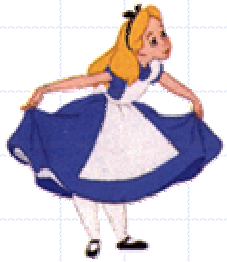

ALICE

F.Carminati/CERN

Root WorkShop

Geneva, October 14

Alice collaboration

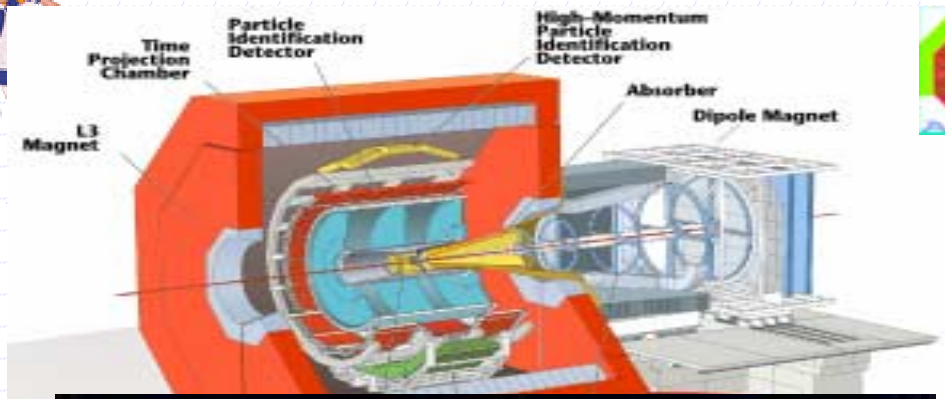

Total weight 2,000t
Overall length 17.3m

online system

multi-level trigger

nd
e

00t
0m

esla

ure
/sec)
proc
GB/
2 -
30
(1.25
data
offl

U. S. S. MISSOURI BB - 63

Total weight 53,000t
Overall length 270.4m

Back to the origin of time

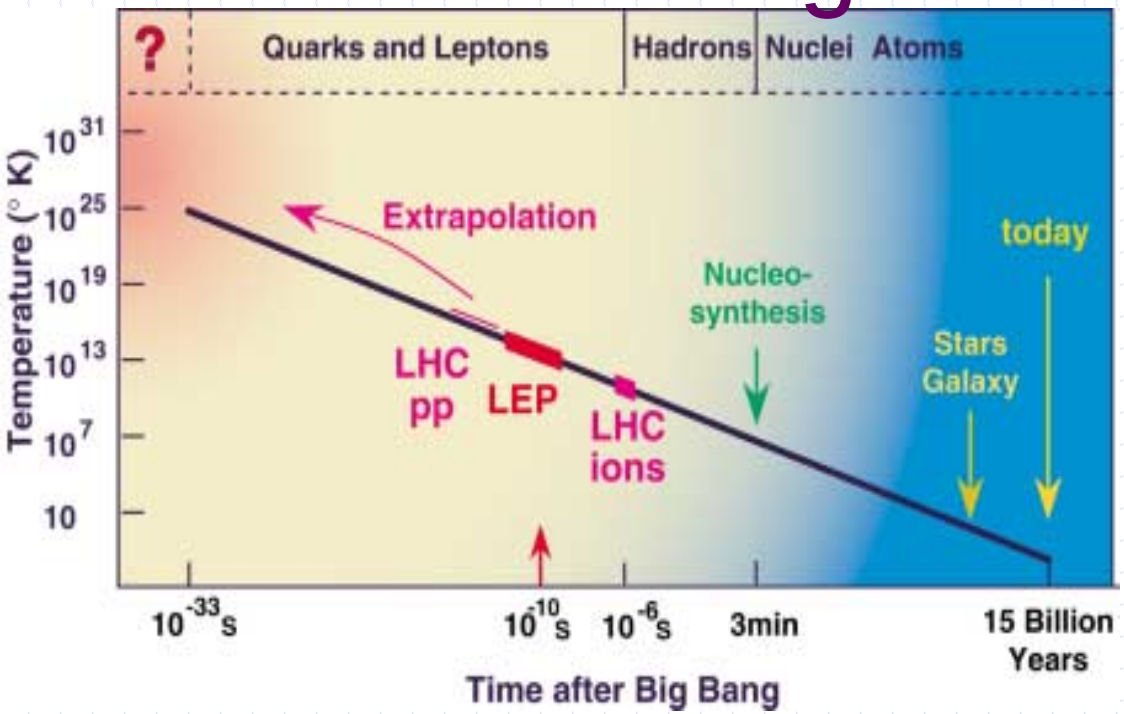

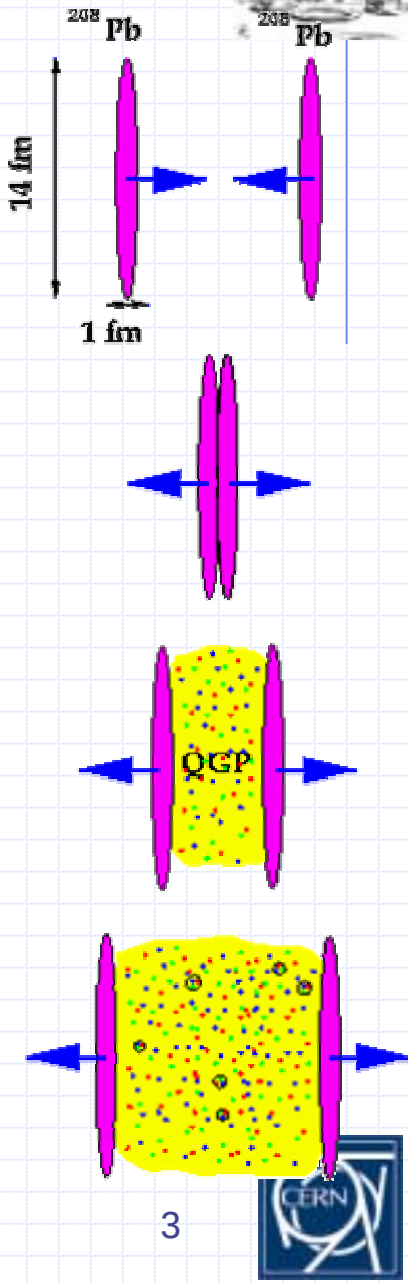

- 10 ms (T = 170 MeV) after the Big Bang quarks and gluons combined in protons and neutrons
- The initial quark and gluons plasma was first recreated at CERN SPS
- At LHC the energy will be 300 times (1144 TeV for Pb-Pb, 5.5 TeV/N)

