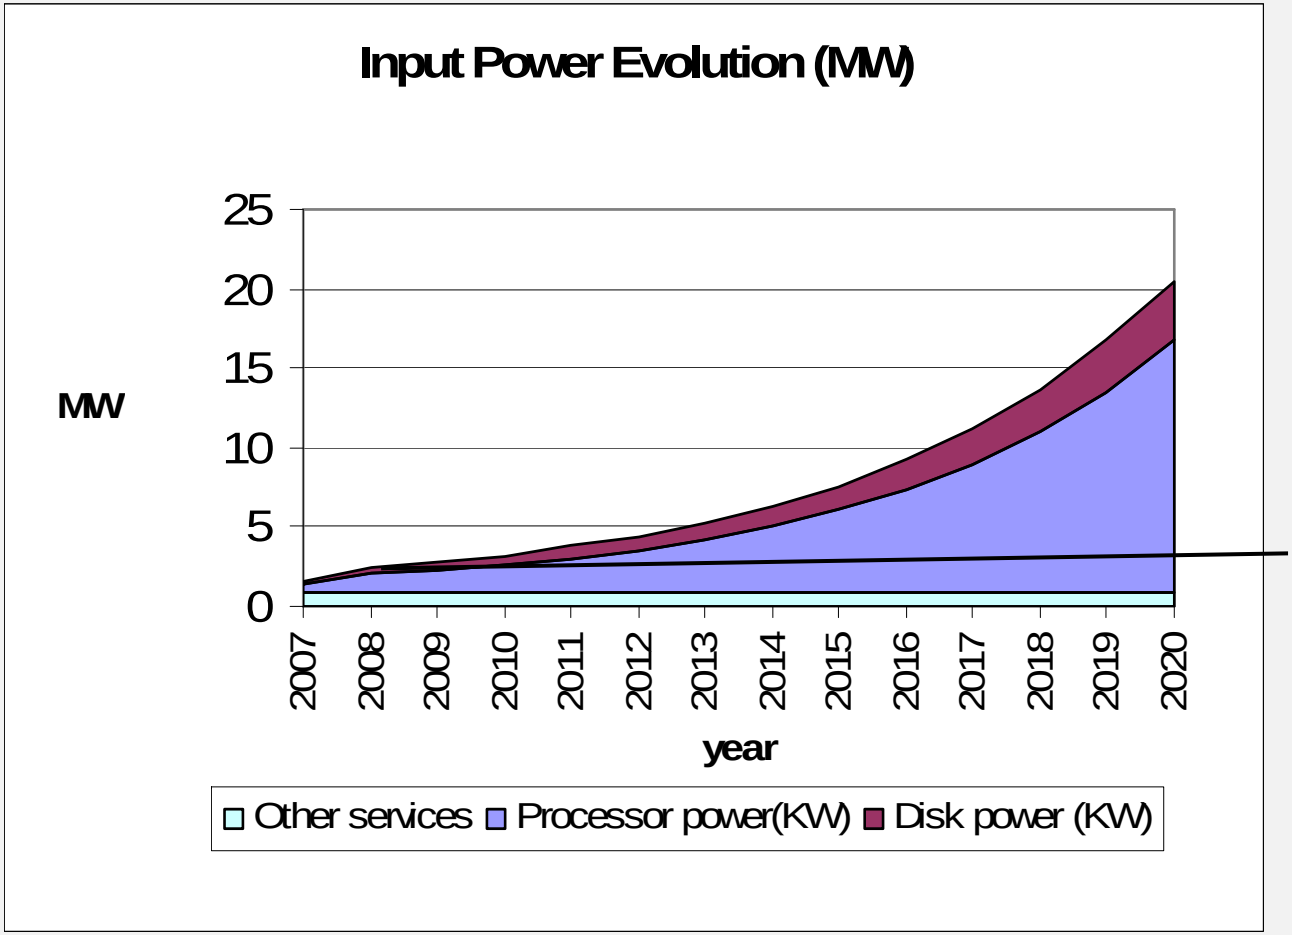

Power limit of present Computer Centre

LHC Computing Grid

- Largest Grid service in the world !

