ALICE Offline

Y. Schutz

CERN, November 17, 2003

Menu

- Organisation of ALICE Off-line
  - CORE Offline
  - Detector projects
- ALICE Framework
  - AliRoot + Root
  - AliEn
- ALICE & EDG/LCG

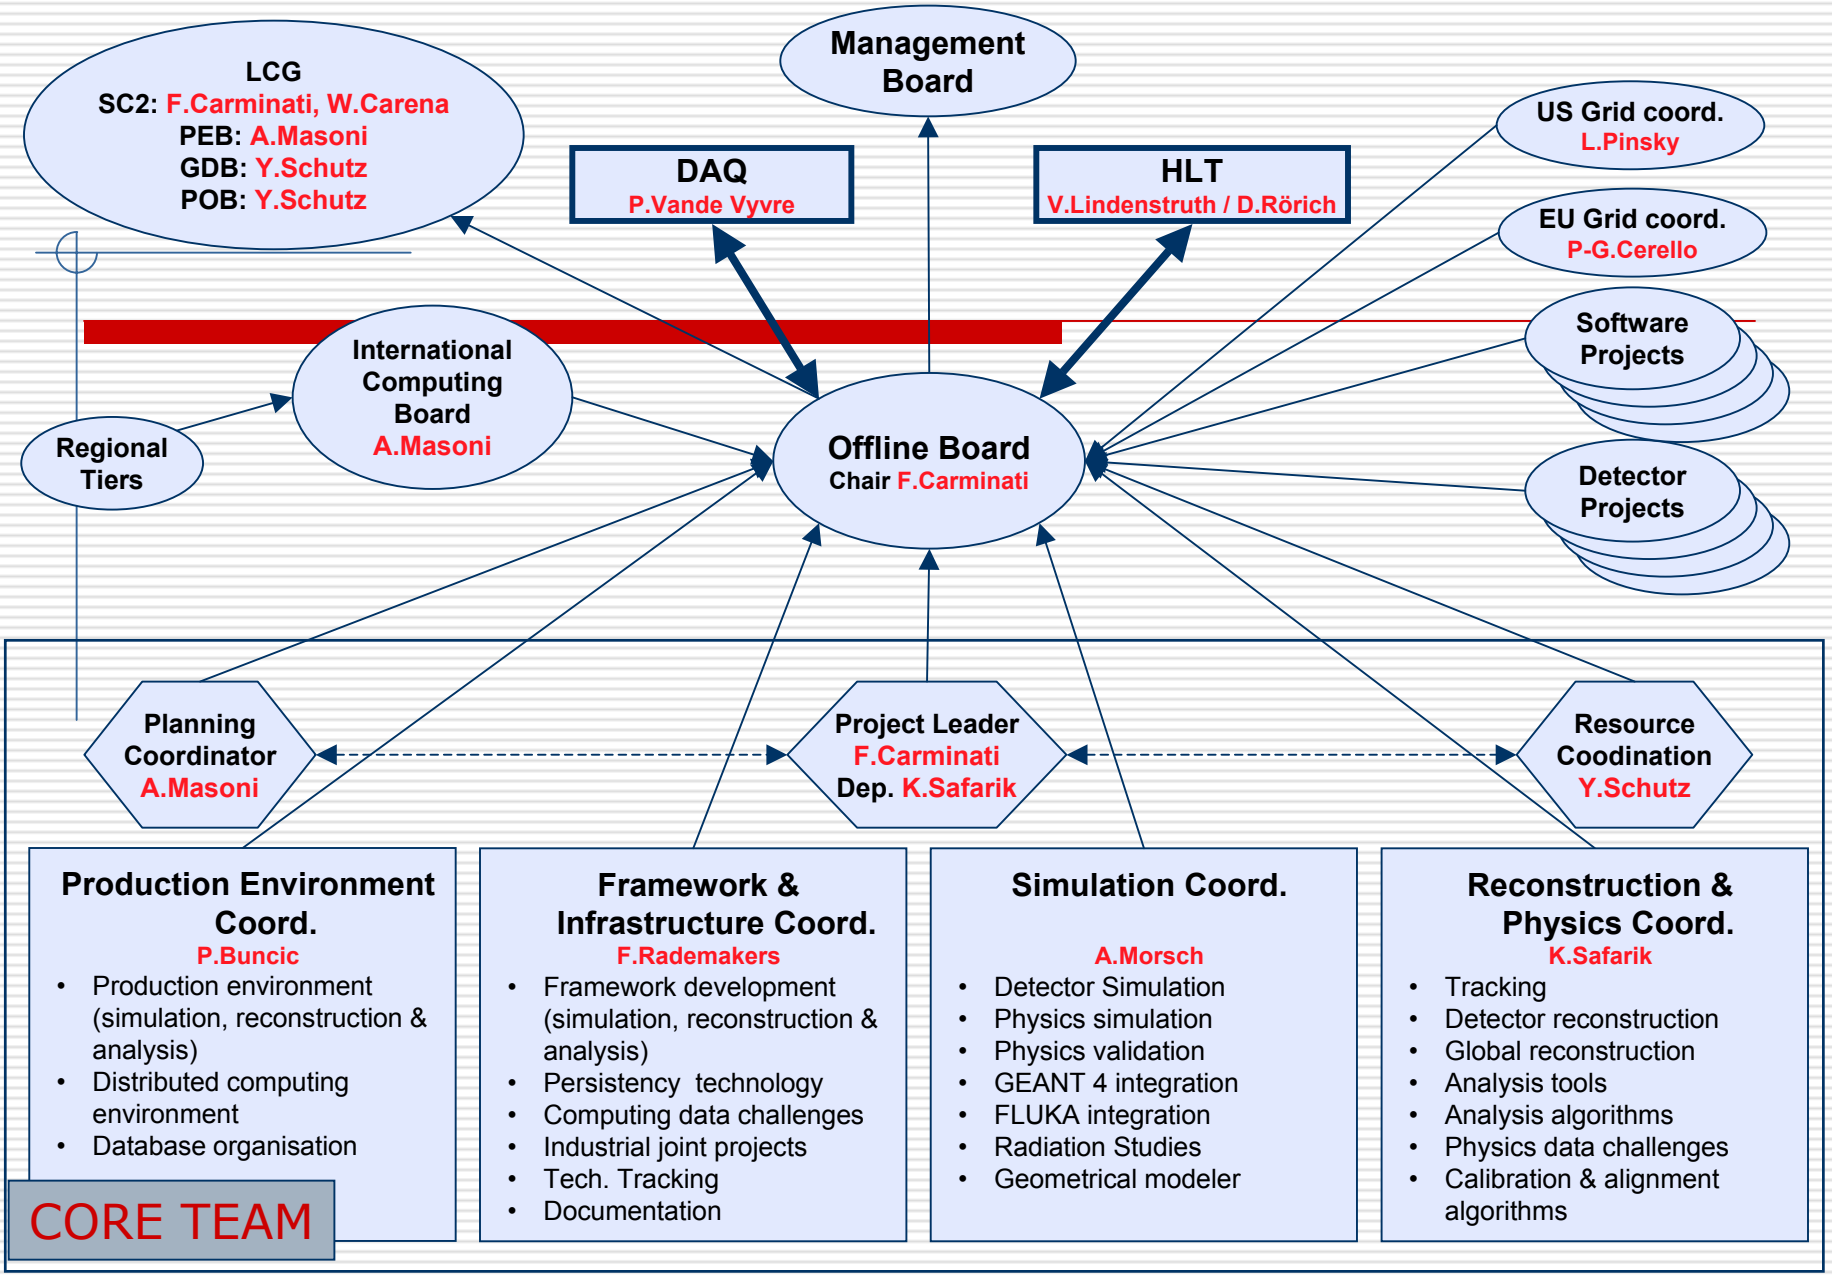
LCG
SC2: F.Carminati, W.Carena
PEB: A.Masoni
GDB: Y.Schutz
POB: Y.Schutz
Management Board
DAQ
P.Vande Vyvre
HLT
V.Lindenstruth / D.Rörich
US Grid coord.
L.Pinsky
EU Grid coord.
P-G.Cerello
Software Projects
Detector Projects
International Computing Board
A.Masoni
Regional Tiers
Offline Board
Chair F.Carminati
Planning Coordinator
A.Masoni
Project Leader
F.Carminati
Dep. K.Safarik
Resource Coodination
Y.Schutz
Production Environment Coord.
P.Buncic
• Production environment (simulation, reconstruction & analysis)
• Distributed computing environment
• Database organisation
Framework & Infrastructure Coord.
F.Rademakers
• Framework development (simulation, reconstruction & analysis)
• Persistency technology
• Computing data challenges
• Industrial joint projects
• Tech. Tracking
• Documentation
Simulation Coord.
A.Morsch
• Detector Simulation
• Physics simulation
• Physics validation
• GEANT 4 integration
• FLUKA integration
• Radiation Studies
• Geometrical modeler
Reconstruction & Physics Coord.
K.Safarik
• Tracking
• Detector reconstruction
• Global reconstruction
• Analysis tools
• Analysis algorithms
• Physics data challenges
• Calibration & alignment algorithms
CORE TEAM

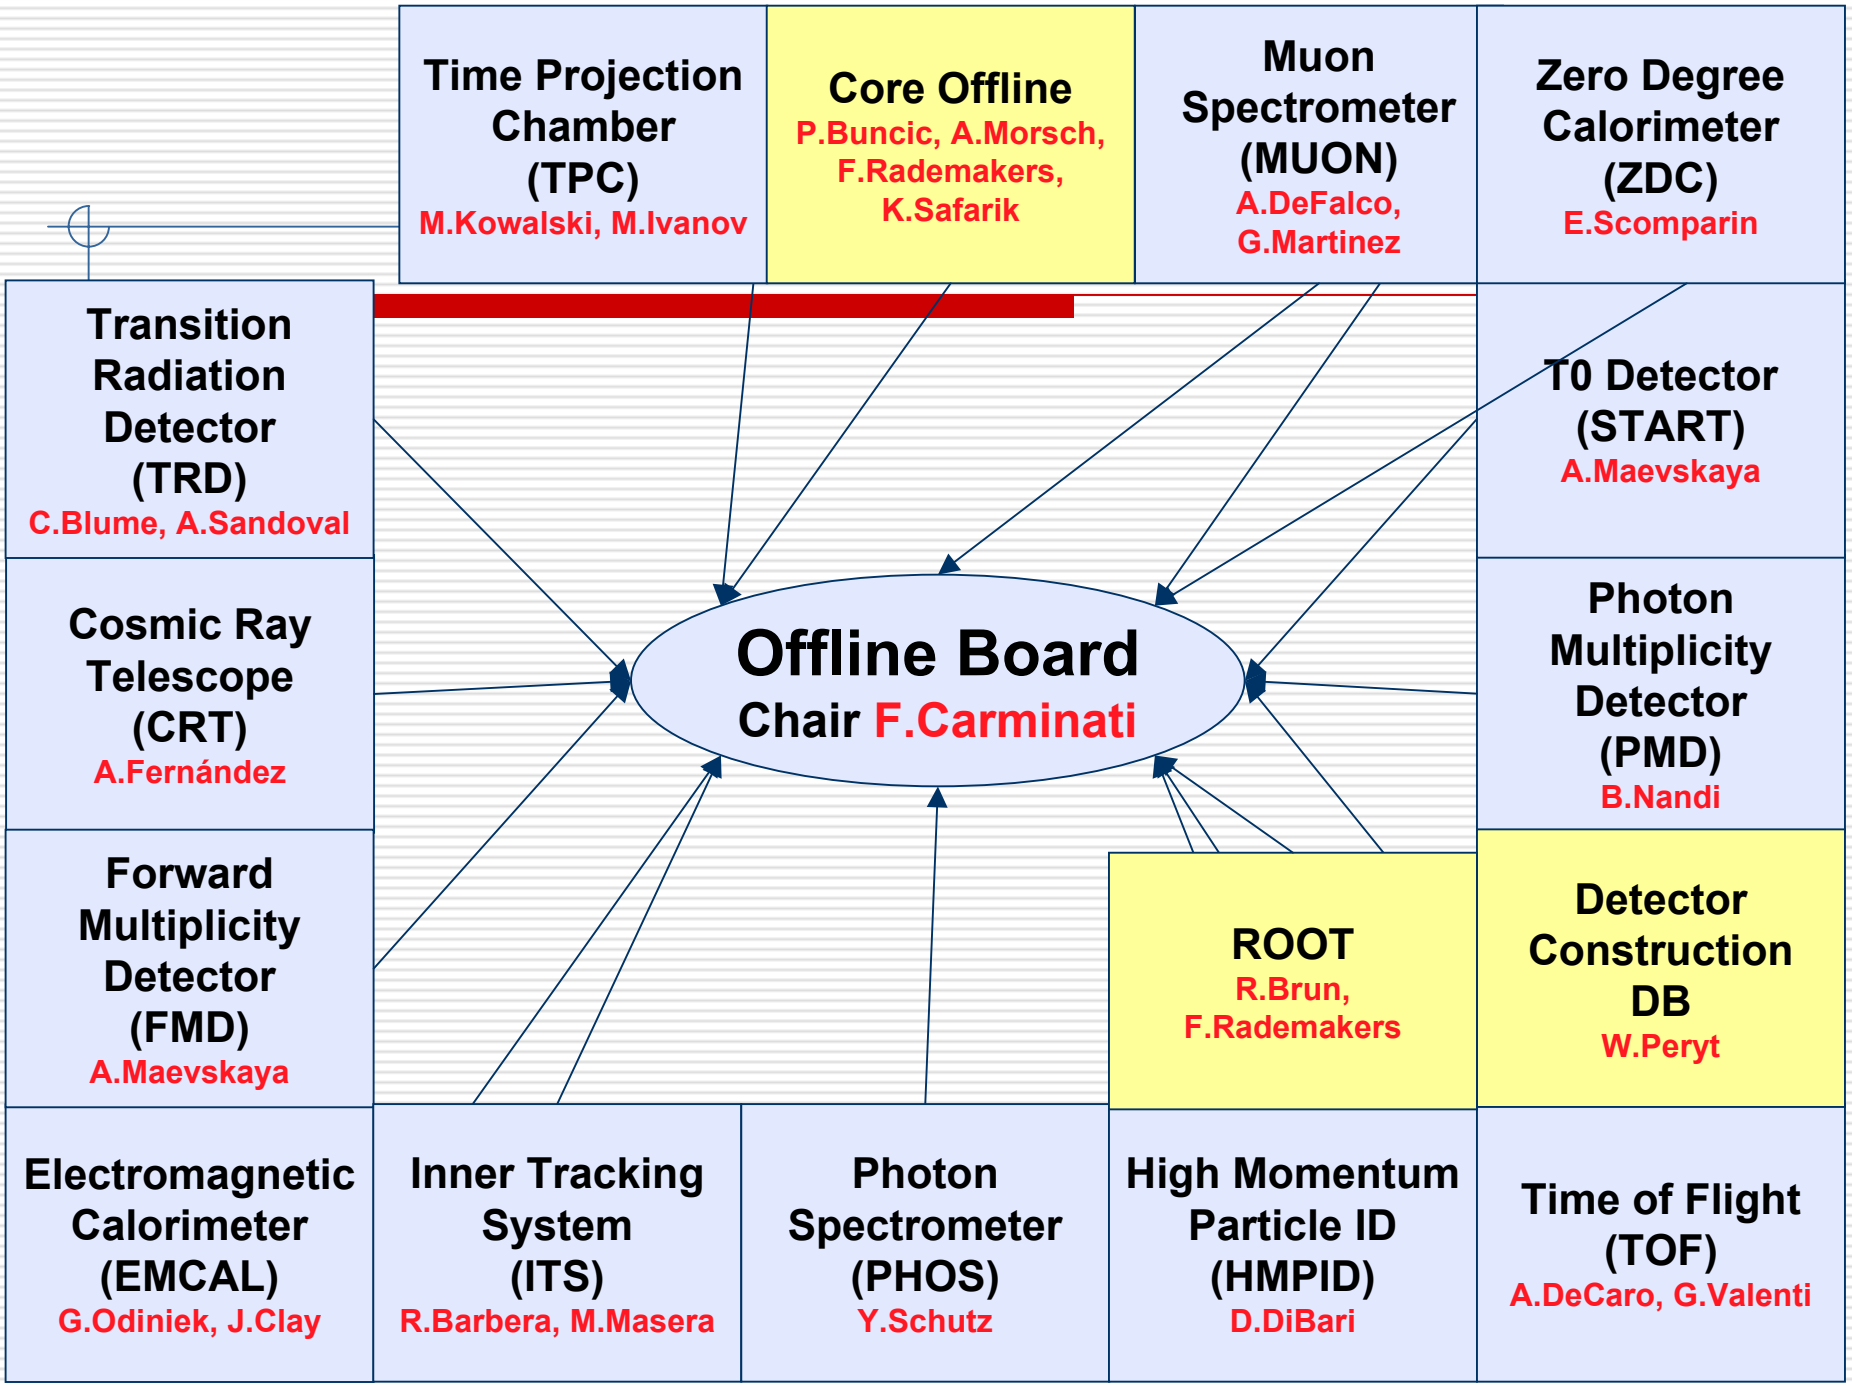
Time Projection Chamber (TPC)
M.Kowalski, M.Ivanov
Core Offline
P.Buncic, A.Morsch, F.Rademakers, K.Safarik
Muon Spectrometer (MUON)
A.DeFalco, G.Martinez
Zero Degree Calorimeter (ZDC)
E.Scomparin
Transition Radiation Detector (TRD)
C.Blume, A.Sandoval
T0 Detector (START)
A.Maevskaya
Cosmic Ray Telescope (CRT)
A.Fernández
Offline Board
Chair F.Carminati
Photon Multiplicity Detector (PMD)
B.Nandi
Forward Multiplicity Detector (FMD)
A.Maevskaya
ROOT
R.Brun, F.Rademakers
Detector Construction DB
W.Peryt
Electromagnetic Calorimeter (EMCAL)
G.Odiniek, J.Clay
Inner Tracking System (ITS)
R.Barbera, M.Masera
Photon Spectrometer (PHOS)
Y.Schutz
High Momentum Particle ID (HMPID)
D.DiBari
Time of Flight (TOF)
A.DeCaro, G.Valenti

