
Applications Area Status

WLCG-LHCC Review
16th February 2010

Andrea Valassi (CERN/IT-ES)

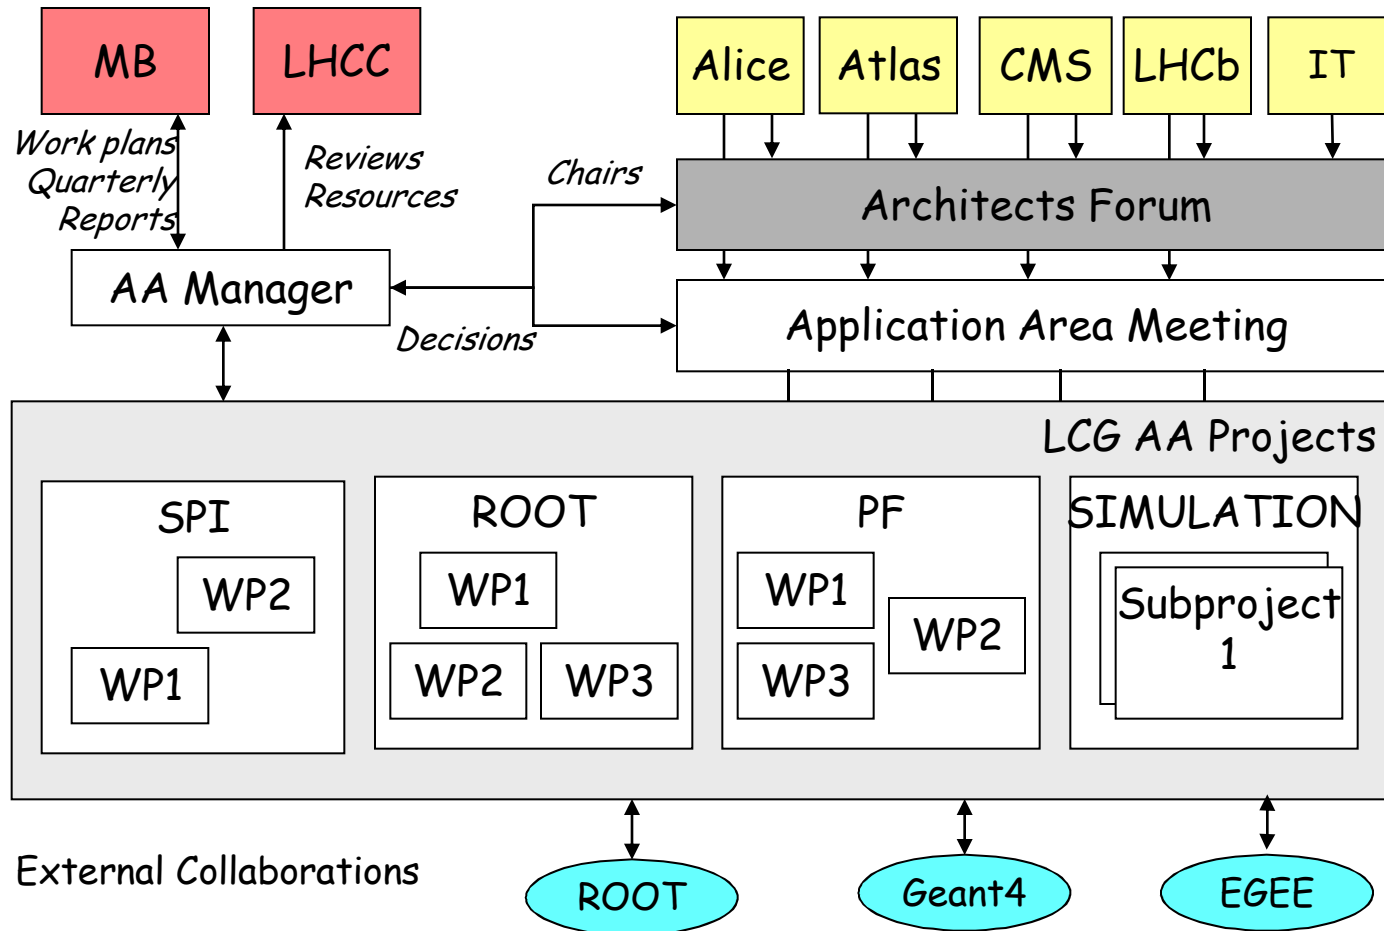
Thanks to Pere Mato and the AA project leaders for their contributions to these slides



Outline

- ◆ Applications Area Overview
- ◆ Status and Plans of the Projects
- ◆ Resources
- ◆ Summary

Applications Area Organization



AA Projects


- ◆ SPI - Software process infrastructure (S. Roiser)
 - Software and development services: external libraries, savannah, software distribution, support for builds, tests, QA, etc.
- ◆ ROOT - Core Libraries and Services (R. Brun)
 - Foundation class libraries, I/O libraries, math libraries, framework services, dictionaries, scripting, GUI, graphics, data analysis, etc.
- ◆ PF - Persistency Framework (A. Valassi)
 - Storage manager, file catalogs, event collections, relational access layer, conditions database, etc.
- ◆ SIMU - Simulation project (G. Cosmo)
 - Simulation framework, physics validation studies, MC event generators, participation in Geant4, Fluka.
- ◆ PH White Paper R&D Projects
 - Exploit **Multi-core** architectures (V. Innocente)
 - Portable Analysis environment using **Virtualization** (P. Buncic)

Unchanged Execution Approach

- ◆ Architects Forum (AF) meetings
 - Decision and action taking meetings
 - » Experiments participate directly in the planning, management, and architectural and technical direction of AA activities
 - Every two weeks, public minutes after internal circulation
 - » <http://lcgapp.cern.ch/project/mgmt/af.html>
- ◆ Application Area meetings
 - Informal forum of exchange of information between the AA projects and experiments, etc.
 - Lower frequency than before
- ◆ Workplans, quarterly reports and reviews
 - Discussed and agreed in the AF