- **Almost 200 sites in 39 countries**
- **100’000 IA processor cores (w/Linux)**
- **Tens of petabytes of storage**

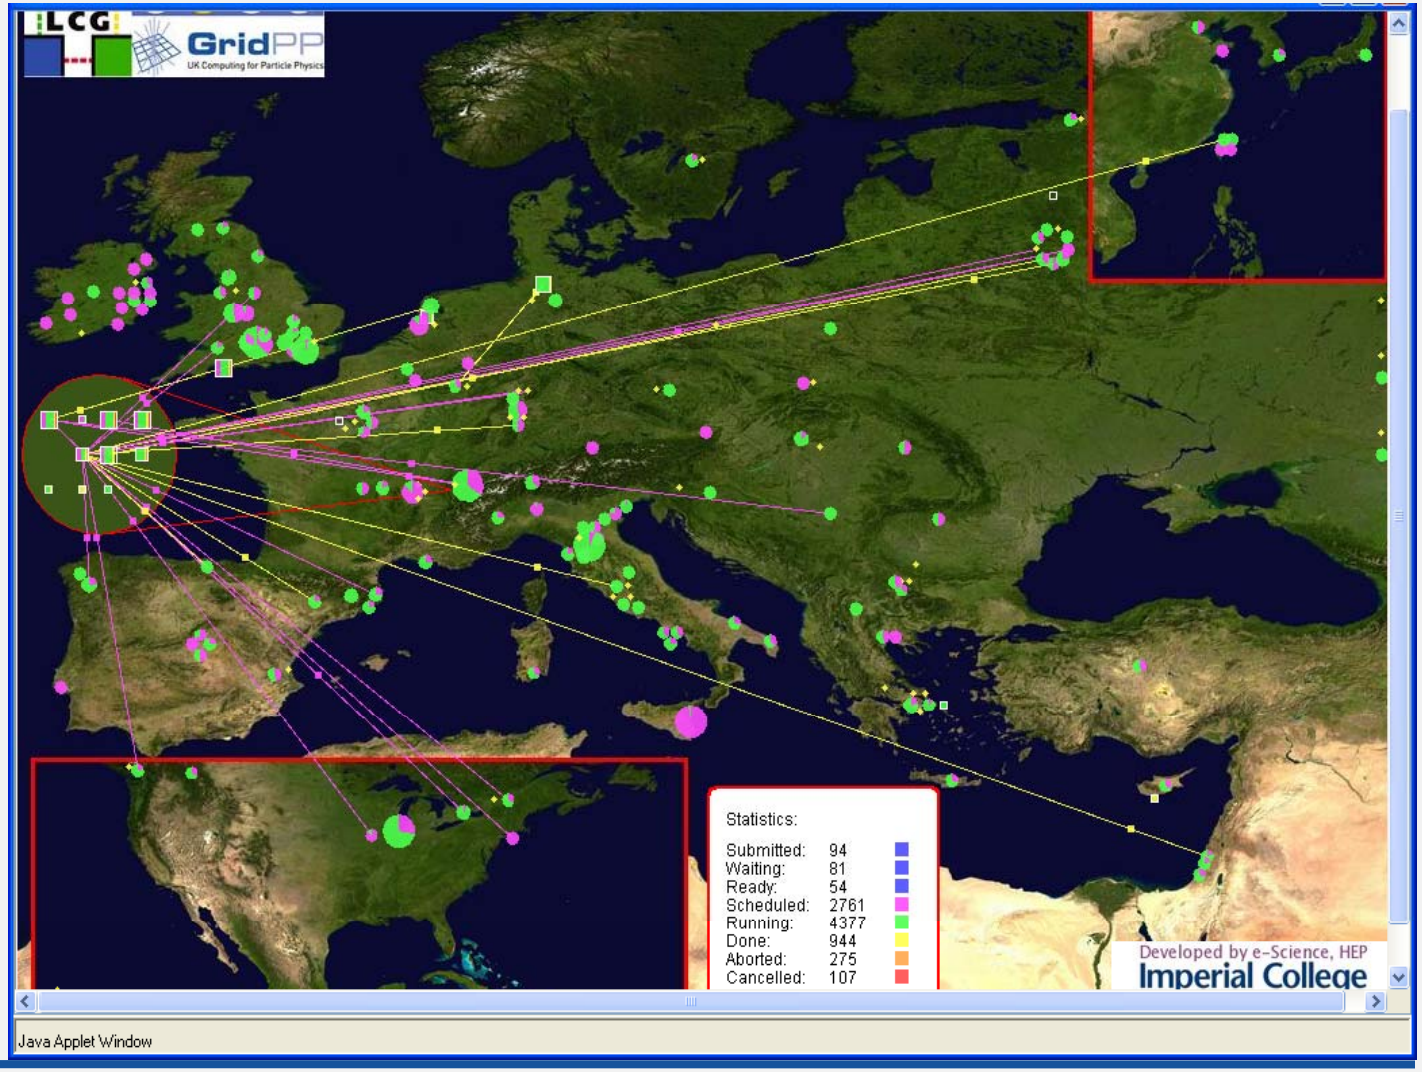

Background to the CERN openlab

- Information Technology has ALWAYS moved at an incredible pace
- During the LEP era (1989 – 2001) CERN changed its computing infrastructure twice:
  - Mainframes (1x) → RISC servers (30x) → PC servers (1000x)
- In openlab, we collaborate to harness the advantages of a continuous set of innovations for improving scientific computing, such as:
  - 10 Gigabit networks, 64-bit computing, Virtualization
  - Performance improvements (Moore's law): HW and SW
  - Many-core throughput increase, Thermal optimization
- We work with a long-term perspective:
  - LHC will operate until at least 2020!

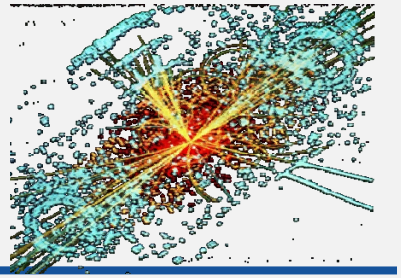

