Technical activity overview

S.Jarp CERN openlab CTO

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

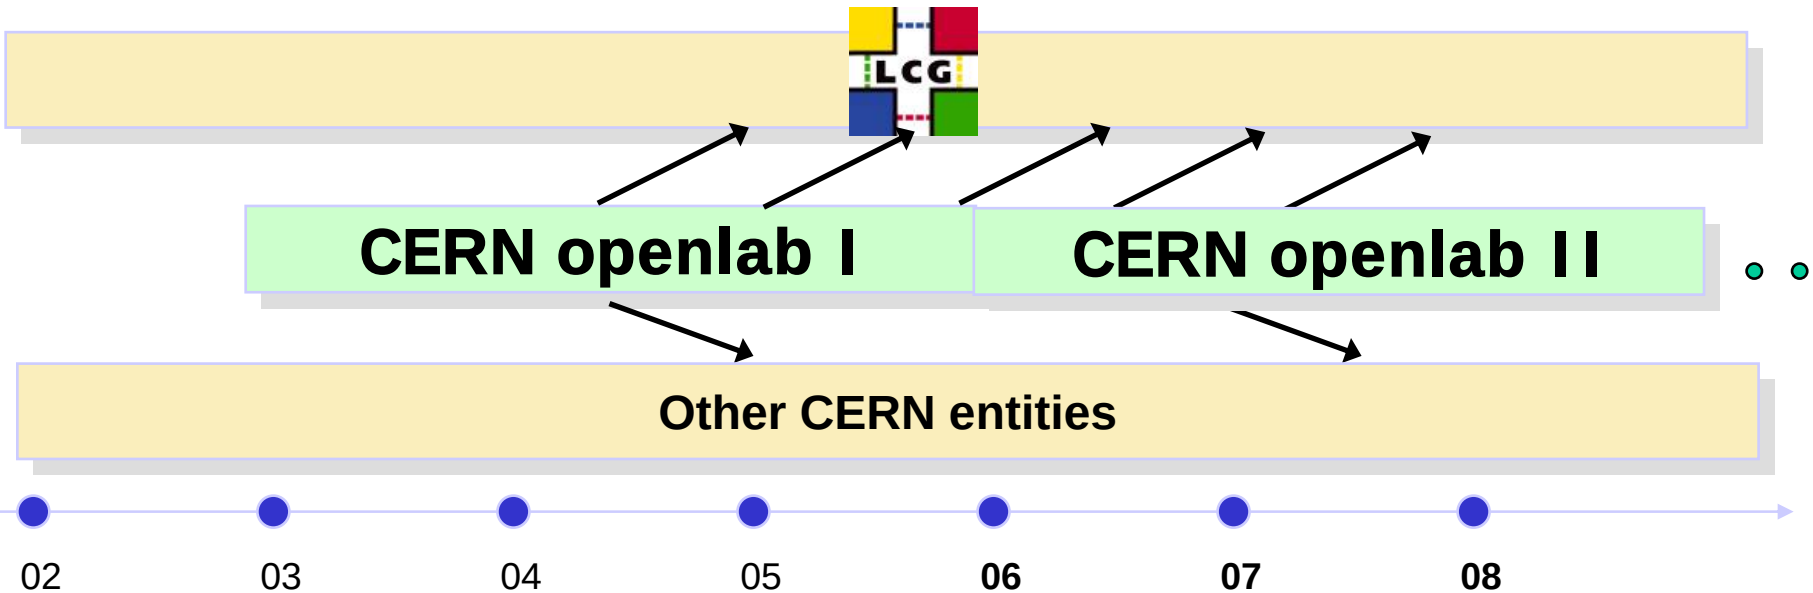

A few highlights from openlab-I

High performance testbed: the opencluster

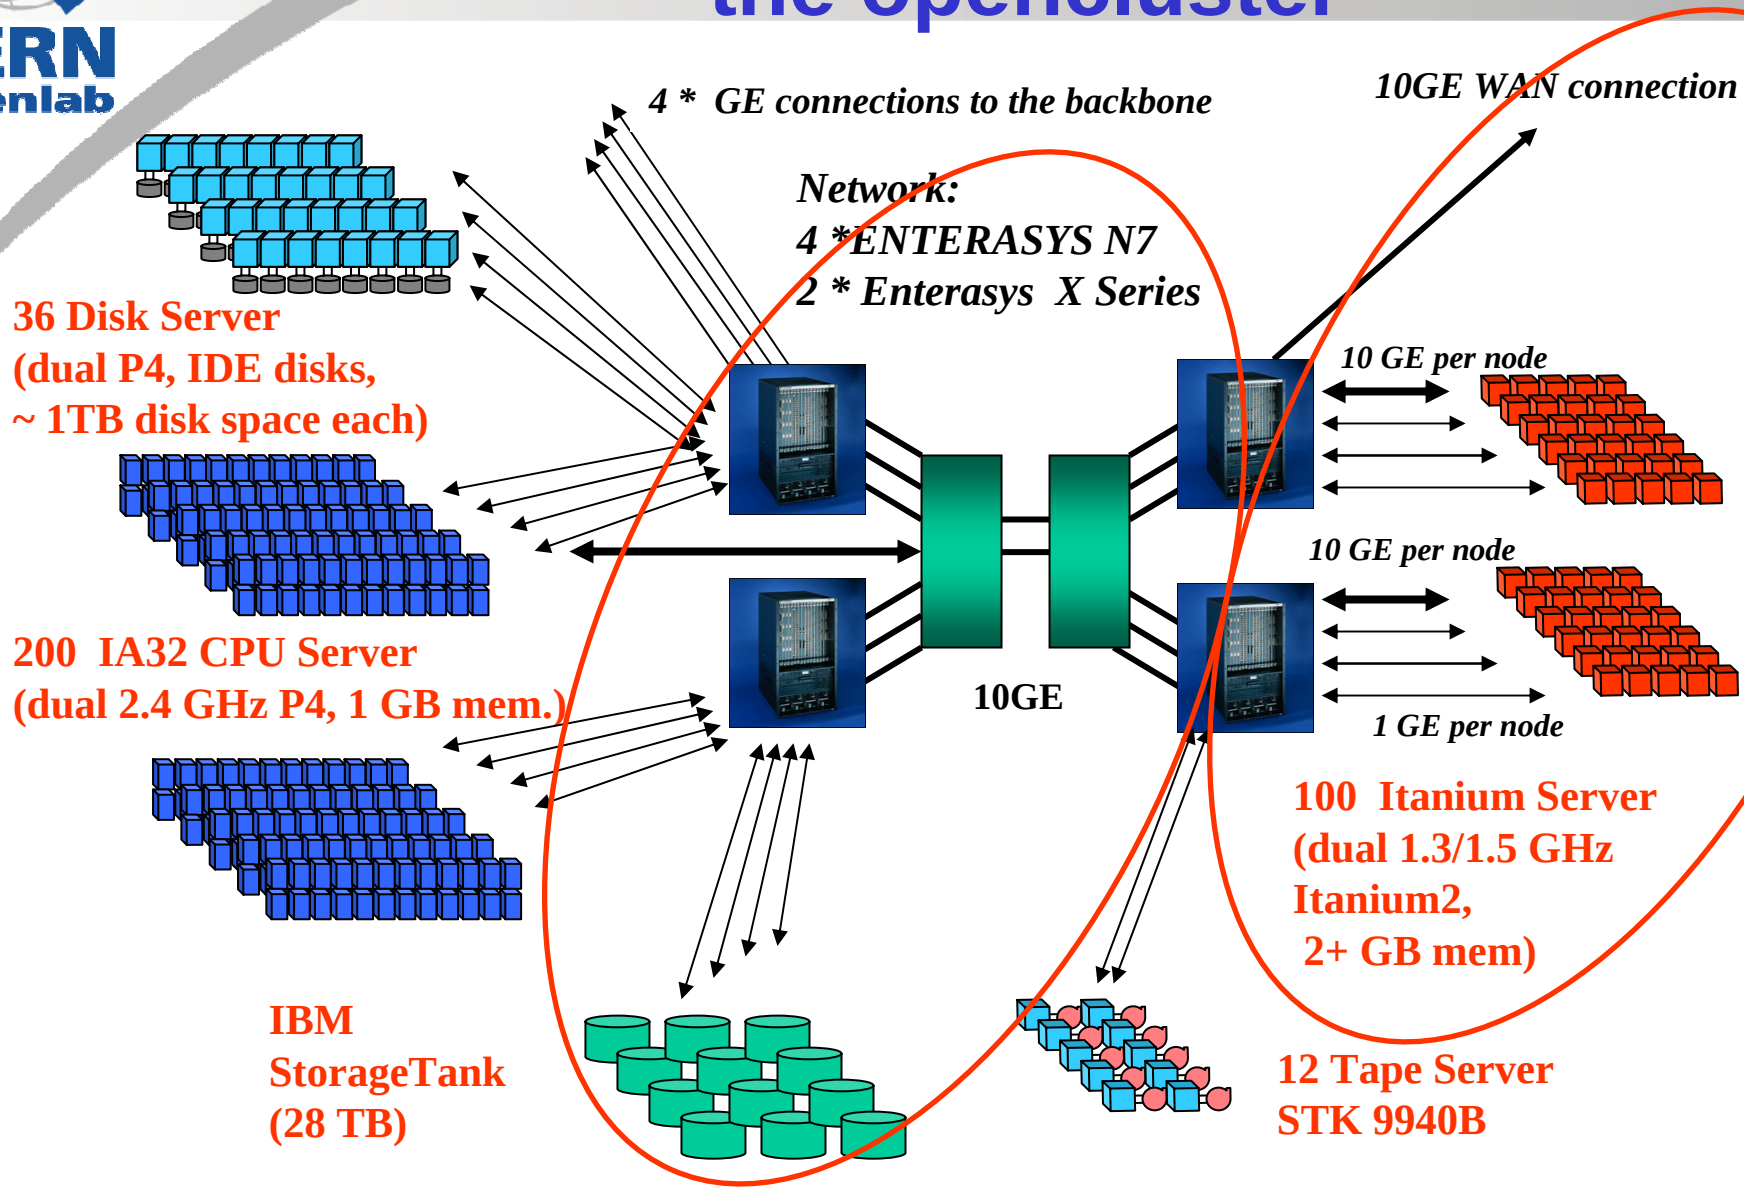

64-bit applications

- **Phase 1 completed (base packages) in collaboration with the physics programmers:**
  - ROOT (Data analysis framework)
  - Geant4 (Physics simulation framework)
  - CLHEP (C++ Class Library)
  - CASTOR (CERN Hierarchical Storage Manager)