Strategic decision in 1998

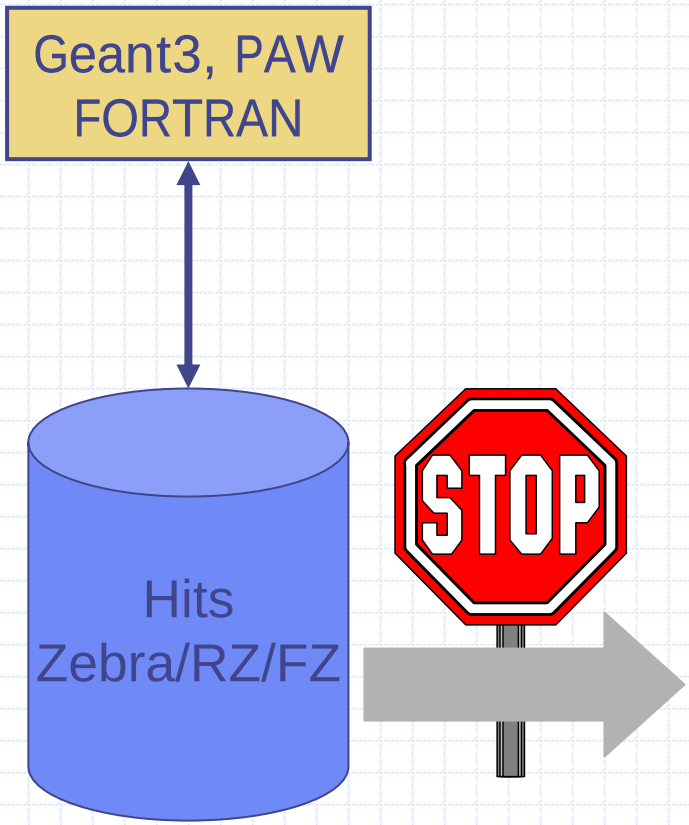

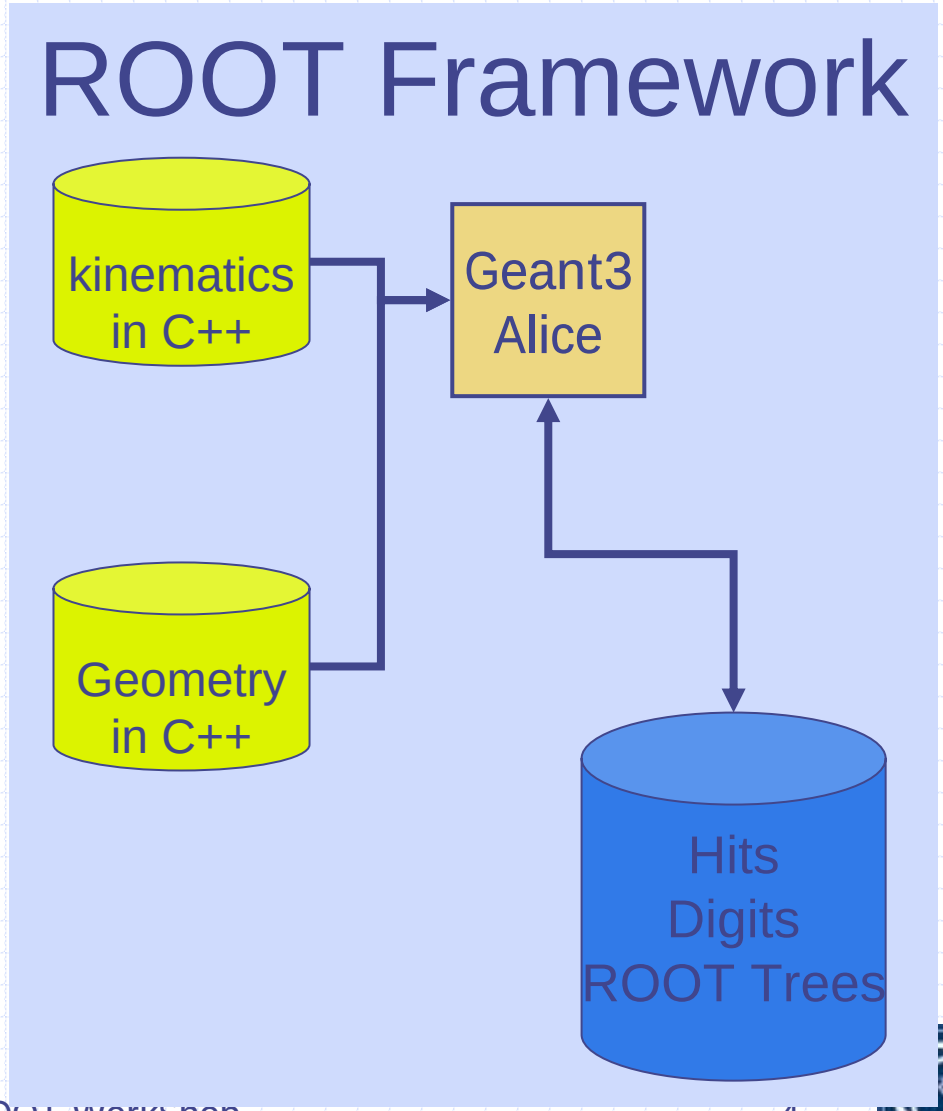

AliRoot evolution schema

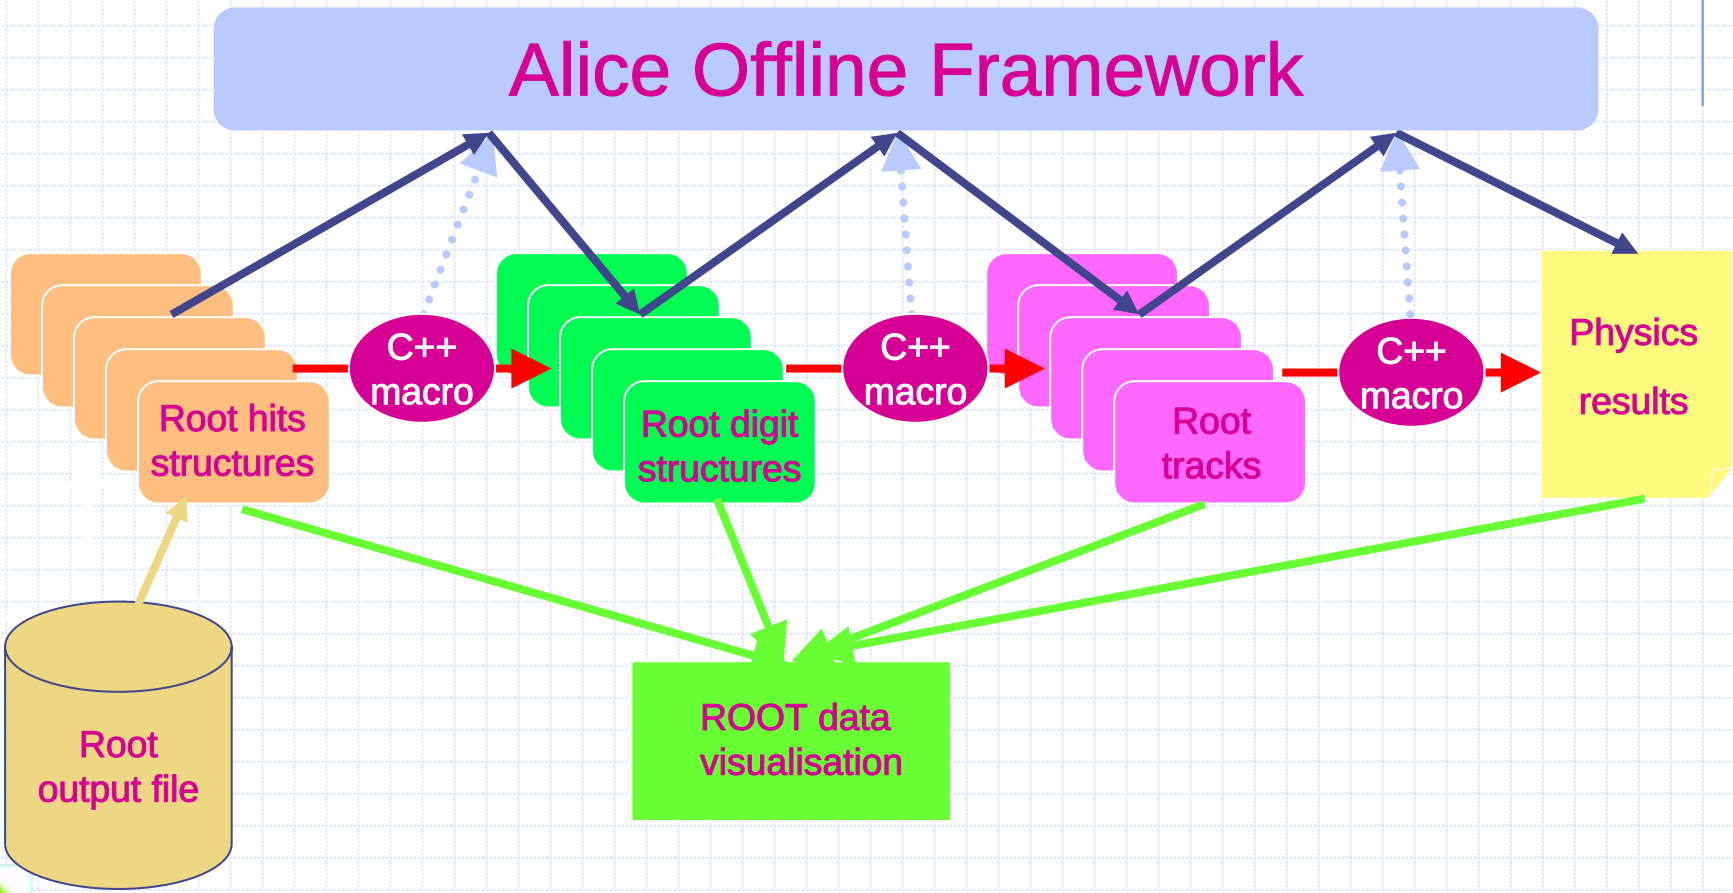

Framework

- AliRoot framework
  - C++: 400kLOC + 225kLOC (generated) + macros: 77kLOC
  - FORTRAN: 13kLOC (ALICE) + 914kLOC (external packages)
  - Maintained on Linux (any version!), HP-UX, DEC Unix, Solaris
  - Works also with Intel icc compiler
- Two packages to install (ROOT+AliRoot)
  - Less that 1 second to link (thanks to 37 shared libs)
  - 1-click-away install: download and make (non-recursive makefile)
- AliEn
  - 25kLOC of PERL5 (ALICE)
  - ~1.5MLOC of PERL5 (opens source components)
- Installed on more than 30 sites by physicists
  - >50 users develop AliRoot from detector groups
  - 70% of code developed outside, 30% by the core Offline team