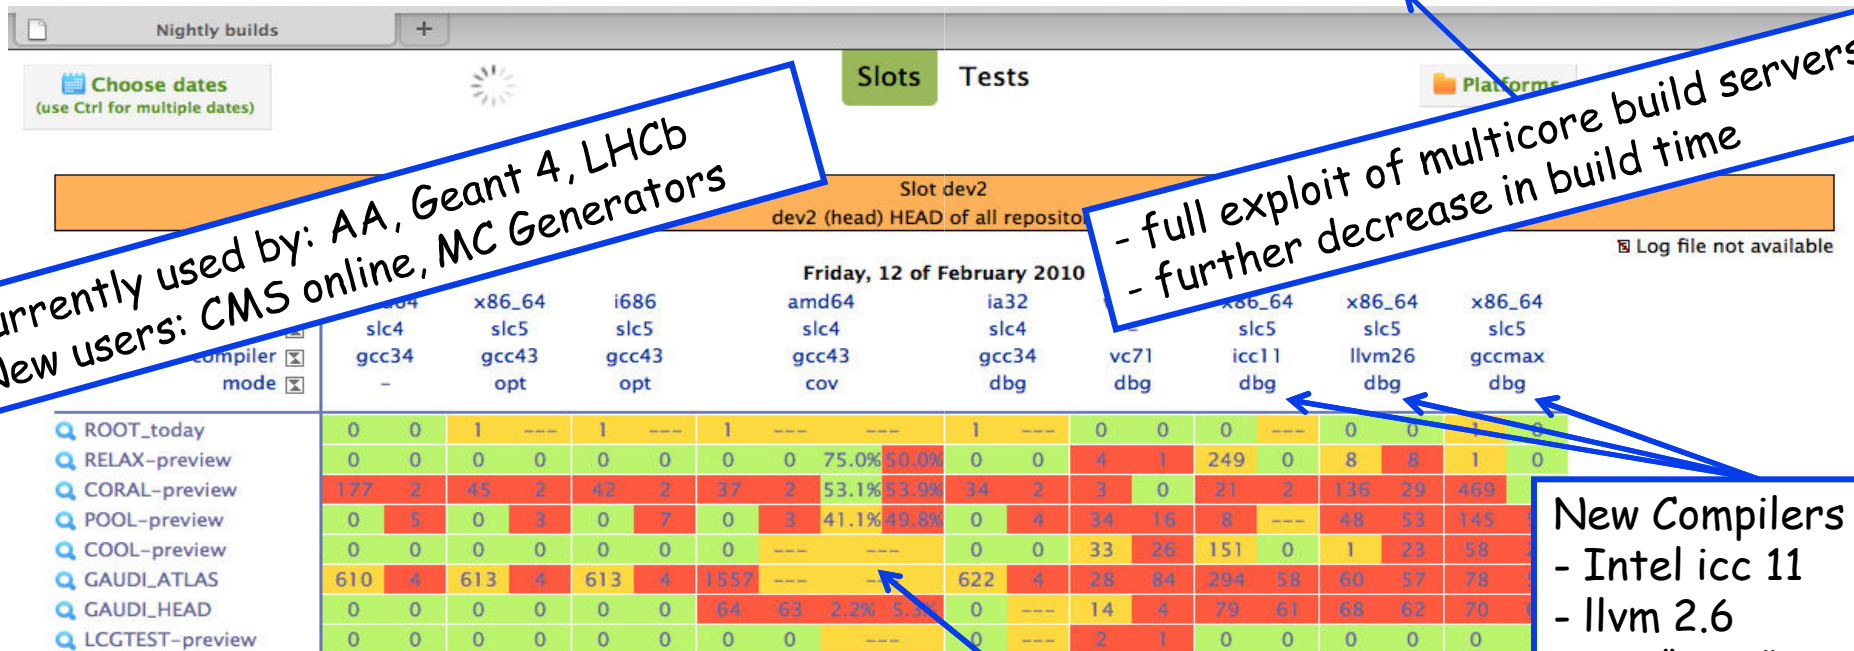
SPI Release Management for LHC

- ◆ Since February 2009
 - 3 major releases LCG 56,57,58 (10 in total) - i.e. major changes in sw stack
 - Overall 123 version upgrades of AA packages
- ◆ Migrated to slc5 / gcc 4.3.2 / i686+x86_64
- ◆ In total 20 different "platforms" are currently supported by SPI
 - OS version * compiler * architectures * optimization
 - Currently introducing MacOSX 10.6 (32 and 64 bit)
- ◆ Optimizing the release process
 - Resulted in significant reduction of release time (~3 hours)



AA Nightly Build Service

- New version of nightly build scripts was put in production
 - Client - Server architecture
 - Multi-threaded builds on several layers (project / package level)



- Currently used by: AA, Geant 4, LHCb
- New users: CMS online, MC Generators

- full exploit of multicore build servers
- further decrease in build time

New overview web page

Code coverage testing
- Provide details which code is tested
- Very much appreciated by AA projects

New Compilers
- Intel icc 11
- llvm 2.6
- gcc "max"
(snapshot from gcc code repository)

Infrastructure

◆ Opengrok

- A new tool for source code cross referencing
- Replacing LXR because of better performance and maintainability

◆ Hypernews phased out

- Instances of AA, Atlas, Alice and LHCb moved to IT provided Sharepoint
- CMS & Totem continue but maintain themselves

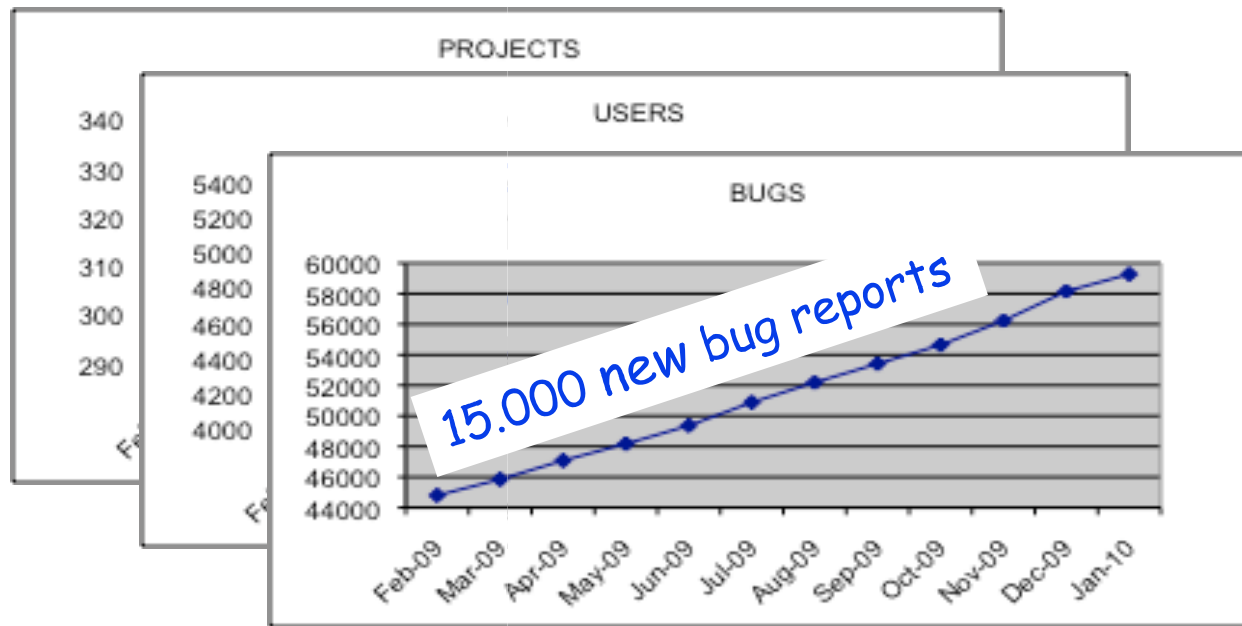
◆ Web infrastructure with Drupal

- Started new effort to provide a common web infrastructure for PH/SFT and LCG/AA projects and work packages