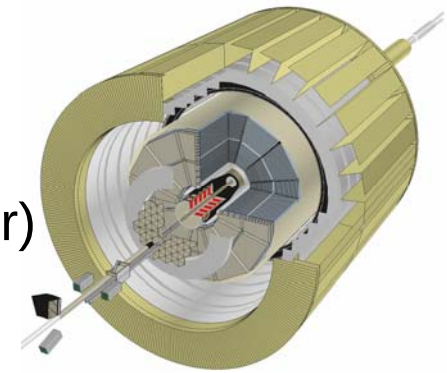

- **Phase 2 underway for the entire sw stacks:**
  - Set of external packages (Boost, etc.): OK
  - Base set of CERN frameworks (as mentioned above): OK
  - Generic HEP packages: OK
  - Specific packages from each experiment: Expected to be ready by Q4

64-bit grid middleware

- **Starting point: The software chosen for LCG had been developed only with IA32 (and specific Red Hat versions) in mind**
- **Two openlab members (working hard for a long time) managed to complete the porting of LCG-2 software to Itanium**
  - Result: All major components made to work on 64-bit Linux:
    - Worker Nodes, Compute Elements, Storage Elements, User Interface, etc.
  - Code, available via Web-site, transferred to HP sites (initially Puerto Rico and Bristol), as well as other interested sites (Poland, Singapore, etc.)
  - Changes fed back to software maintenance teams
  - Porting experience summarized in white paper

A successful demo of a heterogeneous Grid !

