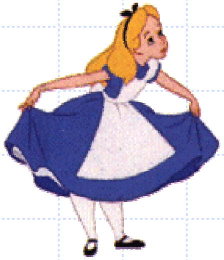

# Plans for ALICE Data Challenges

F.Carminati
Grid Deployment Board, July 8th, 2003
Cosener's House

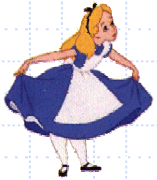

# Summary

- ALICE PDC-3: objectives and strategy
- The framework, AliRoot & AliEn: what's new
- ALICE experience with EDG
- Computing framework needs
- Computing resource needs
- Common tools for PDC-3: LHC Computing Grid (LCG)

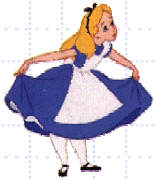

# ALICE Physic Data Challenges

| 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<br>• Simple analysis tools<br>• Digits in ROOT format |
| 01/04-06/04 | 10% | • Complete chain used for trigger studies<br>• Prototype of the analysis tools<br>• Comparison with parameterised MonteCarlo<br>• Simulated raw data |
| 01/06-06/06 | 20% | • Test of the final system for reconstruction and analysis |

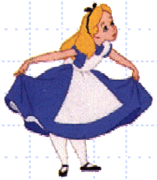

# From PDC 2 to PDC 3

- The ALICE Physics Performance Report will be completed by the end of the year
    - Its emphasis is on detector performance on soft physics
- The strategy makes use of *event mixing*
    - A *signal-free* parameterised *event* is *mixed* with signals
    - *Signal-free events* were reused up to 500 times
- Parameterisations developed from full simulation have been heavily used
- The new phase of physics simulations will be on hard physics, jets and jet quenching
    - At LHC, but also at RHIC, the mini-jet contribution is expected to increase for the production of particles in the central rapidity region
- Event mixing strategy has to be reviewed