AliRoot layout

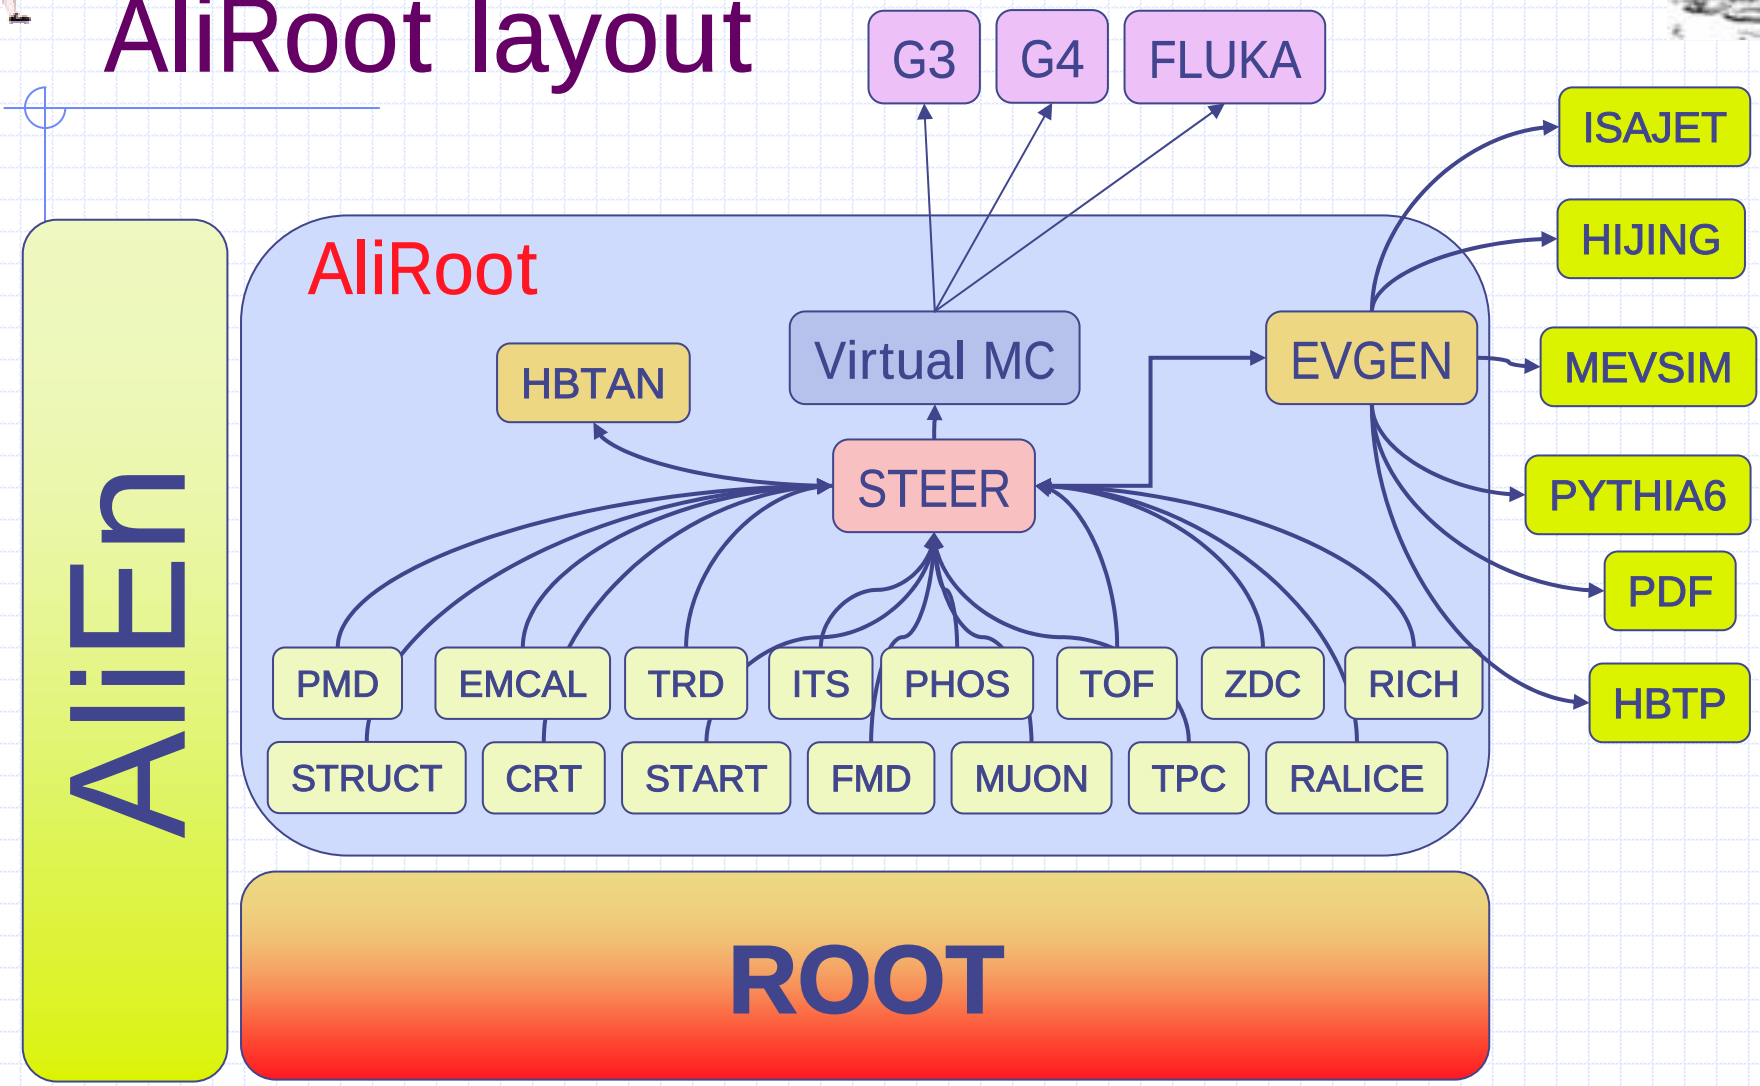

Whiteboard Data Communication

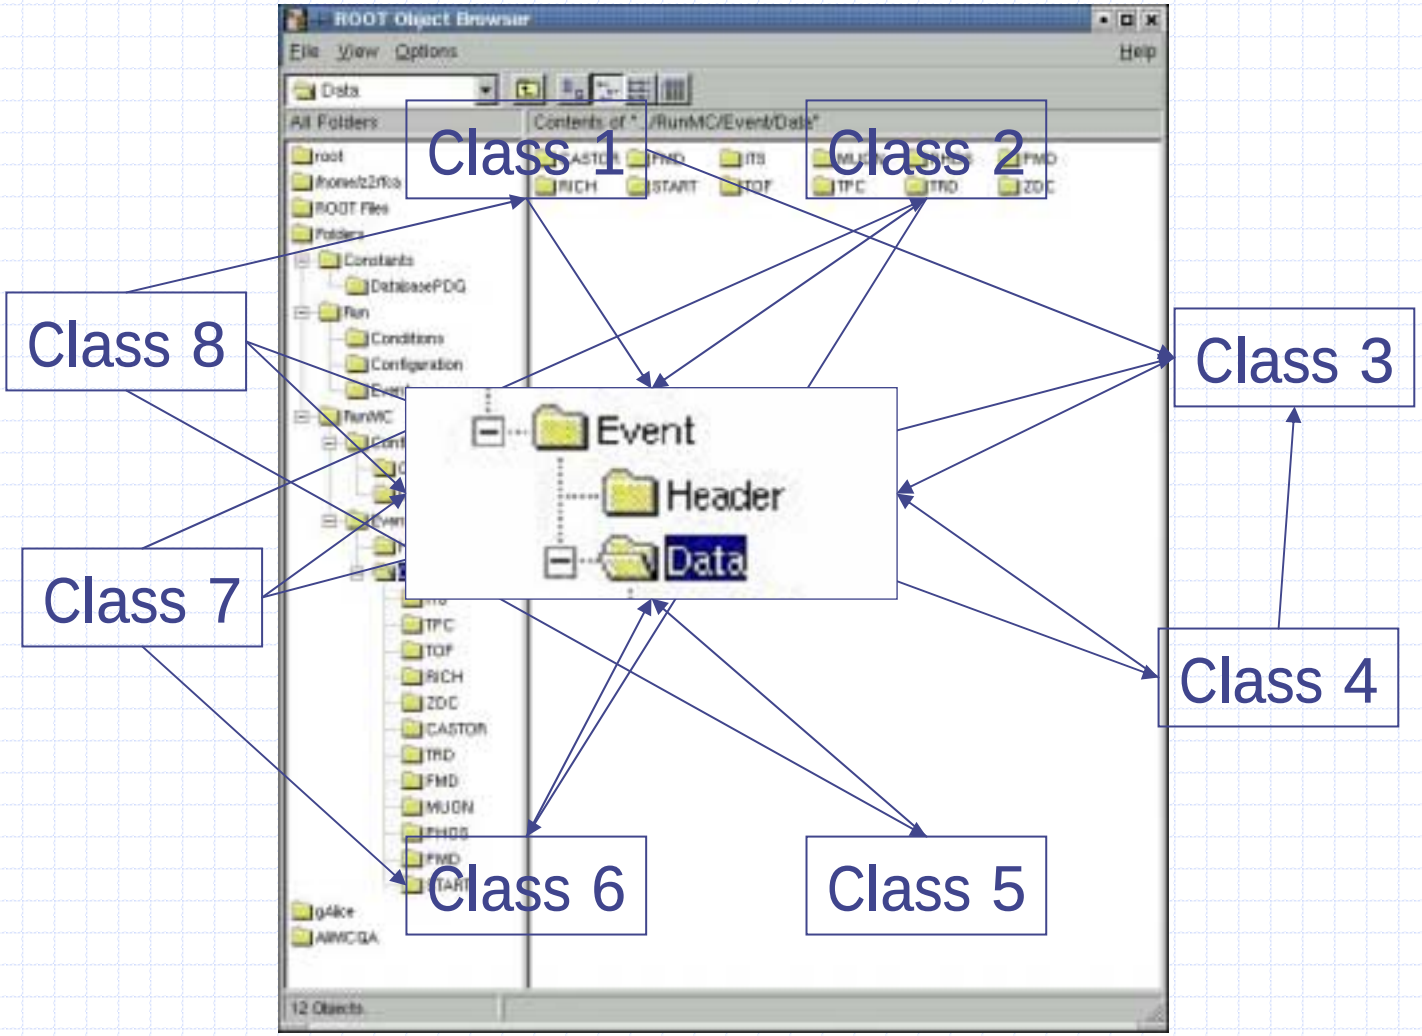

ALICE Event/100

ALICE Detector complexity
similar to ATLAS and CMS

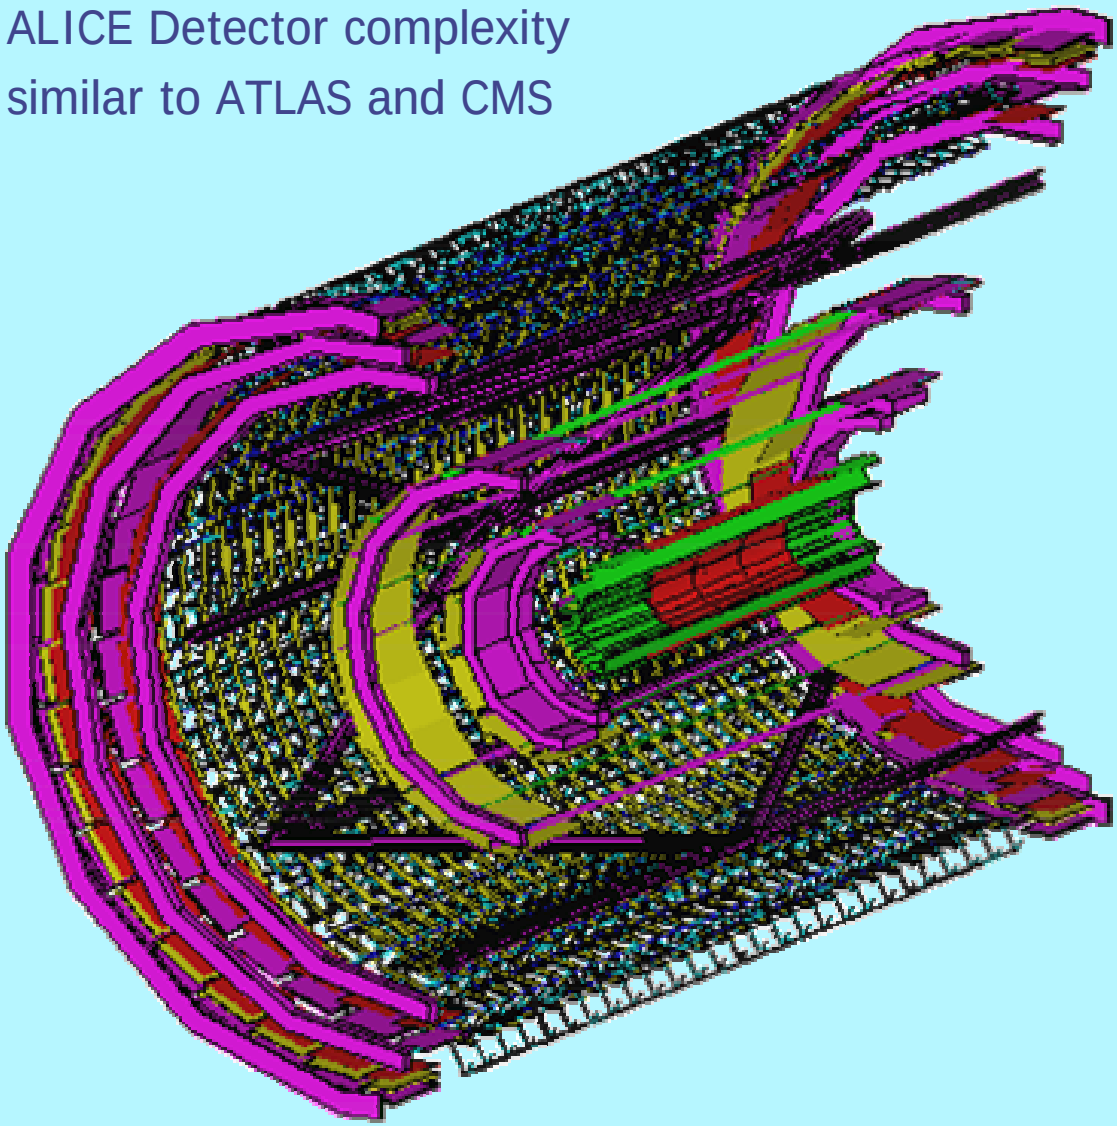

- GEANT3
  - Developed in 1981
  - Still used by the majority of experiments
- Geant4
  - A huge investment
  - Slow penetration in HEP experiments
- FLUKA
  - State of the art for hadronics and neutron physics
  - Difficult to use for full detector simulation

The Virtual MC

Geant3.tar.gz includes an upgraded Geant3 with a C++ interface

Geant4_mc.tar.gz includes the TVirtualMC <--> Geant4 interface classes