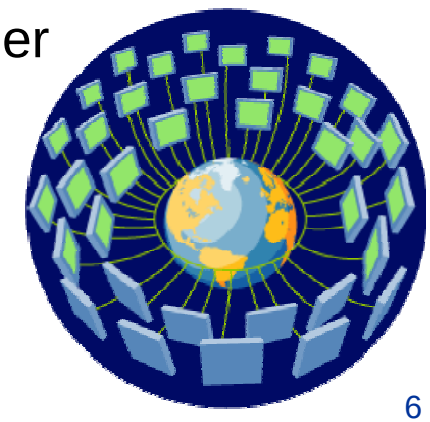

Computational Fluid Dynamics

- **Based on Itanium cluster with Infiniband switches from Voltaire**
- **CFD calculations:**
  - A numerical analysis of fluid flow, heat transfer and associated phenomena in LHC caverns
  - 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

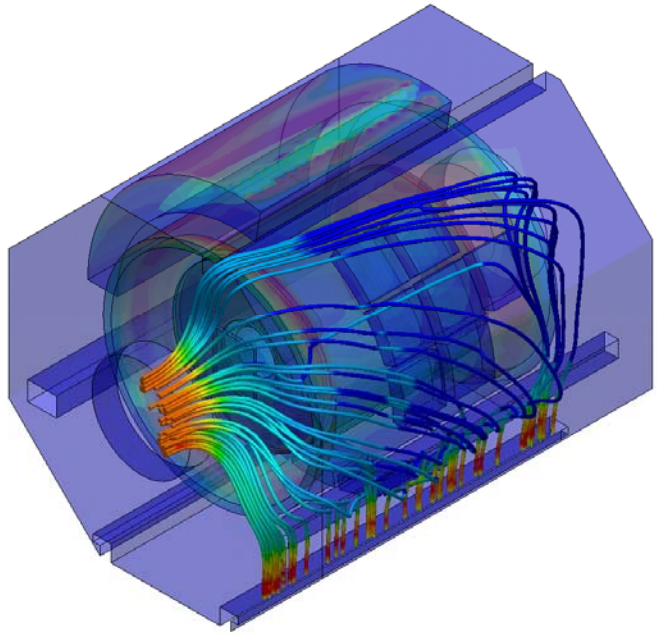

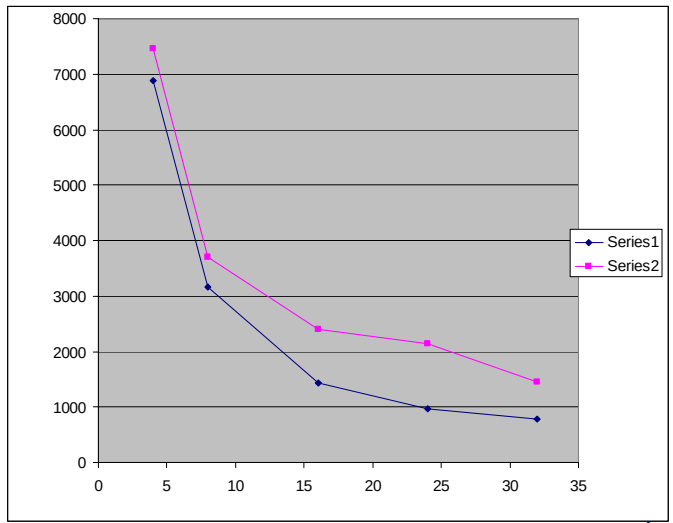

Moving to openlab-II

- **HP-related activities**
  - Procurve (Ethernet switches)
  - SmartFrog
  - Tycoon
- **Intel-related activities**
  - Multicore benchmarking
  - Virtualization
  - Compiler project
- **Oracle-related activities**
  - *Streams* for databased replication
  - Software regression testing