# Example

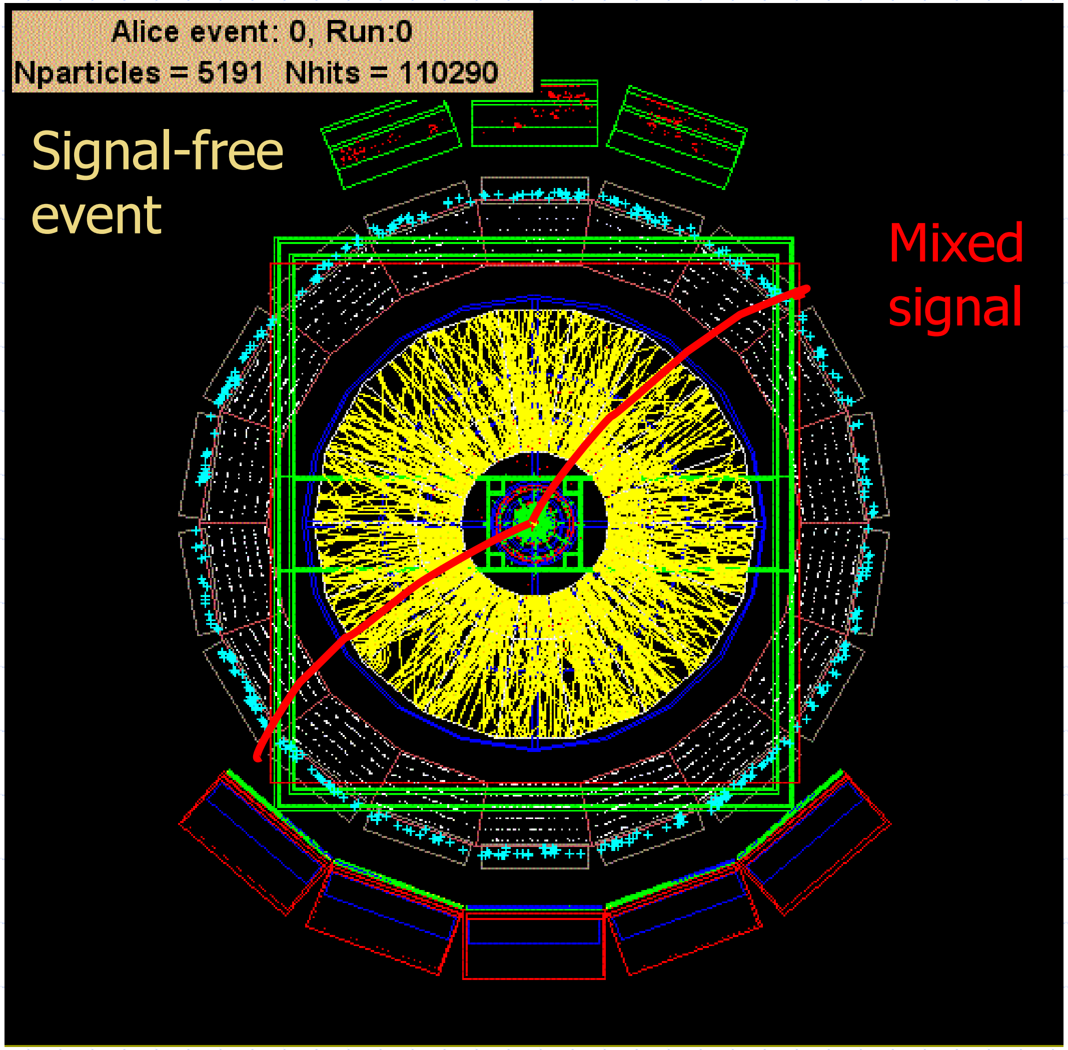

Alice event: 0, Run:0
Nparticles = 5191 Nhits = 110290
Signal-free event
Mixed signal

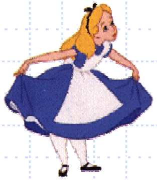

# Objectives of PDC 3

- A simple *signal-free* parameterisation cannot be used due to the absence of jet and mini-jet fluctuations
- We have to use a *real* (!) generator (HIJING) that contains (mini-)jets to simulate the *underlying event* down to the low x-sec signals
- For $p_t$ ~10-20 GeV/c → ~1 jet per event
  - Study via leading particles at high $p_t$ in the barrel (not easy!)
  - Event mixing not possible
- For $p_t$ ~200 GeV/c → 1000 jet per month
  - Event mixing can be used, but with mixing jets with *real events* reused not more than 50 times

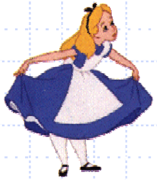

# Objectives of PDC 3

- The estimation of the number of events is essentially defined by the jet study
  - $10^5$ events for jets with $p_t$ ~10-20 GeV/c
  - $10^4$ -- $10^5$ events for studies on particle correlations and simple and double strangeness ($\Lambda$,$\Xi$)
  - $10^6$ events for high $p_t$ jets (→ ~$10^5$ underlying events)
  - $\Omega$, charmonium and bottonium into e+e- → similar statistics, the same underlying events can be reused
  - Centrality
    - 50% central events (b<2-3 fm)
    - 50% minimum bias
- Besides Pb-Pb: $10^7$ events p-p (Pythia) and ~$10^6$ p-A (Hijing)

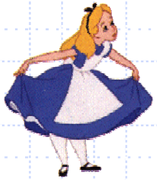

# Objectives of PDC 3

- Verify the (computing) ability to (simulate and) process ~10% of the data of a standard data-taking year
  - Evaluate the consequences for the experiment and middleware frameworks
- See how much data we can simulate, process (including mixing) and analyse in 6 months with LCG and private resources
- Part of the raw data will be shipped to CERN and from there redistributed for reconstruction and analysis
  - Mimic real life when all data come from CERN with a (large) part of the production over a limited time
  - Optimize the reconstruction to get results in time

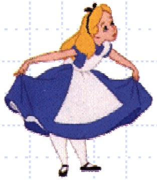

# ALICE ADC’s

| 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 -<br>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<br>Prototype of the final remote data replication<br>(Raw digits for all detectors) |
| 5/2006 | 750 (1250 if possible) | 750 (1250 if possible) | Final test (Final system) |

