CERN openlab for DataGrid applications

Introduction

CERN openlab

- 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

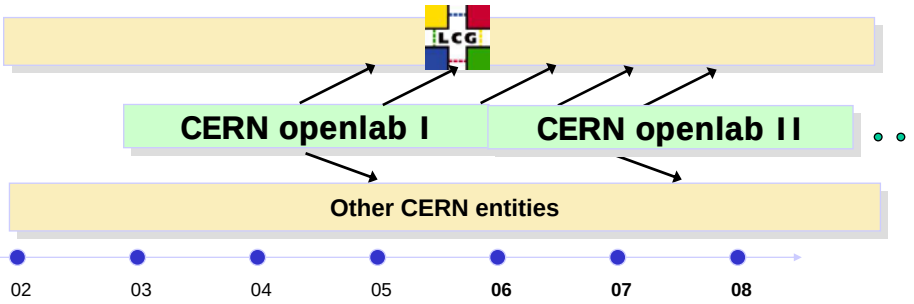

In partnership with

IBM
intel
Today's focus
hp invent
ORACLE
ENTERASYS NETWORKS

Integration with the LCG test bed

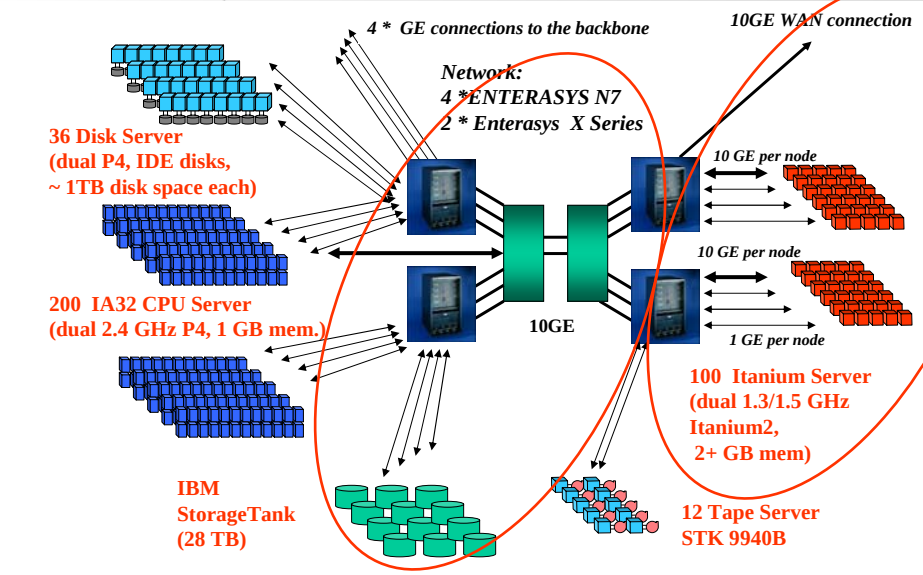

High Througput Prototype (opencluster + LCG prototype)

Examples of certain hardware-oriented activities

Example of early success
(Land Speed records)

- **Internet2 Land Speed record established during Telecom 2003**
  - Collaboration with DataTag project and the ATLAS experiment
  - Speed above 5 Gbits/s was easily reached with single Itanium server
  - Telecom result: 5.44 Gbit/s over 7'067 km
  - Early in 2004: 6.57 Gbits/s over 15'766 km

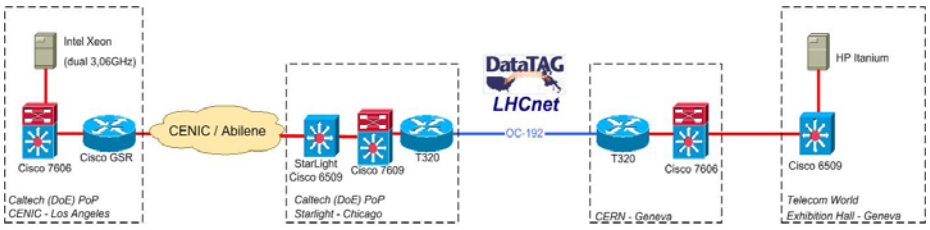

The technical layout for the Land Speed Record during Telecom 2003.

