Workshop on virtualization and multi-core technologies for LHC

Welcome and Introduction

Pere Mato (CERN)

14-16 April 2008

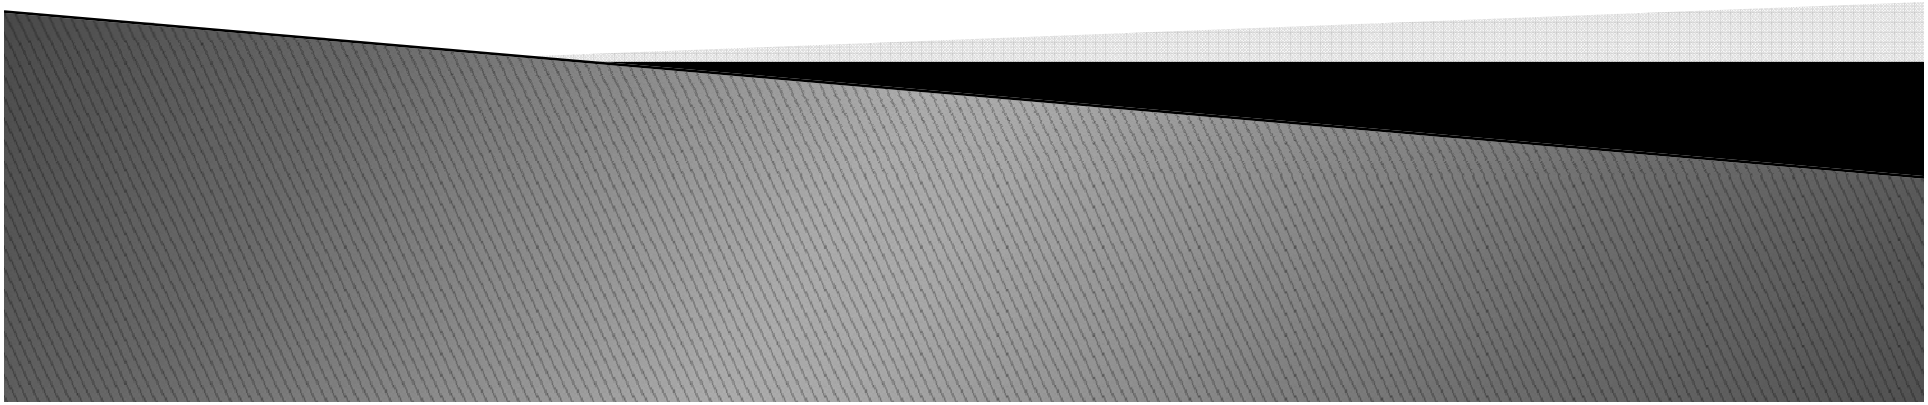

Outline

- Welcome
- History/Origin of R&D Plans
- Motivation and vision for WP8 and WP9
- Goals for this workshop
- Practical Information

Welcome

- Welcome to the kickoff workshop for the new computing R&D projects to exploit the virtualization and multi-core technologies for the particle physics community
- The idea has been to bring together some of the technology vendors, the LHC collaborations and experts on the domain to share their experiences and needs in order to elaborate a possible common program of work for these areas.

History/Origin PH R&D Plan

- During the past few years CERN's resources fully committed to LHC
- In the years 2008-2011 not enough resources for reimbursement of loans, LHC operation and technical R&D
  - PH department R&D request included in CERNCouncil "White Papers", October 2006, covering years **2008-2010**
  - Included in Council/FC documents June 2007, covering years **2008-2011 => additional resources for PH**
- Strategic choice of PH management for initial R&D:
  - "… 9 activities in which a strategic role for CERN has been identified and which concern the development of specific expertise relevant for future projects. Priority has been given to common projects and coordination frameworks, where several experiments will profit from the development…"

Theme 3 R&D work packages

- Microelectronics and optoelectronics
  1. Radiation-hard Technology and Common Building Blocks
  2. On-detector Power Management
  3. Radiation-hard Optical Link for Experiments
- Detector technologies
  4. Radiation Hard Semiconductor Detectors
  5. Micropattern Gas Detectors
  6. Interconnect Technology and Quality Assurance
  7. Facilities and Component Analysis for Detector R&D

Theme 3 R&D work packages (2)

- Physics data analysis, simulation and computing
  8. Parallelization of Software Frameworks to exploit Multi- core Processors
     - Adapt and optimize existing software frameworks in the LHC experiments to run on multi-core processors
  9. Portable Analysis Environment using Virtualization Technology
     - Development of "Virtual Appliances" to provide a complete data analysis environment specialized for each of the LHC experiments

Why these two areas?