User Code

VMC

G3

G4

FLUKA

G3 transport

G4 transport

FLUKA transport

Reconstruction

Visualisation

Generators

Virtual Geometrical Modeller

Geometrical Modeller

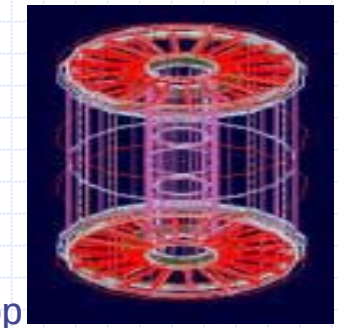

Alice

3 million volumes

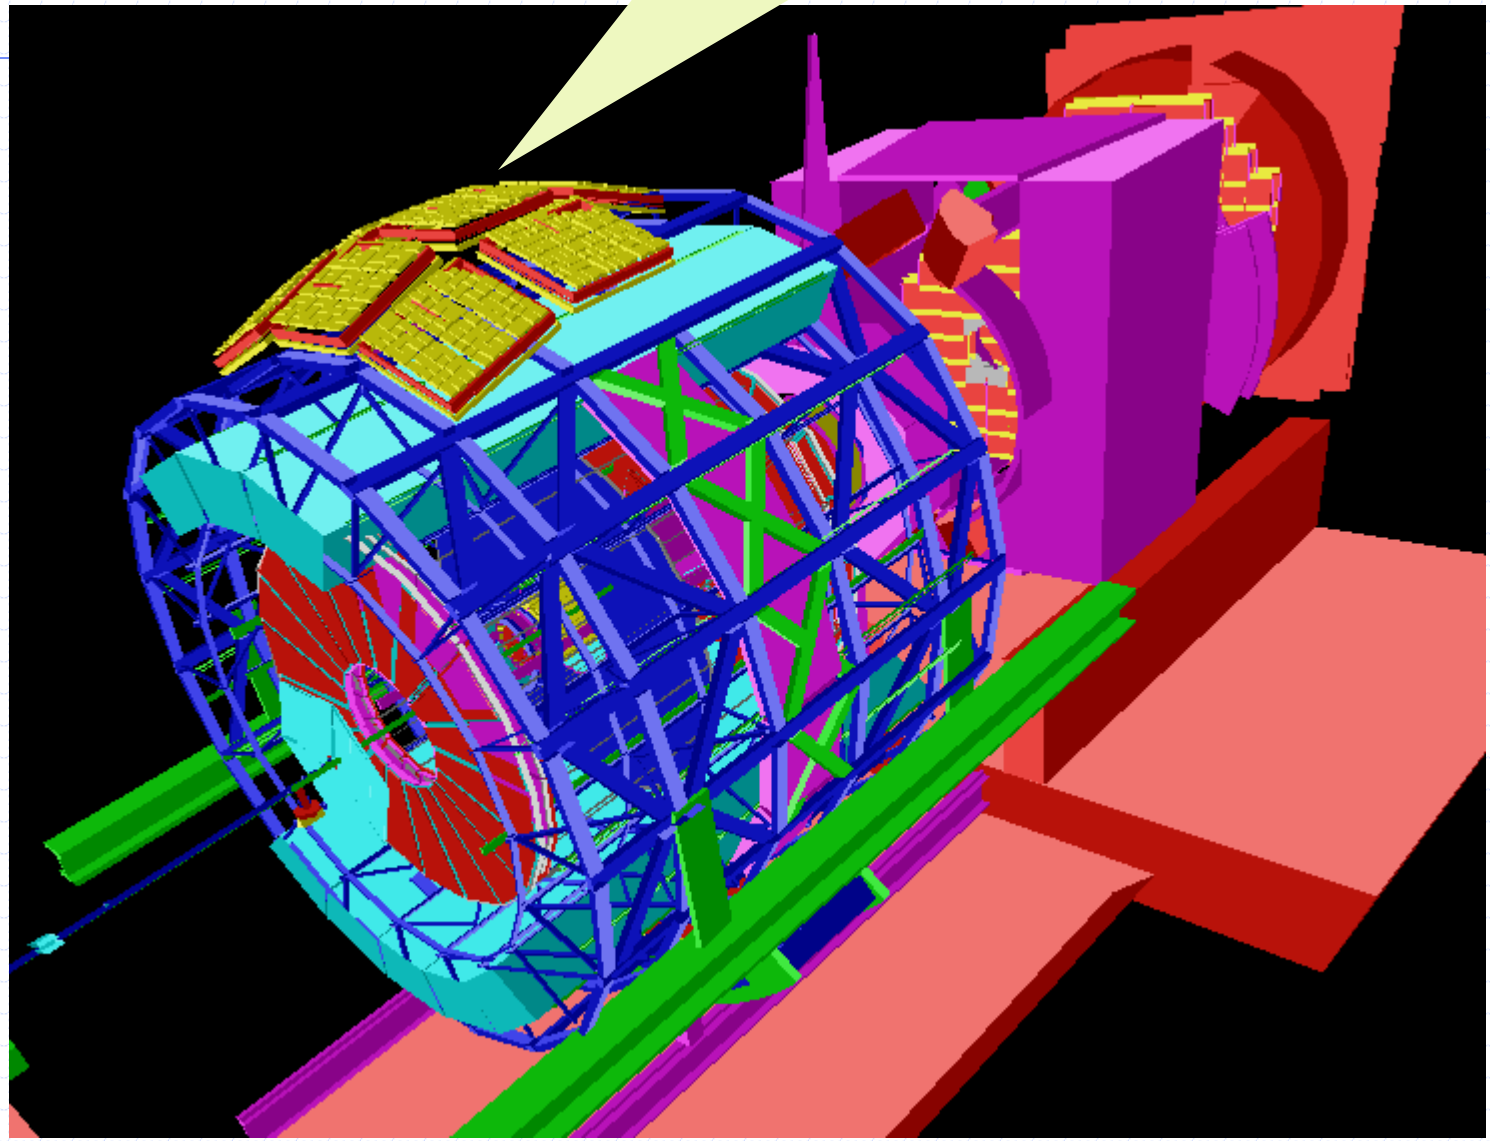

Tracking efficiencies

- TPC+ITS tracking efficiency at dN/dy 8000
  - standard definition of good vs. fake tracks requires all 6 ITS hits to be correct
    - most of 'incorrect' tracks just have one bad point
  - when this definition is relaxed to 5 out of 6 hits the efficiency is better by 7 - 10 % (fake track probability correspondingly lower)

