

# openlab II technical activity overview

Sverre Jarp CERN openlab CTO  
15 November 2007

PARTNERS



ORACLE®

CONTRIBUTORS

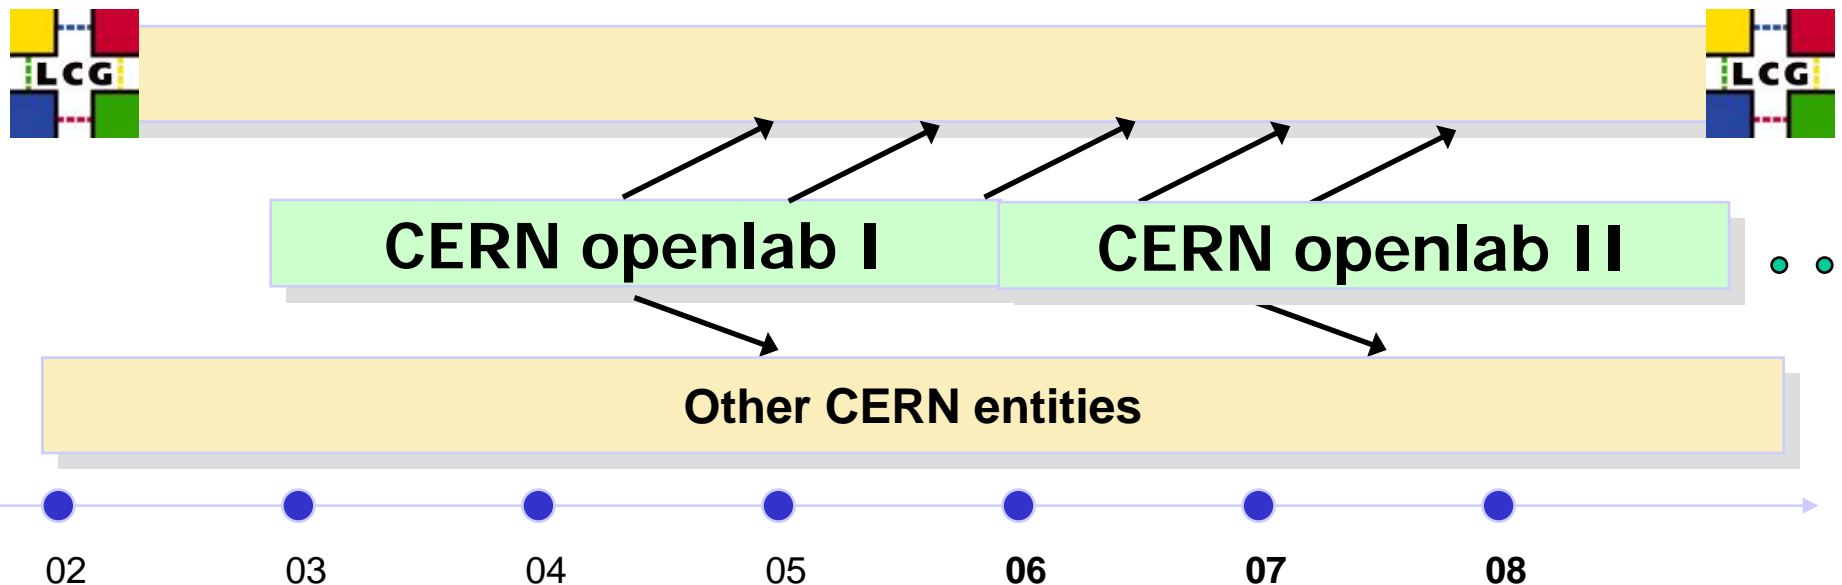


STONESOFT  