- Technologies with potential to be directly applicable to “real” needs of experiments in a short time scale
  - Common to all experiments independently of their software environment
  - Narrowed the scope to make it more effective
- Quite attractive new technologies
  - Attack young people to work in the projects
  - Experience useful in many other domains
- Already some interest in the community
  - Serious possibilities of collaboration

WP8 – Parallelization of Software Frameworks to exploit Multi-core Processors

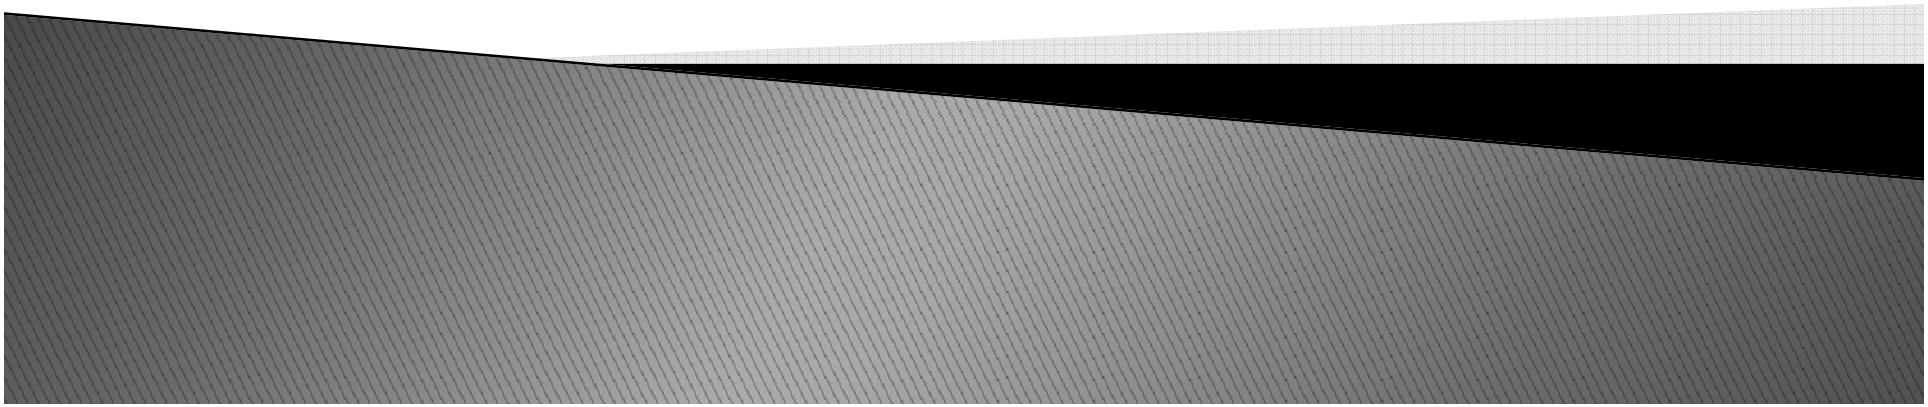

Motivation

- Industry trend to Multi/Many-Core CPU architectures
  - Dual and quad-core already available, 16 cores in the near future
  - Their main motivation is to get a better price/performance/watt
  - Next generation CPUs will be build around tens to hundreds small cores
- Demand for scalable, parallel applications (simulation, reconstruction, analysis) to profit from multi-core CPUs
  - Large architectural impact, different programming model
  - Cannot afford reengineering the totality of LHC experimental software
  - Existing software frameworks would need to be 'adapted'
- Coordinated effort
  - Design patterns and implementation technologies can be applied to common frameworks and/or specific experiment applications
  - Understanding performance issues of physics application software is beneficial to the community