The CERN and Intel collaboration in openlab

Platform Competence Centre

- Intel-related activities:
  - Thermal optimization
    - Servers and entire Computer Centre
  - Virtualization
  - Multi-core performance/throughput improvements
  - 10 Gb networking
  - Sneak peaks at new technologies

Related to today's Xeon-7400 announcement

Thermal management

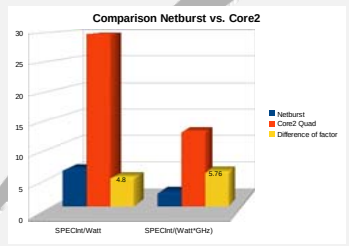

- Joint white paper on Computer Centre efficiency
  - Based on issues with existing building
    - Just issued; see next slide

- Complementary project to understand thermal characteristics of each server component
  - Processors (frequencies and SKUs) ; Memory (type and size); Disks; I/O cards; Power supplies

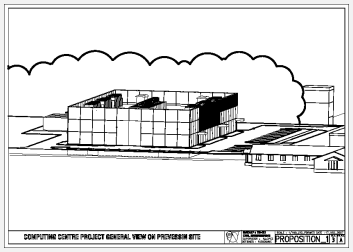

- Project for new Computer Centre
  - Understand all relevant issues (before starting)
  - Aim at 2.5 + 2.5 MW