ALICE Offline framework

World
Interfaces &
Distributed computing
environment
{anythi

AliEn
@GRID

User
Simulation,
Reconstruction,
Calibration,
Analysis
{C++}

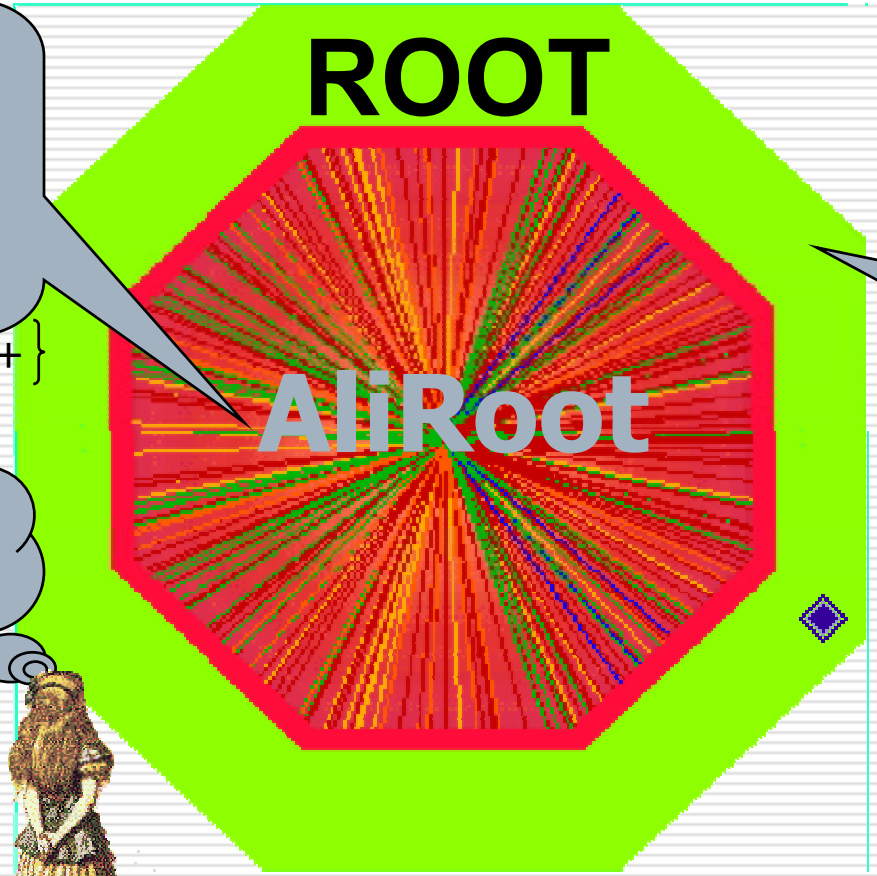

System
GUI
Persistent IO
Utility Libs
{C++}

Nice! I only
have to
learn C++

HEP use cases:

- ✓ Simulation & Reconstruction
- ✓ Event mixing
- ✓ Analysis

Strategic decision in 1998

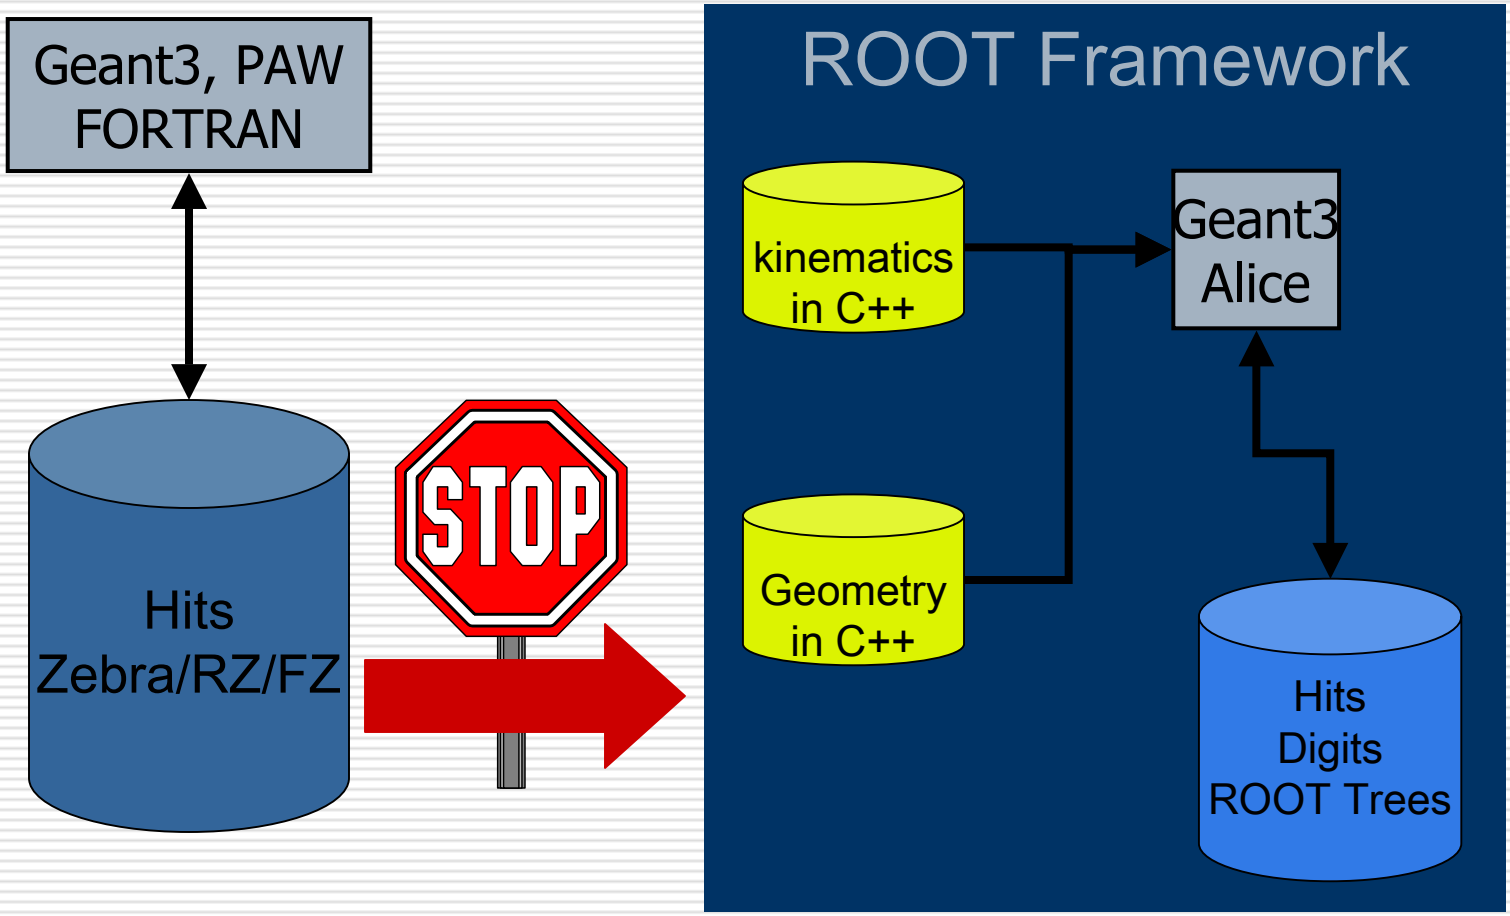

Framework 1/2

- AliRoot framework 597k LOC (+378 generated)
  - C++: 592 (99.17%)
  - fortran: 2.7 (0.44%)
  - sh: 1.3 (0.22%)
  - perl: 0.2 (0.03%)
  - fortran: 903k LOC (external packages)
  - Maintained on Linux (any version!), PC-Pentium, PC-Itanium, HP-UX, DEC Unix, Solaris, Mac OSX
  - Works also with Intel icc compiler

generated using David A. Wheeler's 'SLOCCount'

- Estimates: effort 164 FTE×y, 3.7y schedule, 44 developers, 22 M€ development cost
- Three packages to install (ROOT+GEANT+AliRoot)
  - Less that 1 second to link (thanks to 37 shared libs)
  - 1-click-away install: download and make (non-recursive makefile)

Framework 2/2

- AliEn
  - 30kLOC of PERL5 and C++ (ALICE)
  - ~2MLOC mainly of PERL5 (opens source components)
- Installed on more than 50 sites by physicists
  - ~50 users develop AliRoot from detector groups
  - 70% of code developed outside, 30% by the core Offline team

AliRoot layout

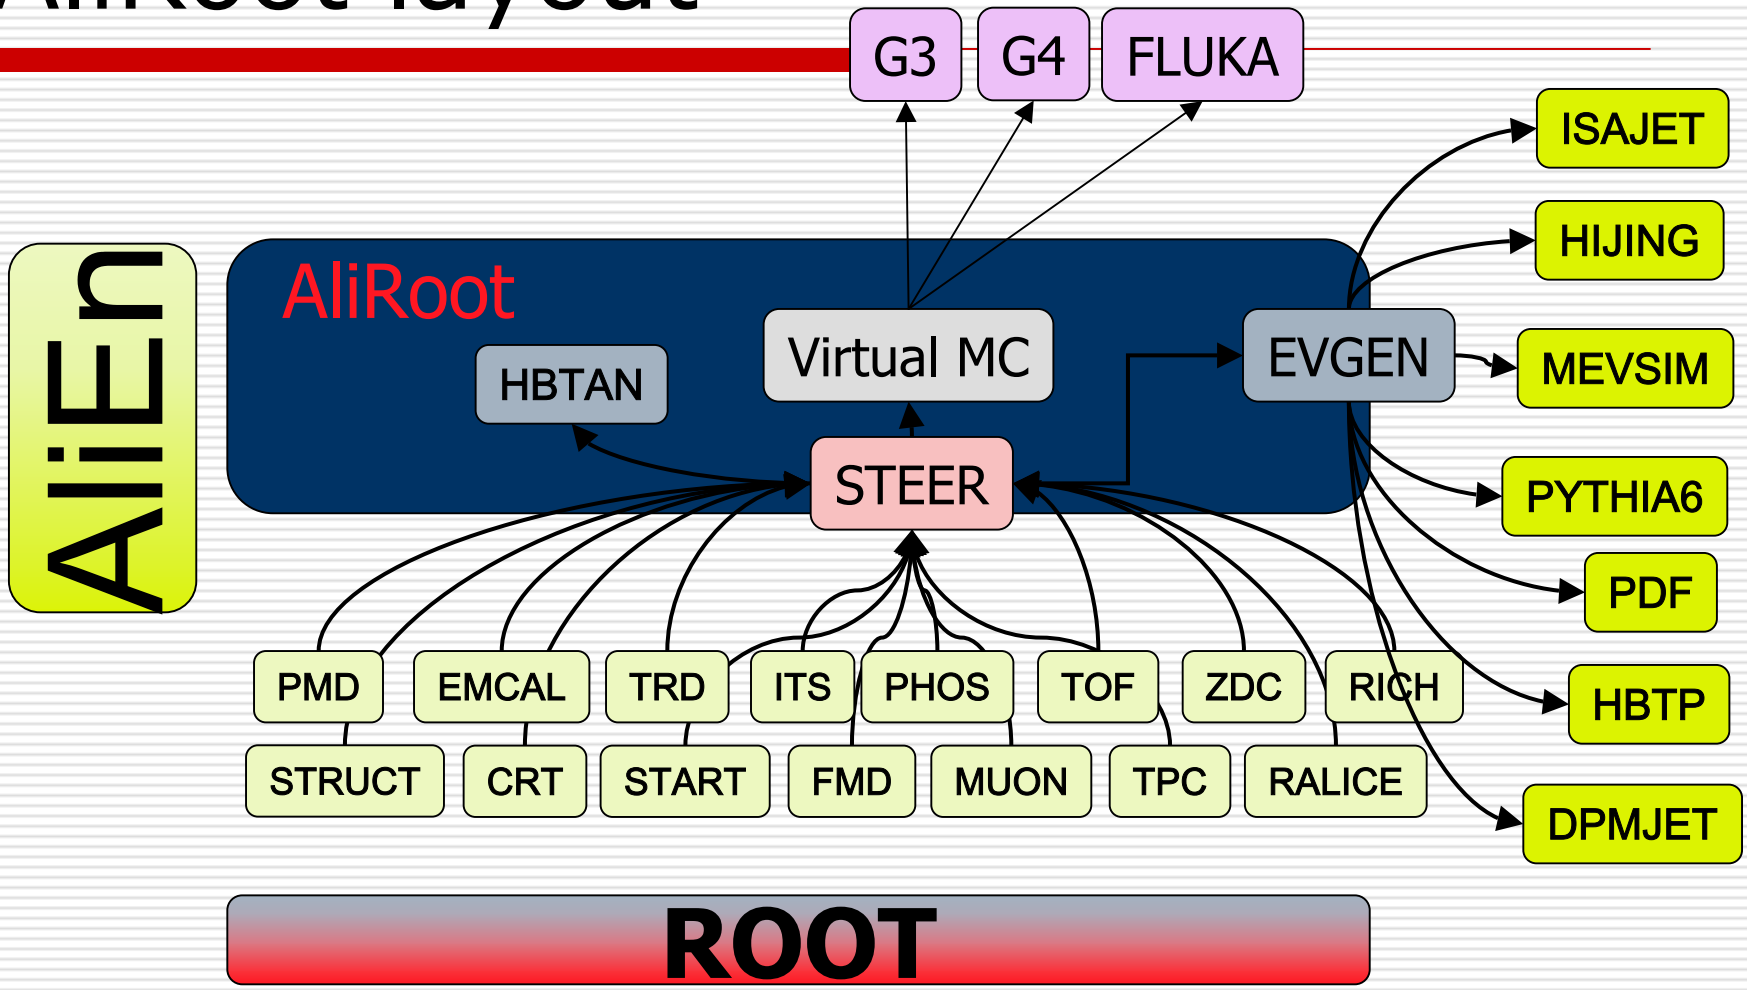

CINT as Interpreter

- CINT used as command line

```
root[0] for (int i = 0; i < 10; i++) printf(“Hello\n”)
root[1] TF1 *f = new TF1(“f”, “sin(x)/x”, 0, 10)
root[2] f->Draw()
```

- And as script interpreter

```
bash$ vi script.C
{
   for (int i = 0; i < 10; i++) printf(“Hello\n”);
   TF1 *f = new TF1(“f”, “sin(x)/x”, 0, 10);
   f->Draw();
}
root[0] .x script.C
```