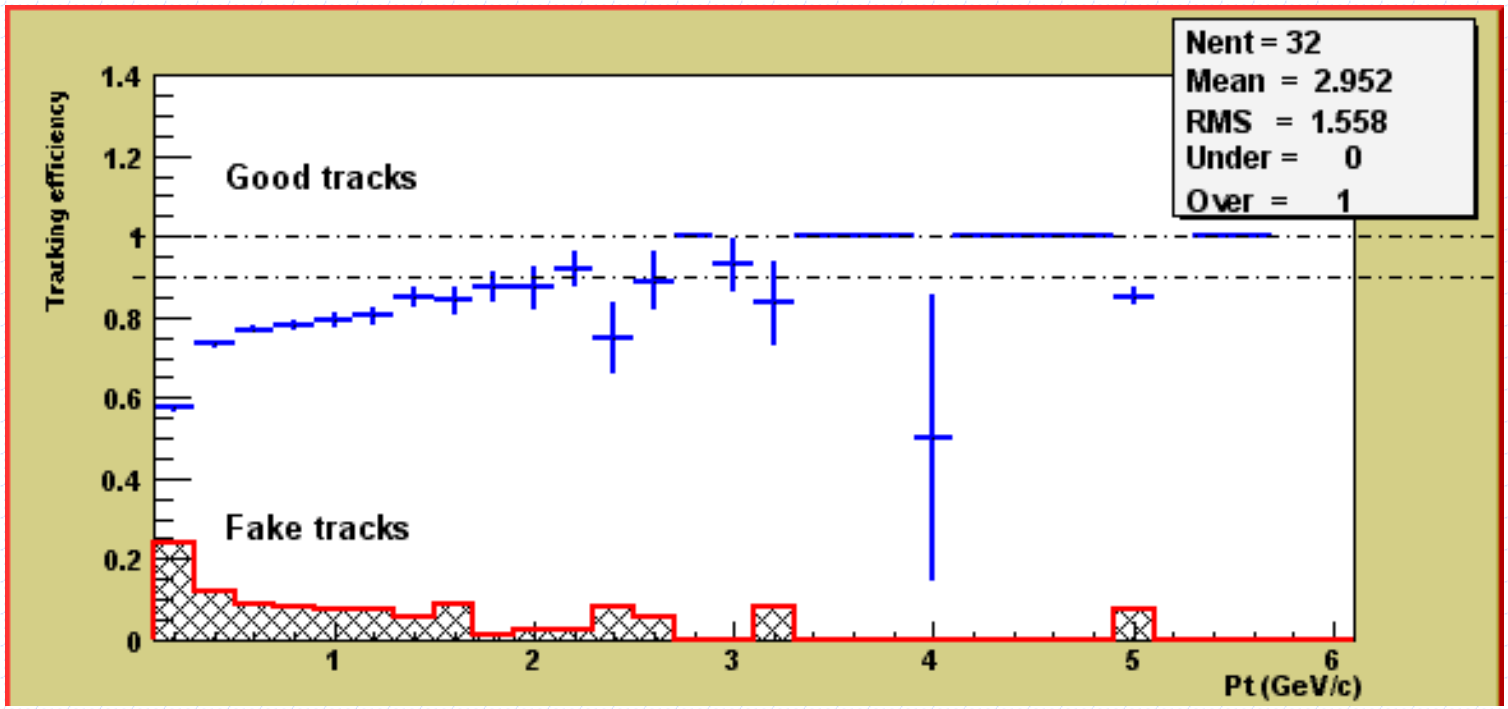

ALICE Physics Performance Report

- 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

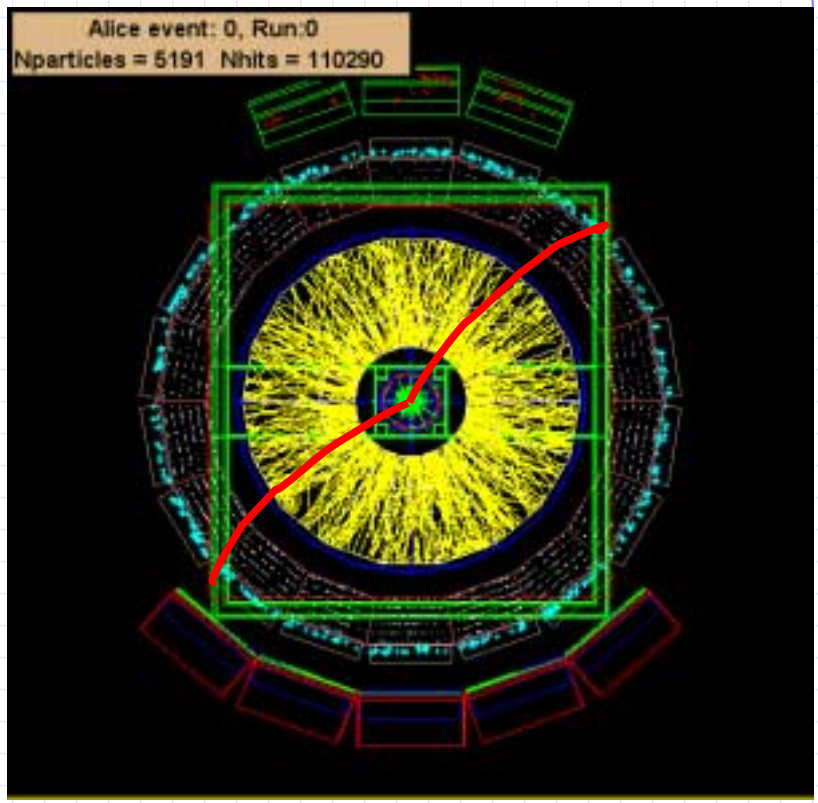

The distributed computing model

- Basic principle
  - Every physicist should have in principle equal access to the data and to the resources
- The system will be extremely complex
  - Number of components in each site
  - Number of sites
  - Different tasks performed in parallel: simulation, reconstruction, scheduled and unscheduled analysis
- Bad news is that the basic tools are missing
  - Distributed resource management
  - Distributed namespace for files and objects
  - Distributed authentication
  - Local resource management of large clusters
  - Data replication and caching
  - WAN/LAN monitoring and logging
- Good news is that we are not alone
  - All the above issues are central to the new developments going on in the US and now in Europe under the collective name of GRID

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

AliEn is used to provide GRID component for MammoGRID, a 3 year, 2M€ EU funded project started in September

AliEn (P.Buncic, P.Saiz, J-E.Revschbak)