The screenshot shows a web browser window displaying the LCG Software Process and Infrastructure website. The address bar shows the URL <http://cern.ch/lcgsource>. The page content includes a search bar, a navigation menu, and a main article titled "New tool for source code cross referencing" with a sub-headline "LCG 58 has been released". A magnifying glass icon is overlaid on the URL in the address bar. The footer includes the Drupal logo and copyright information for 2010.

Savannah

- ◆ Standard tool for issue tracking in LHC
 - Currently being moved to new HW/OS (slc5)
 - » Including integration with CERN SSO (single sign on)
 - » Hot standby machine available in case of incident
- ◆ Last year :

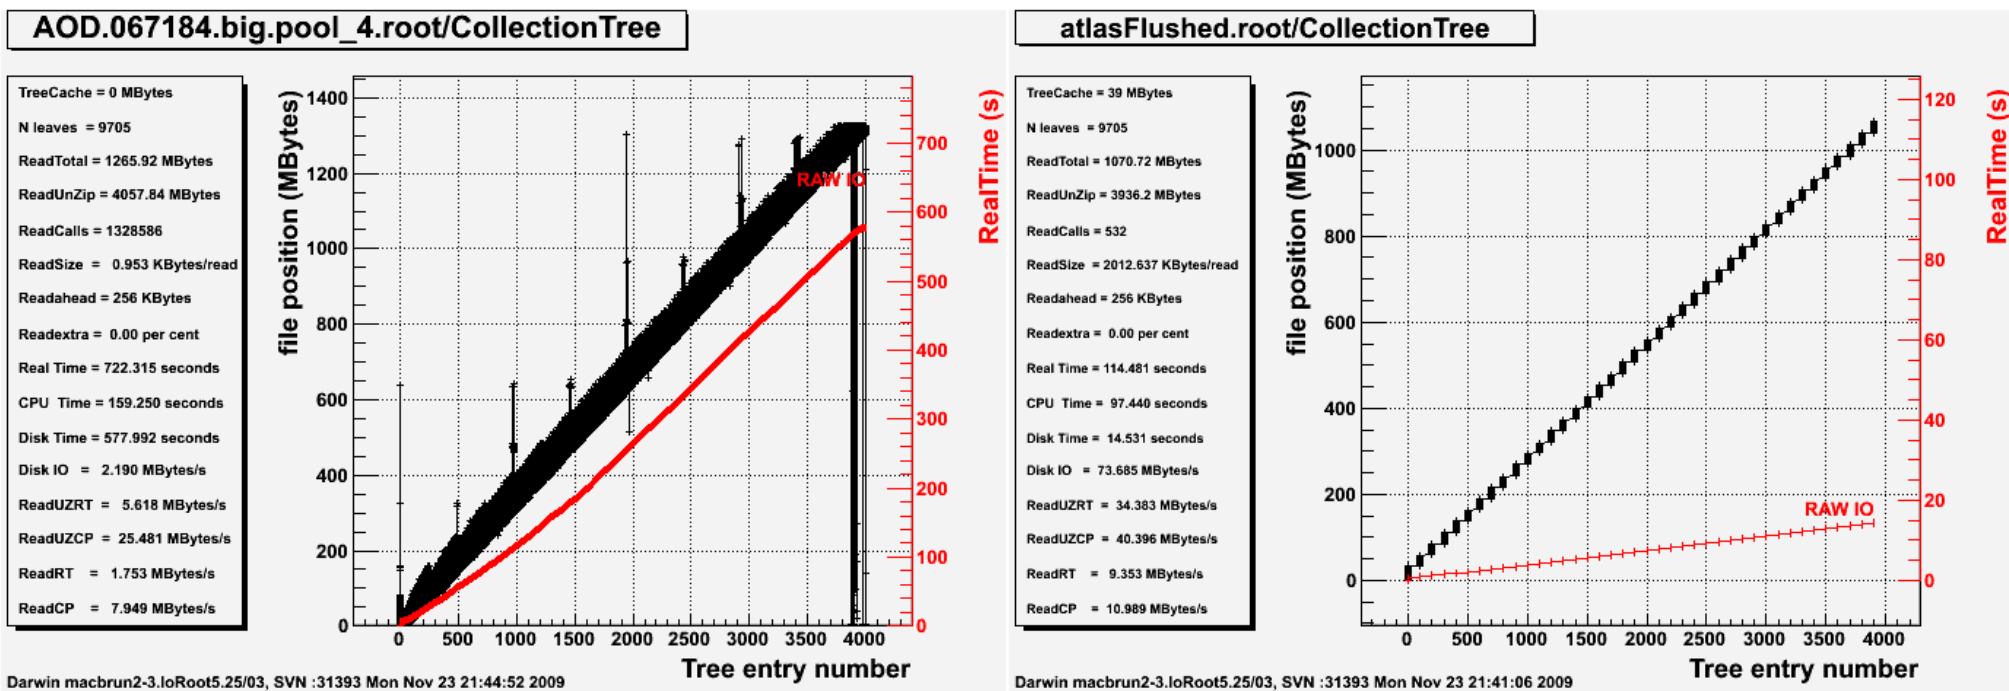


ROOT Highlights

- ◆ Two major releases of ROOT in 2009 (5.24 in June and 5.26 in December).
 - Version 5.26 used by Alice and LHCb in February 2010
 - CMS planning to move from "summer" onwards
- ◆ Version 5.26 includes a long list of new features and improvements (see [release notes](#))
 - It has been ported to new OS (SLC5, Ubuntu, Debian, MACOS10.6, Windows7) with the latest generation of compilers.
- ◆ The test suite including now several thousand test programs grouped in 113 super groups
 - It is now run nightly and is a big help for the general QA of the system

ROOT - I/O Read Performance

- ◆ Improved performance for the data analysis of the LHC experiments
- ◆ New class TTreePerfStats to monitor the various IO parameters



Before (ATLAS AOD):

- Many non-monotonic seeks
- 1.7 MB/s

After several optimizations:

- Only monotonic seeks
- 9.3 MB/s



ROOT - Details (1)

◆ CINT

- Several Improvements in CINT/Reflex/Cintex
 - » delayed loading of dictionaries
 - » gains in memory, e.g. 75 MB gained in case of CMS
- Evaluating the LLVM system (gcc compatible compiler) as a possible replacement of the CINT interpreter with a JIT (Just In Time) compiler with the advantage that it is designed to support the latest C++0x language.