AliRoot evolution schema

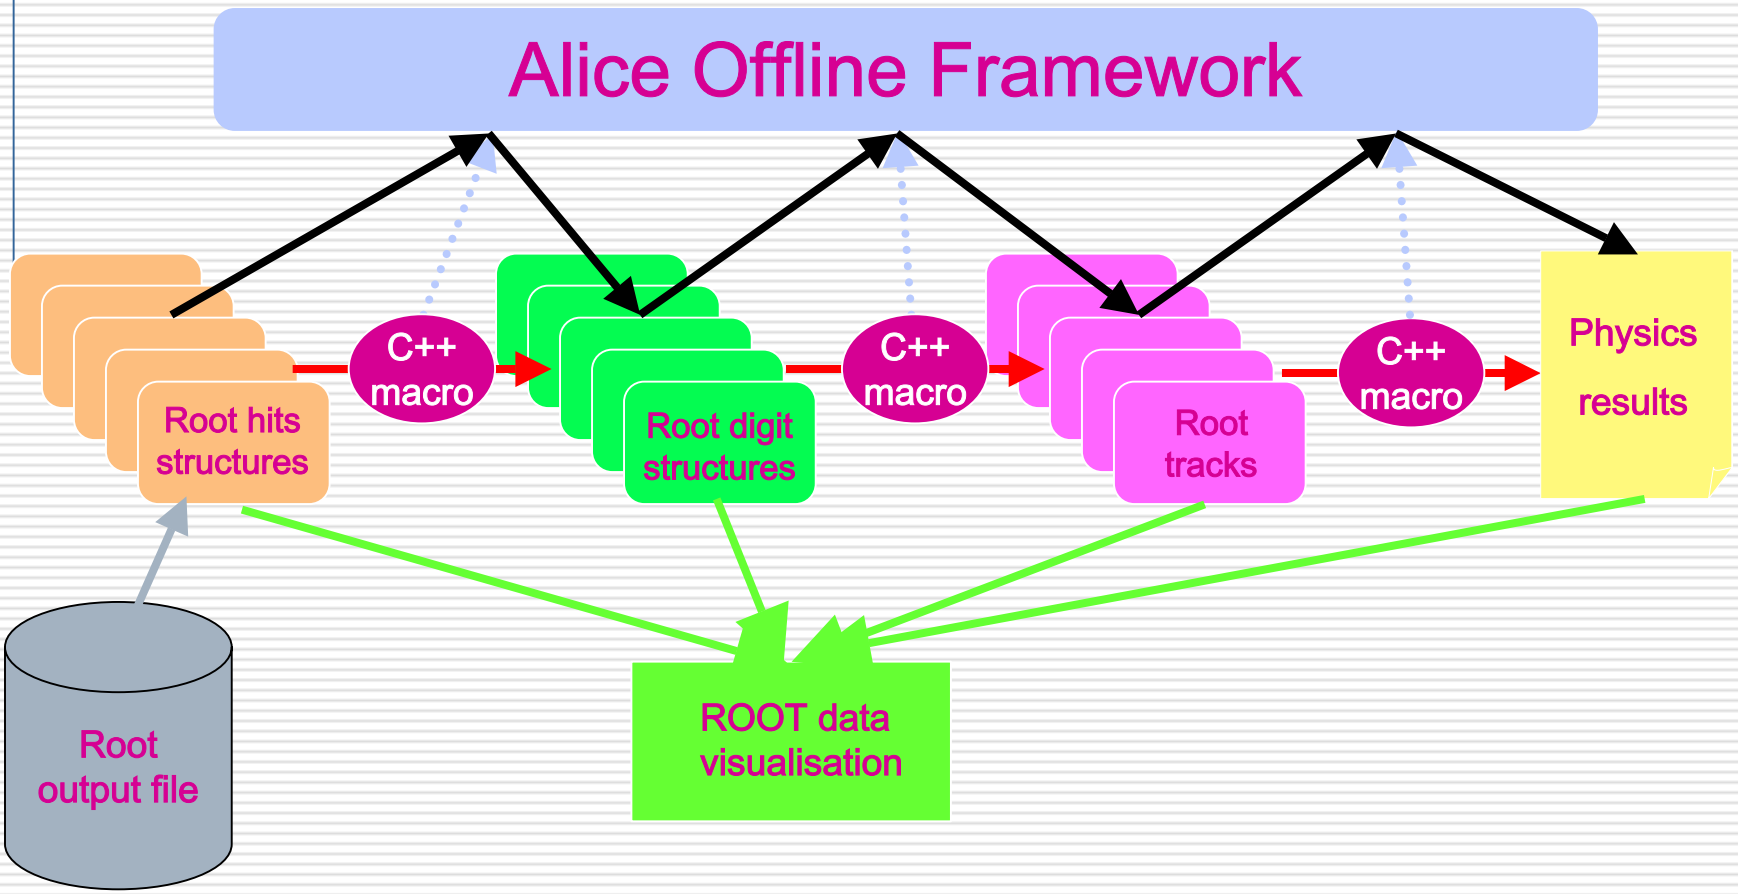

EXplaining Planning in ALICE

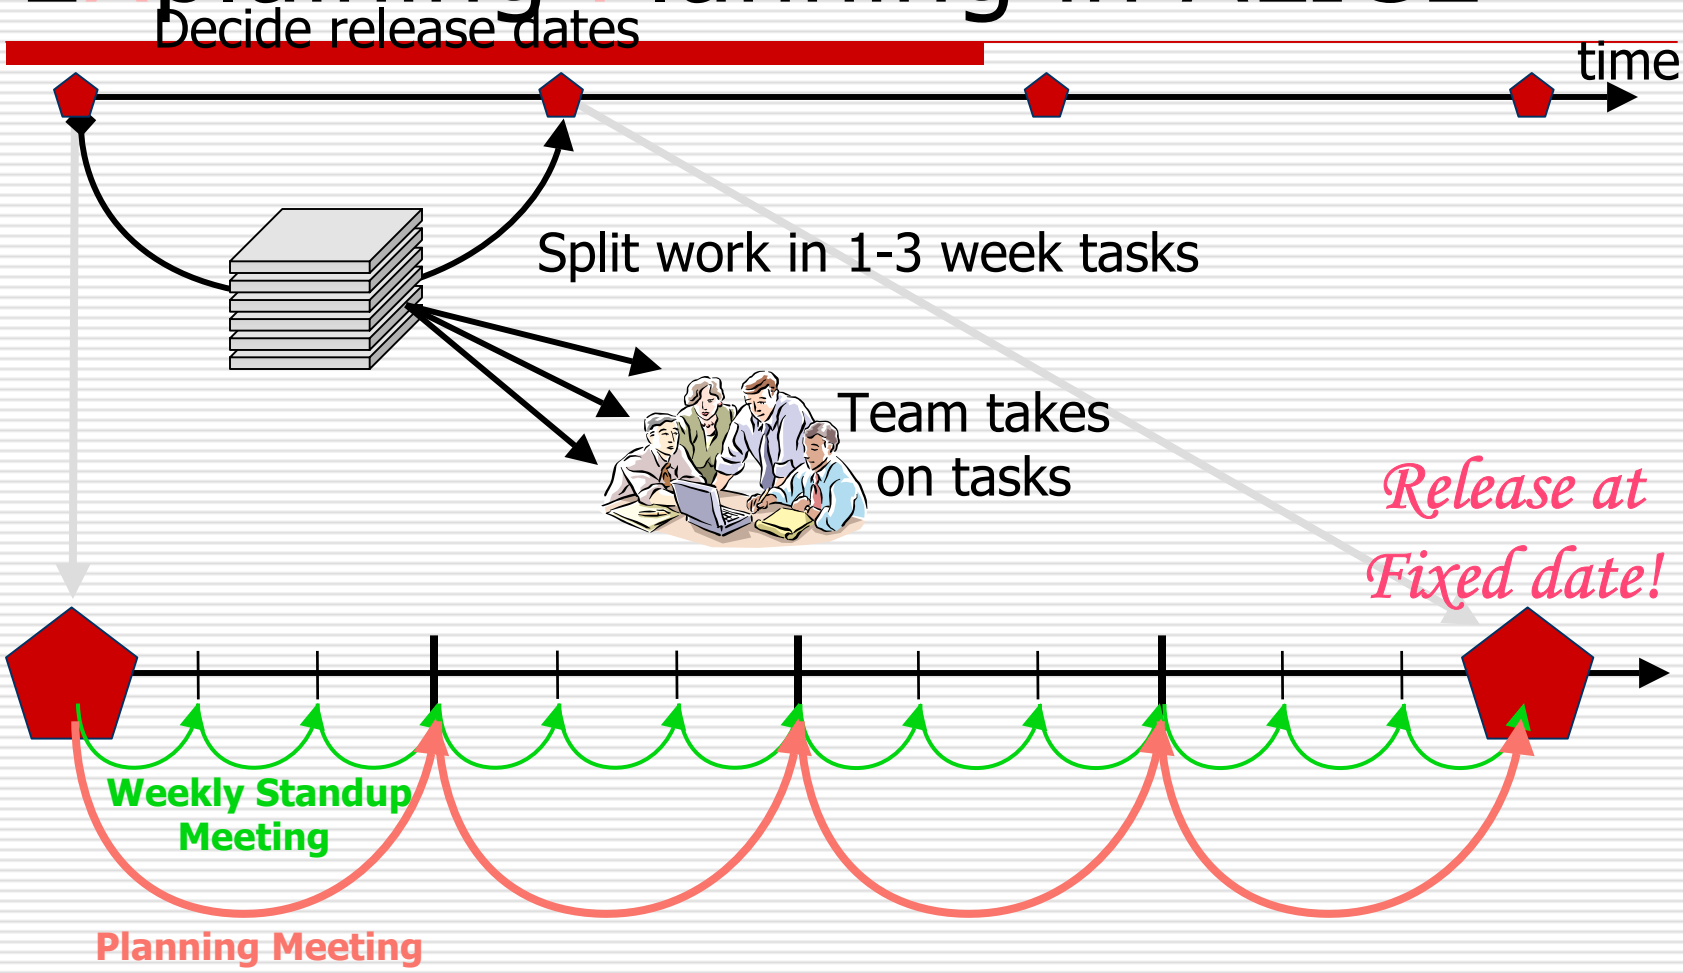

AliRoot maintenance

- Regular release schedule
  - One major release every six months
  - One minor release (tag) every month
- Continuous maintenance and support
  - Very few bugs in the production release because the development version is always available
- Emphasis on delivering production code
  - Corrections, protections, code cleaning, geometry
- Nightly produced UML diagrams, code listing, coding rule violations, build and tests
- One single repository with production and development code

ALICE Event/100

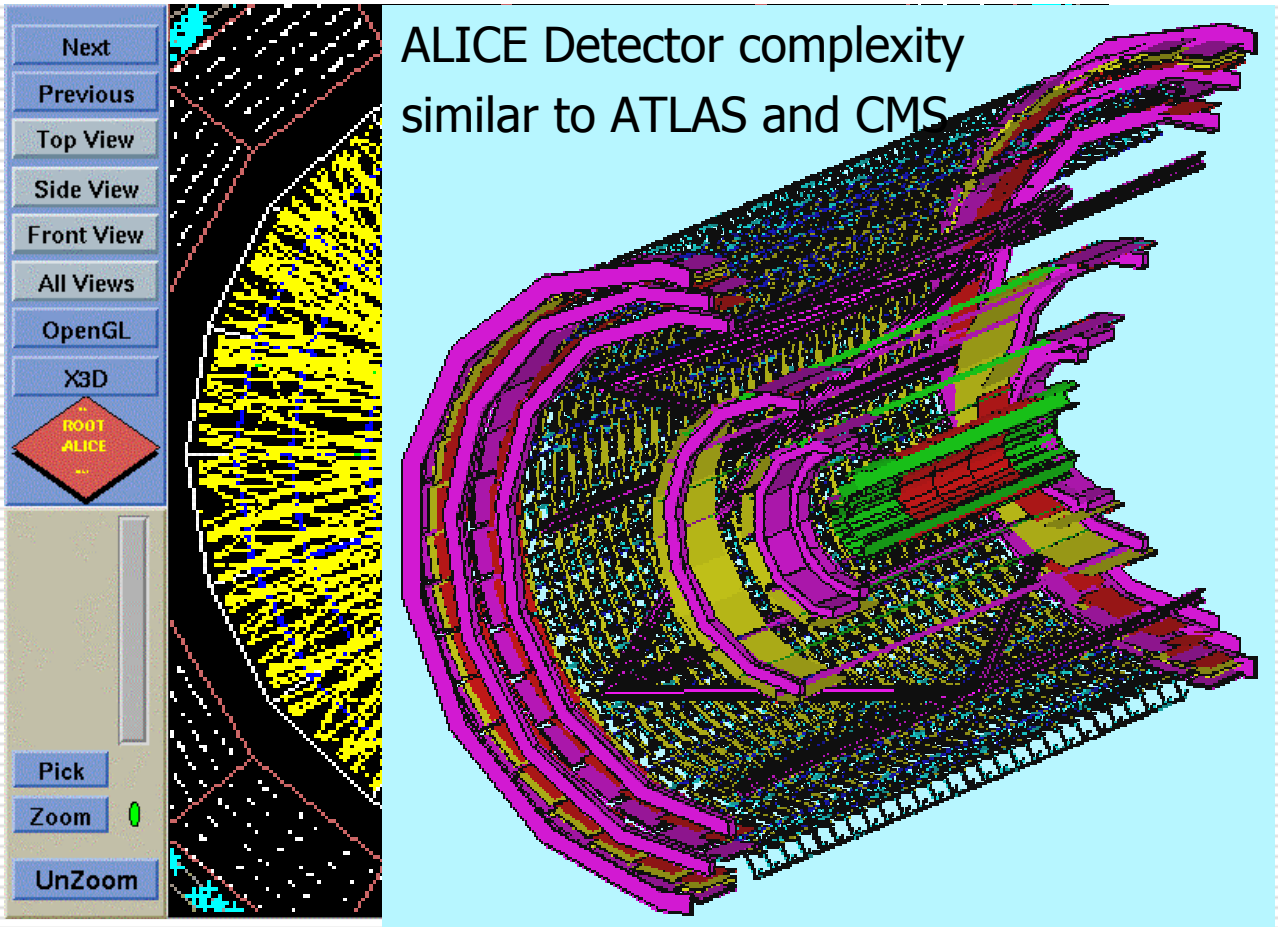

The question of simulation

- Simulation is vital to evaluate the performance of the detector and estimate background

BUT

- Using GEANT 3.21
  - Stay with FORTRAN, old physics and geometry
- Using GEANT 4
  - Not yet validated
- Using FLUKA
  - Good physics but limited user interface