Architecture

Data access

File catalogue

Bookkeeping

Authentication

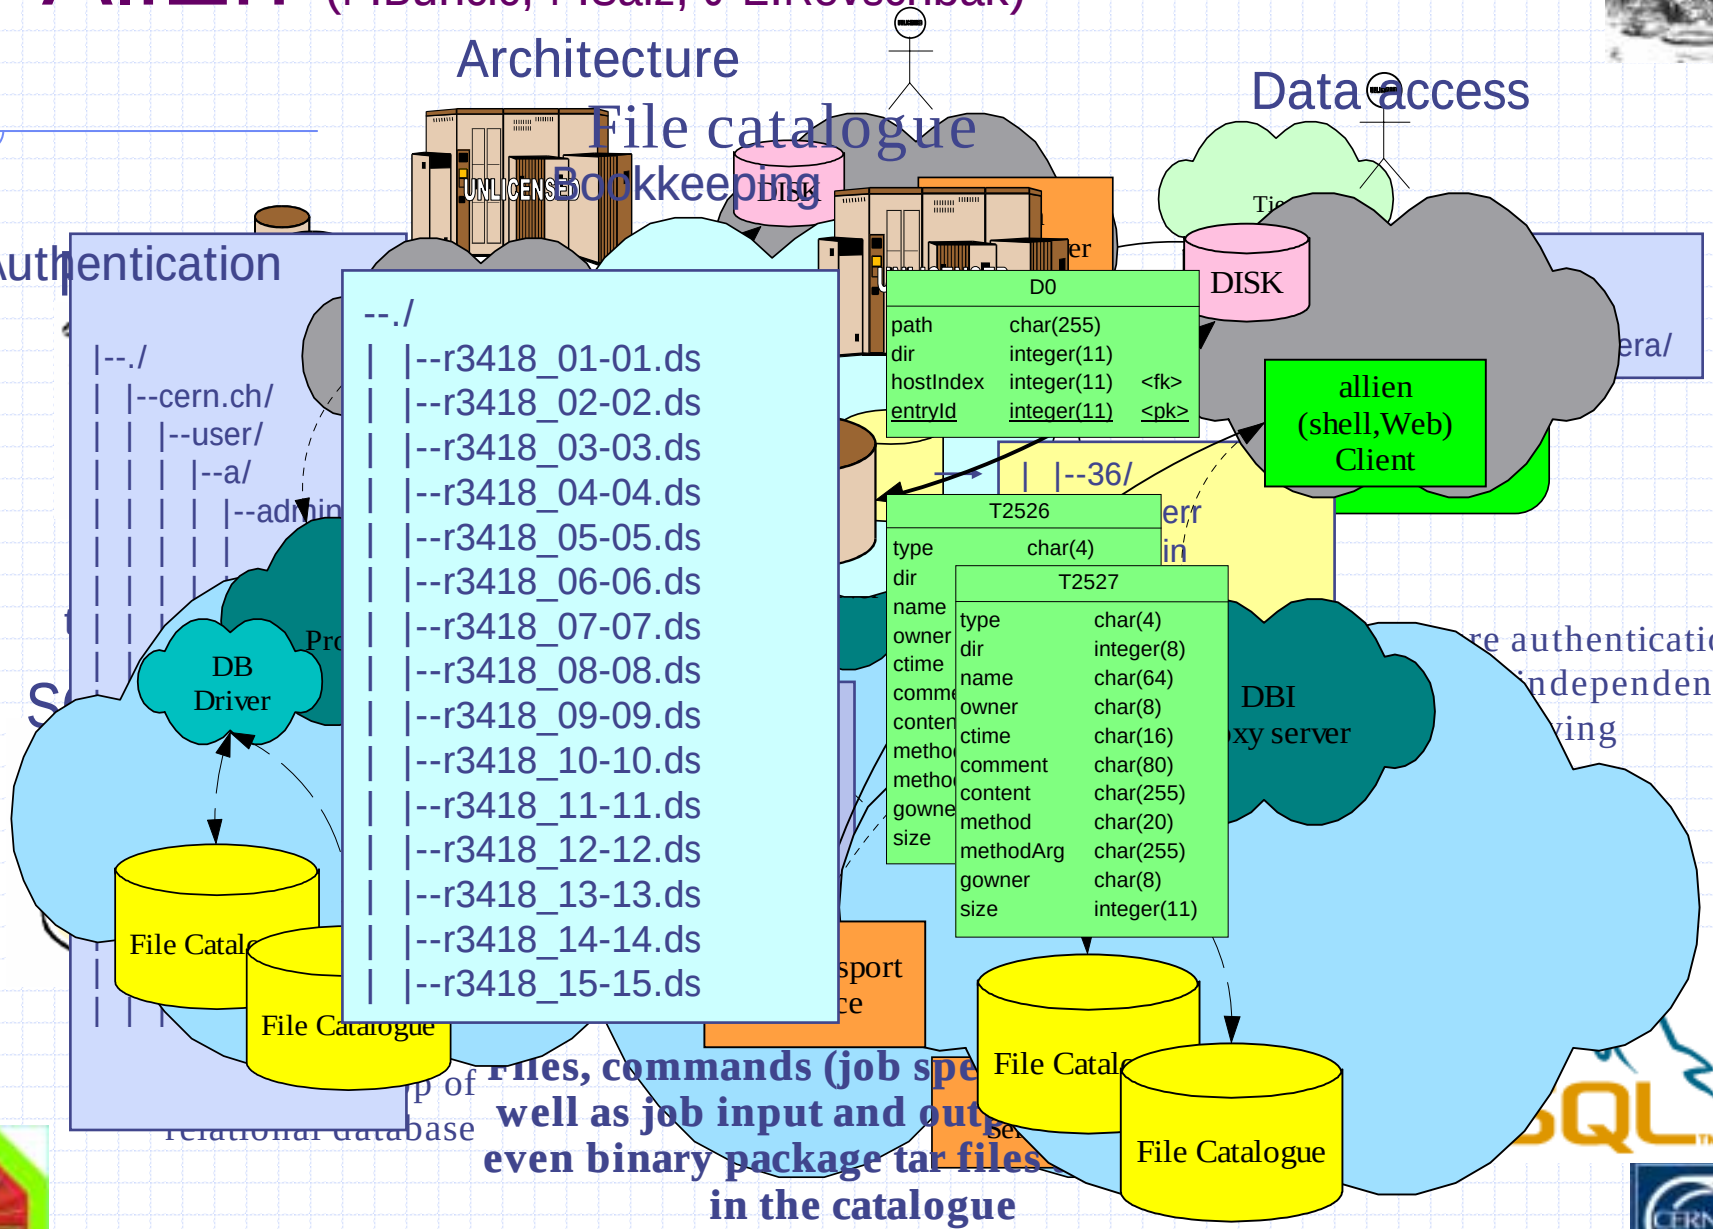

Files, commands (job spe
well as job input and out
even binary package tar files
in the catalogue

Production Summary

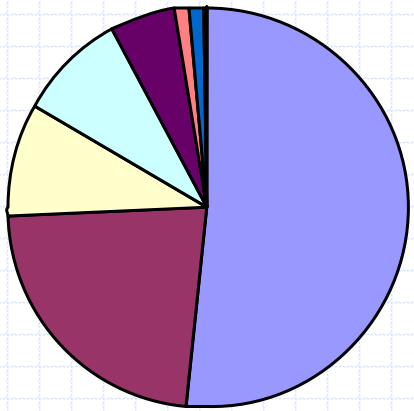

10^5 CPU hours

13 clusters, 9 sites

37 people

21 insitutions

- 5682 events validated, 118 failed (2%)
- Up to 300 concurrently running jobs worldwide (5 weeks)
- 5 TB of data generated and stored at the sites with mass storage capability (CERN 73%,CCIN2P3 14%, LBL, 14%, OSC 1%)
- GSI, Karlsruhe, Dubna, Nantes, Budapest, Bari, Zagreb, Birmingham(?), Calcutta in addition ready by now