◆ PROOF

- With the rapidly increasing number of users more facilities have been developed to support disk and CPU quota management systems
- The performance monitoring has been improved.
- PROOF-LITE has been released. It is designed to take advantage of multi-core systems on desktops/laptops.

ROOT - Details (2)

◆ Math

- Many developments to simplify the user interface for fitting
- New versions of the RooFit and TMVA packages introduced with their long list of improvements and new features

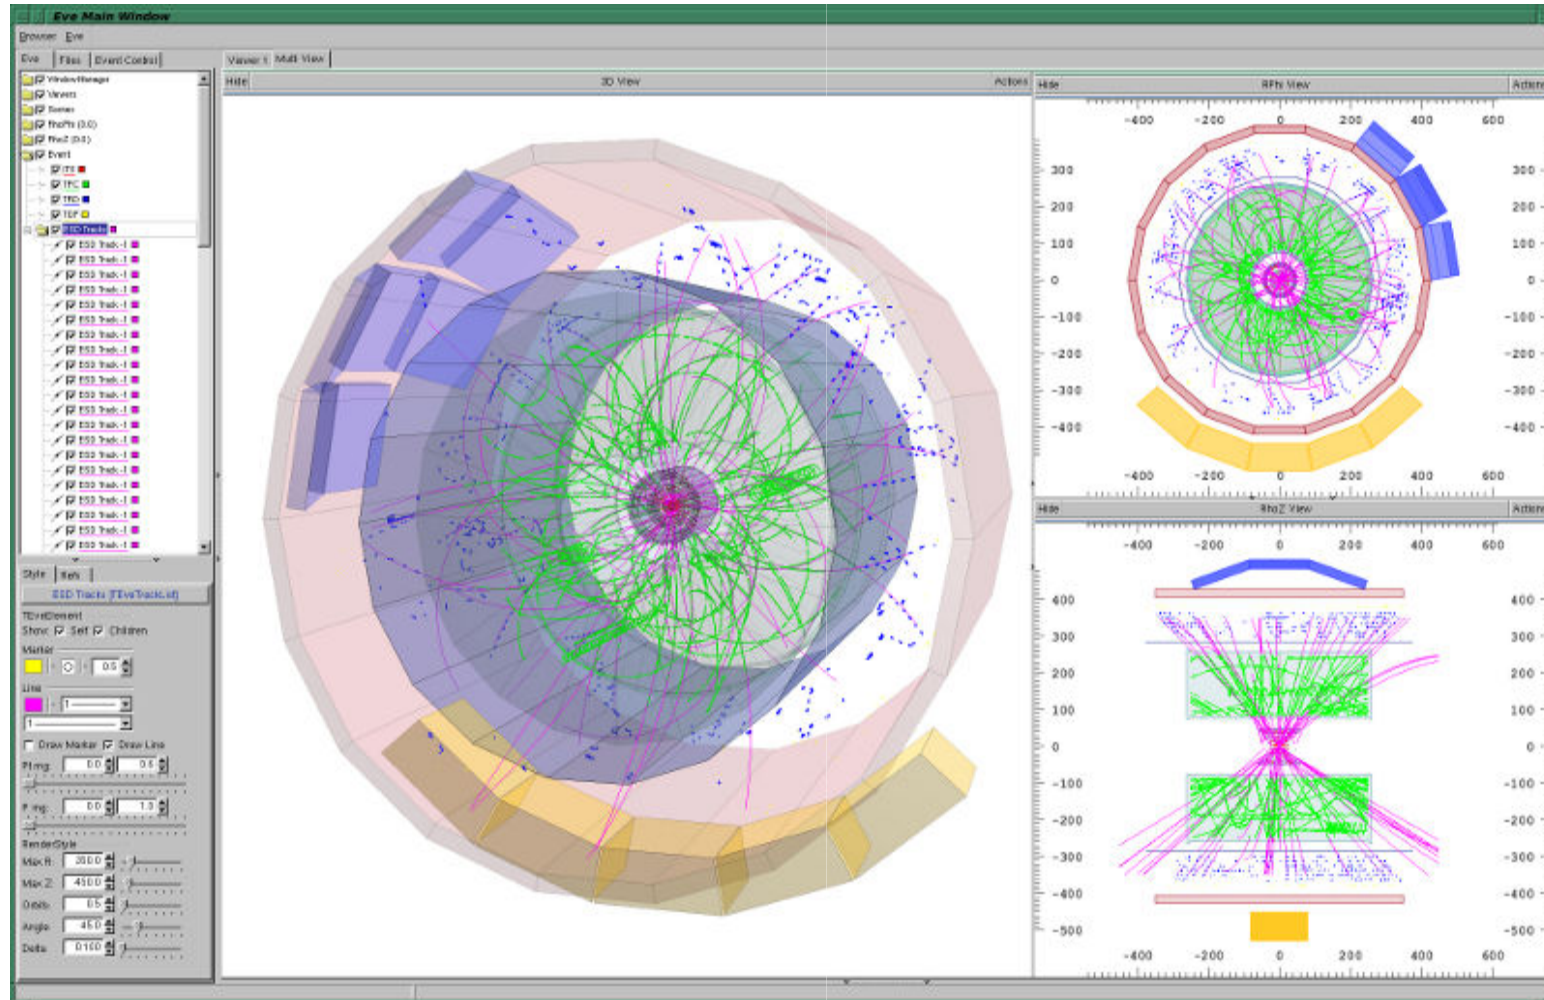
◆ GUI

- The GUI builder has been extended and consolidated
- A GUI recorder implemented to capture an interactive session in a standard ROOT file and replay it later. Used QA now.

◆ Graphics

- A long list of improvements suggested by users to improve the visual representation on the screen or ps, eps, pdf files.
- A new major version of the Event display package EVE is included in the release. Used by event displays of Alice and CMS.