The Virtual MC

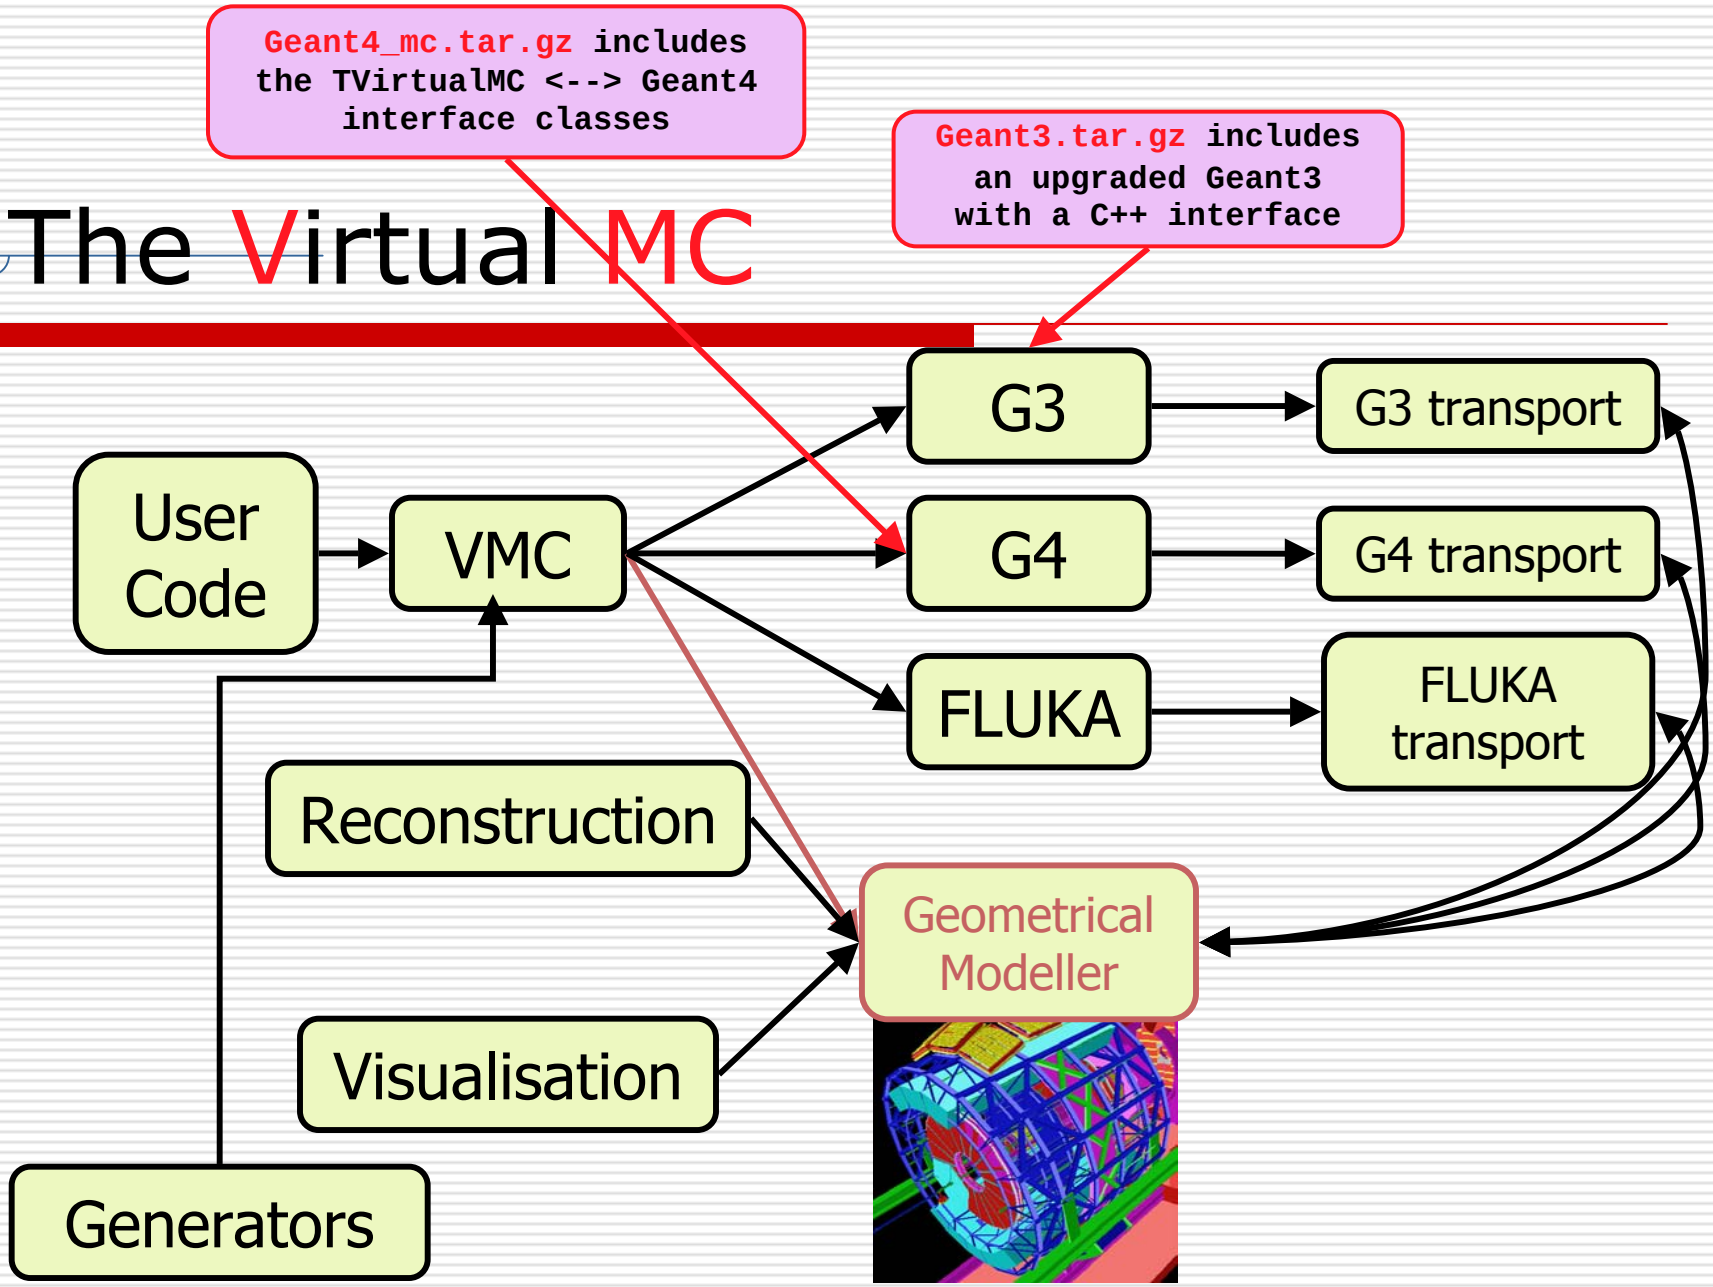

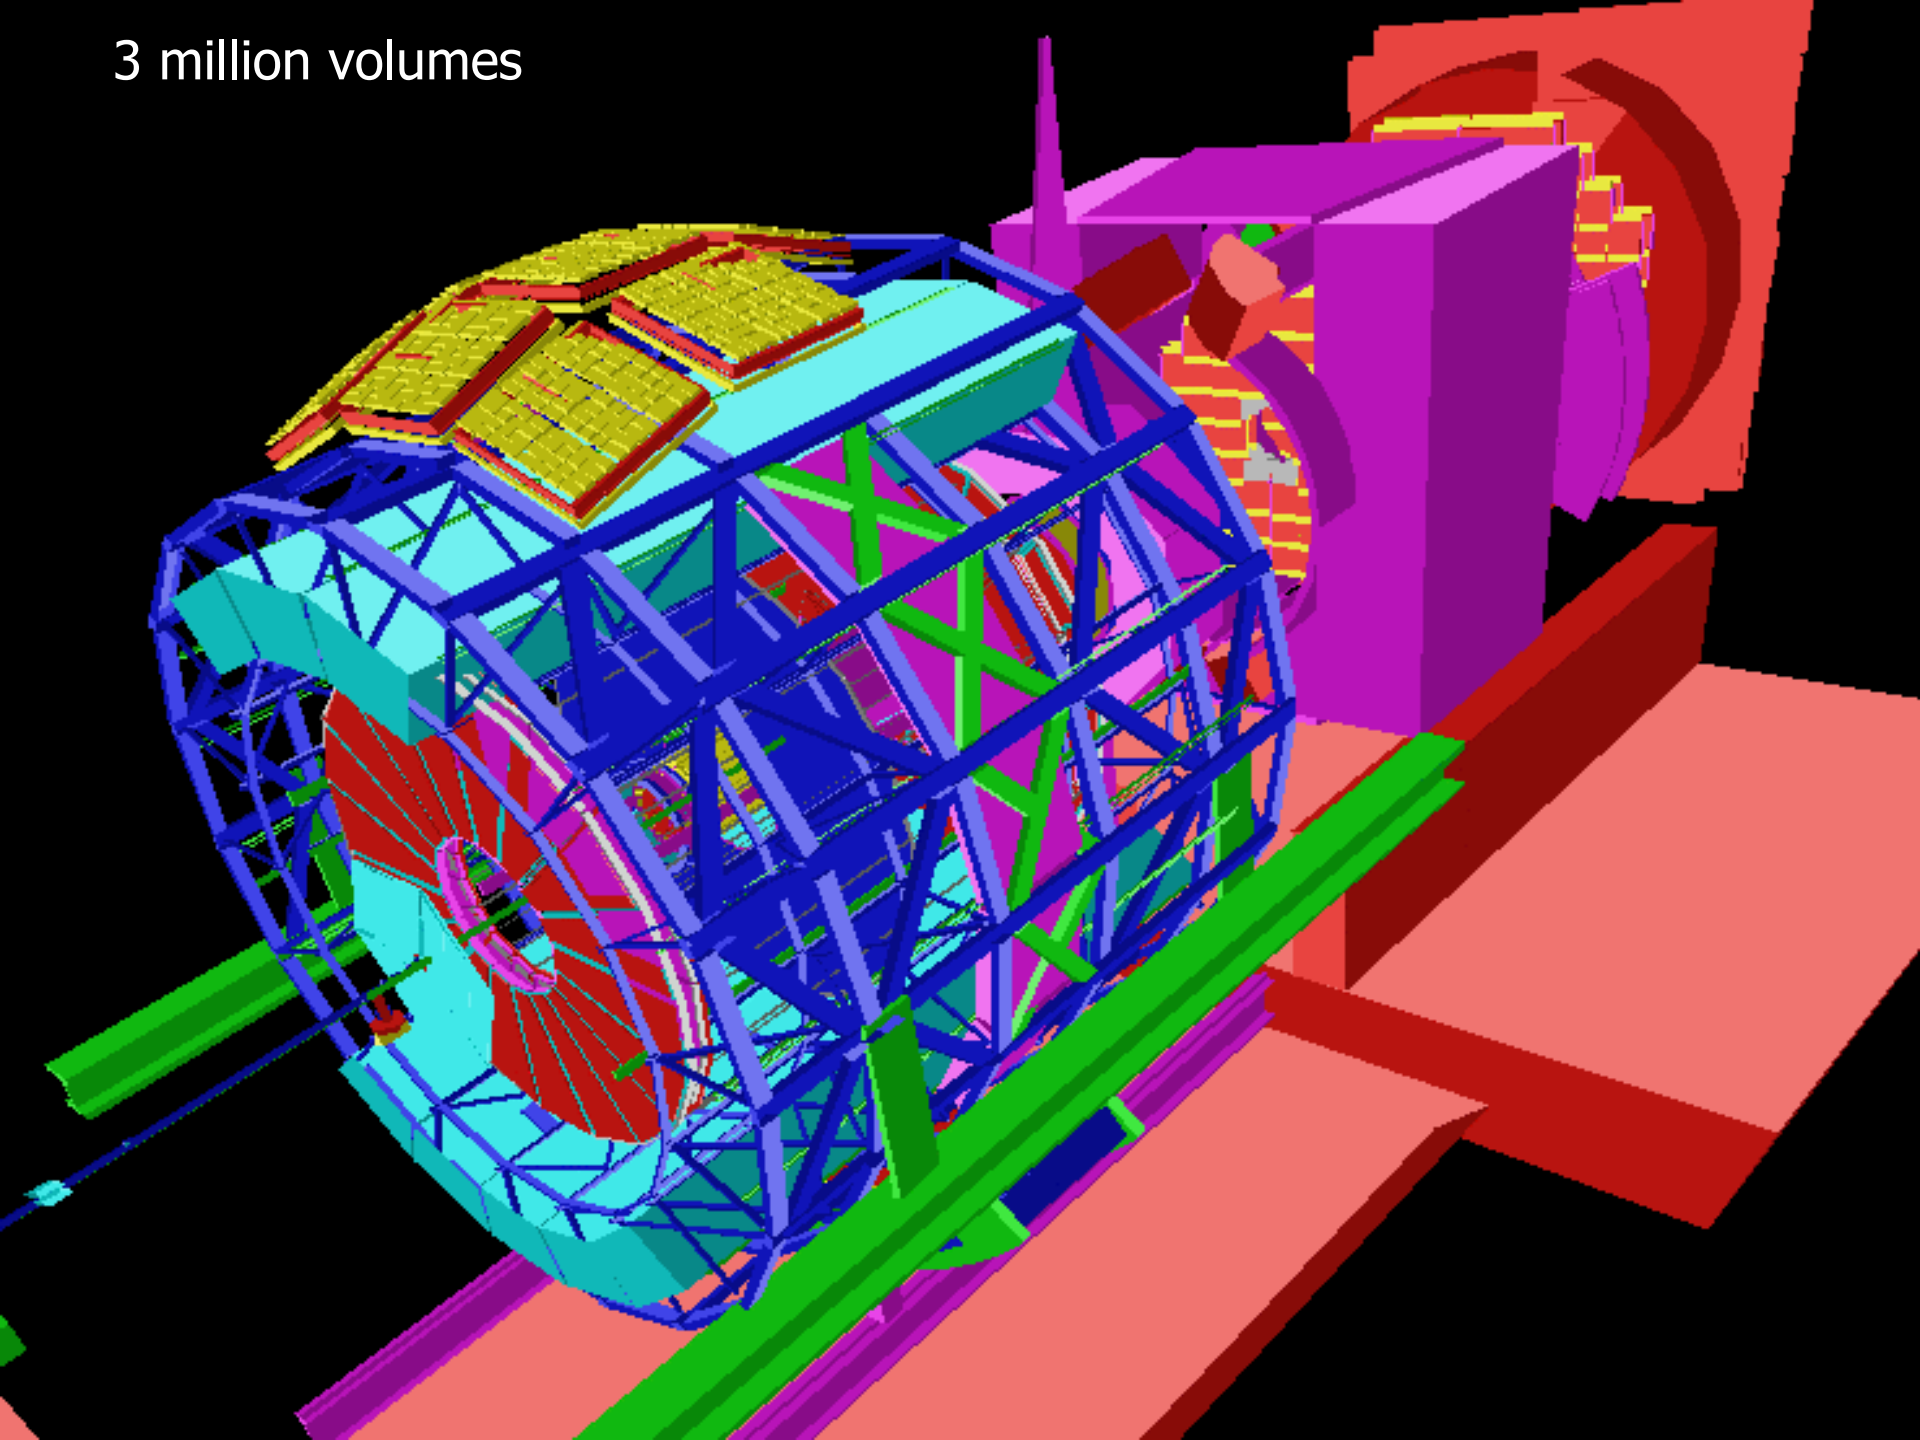
3 million volumes

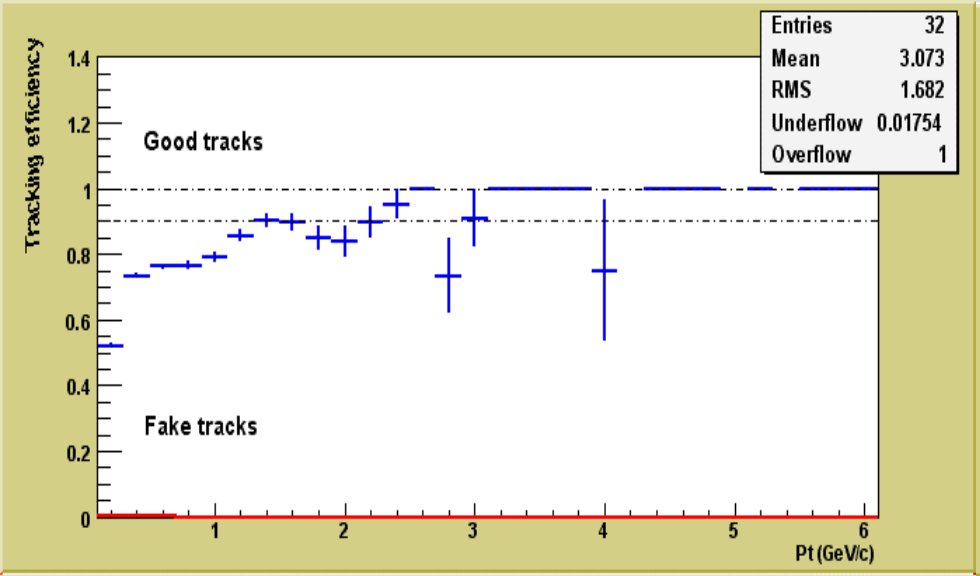
Entries 32
Mean 3.073
RMS 1.682
Underflow 0.01754
Overflow 1
Tracking efficiency
Good tracks
Fake tracks
Pt (GeV/c)

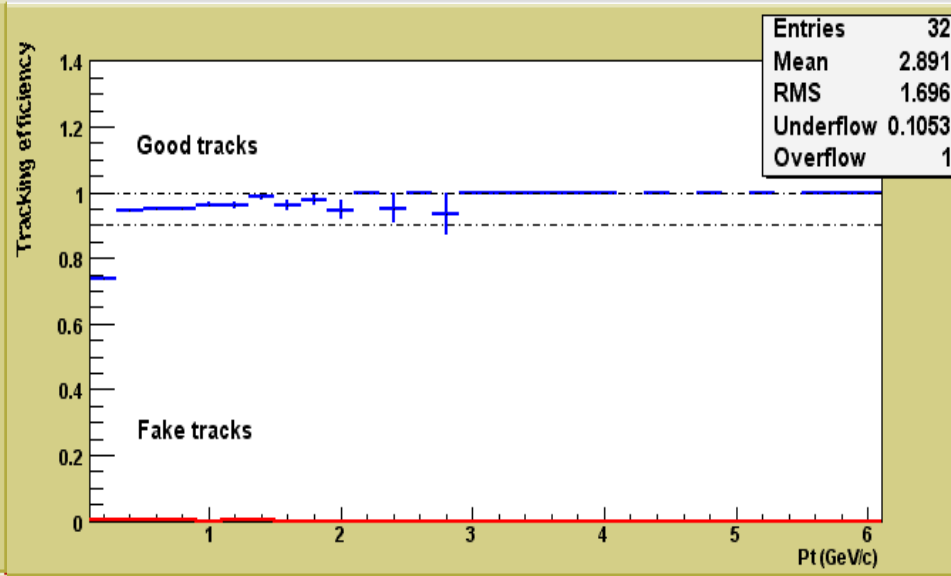
Entries 32
Mean 2.891
RMS 1.696
Underflow 0.1053
Overflow 1
Tracking efficiency
Good tracks
Fake tracks
Pt (GeV/c)

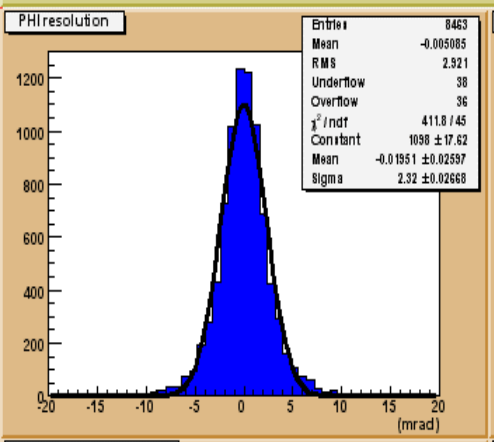
PHI resolution
Entries 8463
Mean -0.005085
RMS 2.921
Underflow 38
Overflow 36
χ²/ndf 411.8 / 45
Constant 1098 ±17.62
Mean -0.01951 ±0.02597
Sigma 2.32 ±0.02668
(mrad)

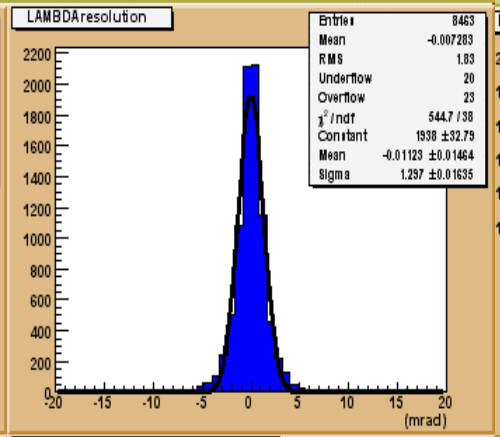
LAMBDA resolution
Entries 8463
Mean -0.007283
RMS 1.83
Underflow 20
Overflow 23
χ²/ndf 544.7 / 38
Constant 1938 ±32.79
Mean -0.01123 ±0.01464
Sigma 1.297 ±0.01635
(mrad)

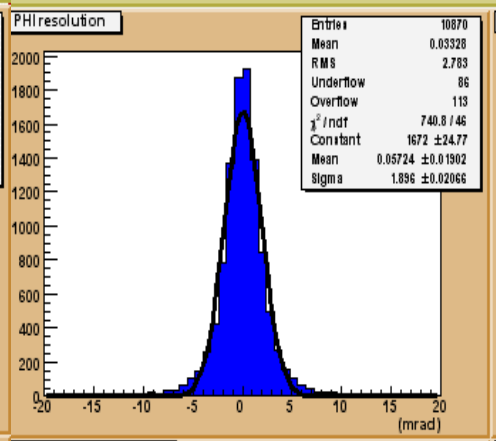
PHI resolution
Entries 10870
Mean 0.03328
RMS 2.783
Underflow 86
Overflow 113
χ²/ndf 740.8 / 46
Constant 1672 ±24.77
Mean 0.05724 ±0.01902
Sigma 1.896 ±0.02066
(mrad)

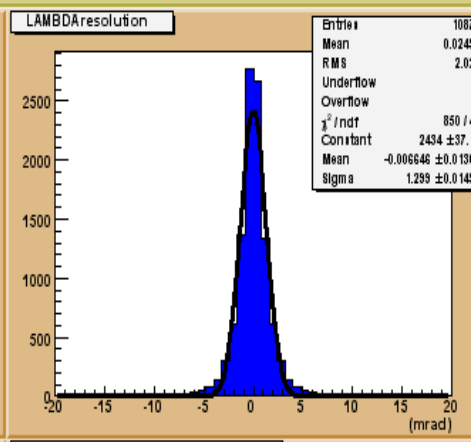
LAMBDA resolution
(mrad)

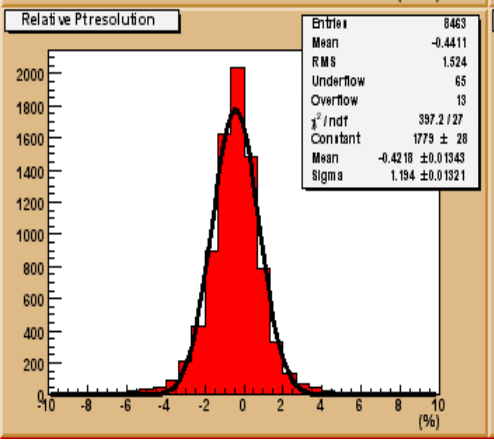
Relative Ptresolution
Entries 8463
Mean -0.4411
RMS 1.524
Underflow 65
Overflow 13
χ²/ndf 397.2 / 27
Constant 1779 ± 28
Mean -0.4218 ±0.01343
Sigma 1.194 ±0.01321
(%)

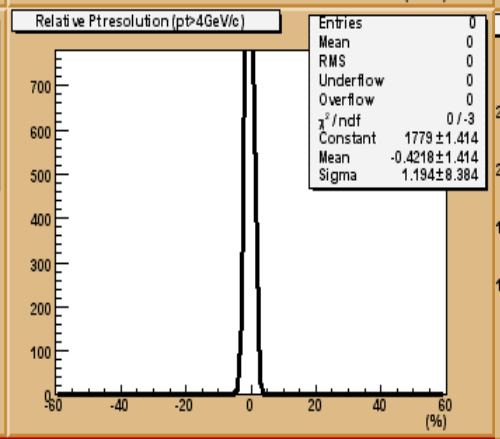
Relative Ptresolution (pt>4GeV/c)
Entries 0
Mean 0
RMS 0
Underflow 0
Overflow 0
χ²/ndf 0 / -3
Constant 1779 ±1.414
Mean -0.4218 ±1.414
Sigma 1.194 ±8.384
(%)

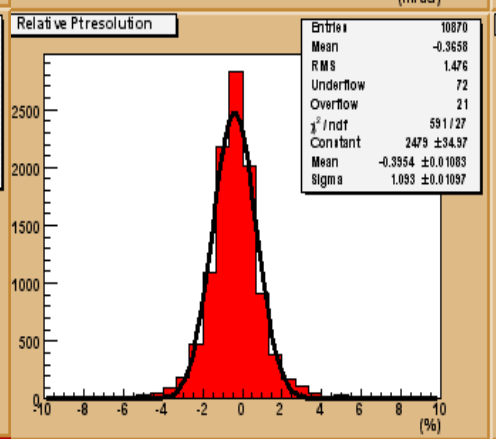
Relative Ptresolution
Entries 10870
Mean -0.3658
RMS 1.476
Underflow 72
Overflow 21
χ²/ndf 591 / 27
Constant 2479 ±34.97
Mean -0.3954 ±0.01083
Sigma 1.093 ±0.01097
(%)