Joint white-paper just issued

- Reducing Data Center Energy Consumption

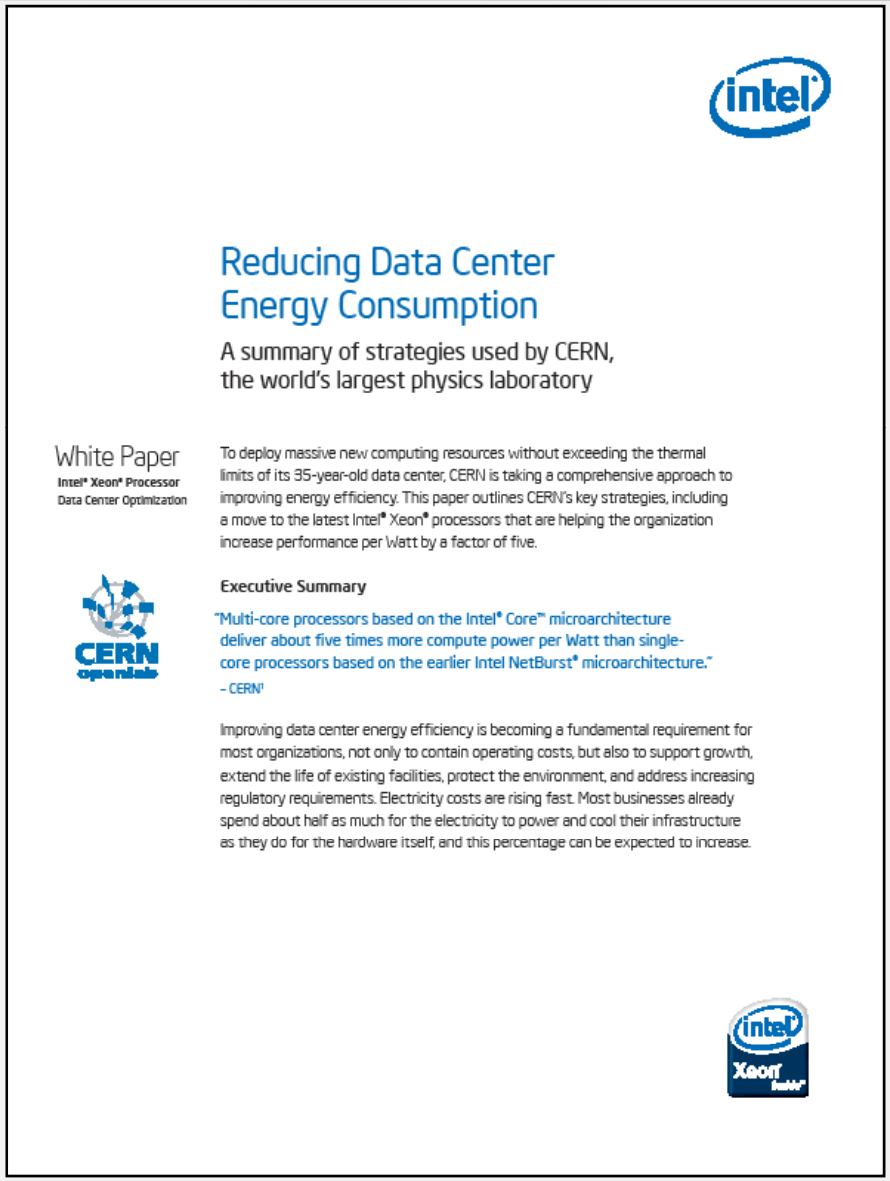

Reducing Data Center Energy Consumption

A summary of strategies used by CERN, the world's largest physics laboratory

White Paper
Intel® Xeon® Processor
Data Center Optimization

To deploy massive new computing resources without exceeding the thermal limits of its 35-year-old data center, CERN is taking a comprehensive approach to improving energy efficiency. This paper outlines CERN's key strategies, including a move to the latest Intel® Xeon® processors that are helping the organization increase performance per Watt by a factor of five.

Executive Summary

"Multi-core processors based on the Intel® Core™ microarchitecture deliver about five times more compute power per Watt than single-core processors based on the earlier Intel NetBurst® microarchitecture."
– CERN[1]

Improving data center energy efficiency is becoming a fundamental requirement for most organizations, not only to contain operating costs, but also to support growth, extend the life of existing facilities, protect the environment, and address increasing regulatory requirements. Electricity costs are rising fast. Most businesses already spend about half as much for the electricity to power and cool their infrastructure as they do for the hardware itself, and this percentage can be expected to increase.

http://download.intel.com/products/processor/xeon5000/CERN_Whitepaper_r04.pdf

Virtualization

- Multiple aims:
  - Server consolidation
  - System testing (used in LCG)
  - Improved flexibility and security
    - also in the grid
  - Personalization of images

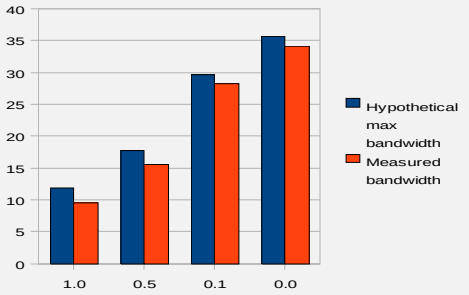

- Also: Development of complementary tools:
  - OSFarm
  - Content-Based Transfer mechanism