During the summer: Strong addition of talented students; accelerating the work in almost all areas

HP Procurve

- **Build on 5-year commercial agreement for edge switches in CERN's network topology**
- **Project aim:**
  - Understand the behaviour of large computing networks (10K+ nodes), such as CERN's:
    - Detect traffic anomalies in the system
    - Enable trend analysis
    - Automate counter measures
    - Provide post-mortem analysis facility
- **Manpower:**
  - Two postdocs (to be hired)
- **Three phases**
  - Identify and exploit relevant data sources
  - Perform data analysis
  - Understand scaling and build prototypes

SmartFrog

- **Understand how such a configuration framework can be used inside Grids**
  - Starting-point:
    - Smartfrog framework is “empty” and needs to be populated
  - Already in openlab-I we did a fair amount of work to enhance the framework for our environment
    - See SmartFrog Web pages (inside CERN’s openlab pages)
  - Current plan:
    - Prototype new components in order to understand the suitability of SmartFrog orchestration inside EGEE testing (ETICS project) and LCG software deployment (new releases of EGEE middleware in 200 sites)
- **Manpower**
  - Xavier Gréhant (PhD student) – since 01/06
    - Decide PhD contents
      - In Sept. together with HP Labs and ENST, Paris
  - One summer student (July - August)

Tycoon

- **Understand Tycoon in the context of e-science grids:**
- **How to deal with swings in demands for computing**
  - Provide complementary services?
    - Based on efficient gateways
  - Provide technology that can be integrated into science grids (such as EGEE)?
    - Analyze Tycoon's features in the context of large-scale centres
- **Manpower**
  - One summer student (now)
  - One postdoc "CERN fellow" (end-October)
- **Tests**
  - Just started with CERN, HP, NGO (Singapore)
  - Soon also others:
    - GRNET (Greek Research and Technology Network)
    - Research communities in INDIA
- **Collaboration**
  - Regular contacts with Kevin Lai (HPL, PA)
  - Summary of experience being written up
    - Tycoon's Wiki pages

Virtualization

- **Rationale:**
  - Grids will be much more flexible and secure when using virtualization
  - All future processors/platforms will have hardware support for virtualization
    - VT-x, VT-I (Hardware processor support)
    - VT-d (Hardware I/O support)
- **Our initial involvement**
  - Xen benchmarked with CERN simulation workload on IA-32
    - Work started by Summer Student 2004
  - Porting of Xen to Itanium
    - Completed last year (Master thesis): Collaboration with HP Labs and Intel
- **Working with CERN's Grid Deployment team**
  - Focused on several requirements:
    - Create an automated server test environment under Xen
    - Allow multiple Linux distributions to be used (in a flexible manner)
      - Image factories
    - Increase security when running "foreign daemons"