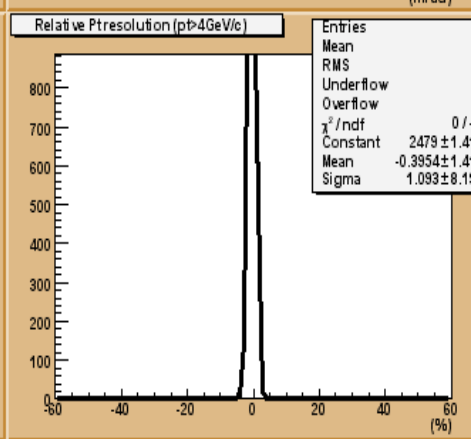
Relative Ptresolution (pt>4GeV/c)
Entries
Mean
RMS
Underflow
Overflow
Constant 2479 ±1.4
Mean -0.3954 ±1.4
Sigma 1.093 ±8.1
(%)

Whiteboard Data Communication

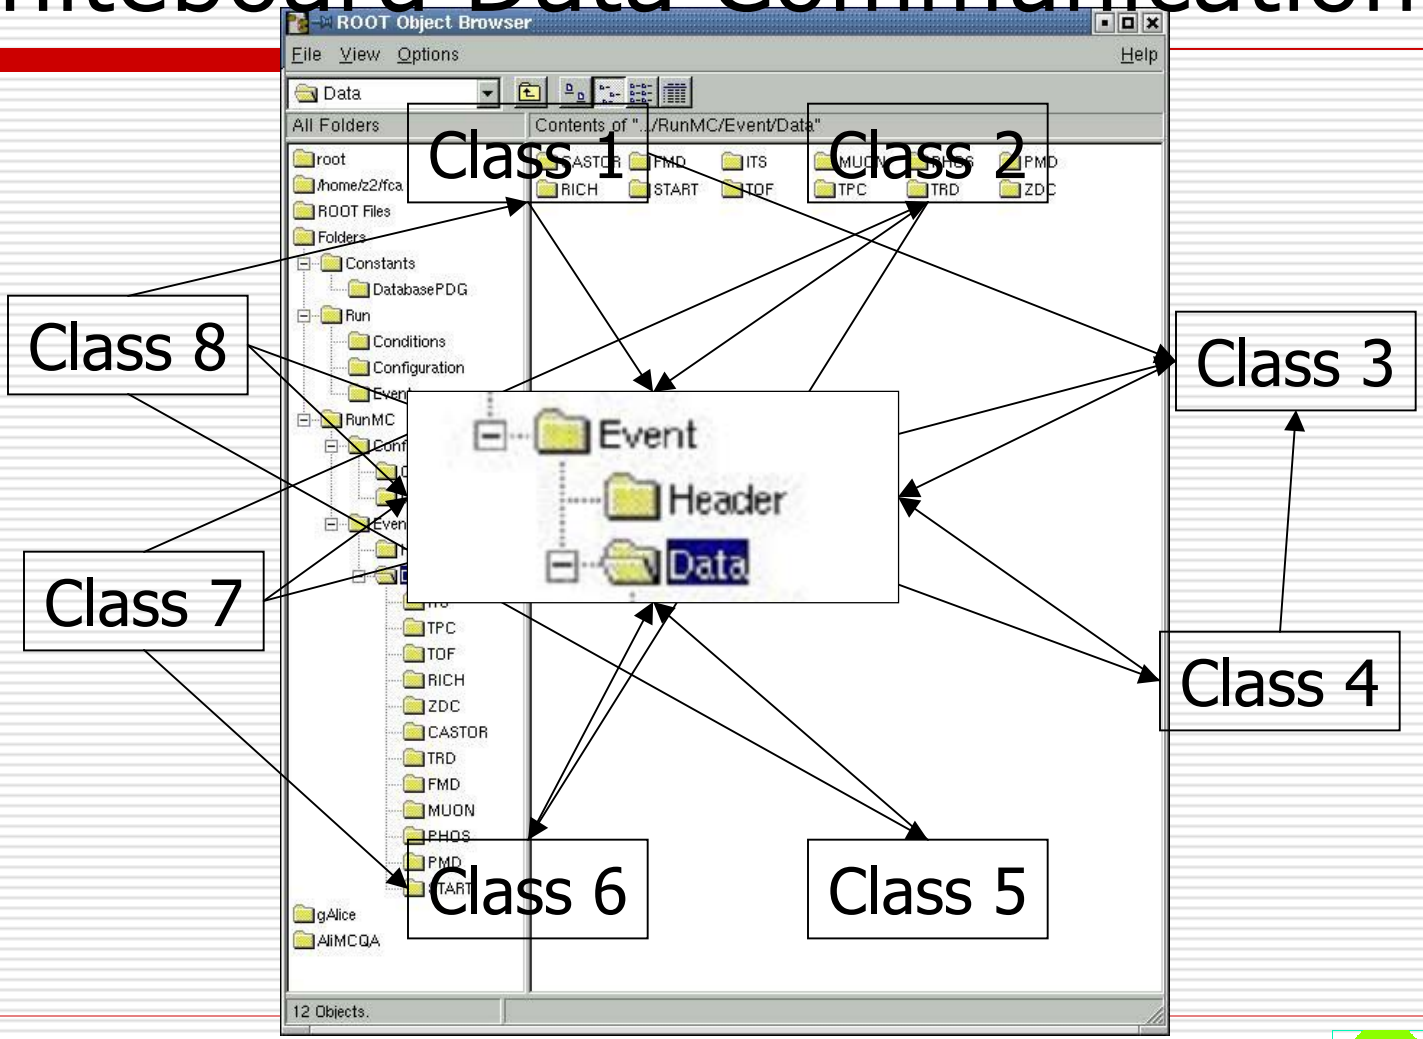

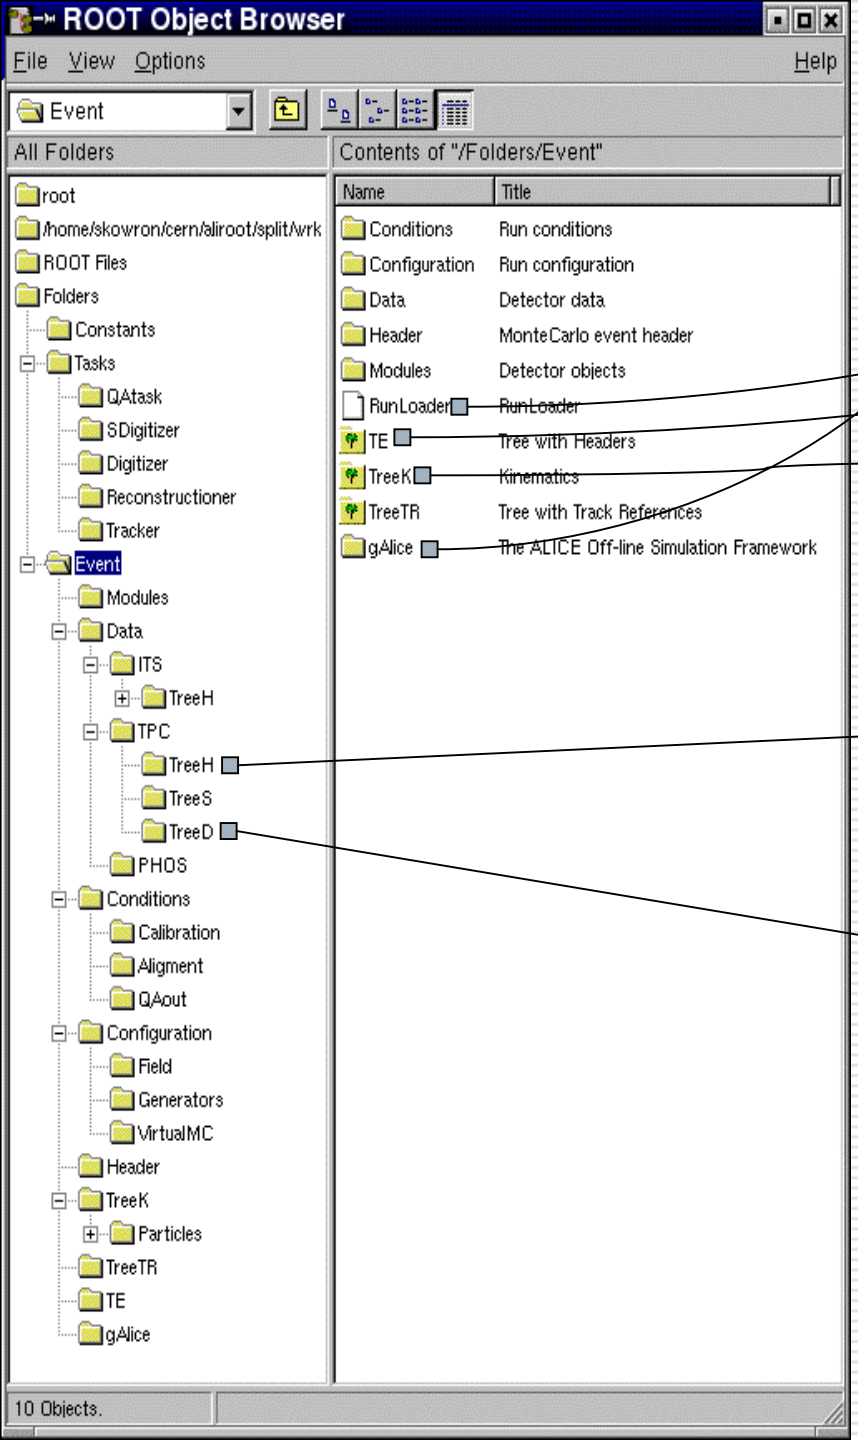
ROOT Object Browser
File View Options Help
Event
All Folders
Contents of "/Folders/Event"
root
/home/skowron/cern/aliroot/split/wrk
ROOT Files
Folders
Constants
Tasks
QAtask
SDigitizer
Digitizer
Reconstructioner
Tracker
Event
Modules
Data
ITS
TreeH
TPC
TreeH
TreeS
TreeD
PHOS
Conditions
Calibration
Aligment
QAout
Configuration
Field
Generators
VirtualMC
Header
TreeK
Particles
TreeTR
TE
gAlice
Name
Title
Conditions
Run conditions
Configuration
Run configuration
Data
Detector data
Header
MonteCarlo event header
Modules
Detector objects
RunLoader
RunLoader
TE
Tree with Headers
TreeK
Kinematics
TreeTR
Tree with Track References
gAlice
The ALICE Off-line Simulation Framework
10 Objects.

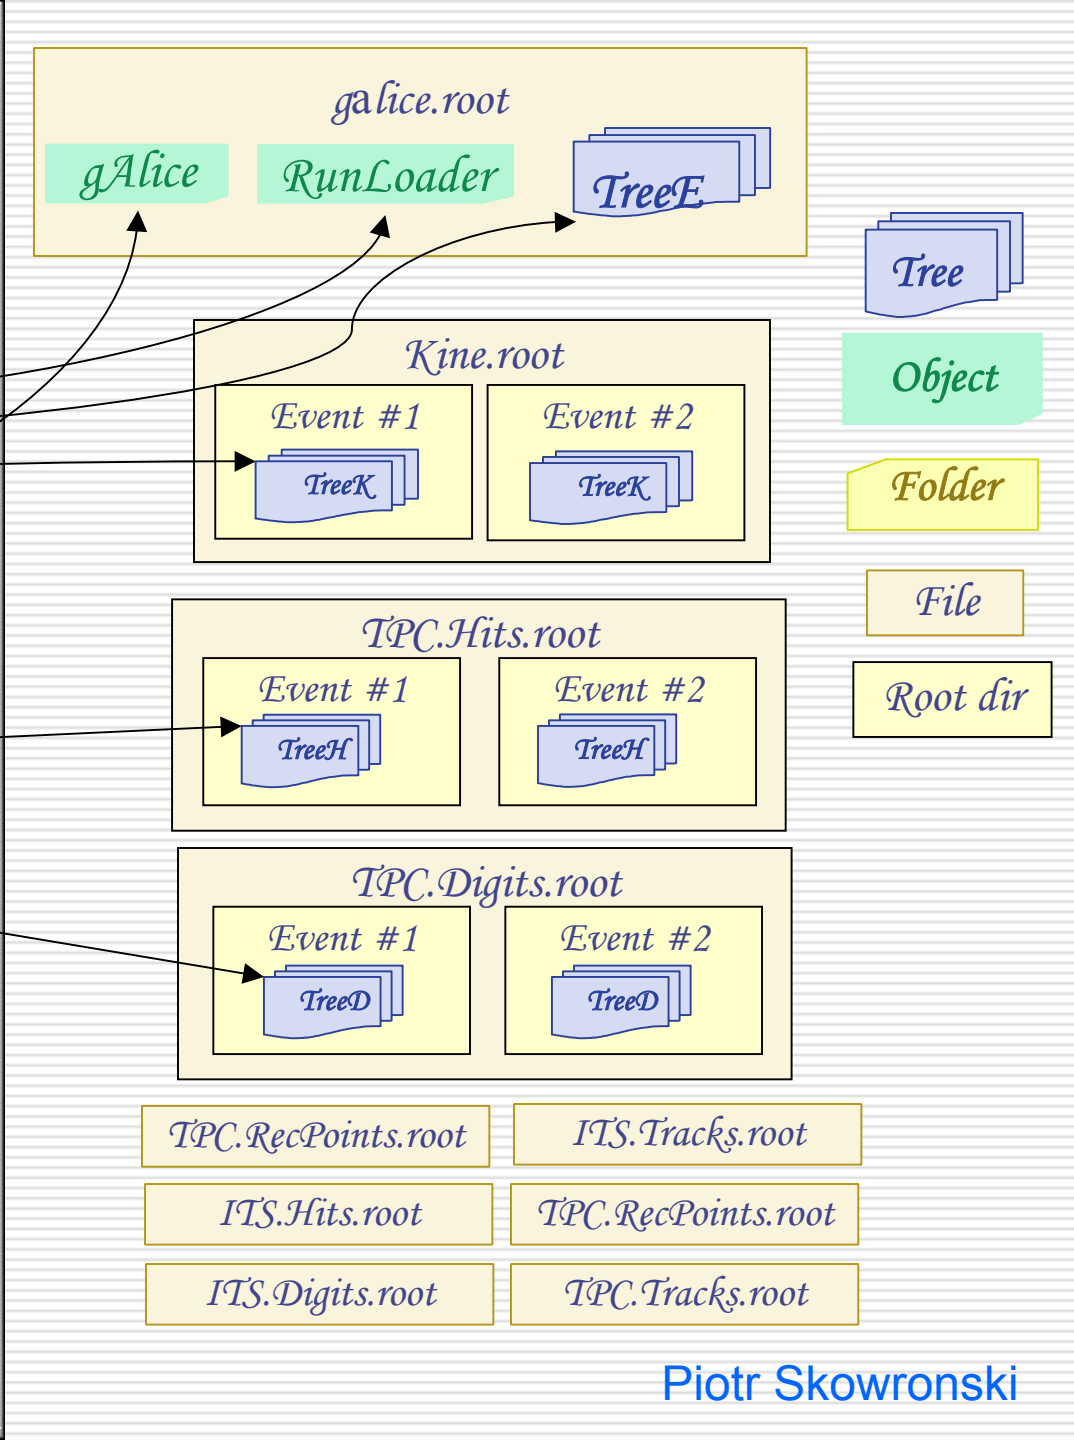
gAlice.root
gAlice
RunLoader
TreeE
Kine.root
Event #1
TreeK
Event #2
TreeK
TPC.Hits.root
Event #1
TreeH
Event #2
TreeH
TPC.Digits.root
Event #1
TreeD
Event #2
TreeD
TPC.RecPoints.root
ITS.Tracks.root
ITS.Hits.root
TPC.RecPoints.root
ITS.Digits.root
TPC.Tracks.root
Tree
Object
Folder
File
Root dir
Piotr Skowronski

ALICE Physics Performance Report

- Detailed evaluation of ALICE performance using the latest simulation tools
  - Acceptance, efficiency, resolution
- Need O(10^7) central HI events
  - 24h@600MHz/ev: 300,000 PC's for one year!
- Need alternative
  - Step1: Generate parametrised background *summable* digits
  - Step2: Generate on the flight the signals and merge
  - Step3: Event analysis
  - O(10^4) background reused O(10^3) times

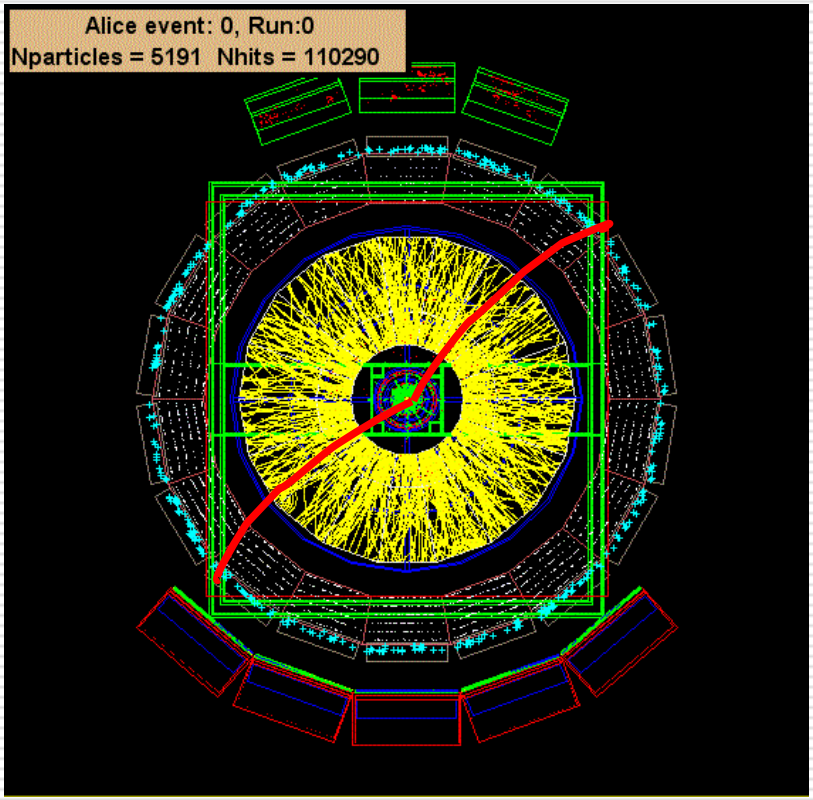

AliEn a lightweight GRID

- AliEn (http://alien.cern.ch) is a lightweight alternative to full blown GRID based on standard components (SOAP, Web services)
  - File Catalogue as a global file system on a RDB
  - TAG Catalogue, as extension
  - Secure Authentication
  - Central Queue Manager ("pull" vs "push" model)
  - Monitoring infrastructure
  - Automatic software installation with AliKit

The Core GRID Functionality !!