ROOT - EVE based Event Display

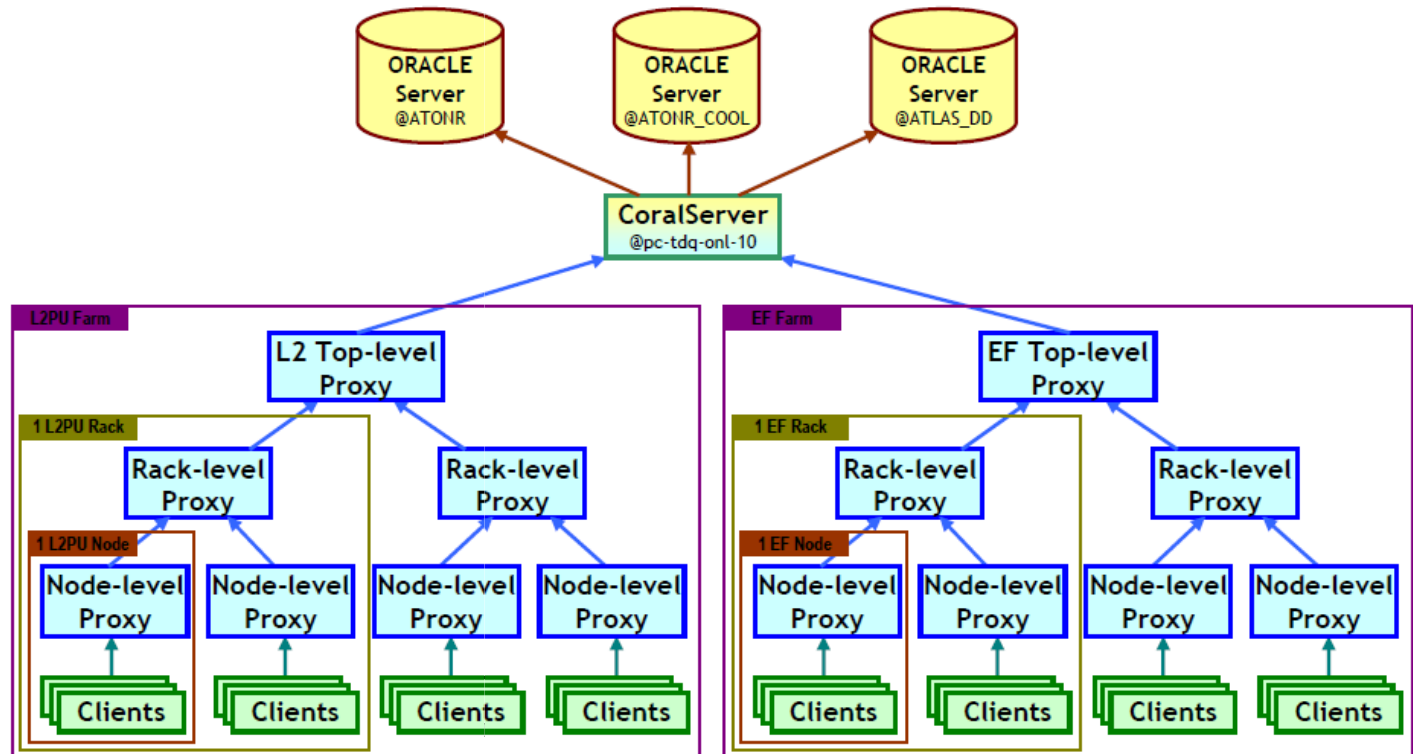


Persistency Framework

- ◆ CORAL relational database access (ATLAS, CMS, LHCb)
 - Most active development area: CORAL server project
- ◆ COOL conditions database (ATLAS, LHCb)
 - Further performance consolidations (11g server validation, CLOBs), partitioning tests, consultancy for Pvss2Cool and ATLAS Frontier
- ◆ POOL object streaming and catalogs (ATLAS, CMS, LHCb)
 - Collection improvements, fast file merge
 - Still active although in "maintenance mode" since summer 2008
- ◆ Large load from user support and software maintenance (50% of project resources)
 - Releases with external upgrades (ROOT) and new platforms (icc)
 - Service operation issues (reconnect after glitch, LFC replica svc...)
 - Oracle client software issues (SELinux, OCI bugs...)

CORAL server

- ◆ In production for ATLAS online since October 2009
 - Only 9 months since restart of development with a new design
 - Deployment and validation at Point1 was extremely smooth
 - Read-only functionality, data caching in CORAL server proxy



Persistency Framework plans for 2010

◆ CORAL and COOL

- CORAL server and CORAL server proxy
 - » Add monitoring for ATLAS online use case
 - » Complete secure authentication and add read-write for offline
- Review read-only transactions (find data added after start of tx)
- Continue support for service operation (user tickets, Oracle client)
 - » Aim for offline CORAL server deployment - also needs support