LHC Computing

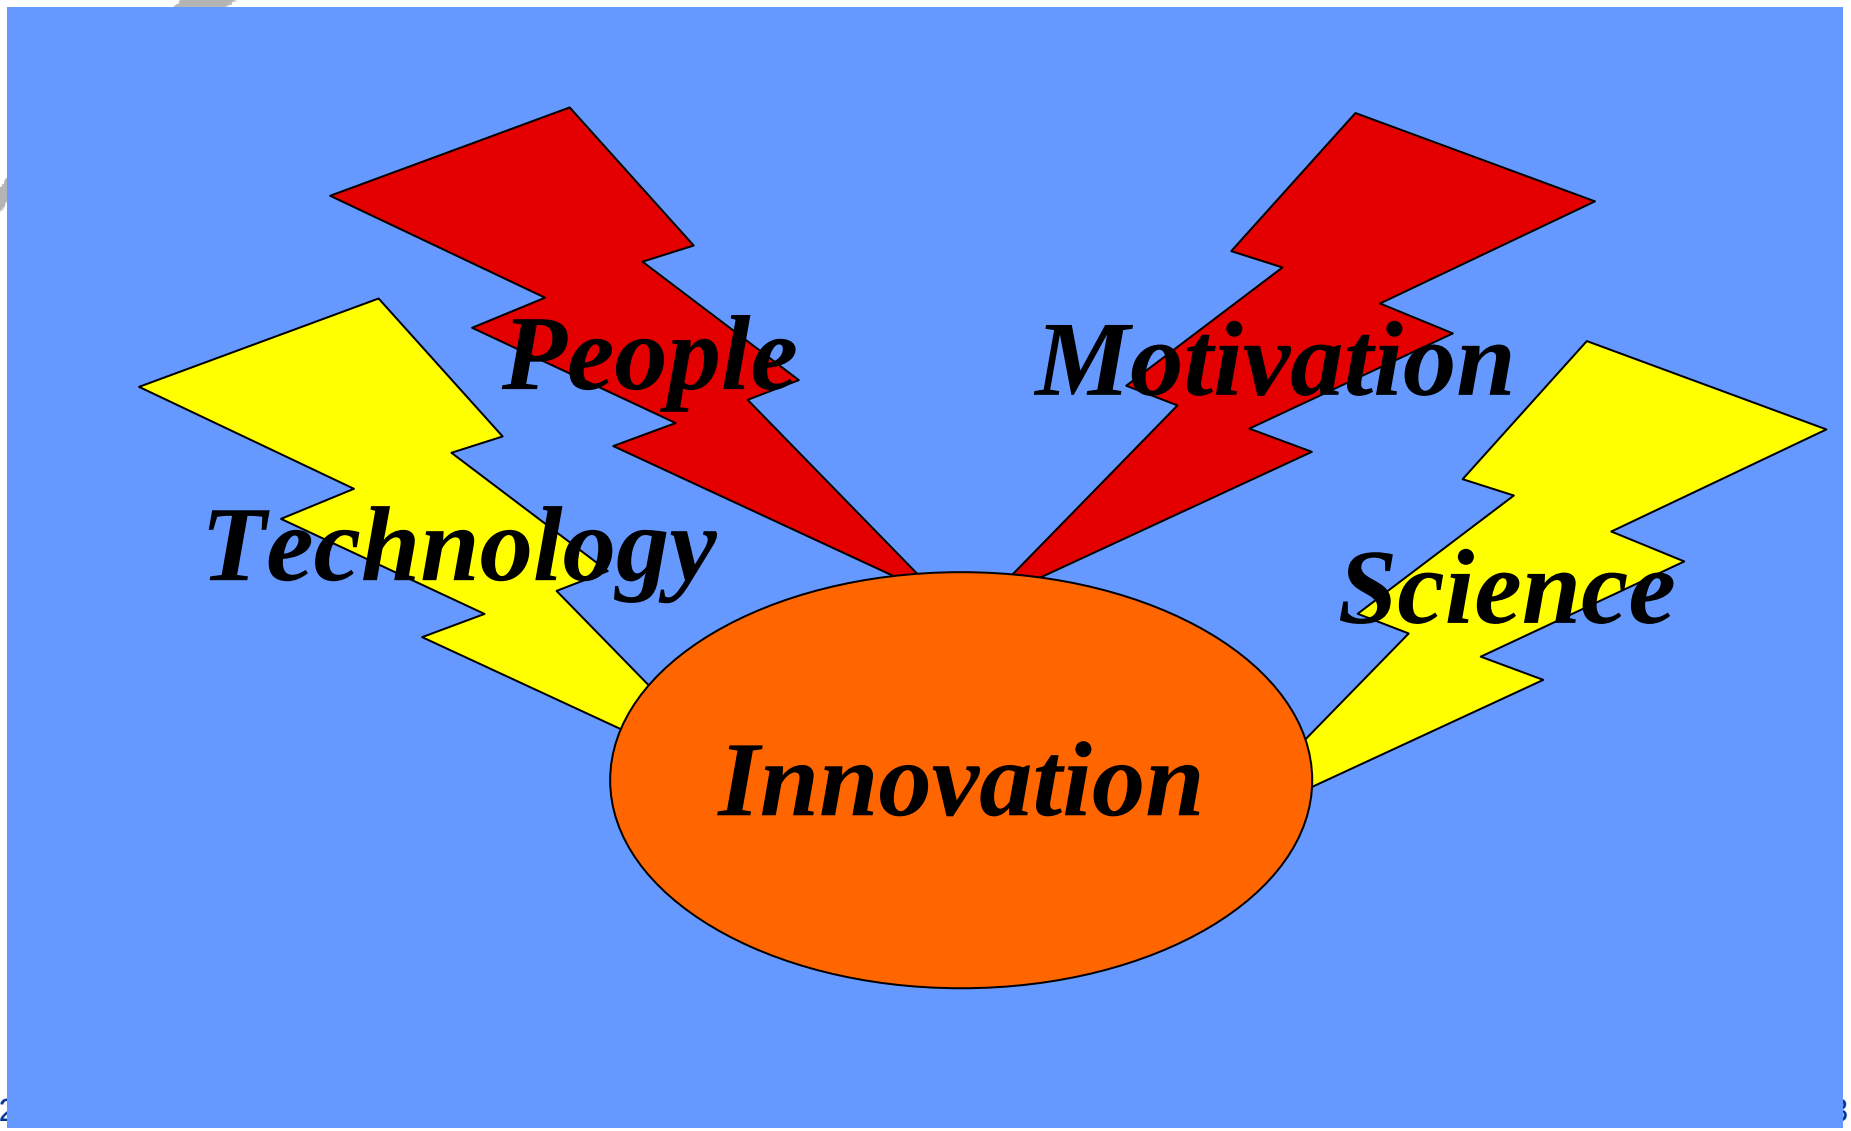

BACKUP

Benchmarking

- **CERN (and our community) will want:**
  - SPECint/USD
    - SPEC2000 with gcc used in acquisitions
  - SPECint/Watt
    - 2.5 MW is our Computing Centre capacity
    - Power meter acquired for accurate testing

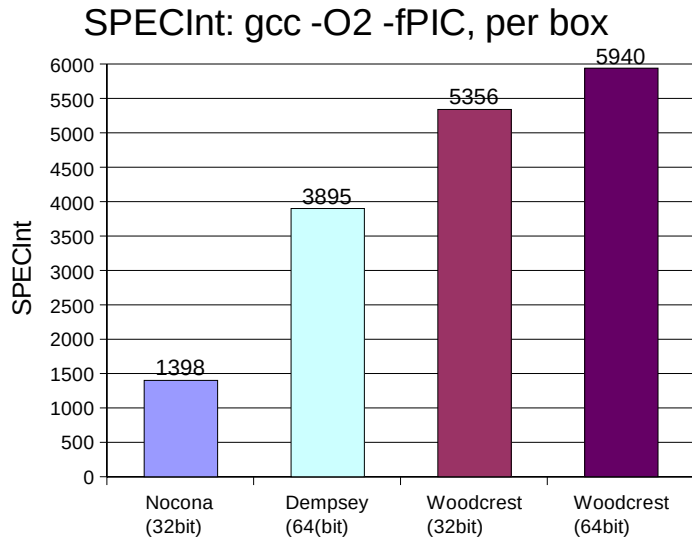

- **Also, effort on optimization of jobs from the LHC collaborations**
  - So far, profiling has been done on:
    - Simulation frameworks (ATLAS and LHCb)
    - Reconstruction framework (ATLAS)

Multicore/Manycore

- **Our "high throughput" computing model is maybe ideally suited:**
  - Independent processes can run on each core, provided
    - Main memory is added
    - Bandwidth to main memory remains reasonable
  - Testing, so far, has been very convincing
    - Dempsey, Sossaman
    - Montecito
    - Woodcrest
- **Active collaboration with Intel**