- AliEn is routinely used in production for Alice PPR

PPR production summary (2001-2002)

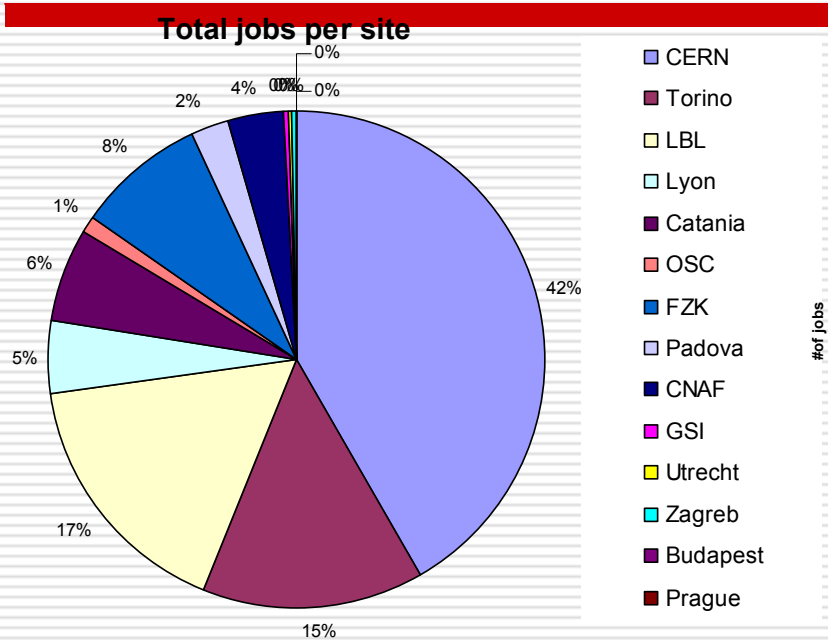

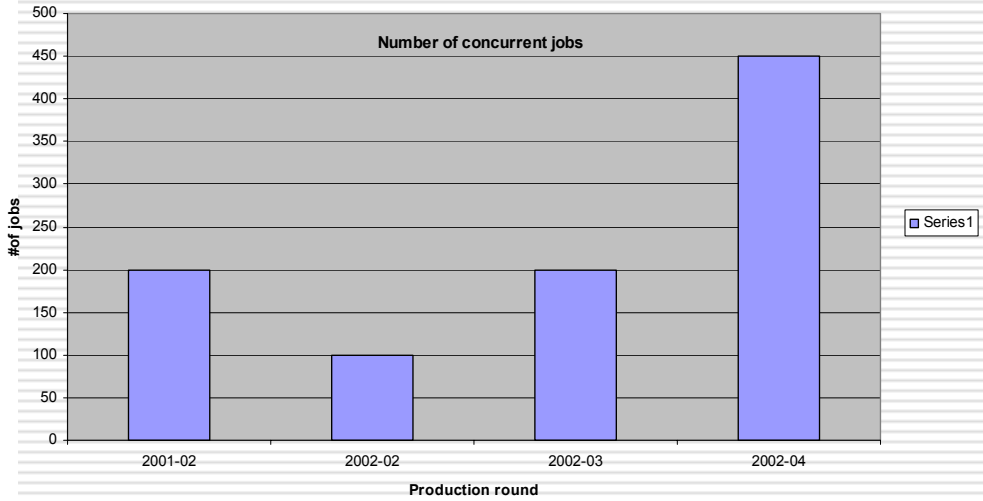

- 32 (was 28) sites configured
- 5 (was 4) sites providing mass storage capability
- 12 production rounds
- 22773 jobs validated, 2428 failed (10%)
- Up to 450 concurrent jobs
- 0.5 operators

ALICE and GRID

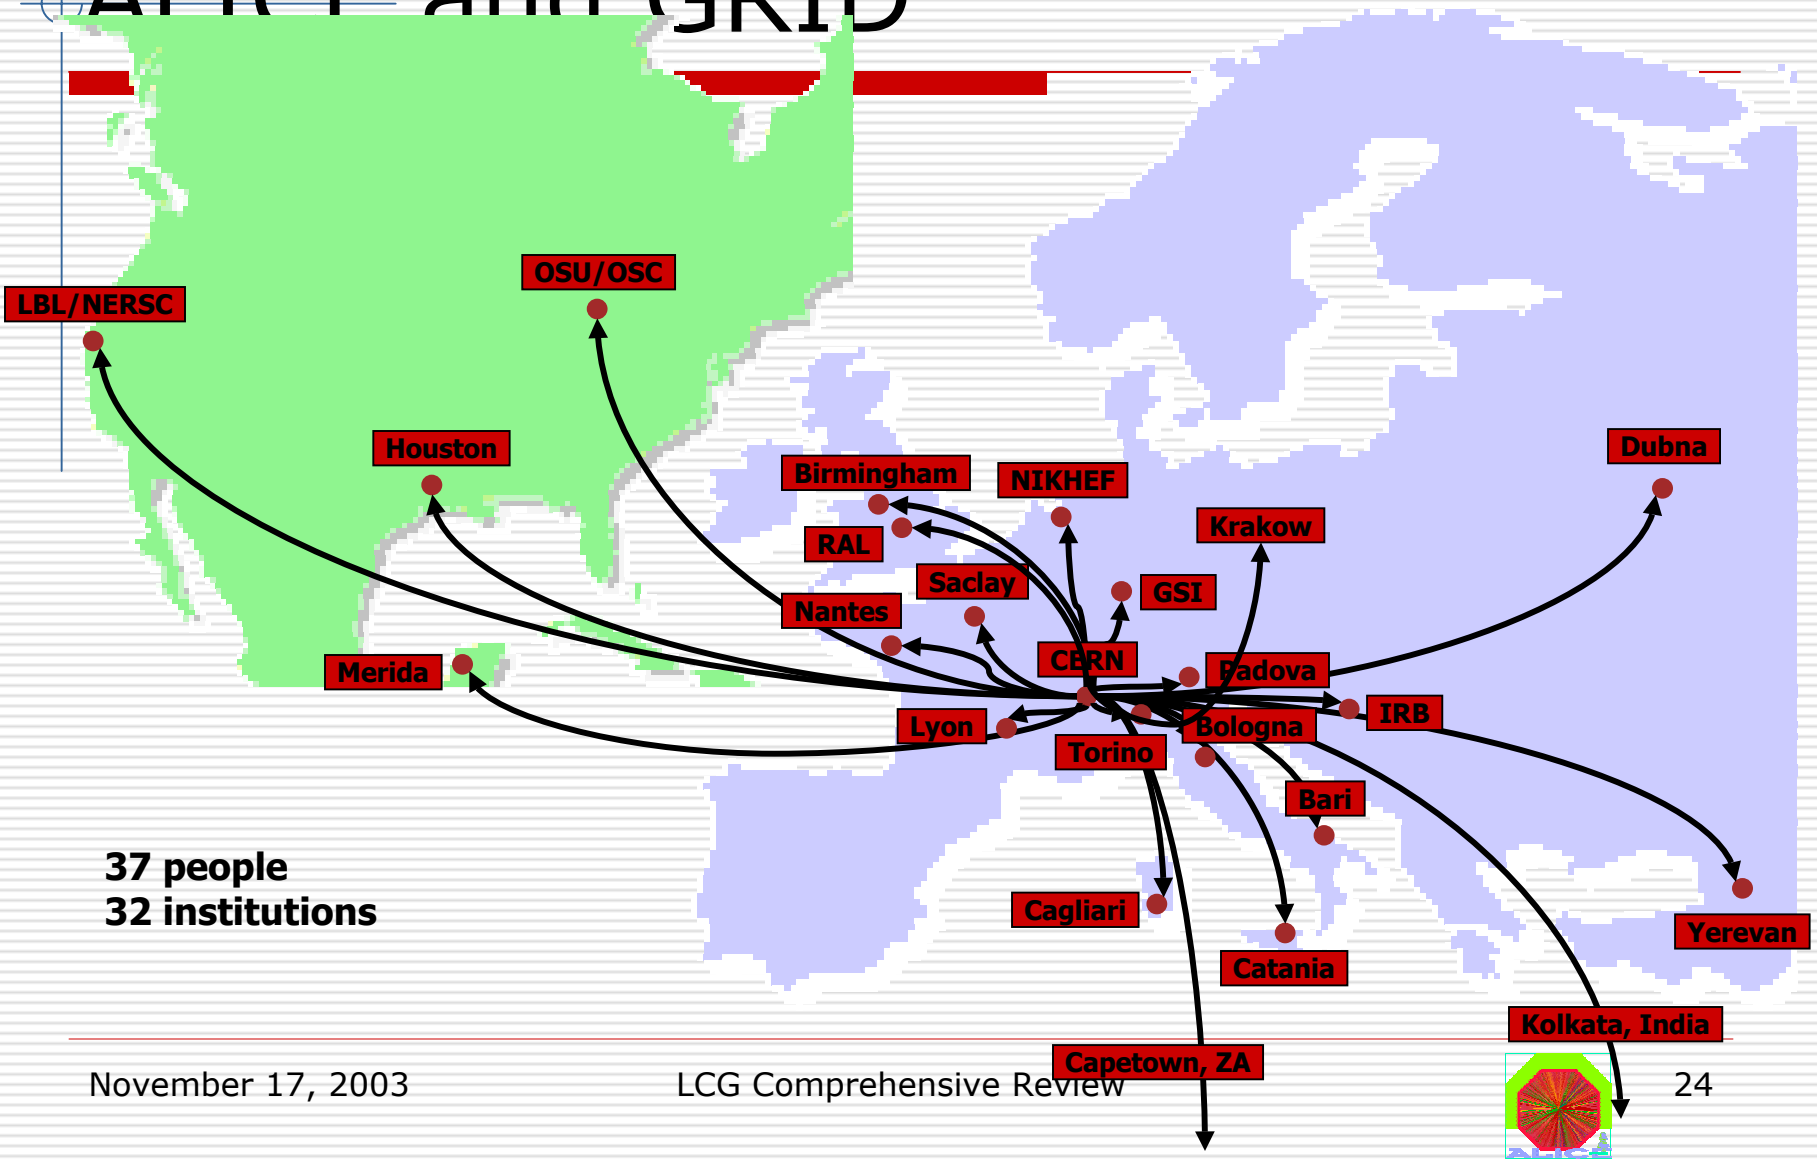

37 people
32 institutions

AliEn as a meta-GRID

AliEn User Interface		
iVDGL stack	AliEn stack	EDG stack

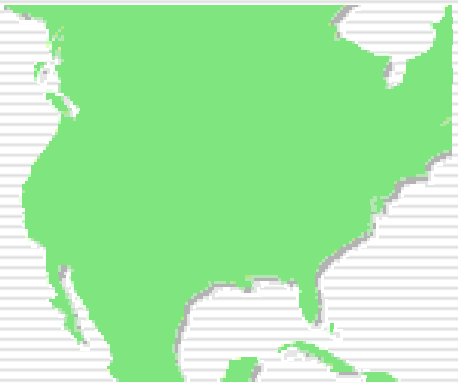

EDG CE's are seen as a single AliEn CE and EDG SE's as a single AliEn SE.

- EDG UI and AliEn suite - ClusterMonitor, CE, SE - active on an interface site
- This interface pulls jobs from the AliEn server, generates JDL requirements ans submits jbs to EDG RB.
- The job launches AliEn processes in the WN for communicatin with the AliEn server.

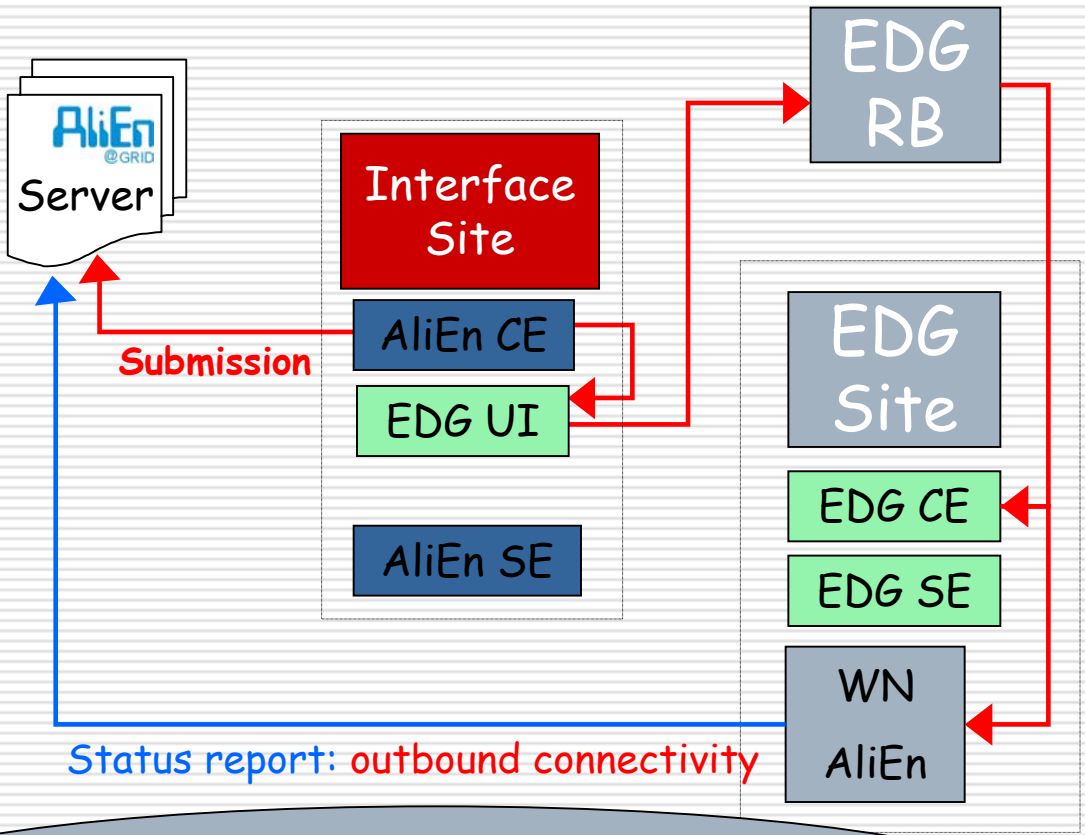

AliEn is part of the ALICE environnement (no daemons in WN)

Communication with EDG

- Outputs from EDG CE saved on an EDG SE. Information is passed to AliEn via .Brokerinfo file.
- Files are recorded in EDG catalog

  Files are also recorded in AliEn catalog (EDG LFN == AliEn PFN)
- Access to files following 2 steps: through AliEn catalog and EDG catalog

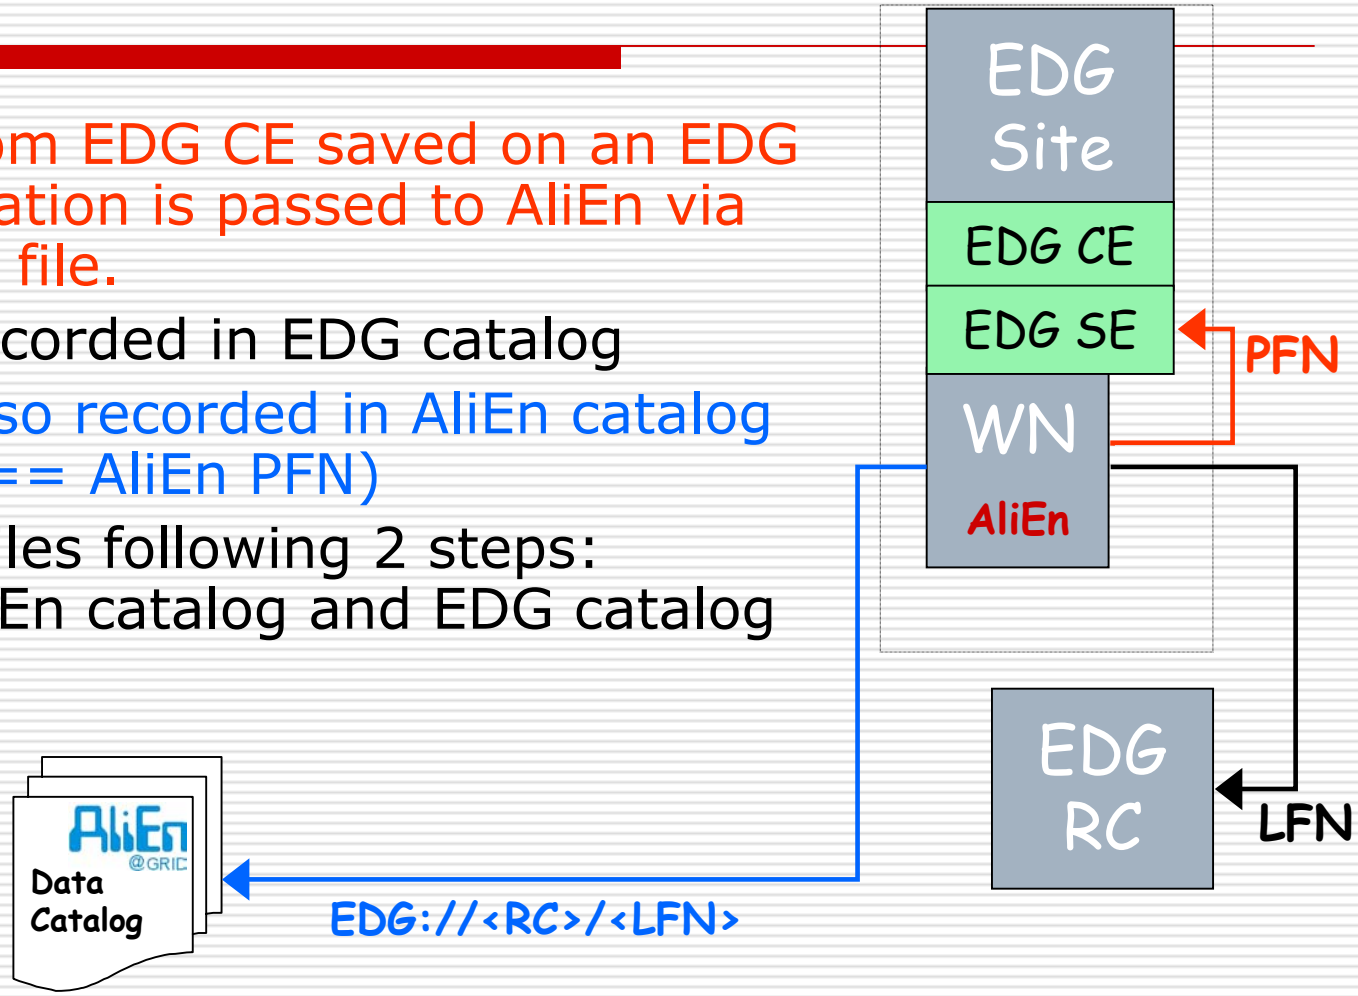

ALICE's experience with EDG

- Produce and analyse 5000 events Pb-Pb (HBT) on EDG 1.4 testbed
  - 1 event/ job, 12-24 CPU hours, 1.8 GB, total of 9TB, ~90,000 CPU h
- Start March 15-03 -- Stop May 15-03
- 450 events produced (6/day)
  - Output on EDG disk SE's (300TB) and MSS SE's (150TB at CNAF and CERN), also registered in AliEn
- No analysis done
- Conclusions
  - Average Efficiency: 35% (more jobs mean lower efficiency)
  - Application Testbed unstable on the scale of job duration (24 h)
  - Most of the jobs failed because of services failures
  - Difficult to track down the errors and to recover
  - EDG 1.4 is far from production status, we are trying now with EDG 2.0