Real World Business Security™

EDS

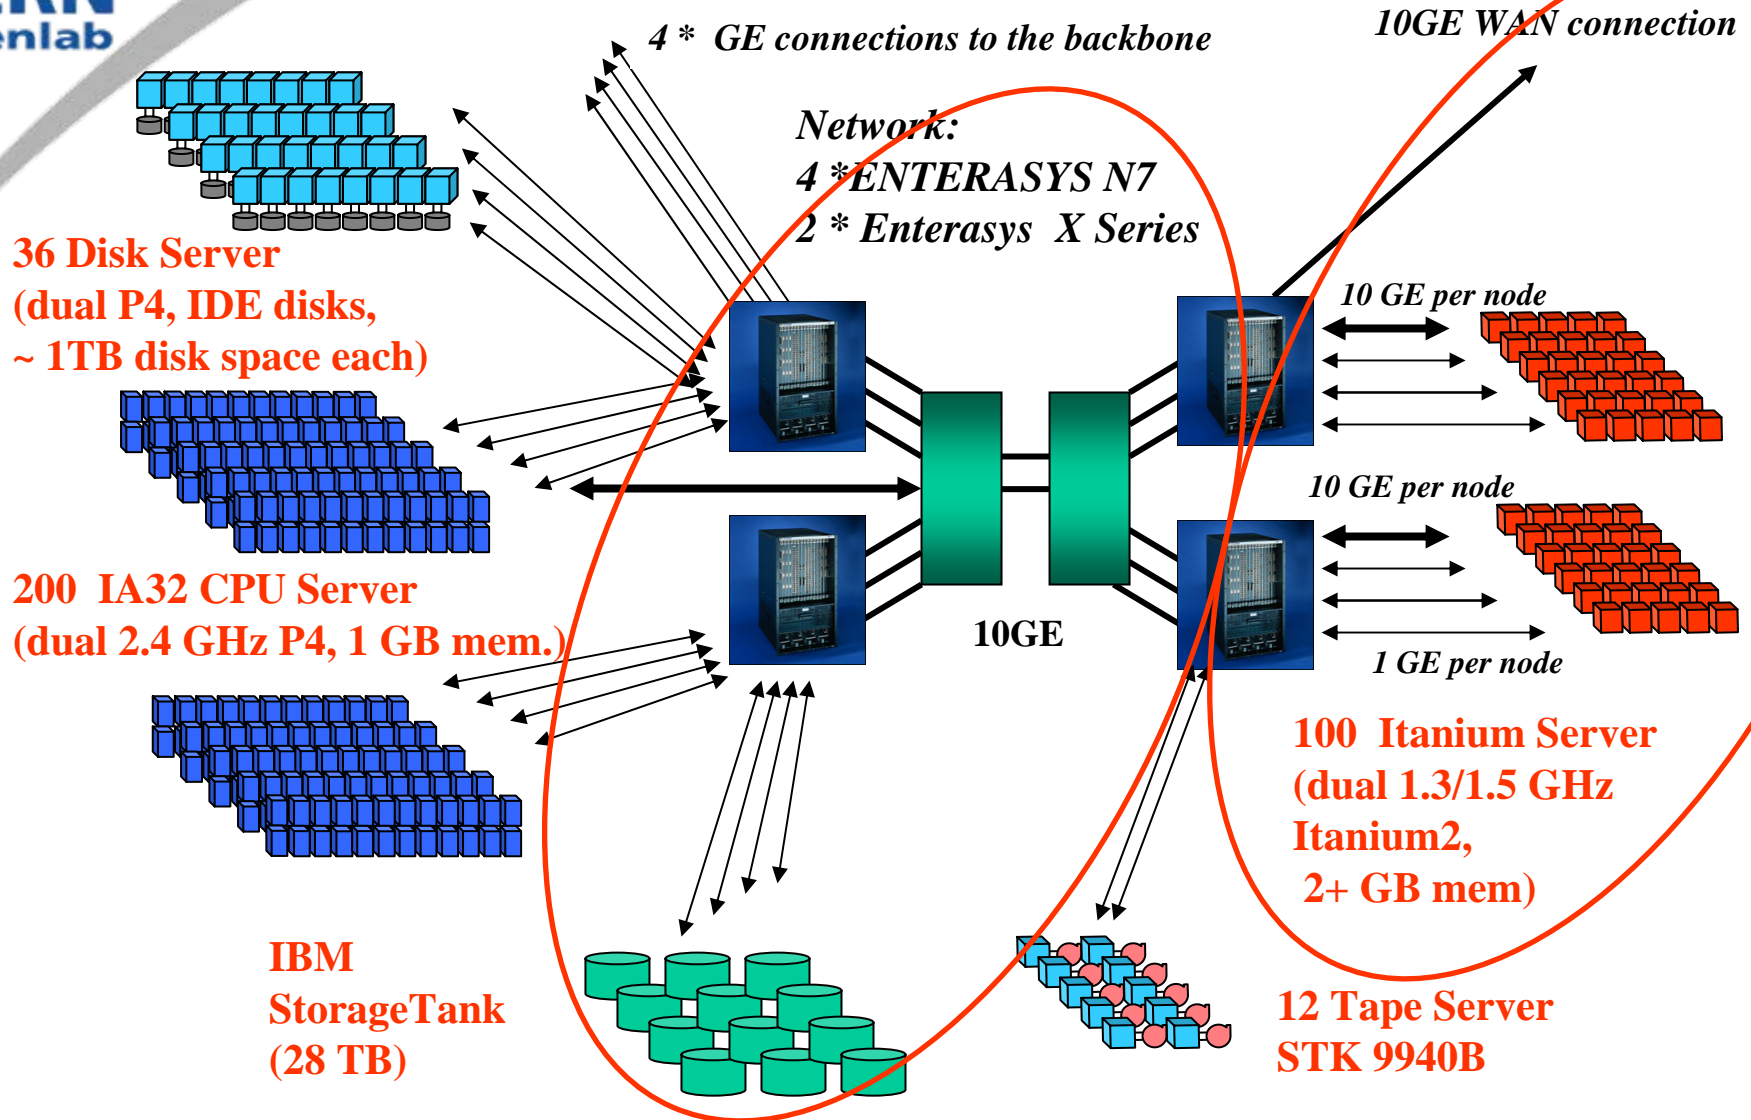
# CERN openlab

- 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:
    - Platform, Grid, databases, Network/Security