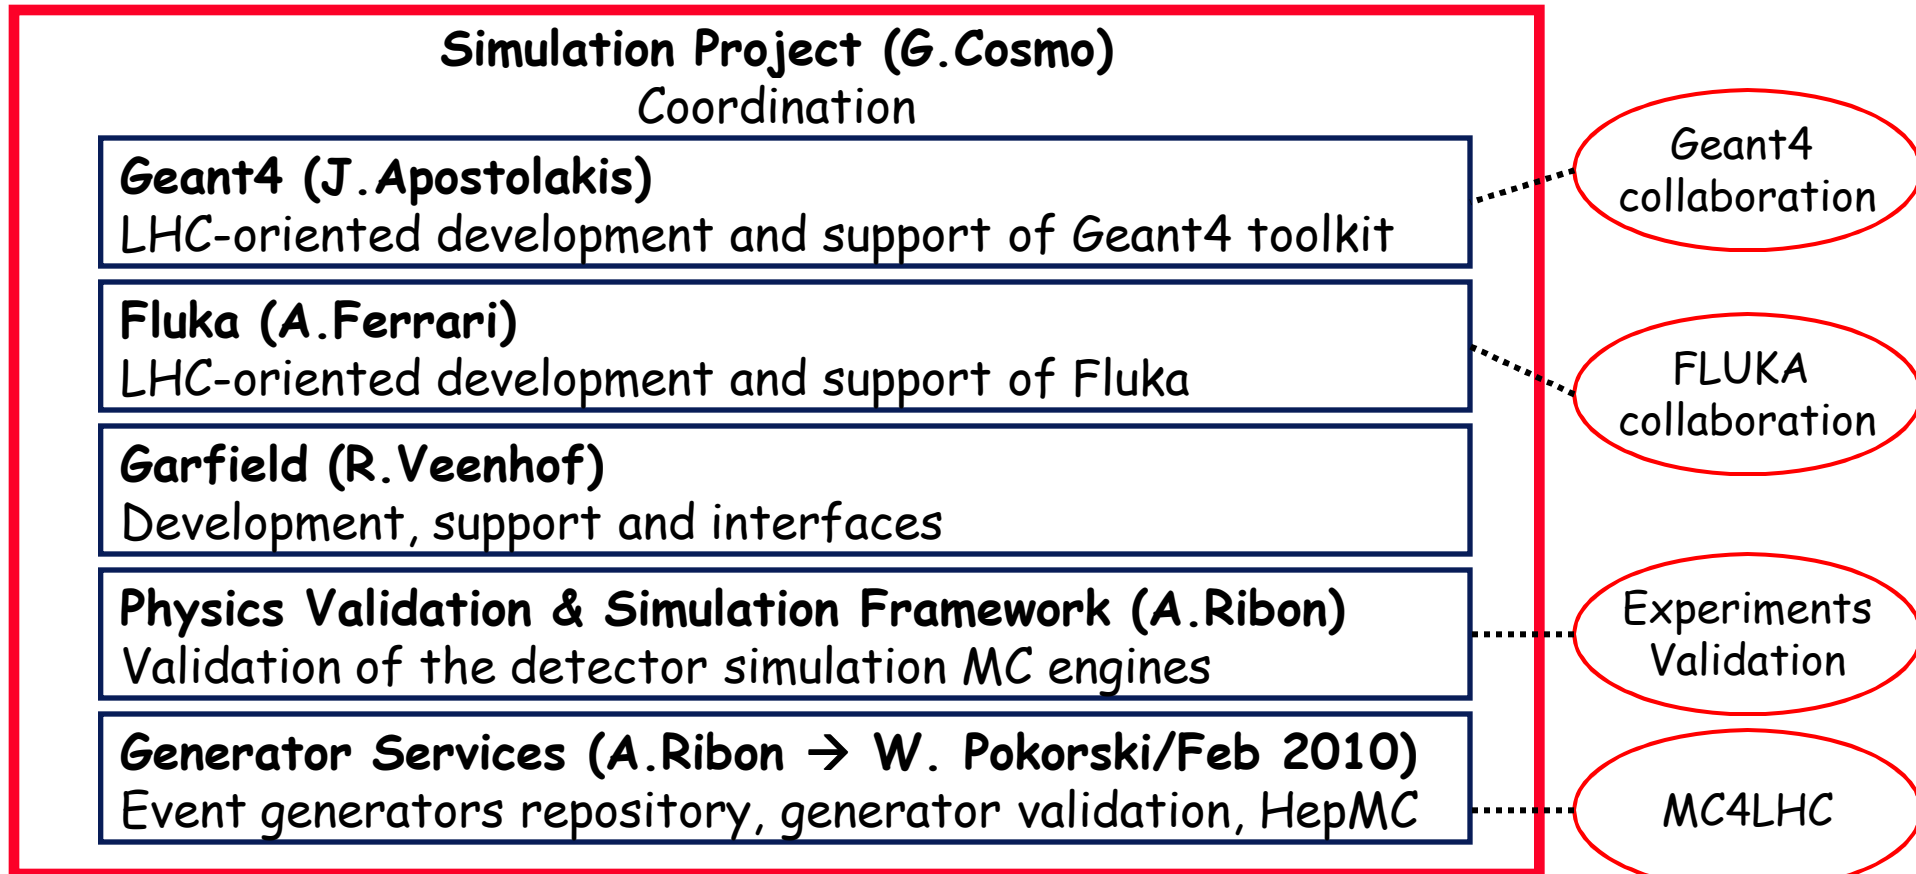
◆ Review support for POOL with IT and the experiments

- Consider splitting POOL into components and moving them out of the Persistency Framework into Gaudi/Athena and CMSSW

◆ All Persistency Framework projects

- Move code repositories from CVS to SVN (postponed from 2009)
- Maintenance with new externals and platforms (llvm, osx106...)

Simulation Project - Structure



Geant4 - Hadronic Physics Progress

- ◆ Release 9.3. (December 2009)
- ◆ Improvement of Fritiof-based FTF model
 - Tuning for energies 3-8 GeV
 - Added Reggeon cascade - as back-end
 - Better matching with Bertini
- ◆ New **CHIPS** hadronic interaction (1st release)
 - Use single model for all Energies => smooth response
 - Improvement and parameter tuning underway
- ◆ Corrected internal cross sections in BERTini
 - At energies below 1 GeV (SLAC Geant4 team)
- ◆ Retuned Multiple Scattering for e^+/e^- (optional)

Geant4 - Performance and Support

◆ Computing Performance

- Reduced number of memory allocations
 - » In geometry, hadronics (reports of ATLAS, CMS)
- Reduced number of costly magnetic field calls
 - » Adapting stepper created by ATLAS