LCG Service Challenges

- **SC1 (Data Export) used 10 Itanium systems**
  - Achieved stability between CERN and Fermilab at 5 Gbps (over a 10 Gbit link)
    - Transfers via GridFTP
    - Limited by receiving side
- **SC2:**
  - Tests with more Itanium systems and more sites (successfully completed in April)
    - Fermilab, Karlsruhe @ 10 Gbps
    - CNAF, RAL, SARA, IN2P3, BNL @ 1 Gbps
    - Sustained 600 MB/s (peaks at 800 MB/s)
    -
- **SC3 is now running:**
  - Demonstrate "production quality" service
  - Involving also Xeon systems
  - With CERN's integrated data management (CASTOR)

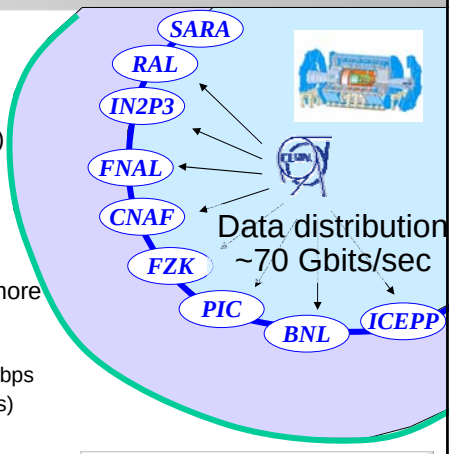

Stable data transfer at ~5Gbit/s to FNAL (US).

Computational Fluid Dynamics

- **Collaboration with dedicated team in TS Department**
- **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 our LHC experiments**
  - and others

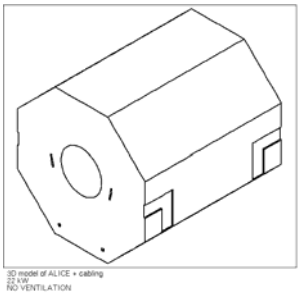

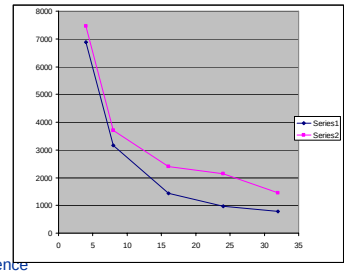

Examples of certain software-oriented activities

64-bit applications

- **Phase 1 completed (base packages) in collaboration with the physics programmers:**
  - ROOT (Data analysis framework)
    - http://root.cern.ch/
  - Geant4 (Physics simulation framework)
    - http://cern.ch/geant4
  - CLHEP (C++ Class Library)
    - http://proj-clhep.web.cern.ch/proj-clhep/
  - CASTOR (CERN Hierarchical Storage Manager)
    - http://cern.ch/castor

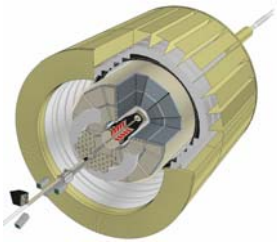

- **Phase 2 underway (entire sw stacks):**
  - Set of external packages (Boost, etc.): OK
  - Base set of CERN frameworks (as mentioned above): OK
  - Generic HEP packages: In progress
  - Specific packages from each experiment: Not yet formalised, but fewer issues expected

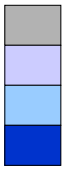

64-bit grid middleware

- **Starting point: The software chosen for LCG had been developed only with IA32 (and specific Red Hat versions) in mind**
- **Two openlab members worked for many months 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
  - Changes given back to software maintenance teams
  - Porting experience summarized in white paper

All of a sudden the Grid was heterogeneous !

