

STONESOFT

# **CERN openlab**

- CERN IT department's main R&D focus
- Framework for collaboration with industry
- Evaluation, integration, validation
  - of cutting-edge technologies that can potentially serve the LHC Computing Grid (LCG)

### Sequence of 3-year agreements

- 2003 2005: the "opencluster" project
- 2006 2008: openlab II with new projects





### Background

- IT Technology continues to move at an incredible pace
- CERN "converted" most of its computing to Intel PCs already in the late 90ies
- Since then, we collaborate closely 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): hardware and software
  - Multicore throughput increases
  - Reduced needs for electric power and air conditioning

### • We need a long-term perspective:

- LHC will operate until at least 2020!



## A few highlights from CERN/Intel collaboration in openlab I (2003 – 2005)



14 November High Througput Prototype (opencluster + LCG testbed)



### Example of early success (Land Speed records)

### Network speed record established during Telecom 2003

- Aim: Transfer the equivalent of a full DVD (5 GB) with a single Itanium server and the Intel 10Gb Network Interface Card (NIC) in less than 10 seconds to California
- Great success!
  - At Telecom: 5.44 Gbit/s over 7'067 km
  - Early in 2004: 6.57 Gbits/s over 15'766 km



- 10GE

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

### 64-bit memory addressing

#### With 32 bits, a PC can address 4 Gigabytes

 This is equivalent to the data that will come from the LHC detectors during a couple of seconds



- With 64 bits, a PC can address 18 Exabytes
  - This is equivalent to the data that would come from the LHC detectors during 100 years!
- CERN has converted all its base software to 64 bits:
  - ROOT (Data analysis framework)
  - Geant4 (Physics simulation framework)
  - CLHEP (C++ Class Library)
  - CASTOR (CERN Hierarchical Storage Manager)
- LHC experiments are now preparing their entire software environment for this exciting capability



### Moving to openlab-ll (2006 – 2008)

### Intel-related activities

- Performance improvements (hw/sw)
  - Compiler improvement project
  - Tuning of physics applications
- Virtualization
- Power measurements
- Multicore benchmarking (See demo)
- Similar list of activities with the other partners in openlab



# Virtualization

#### Virtualization:

- a technique for hiding the physical characteristics of computing resources from the way in which "clients" interact with those resources
- A great feature that CERN and High Energy Physics can exploit in Grid computing
  - Grids will be both more flexible and more secure when using virtualization
  - Large (multi-core) servers can be partitioned easily
  - ... and it will come at a low cost
- CERN openlab has been involved since 2004
  - Using an open source hypervisor (Xen)
- Now working with CERN's Grid Deployment team
  - Increasing security, flexibility, as well as developers' productivity



CERN

# CERN and our community want:

- Best possible price for a given performance
  - Measured as "SPECint/CHF"
- Lowest possible "heat production"
  - Measured as "SPECint/watt"

 With Intel's multi-core systems, we have seen a considerable improvement in both



### **Multicore comparisons**





14 NOVERIDER 2000

CERN

openlab

•



# From Multi to Many

# Our "high throughput" computing model is maybe ideally suited:

- Independent processes can run on each core, provided that:
  - Main memory is added
  - Bandwidth to main memory remains reasonable
- Testing, so far, has been very convincing
- Today, we are happy to be part of the movement to Quad core







### BACKUP



# **Compiler project**

CERN, and High Energy Physics community, write their own computer programs

- Millions of lines of code
- "Compilers" translate these programs into executable entities
- Improved efficiency of this translation results in faster execution for our users
- Openlab is working on several fronts to improve compilers:
  - Focusing on improving performance for C++ (our programming language for developing applications)
  - Helping ensure correctness of compilers
- This project is active since the beginning of openlab I