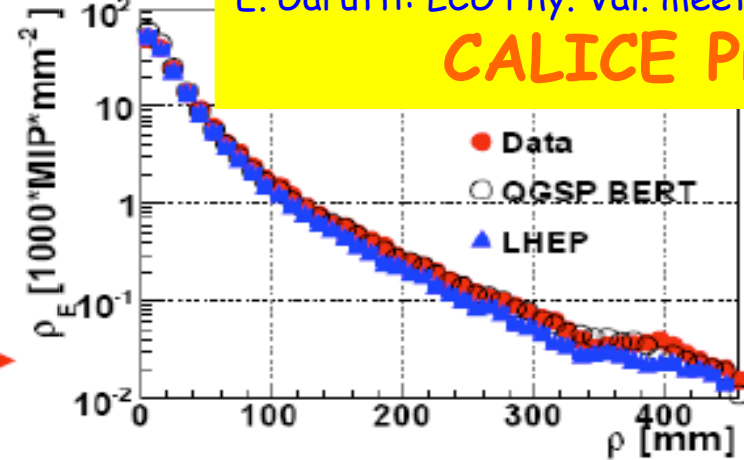
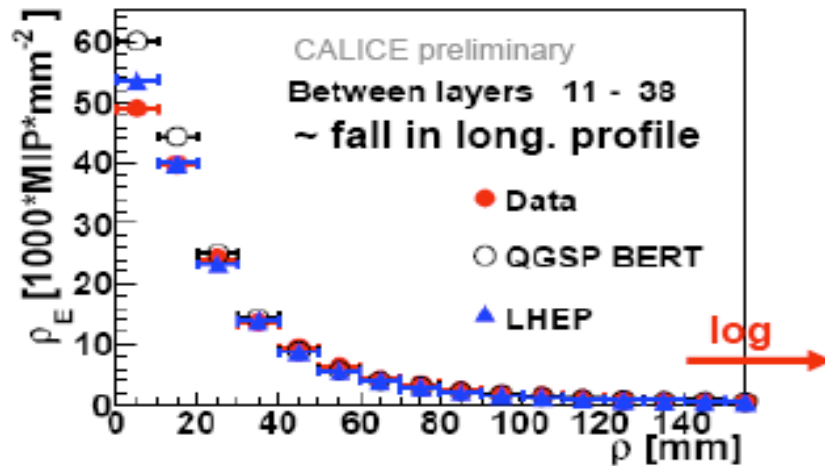
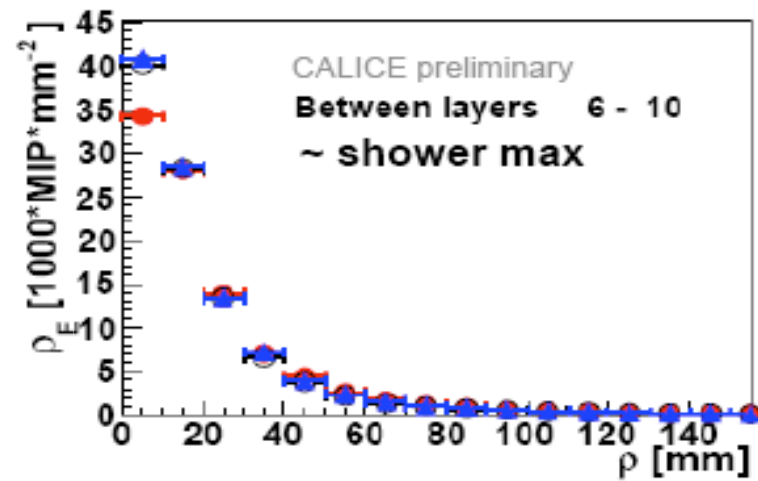
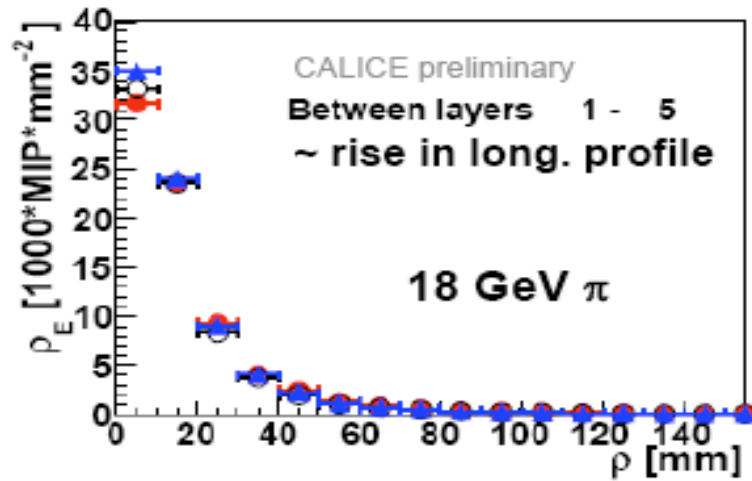
◆ Support and patching

- Support for issues in LHC experiment productions
 - » patches
- Will provide patches for 9.2 until December 2010
 - » Fixes
 - » Some CPU improvements

Physics Validation

- ◆ Main effort to improve the transition between hadronic models in Geant4
 - Physics Lists with different transition regions: e.g. **FTFP_BERT_TRV** (as **FTFP_BERT** but with transition between Fritiof and Bertini between 6 - 8 GeV, instead of 4 - 5 GeV)
 - New combinations of models: e.g. **QGSP_FTFP_BERT** (like **QGSP_BERT**, but with Fritiof replacing LEP, transition 6-8 GeV)
 - Improvement of Geant4 models (next text slide)
- ◆ Regular feedback from LHC test-beam analyses
 - **ATLAS HEC, ATLAS TileCal, CMS HCAL**
- ◆ New validation results from CALICE test-beam
 - **Very granular calorimeter sensitive to details of hadronic showers**

Lateral Shower Profile vs. Depth



E. Garutti: LCG Phy. Val. meet. 10 June 2009
CALICE Preliminary

LCG Generator Services

◆ GENSER

- Structure stable and used by experiments
- 26 generators installed (most with different versions)
- Autotools ready for all generators. "Bootstrap" script ready for building GENSER more easily and everywhere
- Completed migration to SLC5
- Adopted HepMC Analysis Tool for regression testing based on distributions; evaluation of Rivet also for regression testing

◆ HepMC

- New release 2.06.00 in progress, expected in March
- Nightly builds for testing HepMC in preparation

◆ MCDB

- In use for CMS productions

WP8 - Multi core R&D

- ◆ Investigate solutions to parallelize application frameworks and also investigate solutions to parallelize algorithms
- ◆ Status
 - ATLAS and CMS developing multi-process versions of their event-processing frameworks that make use of copy-on-write mechanism
 - A parallel version of GAUDI was developed in the project and will be soon released to the LHCb collaboration
 - In collaboration with IT investigations and tests have started to understand how to best run parallel applications
 - CMS and the Gaudi frameworks instrumented with a "hardware" performance monitor
 - On going work to speed up tracking algorithm in CMS exploiting vectorization technology