AliEn & PROOF

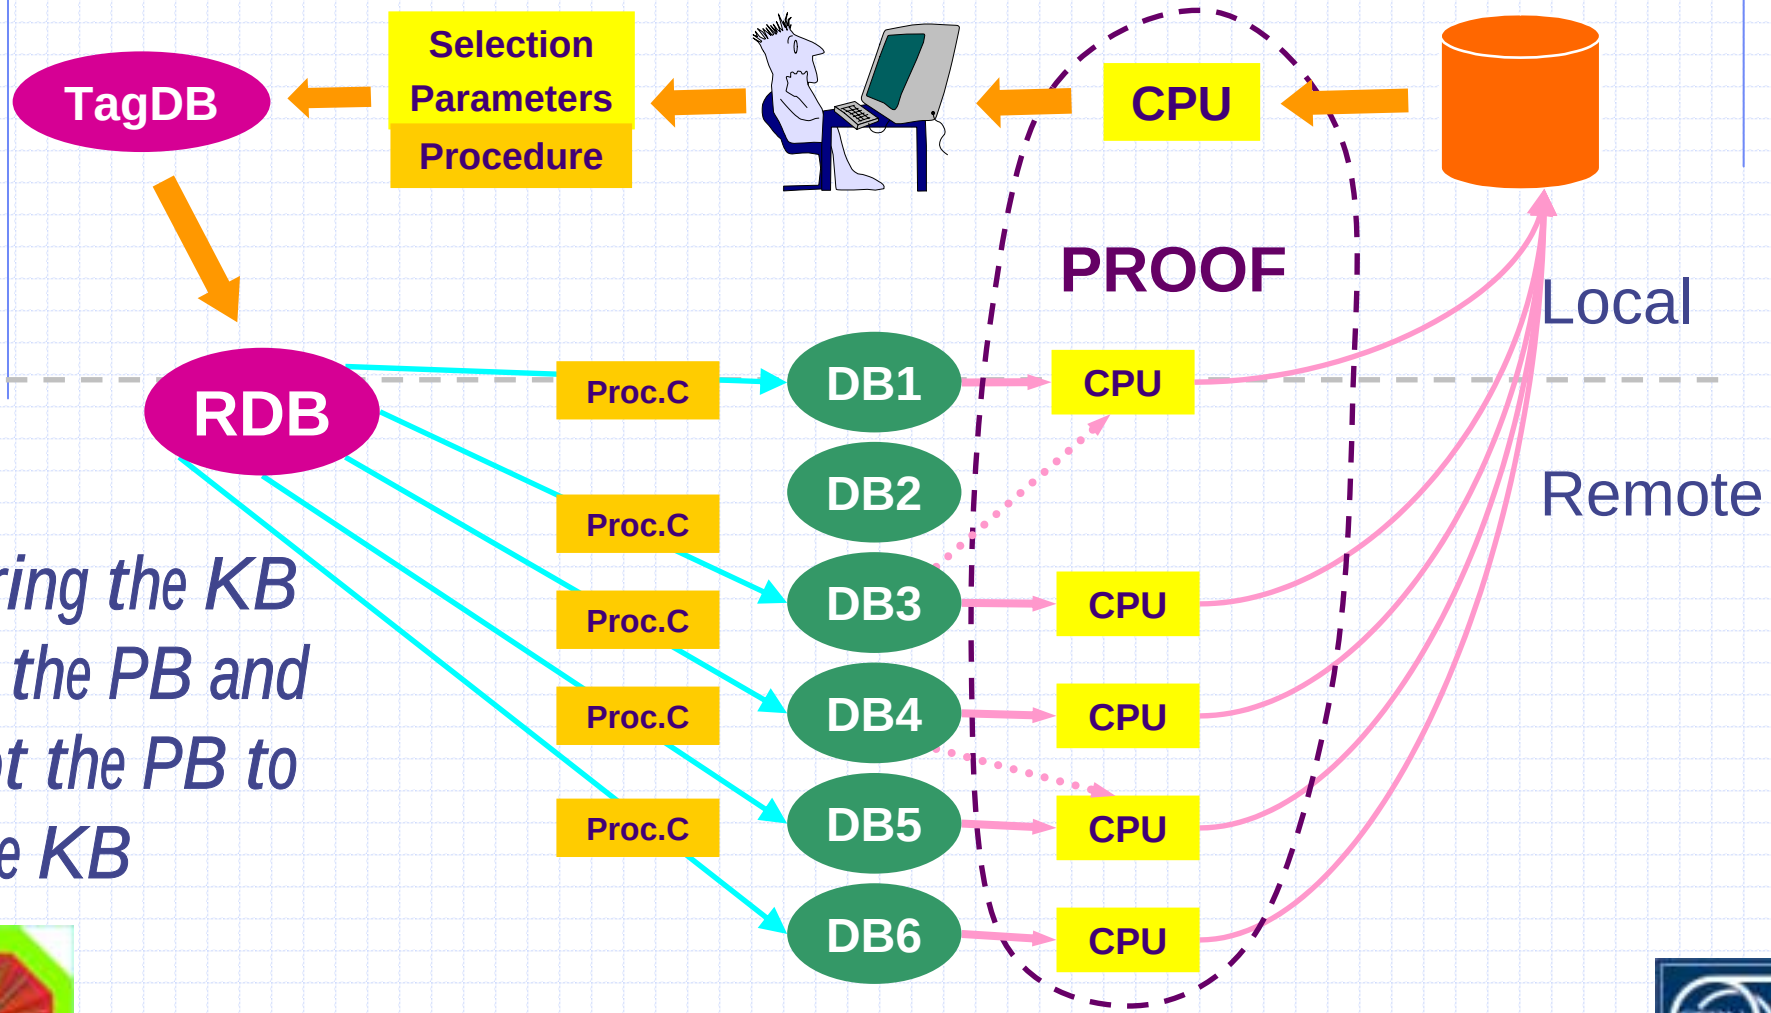

Bring the KB to the PB and not the PB to the KB

ALICE Data challenges

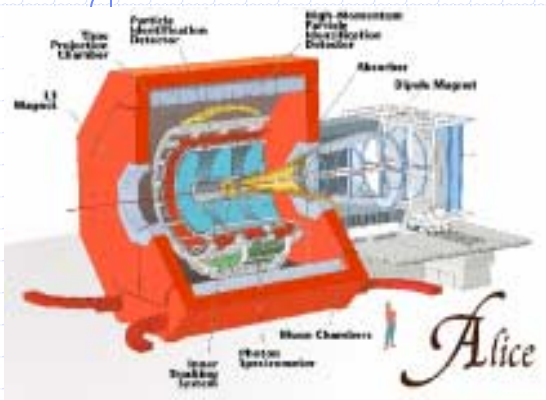

Raw Data

Simulated Data

AliRoot

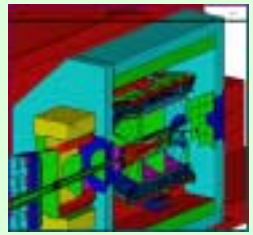

GEANT3
GEANT4
FLUKA

DAQ

ROOT I/O

CERN
TIER 0
TIER 1

ROOT

CASTOR

GRID

Regional
TIER 1
TIER 2

ADC IV Hardware Setup

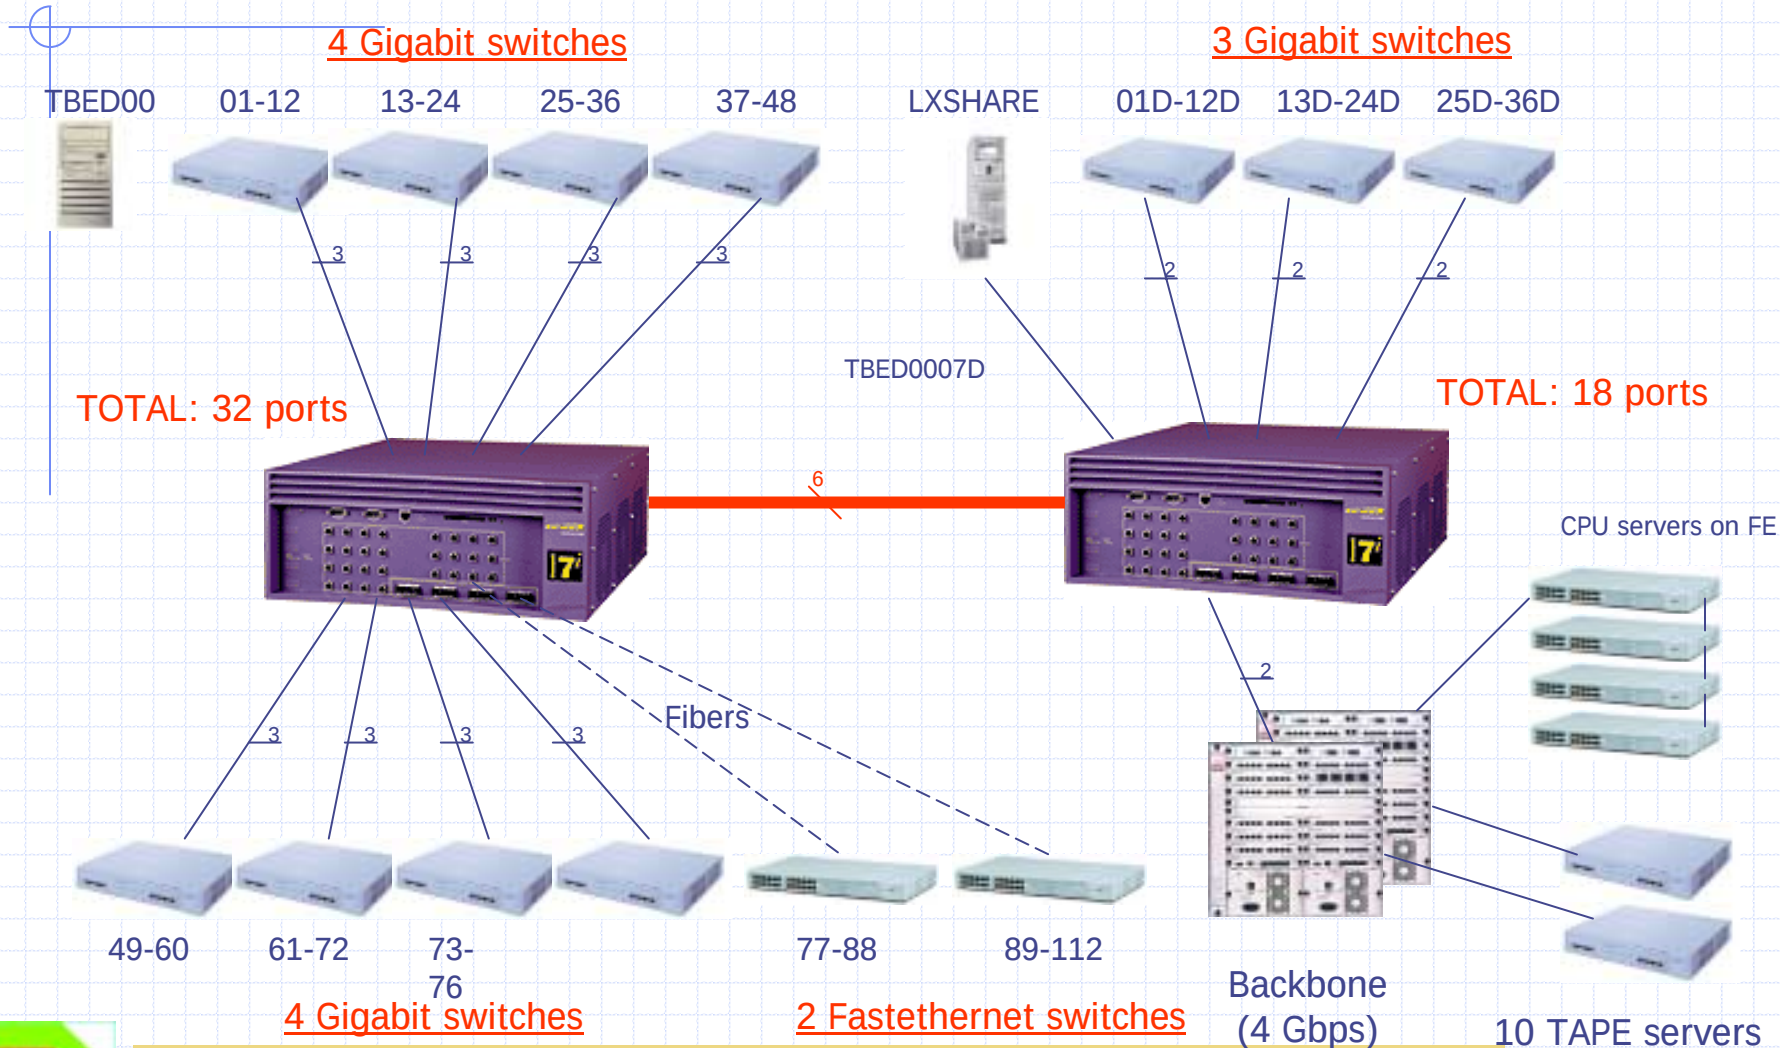

Total: 192 CPU servers (96 on Gbe, 96 on Fe), 36 DISK servers, 10 TAPE servers

Performances

Data generation in LDC, event building, no data recording

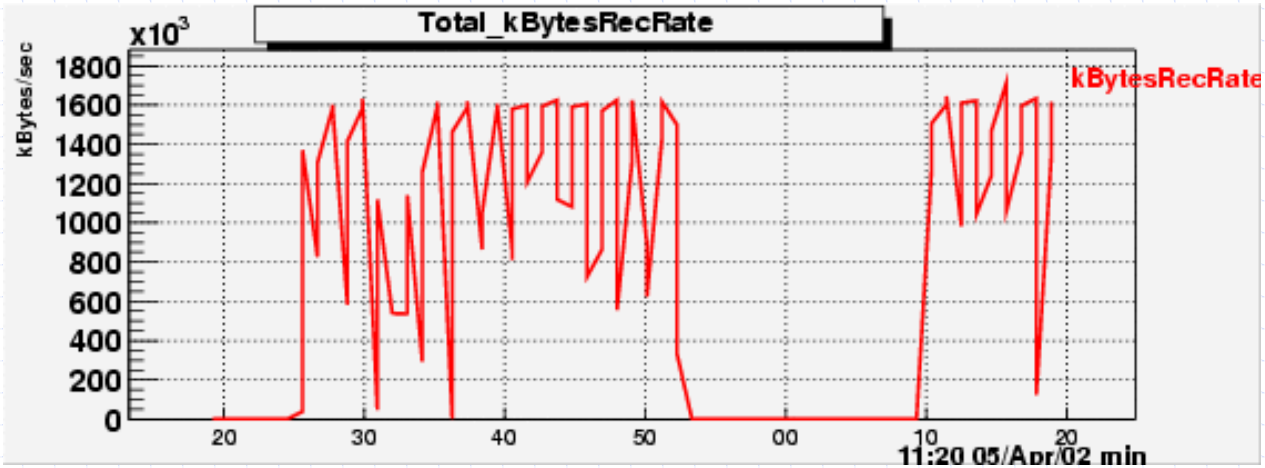

Data generation in LDC, event building, data recording to disk

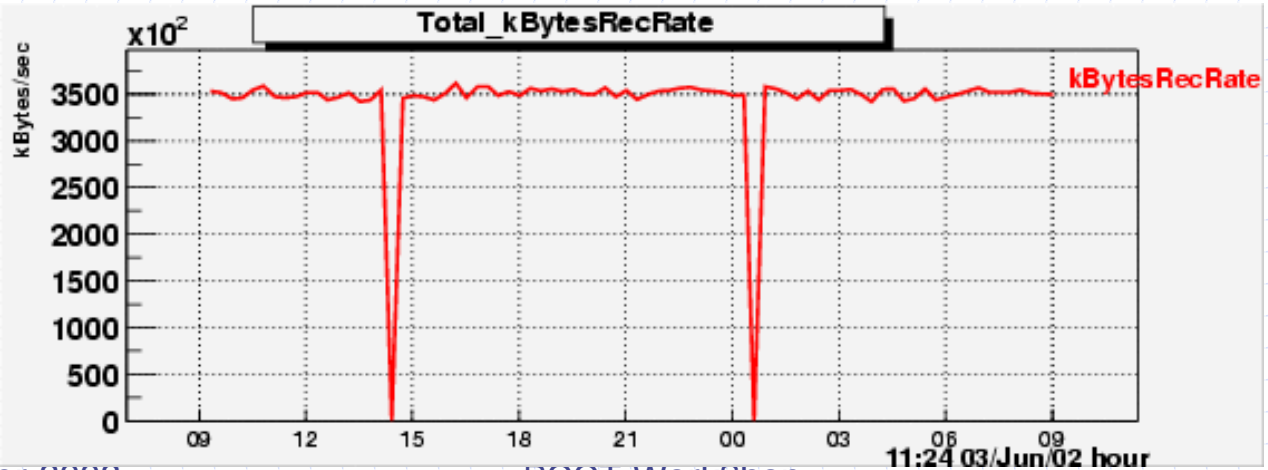

Software development

- In the LEP era the code was 90% written in FORTRAN
  - ~10 instructions, 50 pages!
- In the LHC era the code is in many cooperating languages, mainly C++
  - ~ 50 instructions, 700 pages – nobody understands it completely (B.Sjostrup)
  - But also C#, Java, Perl, Python, php..., Web and GRID
- Users are heterogeneous, sparse and without hierarchical structure
  - From very expert analysts to users, from 5% to 100% of time devoted to computing
- People come and go with a very high rate
  - Programs have to be maintained by people who did not develop them
  - Young physicists need knowledge they can use also outside physics
- And yet HEP software has been largely successful!
  - Experiments have not been hindered by software in their scientific goals
  - GEANT3, PAW and ROOT: in use since 20 years on all architectures and OS
- And yet we (as a community) have not used traditional SE
  - Did we do something right?

Agile Technologies
(aka SE catching up)

- SE response to HCP are the "Agile Methodologies"
  - Adaptive rather than predictive
  - People-oriented rather than process-oriented
  - As simple as possible to be able to react quickly
  - Incremental and iterative, short iterations (weeks)
  - Based on testing and coding rather than on analysis and design

- Uncovering better ways of developing software by valuing:

Individuals and interactions		*processes and tools*
Working software	**OVER**	*huge documentation*
Customer collaboration		*contract negotiation*