- Benchmarking

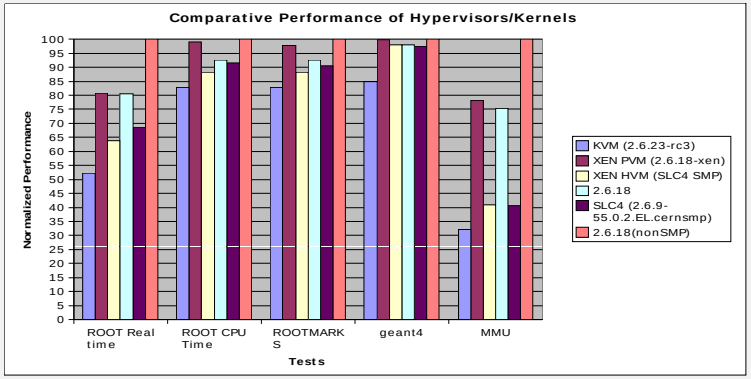

Cores: From Multi to Many

- Our “high throughput” computing is ideally suited:
  - Independent processes can run on each core, provided that:
    - Main memory is added
    - Bandwidth to main memory remains reasonable
  - In openlab, we have had early access to multiple generations:
    - Woodcrest, Clovertown, Harpertown, Dunnington; Montecito
- Already in November 2006, we were proud to be part of Intel’s movement to Quad core
  - All recent acquisitions have been Quad-core systems
  - And we are ready for the next step!