ALICE's experience with LCG

- Results: monitoring of efficiency and stability versus job duration and load
  - Efficiency (algorithm completion): if the system is stable eff ~100%, if any instability eff=0%. Disk space availability on WN.
  - Efficiency(output registration to RC): 100%
  - Geographical job distribution:
    - Few sites accept event until they saturate and then RB looks for other sites
    - By submitting a bunch of jobs and no WN is available, all the jobs enter the Schedule state always on the same CE.
    - Automatic Proxy-renewal: allways OK
- ALICE is carrying out a test production on LCG1:
  - 200 Pb-Pb events
  - 1 job/event → 200 jobs
  - 1.8 GB/job → 288 GB
  - 12-24 hours per job
  - Started on November 14th
  - 17/11 11:00 : 137 done; 31 cancelled; 32 Waiting → 82.2%

AliEn Timeline

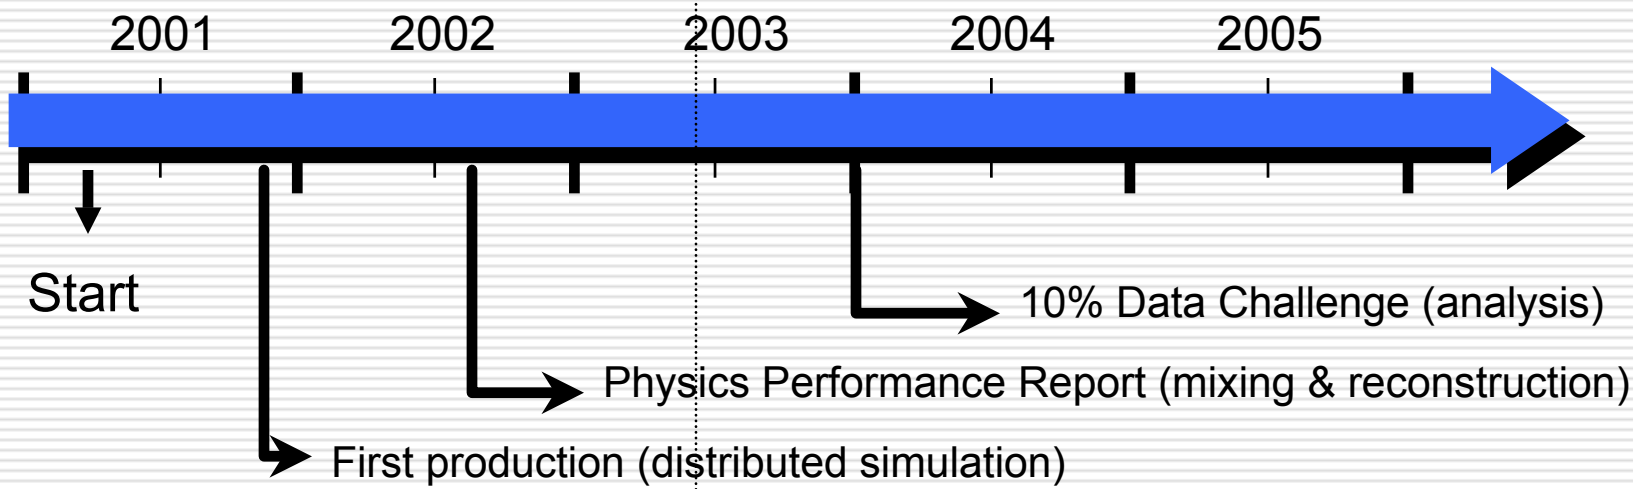

Functionality + Simulation

Interoperability + Reconstruction

Performance, Scalability, Standards + Analysis

Distributed analysis @ GRID

User

Produces
analysis code

.... and wants to run it on all data which are
tagged as <X> and belong to run <Y>
....

.... but data are distributed world wide

- Asynchronous
  - Pseudo-interactive
  - Job-splitting+batch processing (transparent to user))
  - Can be done using existing tools (AliEn + ROOT)
- True Interactive
  - Highly synchronous system
  - Instantenuous analysis results
  - Needs:
    - New functionality (AliEn + PROOF)
    - High system availability

Distributed analysis: job splitting

- ROOT/AliEn prototype demonstrated
- From an analysis macro as many jobs are generated as there are files to be analysed
- Jobs are executed in parallel
- Results in the form of TTree, TNtuple and histograms are merged
- Partial results are accessible

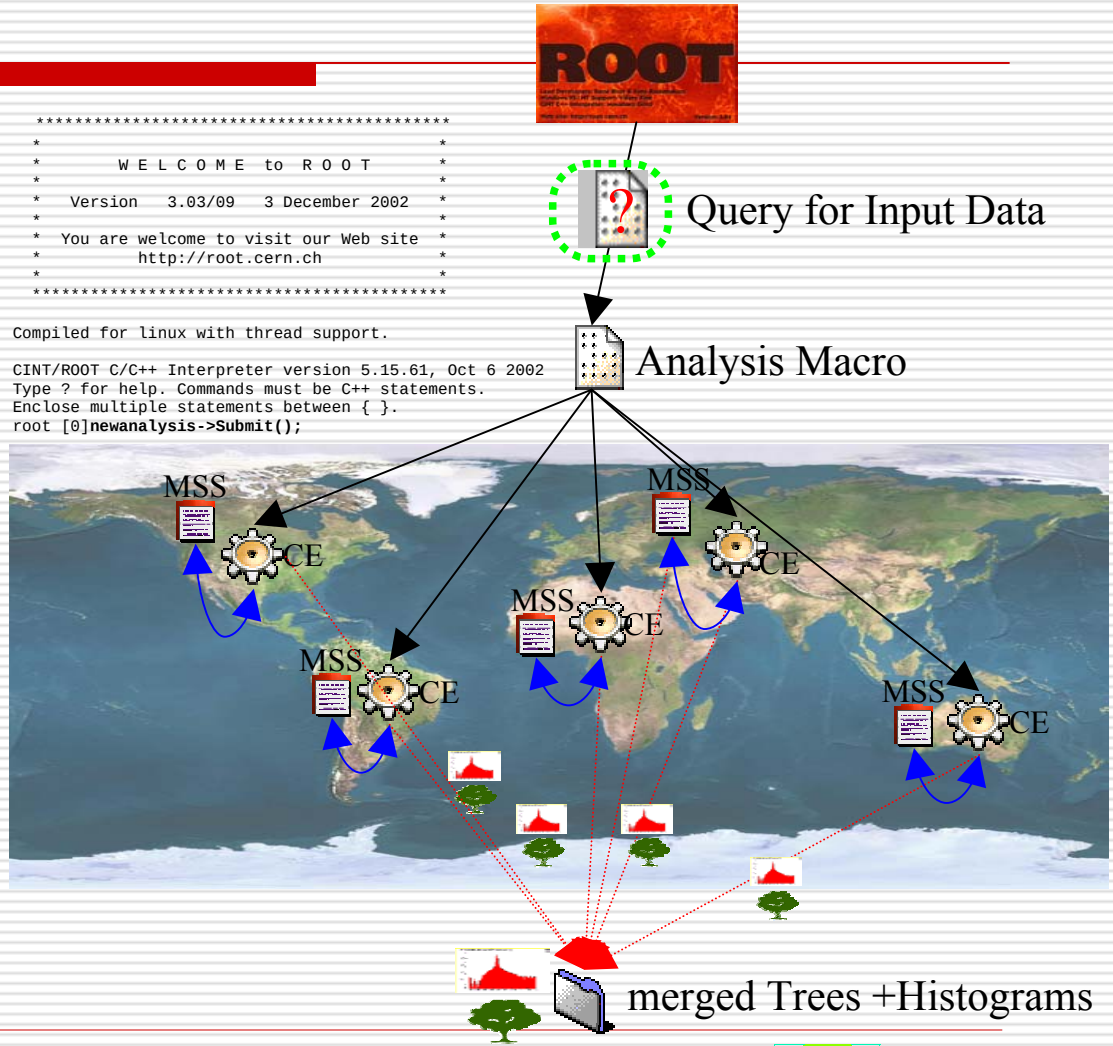


```
*******************************************
*                                         *
*       W E L C O M E  to  R O O T        *
*                                         *
*   Version   3.03/09   3 December 2002   *
*                                         *
*  You are welcome to visit our Web site  *
*          http://root.cern.ch            *
*                                         *
*******************************************

Compiled for linux with thread support.

CINT/ROOT C/C++ Interpreter version 5.15.61, Oct 6 2002
Type ? for help. Commands must be C++ statements.
Enclose multiple statements between { }.
root [0]newanalysis->Submit();
```

AliEn + ROOT

Demo at Supercomputing 2003

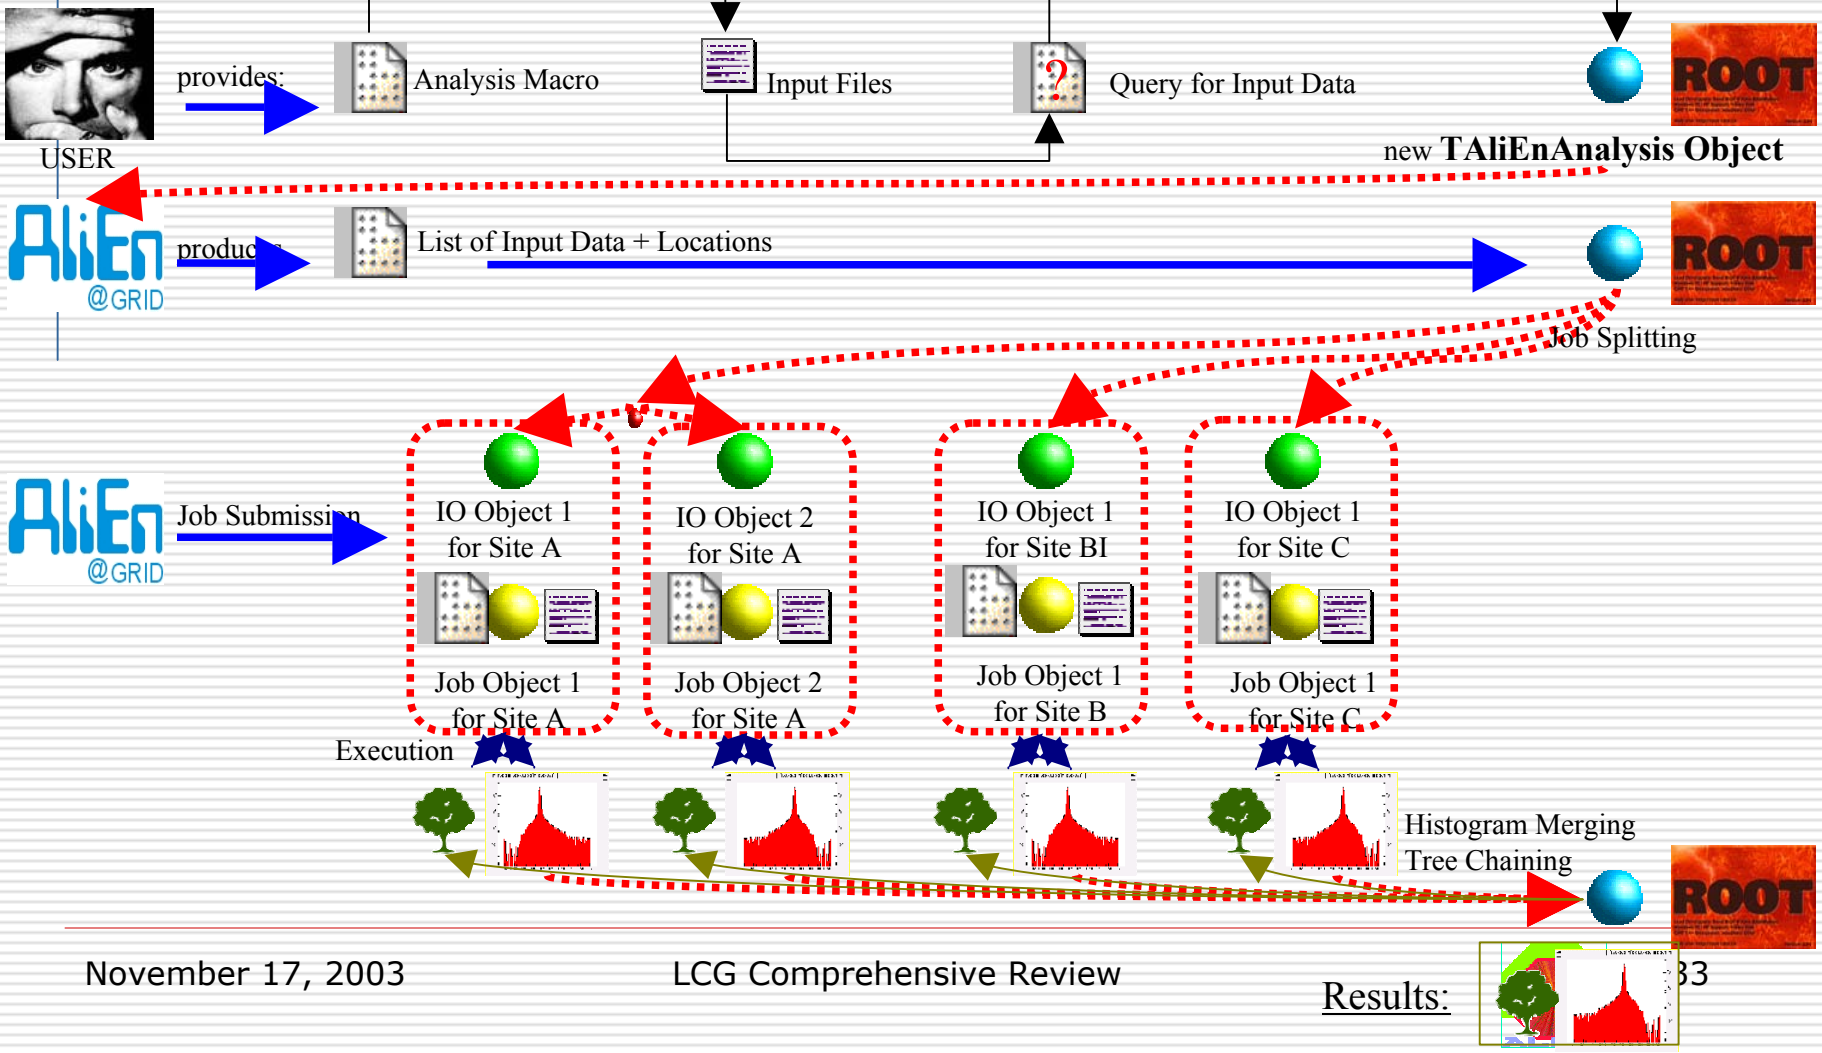

Federico Carminati — Alice Grid Proposal

13/4/2000 — Draft Revision 04

The ALICE Grid Proposal for Using Computational Grids

F.Carminati

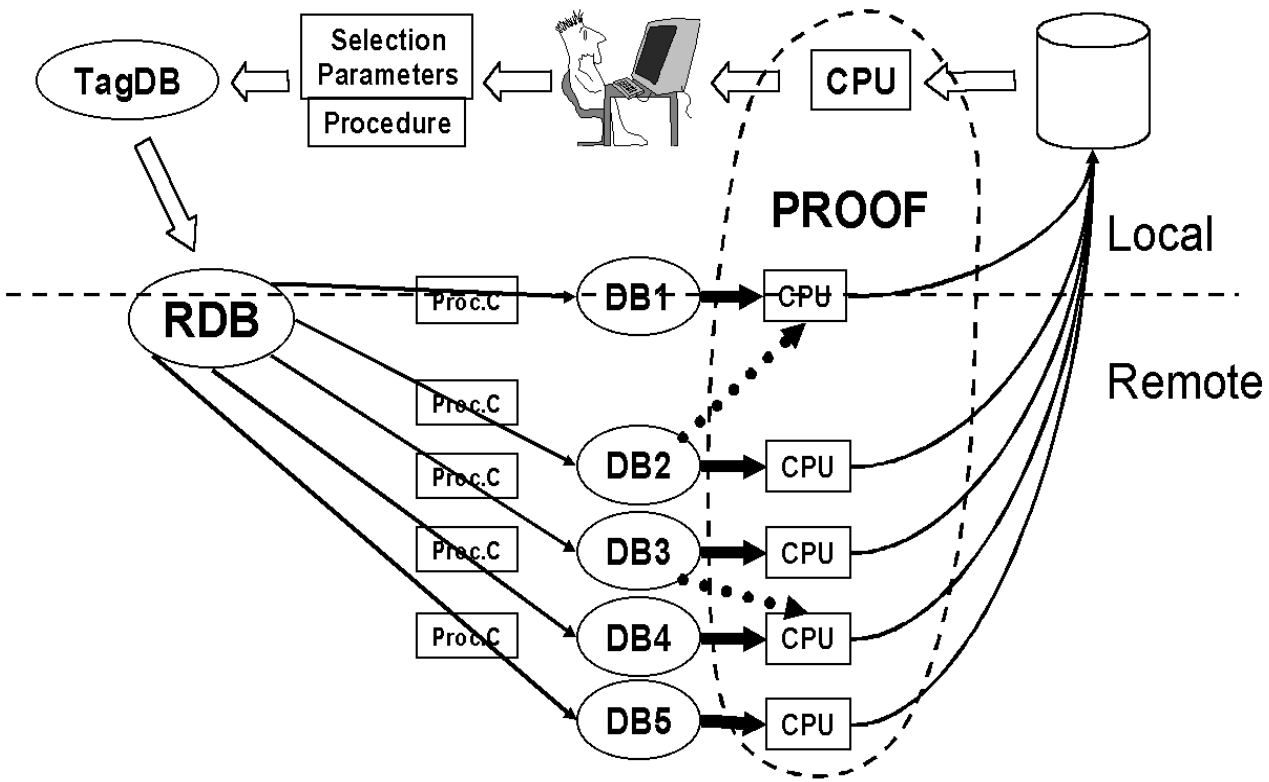

Figure 1: Schematic view of the analysis via PROOF facility

ALICE Physic Data Challenges

Period (milestone)	Fraction of the final capacity (%)	Physics Objective
06/01-12/01	1%	pp studies, reconstruction of TPC and ITS
06/02-12/02	5%	☐ First test of the complete chain from simulation to reconstruction for the PPR ☐ Simple analysis tools ☐ Digits in ROOT format
01/04-06/04	10%	☐ Complete chain used for trigger studies ☐ Prototype of the analysis tools ☐ Comparison with parameterised MonteCarlo ☐ Simulated raw data
01/06-06/06	20%	• Test of the final system for reconstruction and analysis