Compiler project

- **CERN needs an unlimited amount of CPU cycles**
  - Research knows no upper bounds
- **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**
  - Add programs, such as ROOT, to the regression testing of compilers
    - Avoid bugs when new compiler versions come out
  - 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 should come with SPEC2006
  - Work with compiler developers (mainly Intel, so far) to improve generated code for key sequences inside HEP programs
    - Example: Random number generators, Geometric rotation routines, etc.
- **In openlab I the focus was mainly on the Itanium processor**
- **In openlab II the focus will be on both Itanium and Xeon**

Virtualization

- **Important fact:**
  - All future processors (Itanium, Pentium 4, Opteron, etc.) will have hardware support for virtualization
- **Our involvement**
  - Xen benchmarked with CERN simulation workload on IA-32
    - Work started by Summer Intern 2004
  - Porting of Xen to Itanium
    - Completed in July (Master thesis): Collaboration with HP Labs
  - New efforts:
    - Demonstrate I/O intensive workloads (ROOT analysis, etc.)
- **Our view:**
  - Grids will be much more flexible and secure when using virtualization
- **Three requirements have come from LCG:**
  - Create a simple server test environment under Xen — OK
  - Allow multiple Linux distributions to be used — OK
  - Allow Virtual Organizations to run “foreign daemons” in a VO box — In progress

Plans and Conclusion

Plans for openlab II

- **Initially two competence centres and plans are still being worked out**
  - Platform Competence Centre
    - Optimization of processors, memory, I/O equipment
    - Exploitation of 64-bit addressing, virtualization and multicore

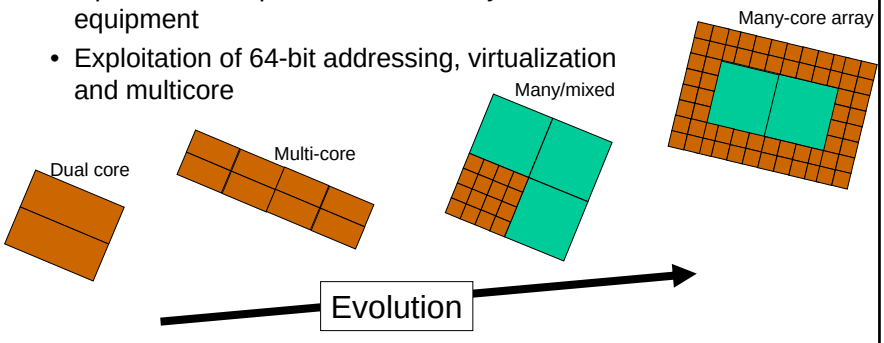

Plans for openlab II (cont'd)

- **In collaboration with EGEE II**
  - Grid Interoperability Centre
    - Test and integration of new software components
    - Interoperability tests
    - Testing of new grid components
      - For instance, improved security

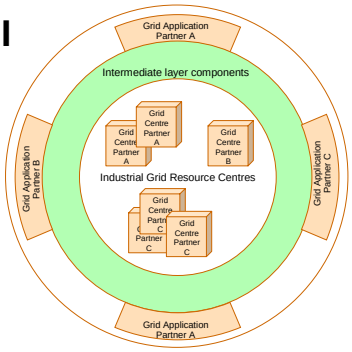

- **Thirdly**
  - Continued set of network-related activities

Conclusions

- **CERN openlab:**
  - Solid collaboration with our industrial partners
    - Encouraging results in multiple domains
  - We believe partners are getting good "ROI"
    - But only they can really confirm it → so ask them
  - No risk of running short of R&D
    - IT Technology is still moving at an incredible pace
  - Vital for LCG that the "right" pieces of technology are available for deployment
    - Performance, cost, but also reliability and resilience, etc.
  - Likely ingredients for LCG (so far):
    - 64-bit "clean" software, improved C++ compilation technology, next generation I/O (10Gbit Ethernet, Infiniband, etc.), new processor features (virtualization, etc.)