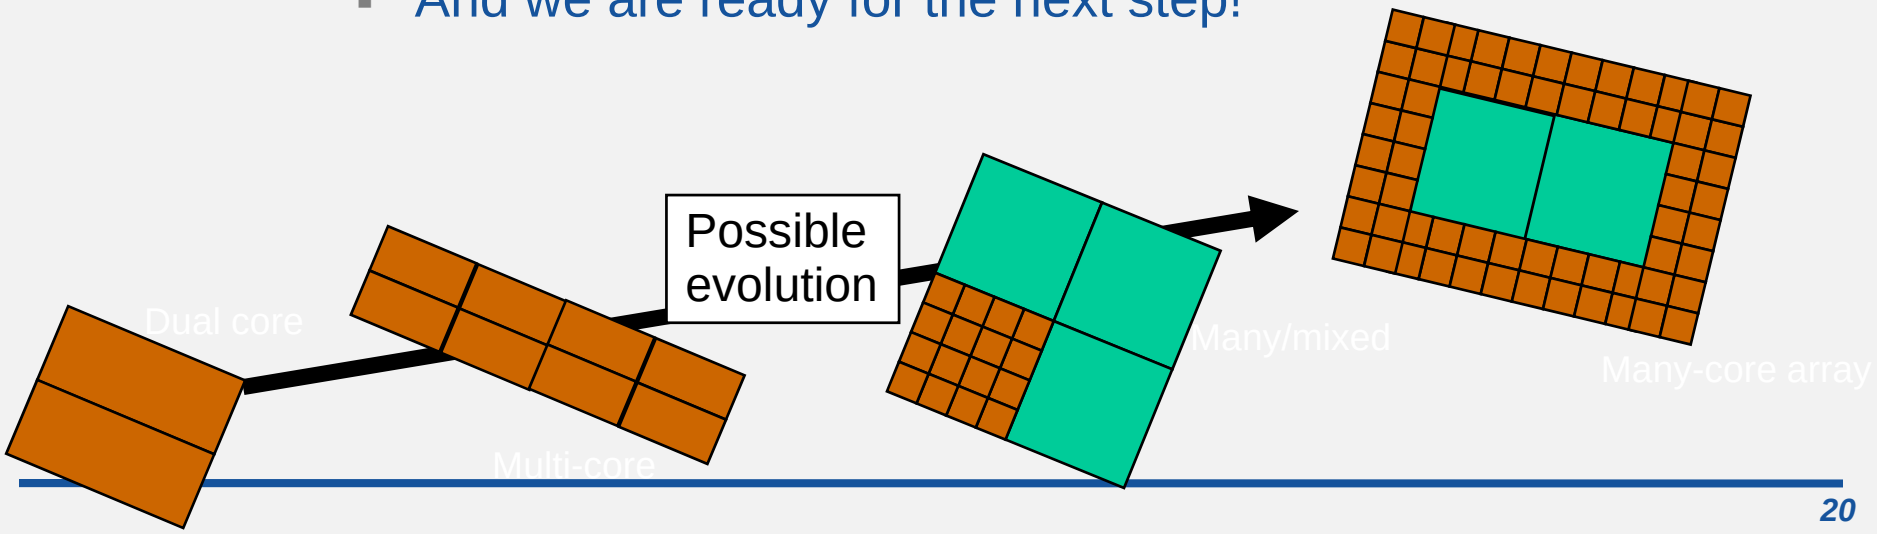

Multicore comparisons

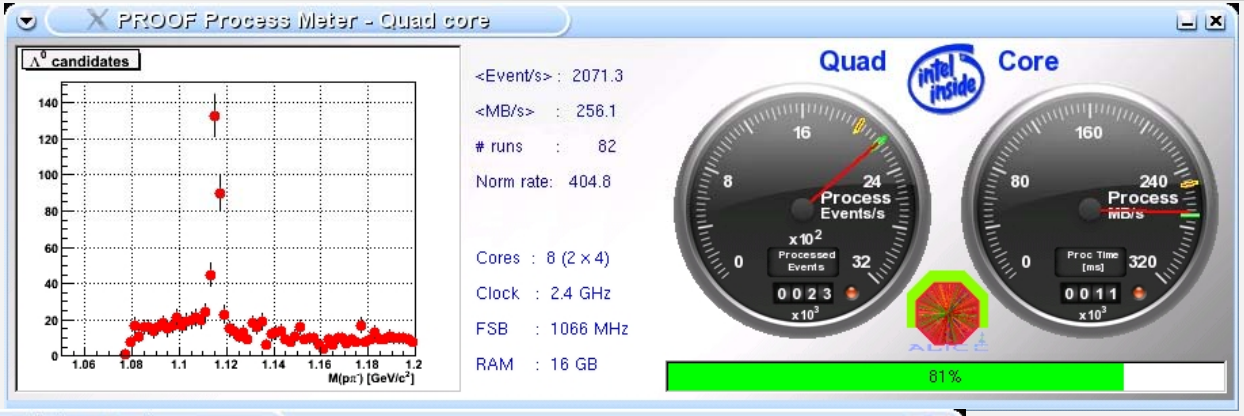

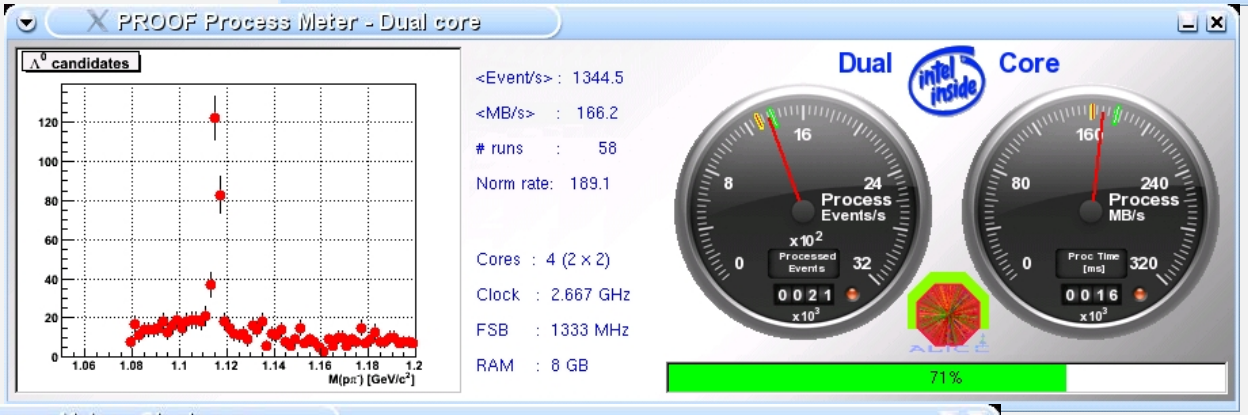

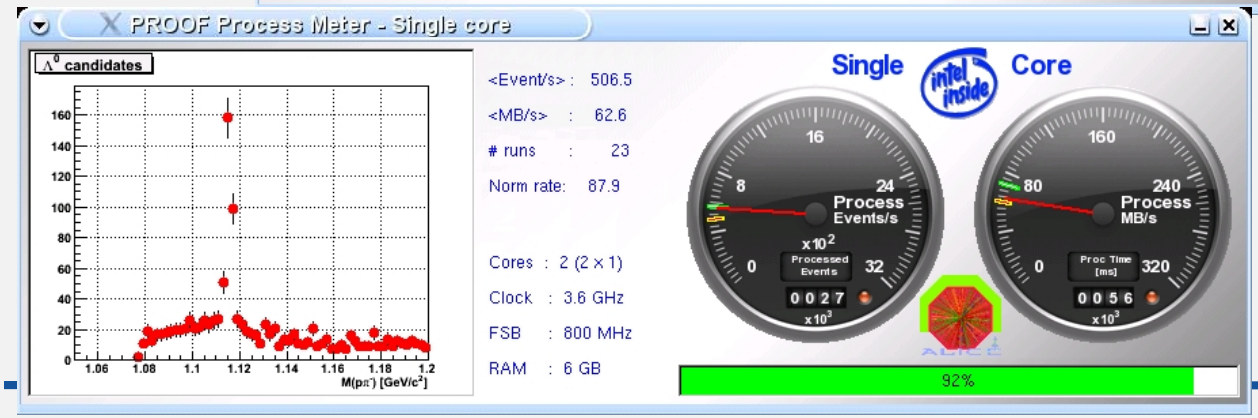

Conclusions

- As a flagship scientific instrument, the Large Hadron Collider and the corresponding Computing Grid will be around for 15 years
- It will rely on continued innovation in Information Technology:
  - Thermal improvements
  - Multicore throughput improvements
  - Virtualization improvements
- Today's Xeon announcement is a great step in this direction!