Vision

- Optimal point between 'Task' versus 'Data' parallelization
  - Task parallelization is probably not sufficiently scalable in our environment
  - Data parallelization scales better in our domain
    - E.g. processing in parallel several 'particles' or 'events' or 'event collections (files)'
- Should be possible to provide changes and additions to the LHC application's frameworks such that the algorithmic code is developed as if it was single threaded
  - Hide the complexity of developing and debugging multi-threaded applications
  - Devise mechanisms to keep the **memory requirements** under control
- What we learn for one framework should be applicable to other frameworks
  - E.g. design patterns in CMSSW should be applicable to Gaudi/Athena

WP9 – Portable Analysis Environment using Virtualization Technology

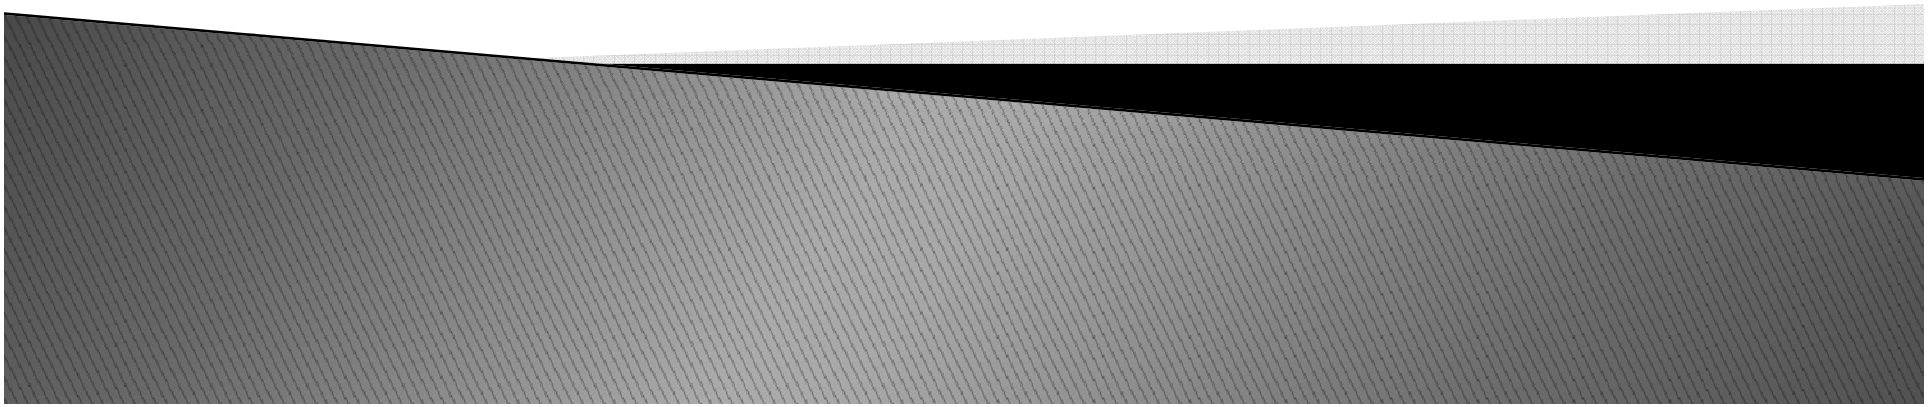

Motivation

- Installing the "complete software environment" in the Physicist's desktop/laptop to be able to do data analysis for any of the LHC experiments is complex and manpower intensive
  - In some cases not even possible if the desktop/laptop OS does not match any of the supported platforms
  - Application software versions change often
  - Only a tiny fraction of the installed software is actually used on most of the cases
- High cost to support large number of compiler-platform combinations
- The system infrastructure cannot evolve independently from the evolution of the application
  - The coupling between OS and application is very strong nowadays and is constraining very much
- These problems are mitigated by using remote login services such as LXPLUS and distributed file systems such as AFS
  - But the price to pay is a stricter usage rules, wasted local resources, competing for common resources, work only while connected, etc.