Responding to change		*following a plan*

That is, while there is value in the items on the right, we value the items on the left more.

Software Development Process

- ALICE opted for a light core CERN offline team...
  - Concentrate on framework, software distribution and maintenance
- ...plus 10-15 people from the collaboration
  - GRID coordination (Torino), World Computing Model (Nantes), Detector Database (Warsaw)
- A development cycle adapted to ALICE has been elaborated
  - Developers work on the most important feature at any moment
  - A stable production version exists
  - Collective ownership of the code
  - Flexible release cycle and simple packaging and installation
- Micro-cycles happen continuously
- 2-3 macro-cycles per year
  - Discussed & implemented at Off-line meetings and Code Reviews
  - Corresponding to major code releases
- We have high-level milestones for technology and physics
  - Computing Data Challenges test technology and integration DAQ – Off-line
  - Physics Data Challenges test the Off-line from the user viewpoint

EXplaining Planning in ALICE

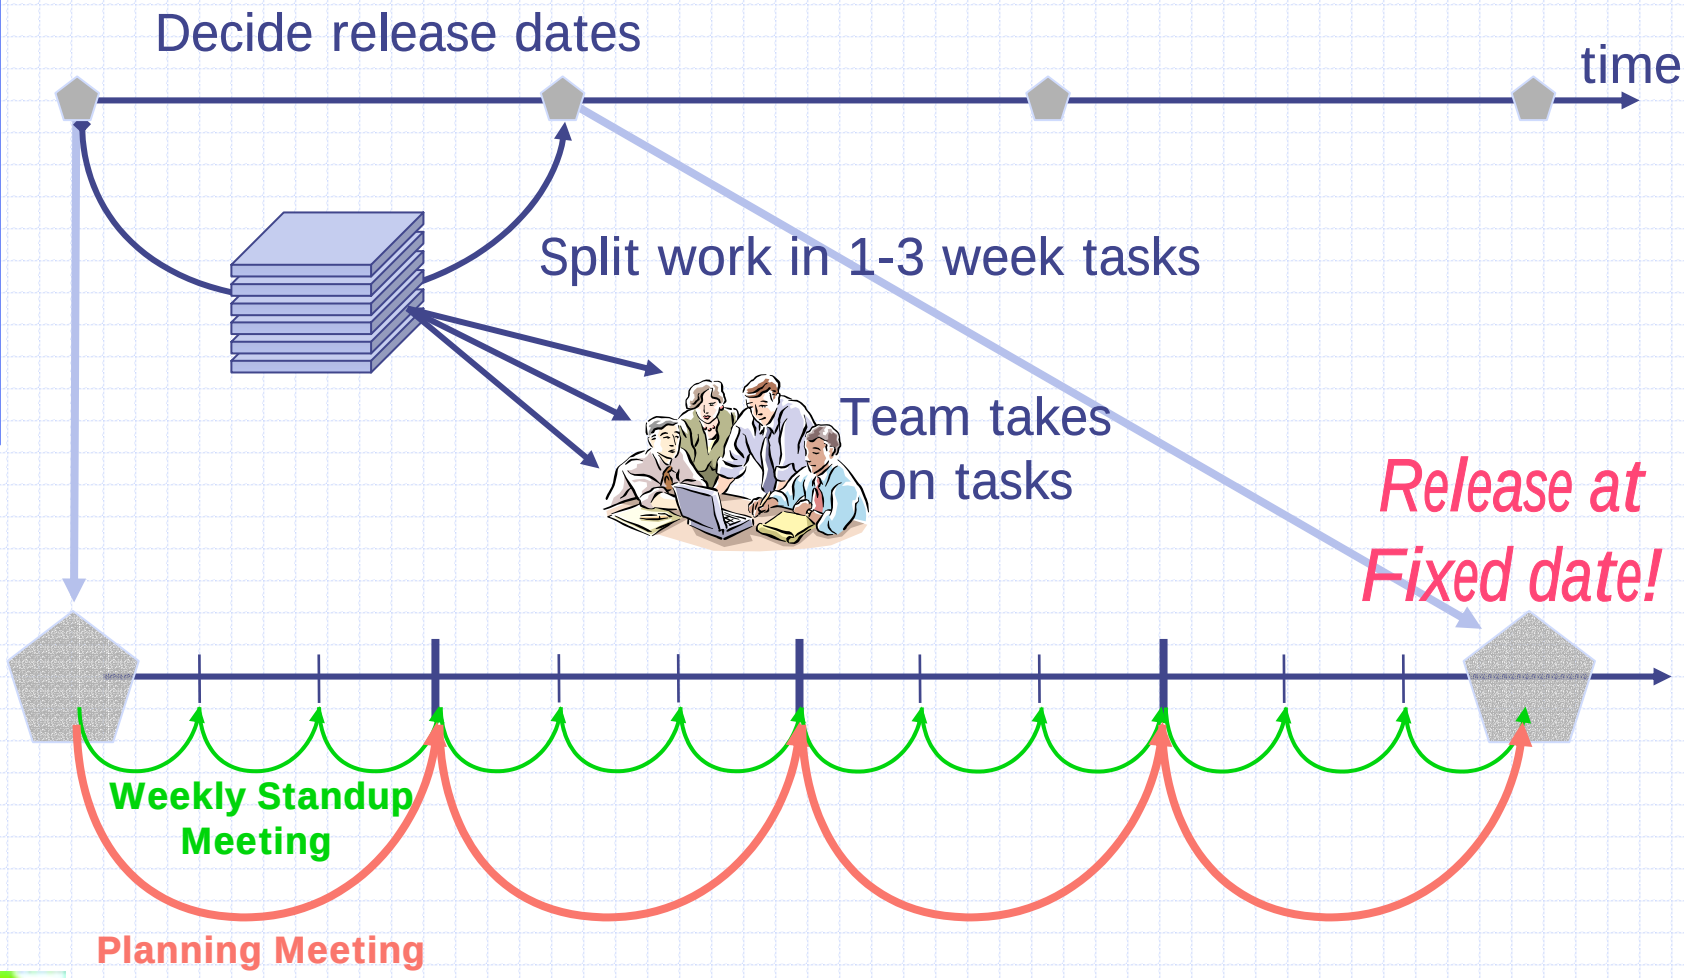

AliRoot maintenance

- 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

ROOT, ALICE & LCG

- LCG is the LHC Computing Project
  - The objective is to build the computing environment for LHC
- ALICE has lobbied strongly to base the LCG project on ROOT and AliEn
- The choice of the other experiments is to establish a *client-provider* relationship with ROOT
  - While developing alternatives for some existing ROOT elements or *hiding* them behind abstract interfaces
- And to use the result of GRID MiddleWare projects
- ALICE has expressed its worries
  - Little time to develop and deploy a new system
  - Duplication and dispersion of efforts
  - Divergence with the rest of the HEP community
- ALICE will continue to develop its system
  - And to provide basic technology, i.e. VMC and geometrical modeller
- ... and it will try to collaborate with LCG wherever possible

nemo profeta in patria

the names have been changed to protect the innocent

parental advice is suggested before visioning these slides in the presence of little children

no animal has been hurt during the preparation of this presentation