CERN and Intel in openlab

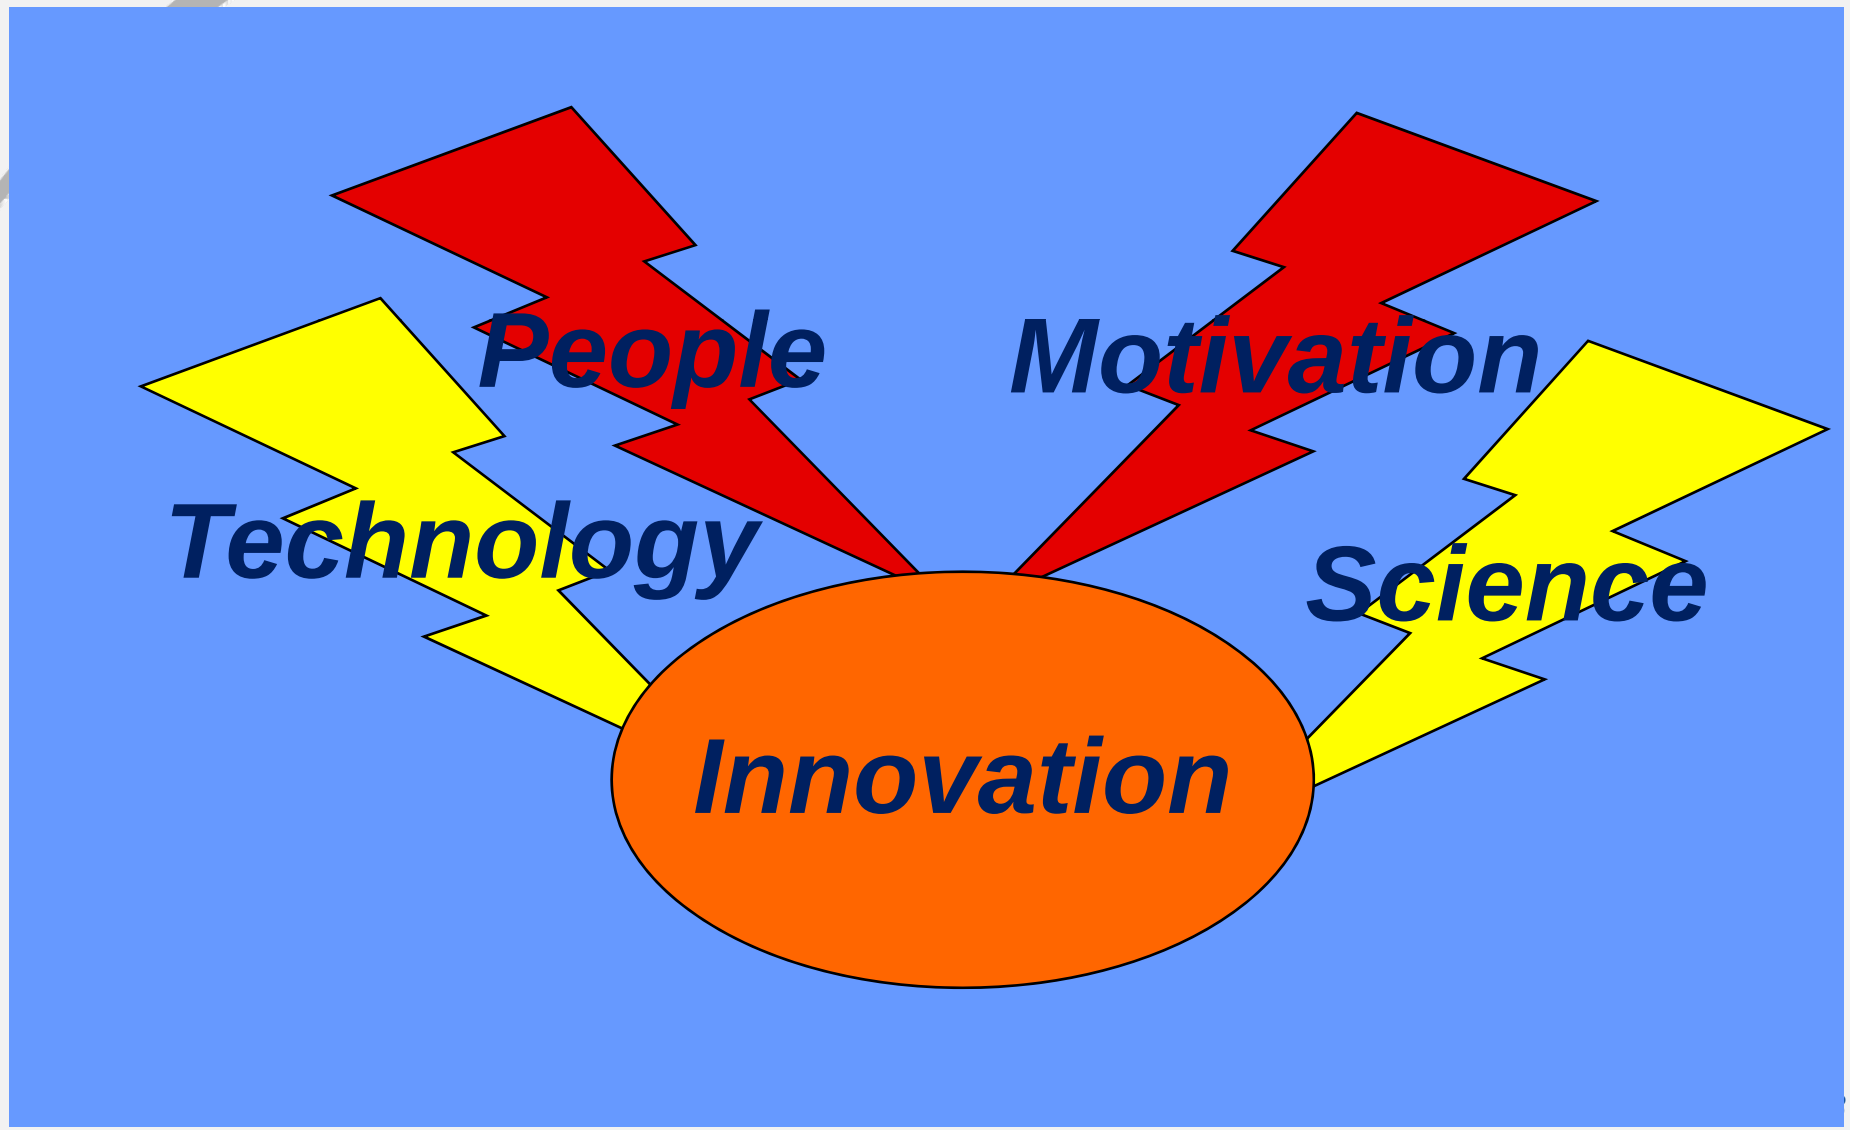

Backup