Vision

- Provide a minimal OS (<100 MB) and a set of basic libraries and tools packaged as a "virtual appliance" with sufficient functionality to support any analysis activity on any host platform (Linux, Mac, Windows)
  - Easy to install
  - Specialized to each LHC experiment
  - Edit, compile, run, grid-submit, interactive analysis, etc
- Fast changing software could be accessed transparently by means of a "read-only network file system" with aggressive caching schema
  - Only needed files would be copied
  - Continue to work offline

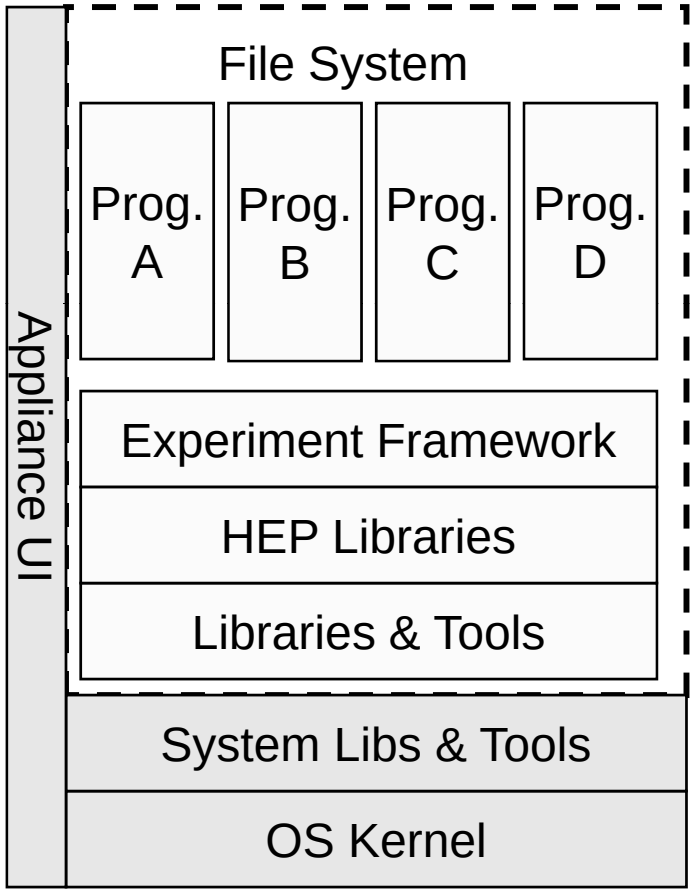

Resources & Organization

- Multi-Core
  - Coordinated by VincenzoInnocente
  - From PH: 6 FTE-year (staff) + 6.5 FTE-year (fellow) over 3-4 years
- Virtualization
  - Coordinated by Predrag Buncic
  - From PH: 7.5 FTE-year (staff) + 7 FTE-year (fellow) over 3-4 years
- Both projects are hosted in PH-SFT with participation of LHC experiments
  - Open participation to anybody interested
- Regular monitoring and control by the LHC experiments via the Architect's Forum (AF)
  - Same channel that is used for the LCG-AA projects (ROOT, Geant4, etc.)
- Consultation and collaboration with IT department to ensure integration with IT services
  - Privileged relations with OpenLab and other groups (GD, CS,…)

Aims for this workshop

- Make these R&D plans known
  - Share existing similar initiatives in the community
  - Identify possible collaborations between individuals/teams in experiments, vendors and CERN groups
- Collect feedback/suggestions from LHC experiments
  - Address real problems
  - Use available resources at best
- Define initial program of work for both projects
  - Better define their scope and a possible set of initial deliverables

Agenda of the Workshop

- Monday 14 April 2008
  - 14:00->18:00 Welcome and Technology session
- Tuesday 15 April 2008
  - 09:00->13:00 Experience and plans on adapting software for multi-core
  - 14:00->18:00 Virtualization in Experiments
- Wednesday 16 April 2008
  - 09:00->13:00 Hands on and Demo session
  - 14:00->18:00 Discussion and Sum up session

Practical Information

- Presentation slides
  - Uploaded into agenda
  - Speakers please send them to one of the chairpersons
- EVO conference
  - In "Universe" community. Possible to phone in.
- Coffee breaks
  - Coffees will be available in the B. 40 main hall
- Dinner Tuesday evening
  - Possibility to organize informal dinner in a nearby Restaurant (Meyrin)
  - If interested send mail to Nathalie.Knoors@cern.ch today