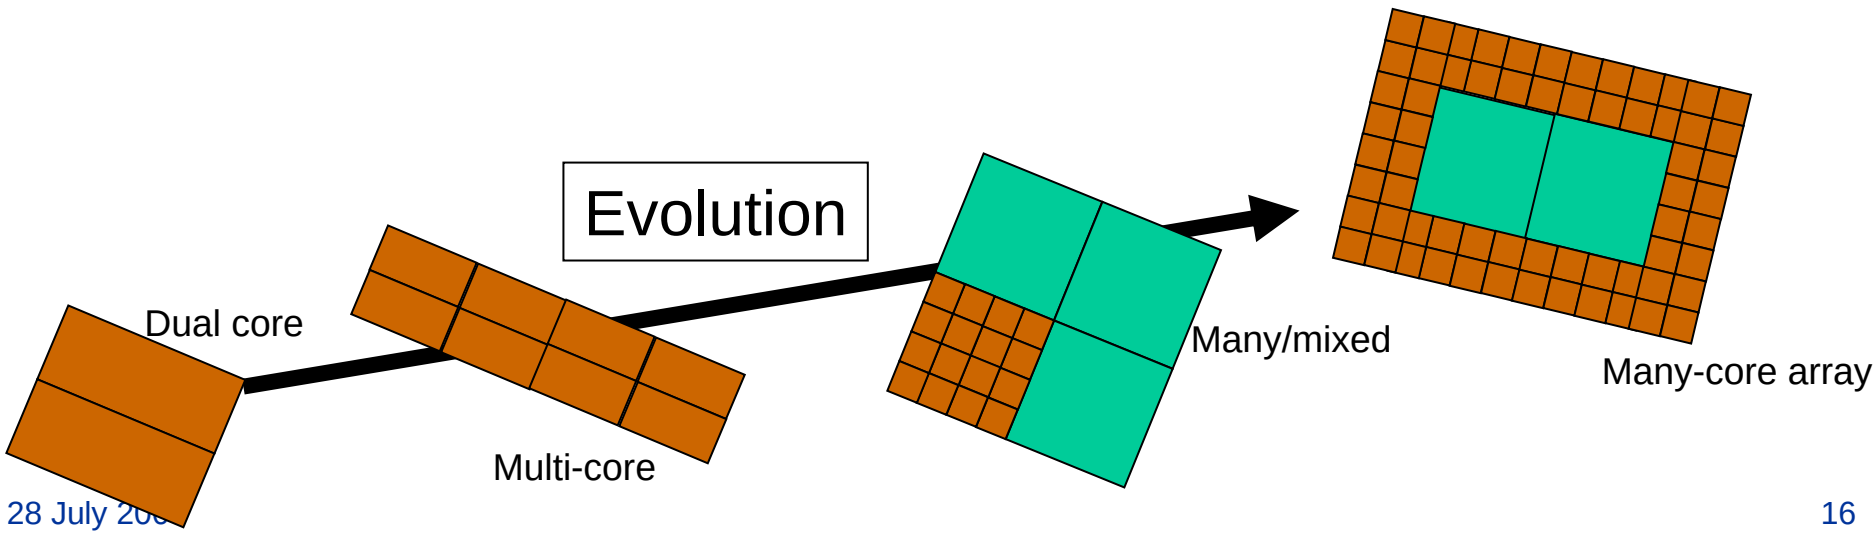

Compiler project

- **Since most High Energy Physics programs are written in-house, compiler optimization translates directly into reduced cycle consumption per job**

- **Openlab has worked on several fronts**
  - Get the compiler writers to take large C++ jobs more seriously
    - The world was (still is?) too dominated by the C and FORTRAN languages
      - More emphasis will come with SPEC2006
  - Add programs, such as ROOT, to the regression testing of compilers
  - Work with compiler developers (Intel, HP, GNU) to improve generated code for key sequences inside HEP programs
    - Example: Random number generators, Geometric navigation routines, etc.

- **In openlab I the focus was mainly on the Itanium processor**

- **In openlab II the focus is on both Itanium and Xeon 64-bit**