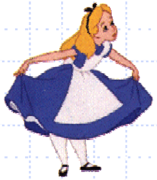

# ALICE DC – BW

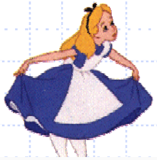

# ADC IV performances

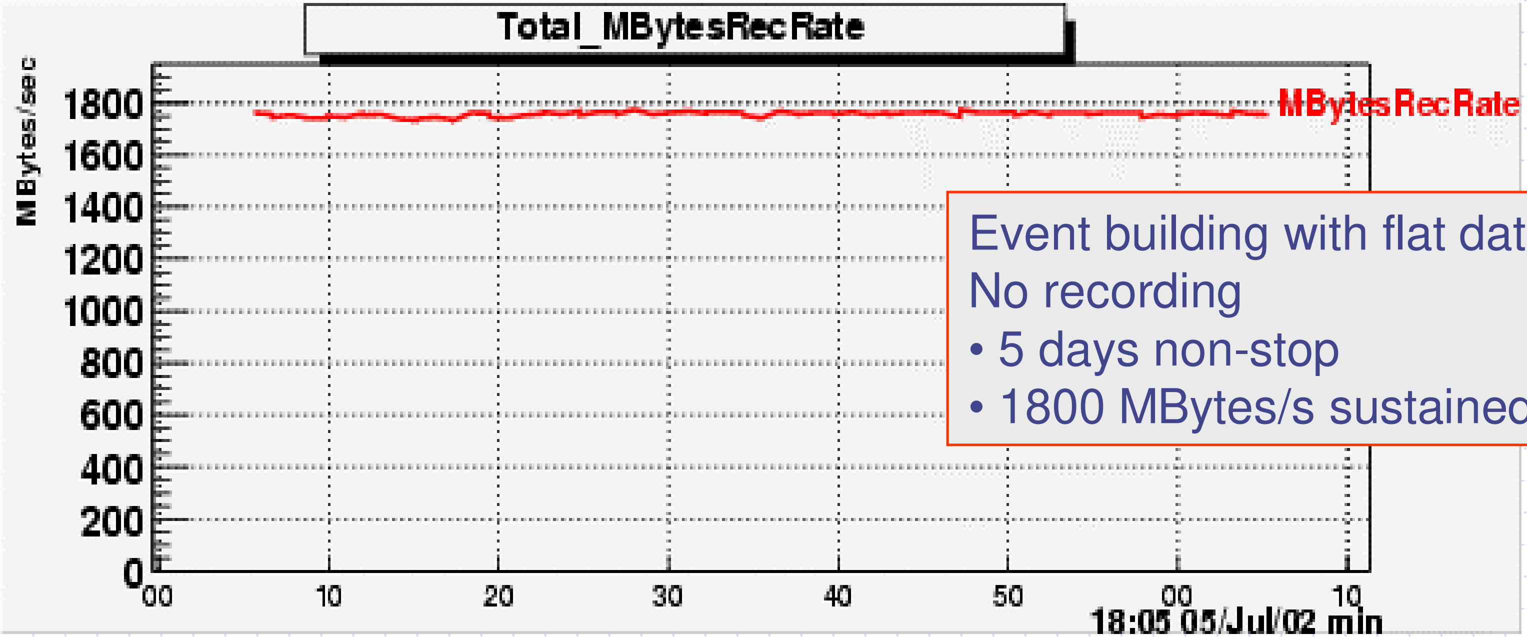

Event building with flat data traffic

No recording

- 5 days non-stop
- 1800 MBytes/s sustained