CERN openlab II

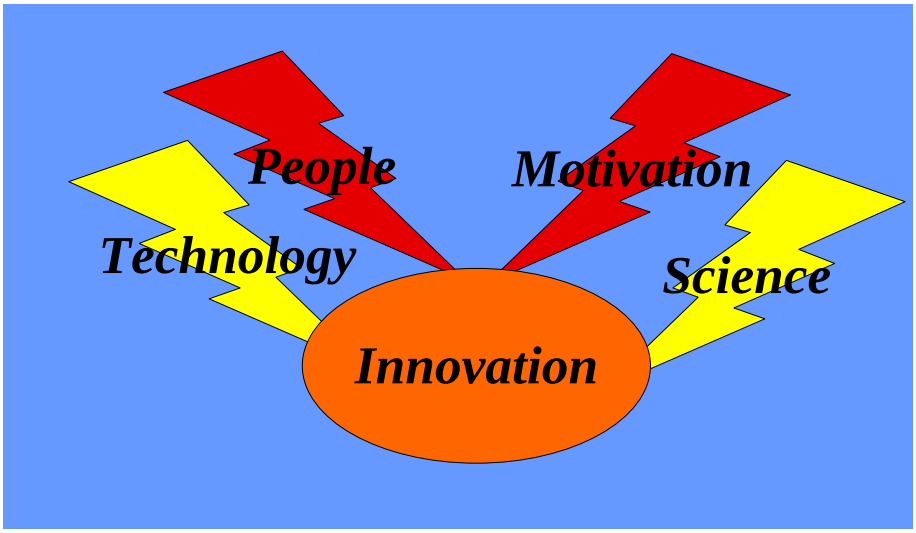

BACKUP

Parallel ROOT Scalability
(Presented at CHEP2003)

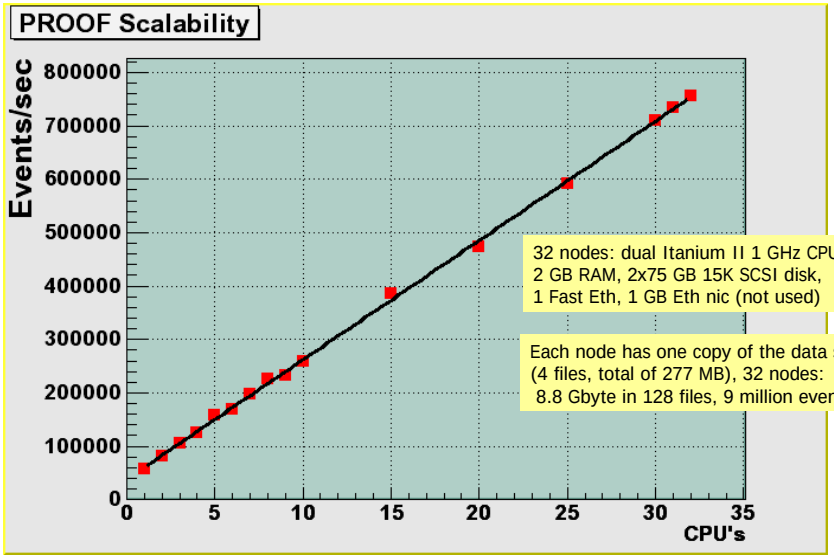

Next generation disk servers

- **Based on excellent equipment:**
  - Two 4-way Itanium servers (RX4640)
    - Two full-speed PCI-X slots
      - 10 GbE and/or Infiniband
    - Two sets of RAID controllers
    - 24 * S-ATA disks with 74 GB
      - WD740 "Raptor" @ 10k rpm
      - Sustained R/W speed of 55 MB/s

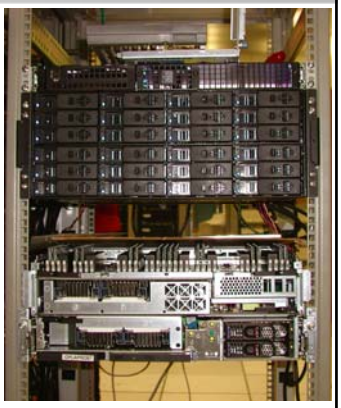

	Three 3-ware controllers	Three Areca controllers
Read speed	**~ 1100 MB/s**	**~1100 MB/s**
Write speed	**~550 MB/s**	**~900 MB/s**