Objectives of PDC 3

- We have to use a *real* (!) generator (HIJING) that contains (mini-)jets to simulate the *underlying event*
  - For p_t ~ 20 GeV/c → ~1 jet per event
  - For p_t ~200 GeV/c → 1000 jet/month (reuse events 50 times)
- Verify the ability to (simulate and) process ~10% of the data of a standard data-taking year
  - Evaluate consequences for the MW and offline frameworks
- See how much data we can simulate, process (including mixing) and analyse in 6 months on LCG and private resources
- Part of the raw data will be shipped to CERN and redistributed for reconstruction and analysis
  - Mimic when all data come from CERN
  - Optimize the reconstruction to get results in time

PDC3 Strategy

- RAW Production
- Transfert RAW to CERN
- RAW Reconstruction
- Analyse

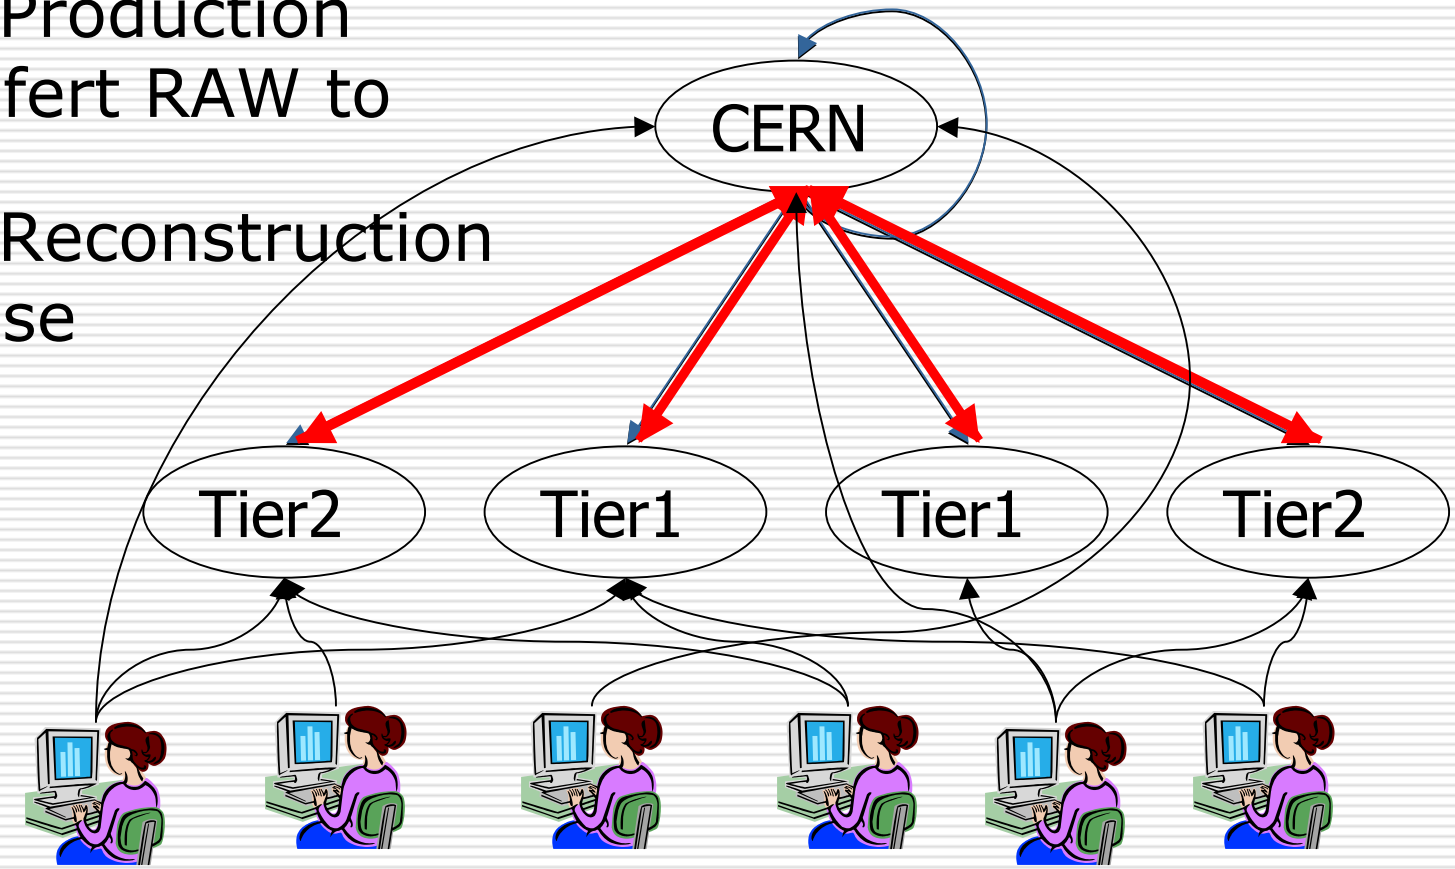

Resources for PDC3

- **Simulation**
  - 10^5 Pb-Pb + 10^7 p-p
  - Distributed production, data replication at CERN
- **Reconstruction and analyse**
  - 5×10^6 Pb-Pb + 10^7 p-p
  - Reconstruction at CERN and outside according available resources
  - CERN is the data source
- **Resources (CPU and Storage)**
  - 2004 Q1: 1354 KSI2K and 165 TB
  - 2004 Q2: 1400 KSI2K and 301 TB
- **Band width**
  - Simulation in 2004 Q1
  - ~90 TB to be shipped to CERN in 2 months (~10 days using 10% of CERN bandwidth).

Computing needs for PDC III

Flexibility of distributed computing model

Alternative scenarios

Year,Quarter	04Q1	04Q2
CPU	ksi95 month	ksi95 month
Total Requirement for Simulation	65	
Total Requirement for reconstruction	124	156
Total Requirement for analysis	249	311
Total CPU Requirement	438	467
STORAGE	TB	TB
Total Requirement for Simulation	21	21
Total Requirement for reconstruction	63	187
Total Requirement for analysis		13
Total Storage requirements	84	220
LOAD BALANCE CPU	ksi95 month	ksi95 month
CERN	131	140
TIER-1s OUT CERN	175	187
Tier-2s	131	140
Total CPU Requirement (check)	438	467
LOAD BALANCE STORAGE	TB	TB
CERN	25	66
TIER-1s OUT CERN	34	88
Tier-2s	25	66
Total Storage requirements (check)	84	220

Year,Quarter	04Q1	04Q2
LOAD BALANCE CPU	ksi95 month	ksi95 month
CERN	66	70
TIER-1s OUT CERN	241	257
Tier-2s	131	140
Total CPU Requirement (check)	438	467
LOAD BALANCE STORAGE	TB	TB
CERN	13	33
TIER-1s OUT CERN	46	121
Tier-2s	25	66
Total Storage requirements (check)	84	220

Resources available

		Q4Q1	Q4Q2
CPU Requirements kSI2k		1354	1400
LCG Declared Capacity for ALICE		941	941
Storage Requirements - total TB active data		165	301
LCG Declared Capacity for ALICE	Disk	192	192
	Tapes	578	578
	Total	770	770

- 200TB needed after PDC3 during 2004
- Available resources estimated considering simultaneous use by 4 experiments.

ALICE PDC3 & LCG

- All the production will be started via AliEn, the analysis will be done via Root/Proof/AliEn
- LCG-1 will be one CE element of AliEn, which will integrate seamlessly LCG and non LCG resources
- If LCG-1 works well, it will *suck* a large amount of jobs, and it will be used heavily
- If LCG-1 does not work well, AliEn will privilege other resources, and it will be less used
- In all cases we will use LCG-1 as much as possible
- We will not take any *decision*: the system performance will decide for us.

ALICE ADC's

Date	MBytes/s	Tbytes to MSS	Offline milestone
10/2002	200	200	Rootification of raw data -Raw data for TPC and ITS
9/2003	300	300	Integration of single detector HLT, at least for TPC and ITS - Quasi on-line reconstruction at CERN - Partial data replication to remote centres
5/2004	450	450	HLT prototype for all detectors that plan to use it - Remote reconstruction of partial data streams - Raw digits for barrel and MUON
5/2005	750	750	Prototype of the final HLT software Prototype of the final remote data replication (Raw digits for all detectors)
5/2006	750 (1250 if possible)	750 (1250 if possible)	Final test (Final system)

ALICE Data challenges

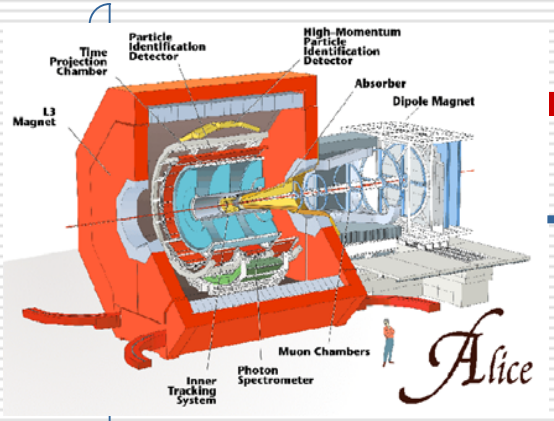

Raw Data

Simulated Data

AliRoot

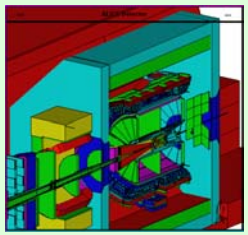

GEANT3
GEANT4
FLUKA

DAQ

ROOT I/O

CERN
TIER 0
TIER 1

ROOT

CASTOR

GRID

Regional
TIER 1
TIER 2

ADC IV performances

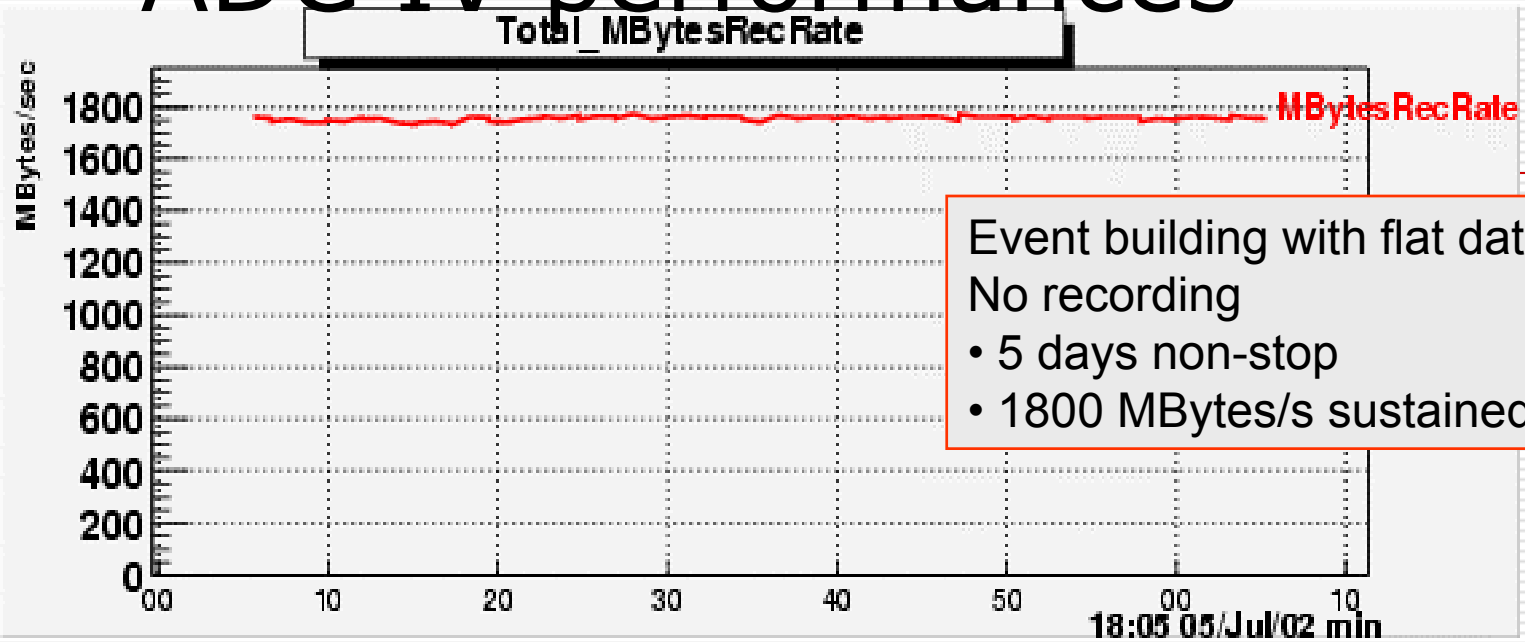

Event building with flat data traffic
No recording
- 5 days non-stop
- 1800 MBytes/s sustained

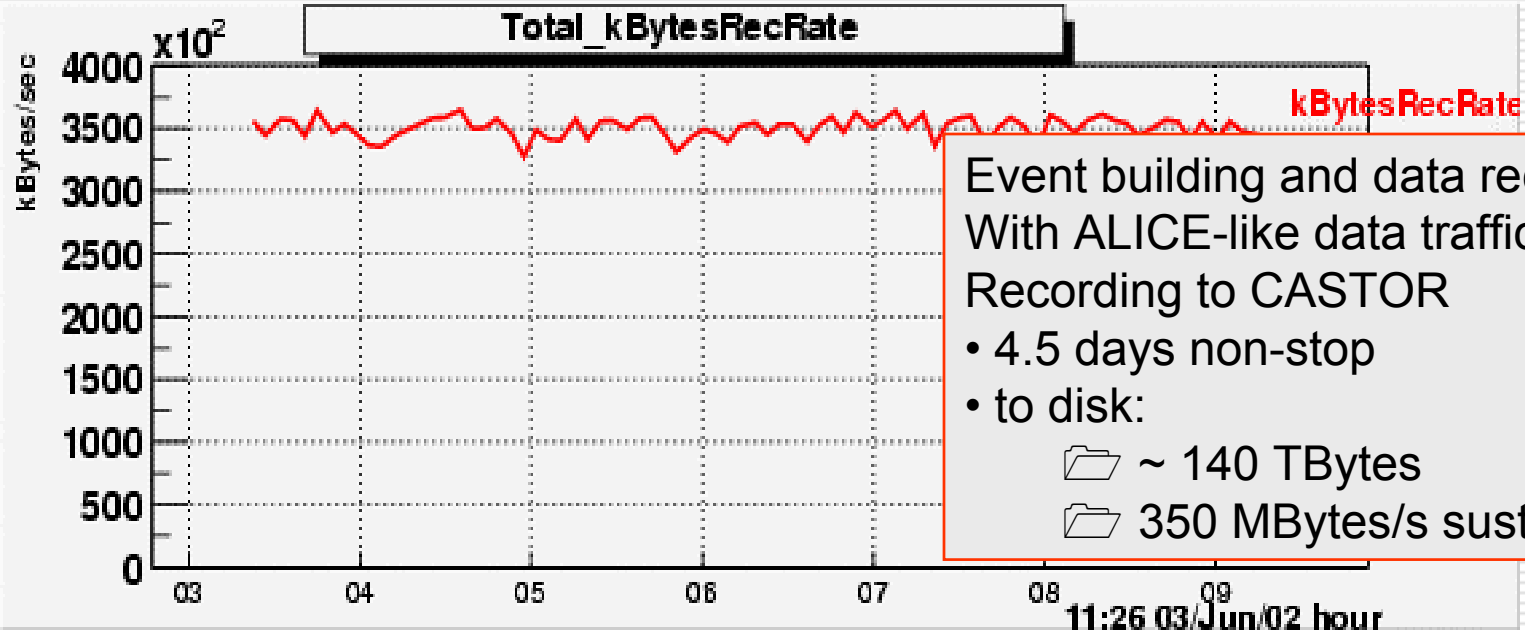

Event building and data recording
With ALICE-like data traffic
Recording to CASTOR
- 4.5 days non-stop
- to disk:
  - ~ 140 TBytes
  - 350 MBytes/s sustained

Relations with HLT

- In ALICE Offline, HLT and DAQ are three distinct projects
- Cooperation between HLT and Offline
  - HLT is using the common AliRoot framework to do simulation
  - Some of the HLT algorithms are integrated in AliRoot for testing
- As the Offline, HLT coordinates the activities of the different subdetector projects
- HLT main trust is the definition of the HLT architecture (HW and SW) and some seminal work on algorithms
- More work on algorithms is done in collaboration with the subdetector projects
- Integration of the three projects is also tested during DCs

Conclusions

- ALICE has solutions that are evolving into a solid computing infrastructure
- Major decisions have been taken and users have adopted them
- Collaboration between physicists and computer scientists is excellent
- The tight integration with ROOT allows a fast prototyping and development cycle
- AliEn provides a complete GRID solution adapted to HEP needs and it allowed us large productions with very few people “in charge”
- Many ALICE-developed solutions have a high potential to be adopted by other experiments and become “common solutions”