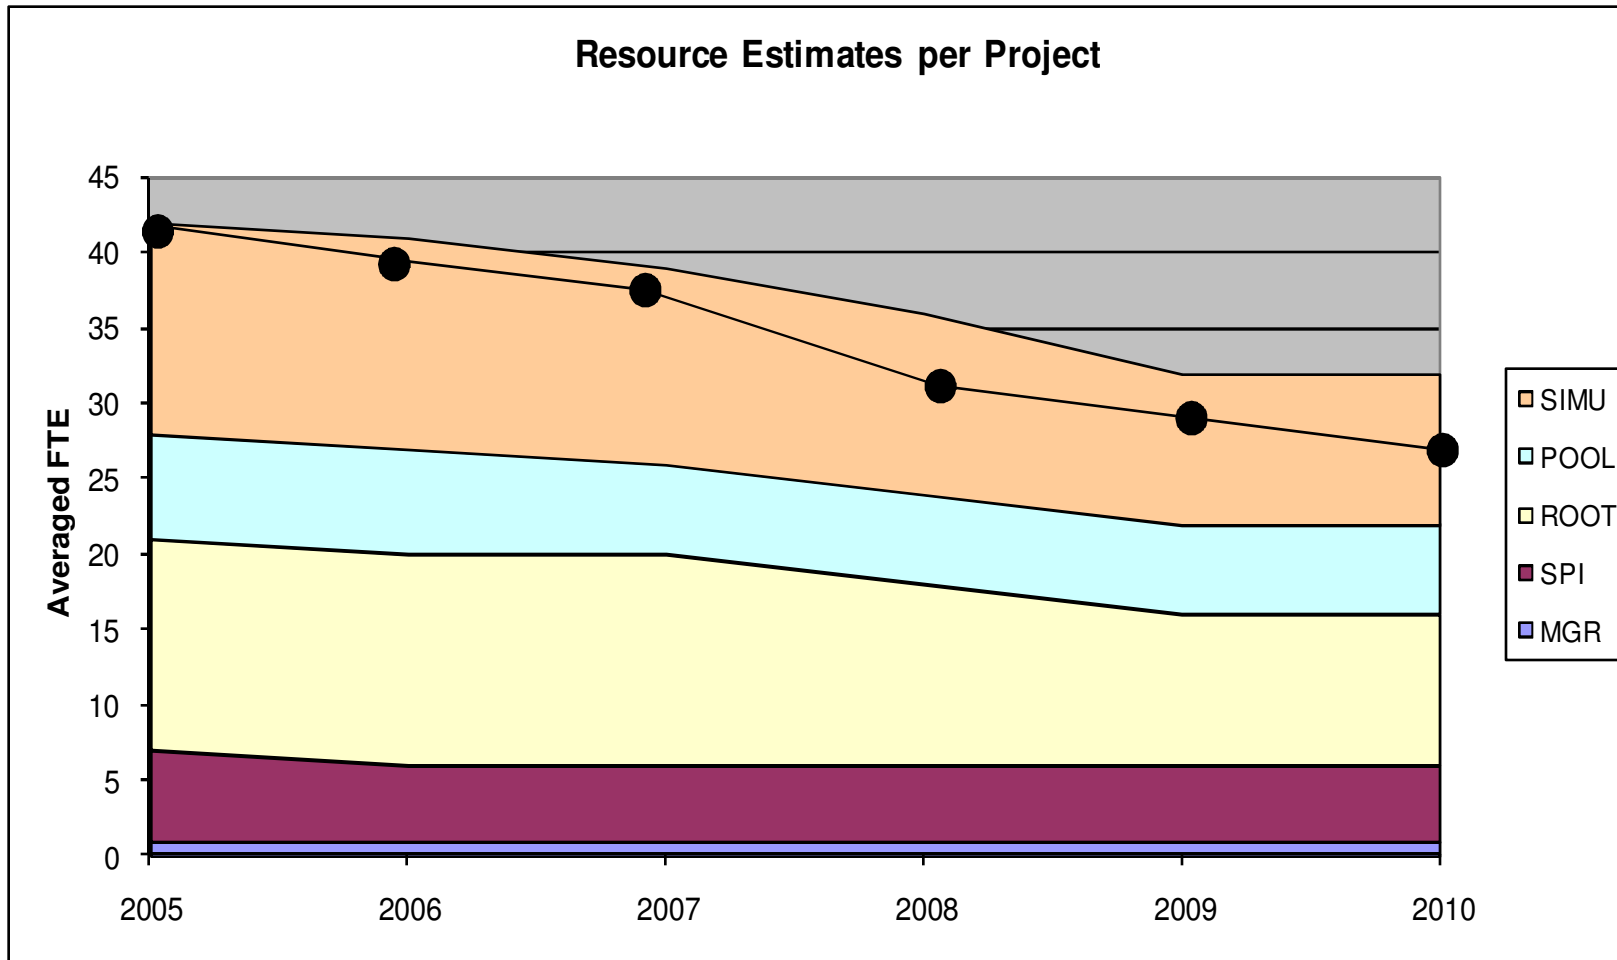
WP9 - Virtualization R&D

- ◆ Project delivers *CernVM*, a Virtual Software Appliance that provides a complete user environment for developing and running LHC applications
 - *CernVM* consists of minimal operating system and dedicated file system for distribution of experiment software (CVMFS)
- ◆ Status
 - Developing *CernVM* to meet requirements of the LHC experiments and assure scalability and reliability of the software distribution
 - Investigating possible deployment of *CernVM* on managed cloud infrastructure (e.g. Amazon EC2) as well as unmanaged clouds (e.g. BOINC for computing using volunteered resources)
 - Two production releases and many development releases
 - » Most of users from ATLAS and LHCb but also support ALICE, CMS, LCD and NA61

Manpower Table (excluding R&D)

Project	Sub-Project	ALICE	ATLAS	CERN	CMS	FNAL	LCG- Russia	LHCb	Grand Total
MGR	(blank)			0.9					0.9
MGR Total				0.9					0.9
PF	Cool		0.2	0.6				0.2	1.0
	Coral		0.5	1.3	0.2				2.0
	Pool		0.4	0	0.2			0.1	0.7
PF Total			1.1	1.9	0.4			0.3	3.7
ROOT	Base			0.6		0.1			0.7
	Dictionary		0.1	1		0.2			1.3
	Geom	0.2							0.2
	Graf	0.1		1.5					1.6
	Gui			0.6					0.6
	I/O					0.7			0.7
	Math			1					1
	Mgr			1					1
	Proof			1.4					1.4
ROOT Total		0.3	0.1	7.1		1			8.5
SIMU	Garfield			0.2					0.2
	Geant4			5.7					5.7
	Genser			0.7		0.1	1.6		2.4
	Mgr			0.25					0.25
	Validation			0.8					0.8
SIMU Total				7.65		0.1	1.6		9.35
SPI	(blank)			3.1					3.1
SPI Total				3.1					3.1
Grand Total		0.3	1.2	20.65	0.4	1.1	1.6	0.3	25.55

Estimated Resource Needs



Staffing Remarks

- ◆ Total reduction in 2010 is more than expected
 - 2009 was about 29 FTE, now is about 25.5 FTE
 - A couple FTE could be recovered later in the year (Fellow + PJAS)
- ◆ Persistency Framework
 - POOL coordinator in IT has moved and has not yet been replaced
- ◆ SPI
 - The complete team (3 FTE) will be renewed in 2010 (one retirement, departures of one fellow and one LD staff)

Summary

- ◆ Applications Area has continued to provide the software and services required by the experiments
 - The provided software did work reasonable well for the first data
- ◆ Well oiled software release procedures
 - Good response time to requests from experiments
 - Software releases (patches) done very quickly (hours)
- ◆ Entering a new phase
 - The challenge of evolving the products while keeping them stable
- ◆ Available manpower getting lower
 - Indication that some projects are entering maintenance mode
- ◆ The two R&D work packages producing nice results
 - Good collaboration between experiments