# A few highlights from openlab-I


# High performance testbed: the opencluster



High Throughput Prototype (opencluster + LCG testbed)



# 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 production software environment (hundreds of libraries) for this exciting capability**



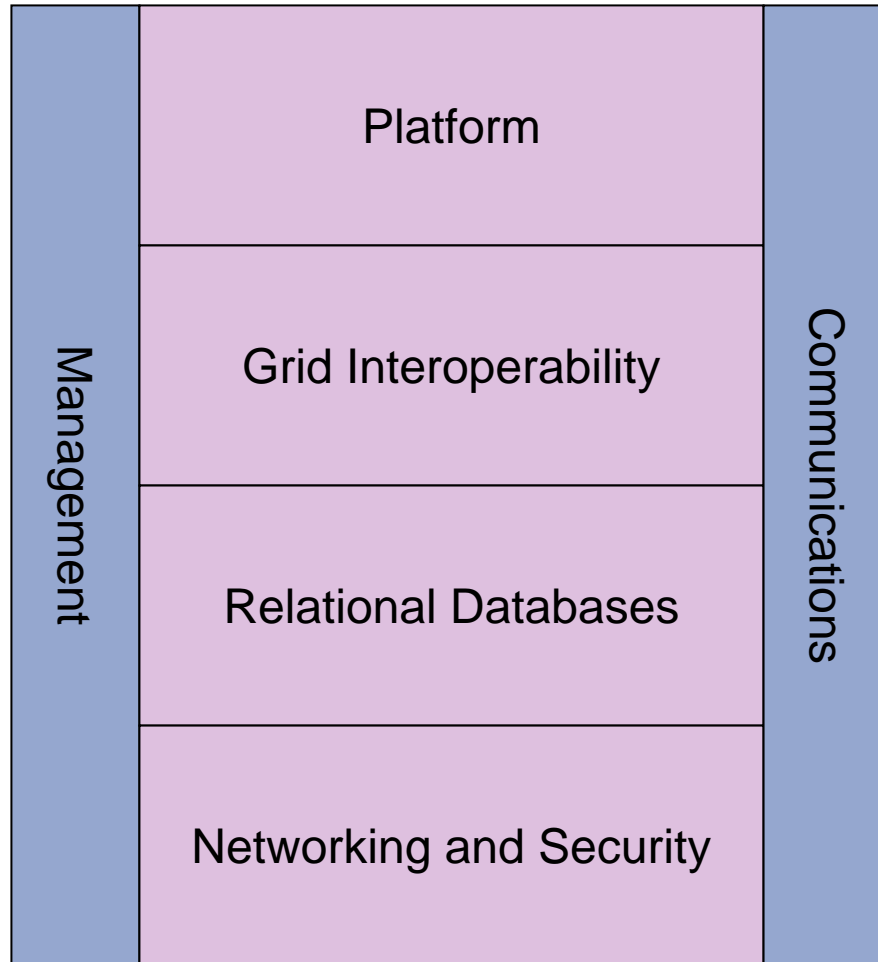
# Computational Fluid Dynamics

- **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



# Benchmarking

- **CERN and our community want:**
  - Best price for a given performance
    - Measured as “SPECint/CHF”
  - Lowest “heat production”
    - Measured as “SPECint/watt”
- **With Intel’s multi-core systems, we have seen a considerable improvement in both**
- **Also, strong effort on optimization of jobs from the LHC collaborations**
  - So far, profiling has been done on:
    - Simulation and reconstruction framework
  - Work on Event Filter in preparation






# Compiler project

- **Since most High Energy Physics programs are written in-house, compiler optimization translates directly into reduced cycle consumption per job**
  - Millions of lines of source code
- **openlab is working on several fronts to improve compilers:**
  - Focusing on improving performance for C++ (preferred programming language for developing applications)
  - Helping ensure correctness of compilers
- **This project is active since the beginning of openlab I**
  - Both on Itanium (IA-64) and Intel-64 (x86-64)
  - Currently working on Intel C++ Compiler V.10
  - Also working with the GNU compiler community

# Multicore/Manycore

- 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
    - Initially on Dual Core systems (Dempsey, Woodcrest, Montecito, etc.)
- Last November we were happy to be part of the first move to Quad core
  - And we are ready to go further!



# 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
- **Local developments:**
  - Osfarm and Content Based Transfers (CBT)

# CERN and Intel




- **CERN is strongly committed to the openlab collaboration:**
  - VIP visits
  - Customer visits, testimonials
  - Strong technical collaboration
    - Multi-core evolution, multithreaded programming
    - Thermal evolution (servers and data centres)
    - Benchmarking
    - Optimization
    - Compilers
    - Virtualization
- **We are looking forward to continued innovation from Intel during the “decade of the LHC”**
  - From now until 2017!



**CERN**  
openlab

# LHC Computing



# BACKUP