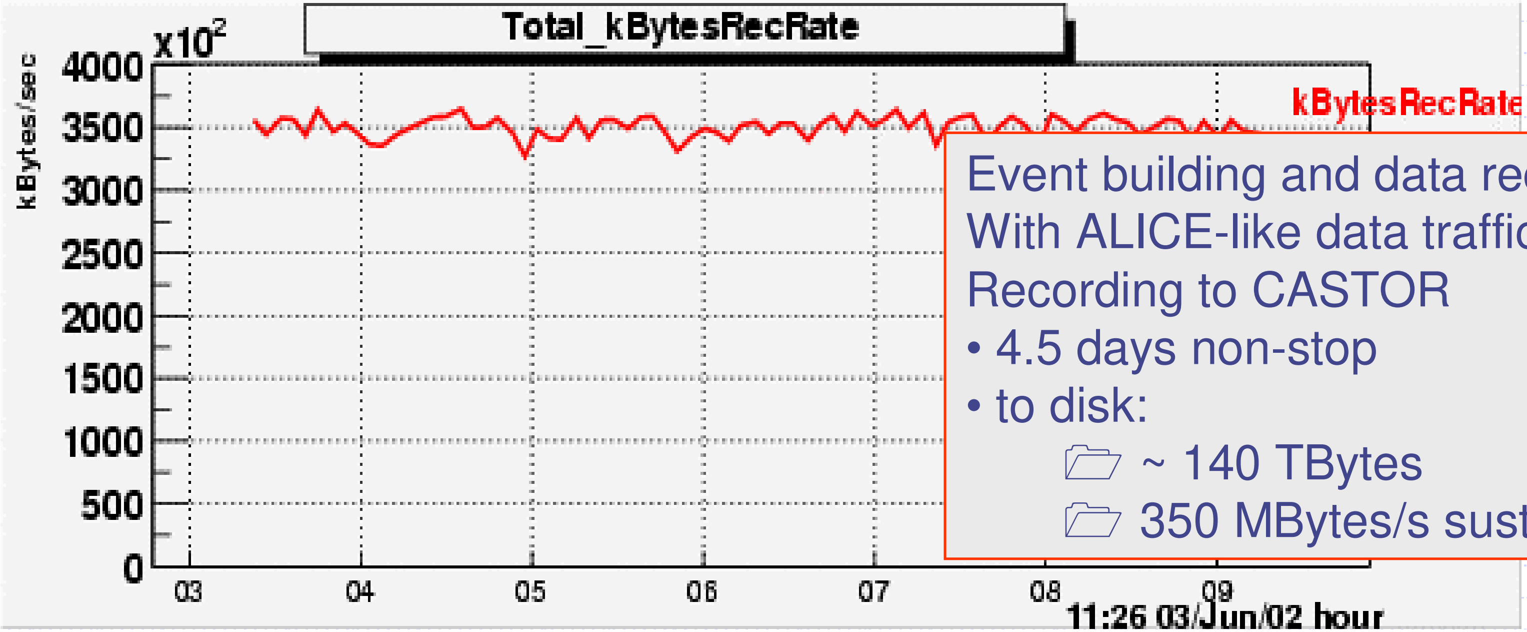

Event building and data recording

With ALICE-like data traffic

Recording to CASTOR

- 4.5 days non-stop
- to disk:
  - ~ 140 TBytes
  - 350 MBytes/s sustained

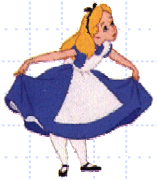

# The framework

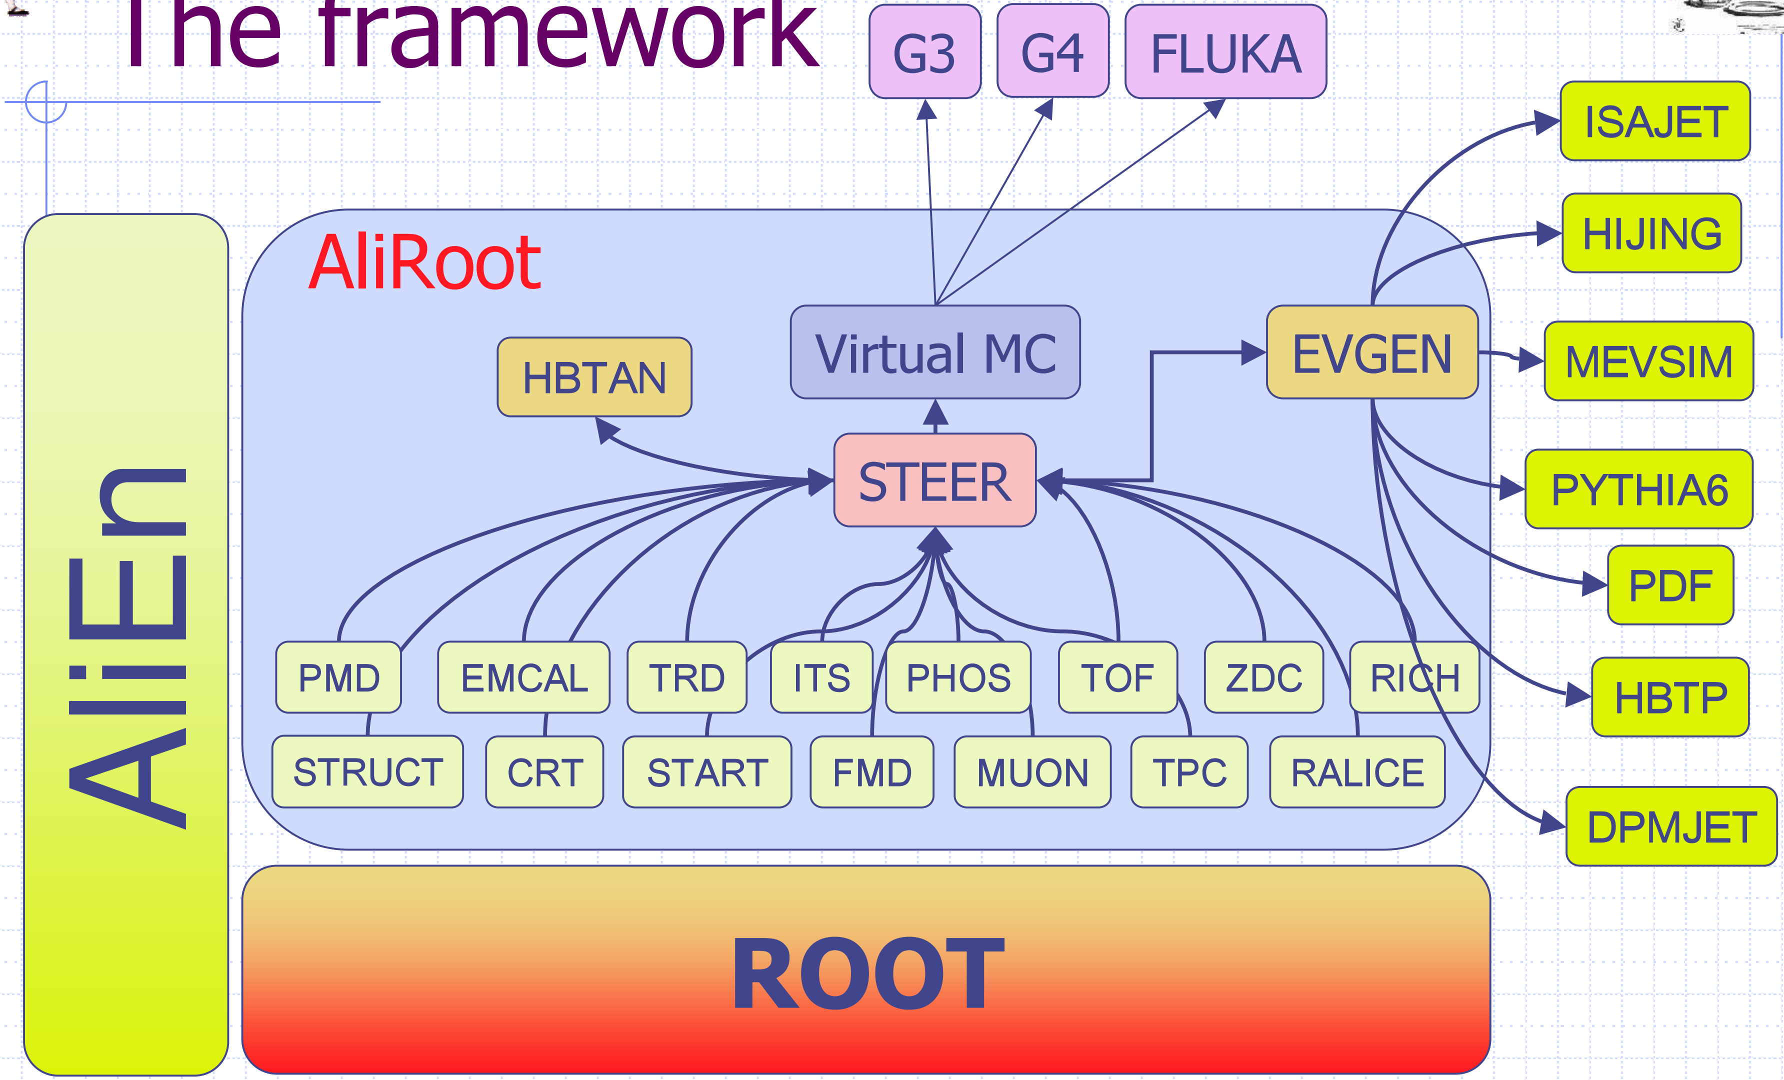
G3
G4
FLUKA
ISAJET
HIJING
AliEn
AliRoot
HBTAN
Virtual MC
EVGEN
MEVSIM
STEER
PYTHIA6
PDF
PMD
EMCAL
TRD
ITS
PHOS
TOF
ZDC
RICH
HBTP
STRUCT
CRT
START
FMD
MUON
TPC
RALICE
DPMJET
ROOT

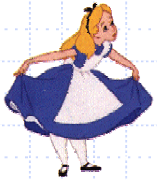

# AliRoot + AliEn @GRID

- AliRoot
  - C++: 400kLOC + 225kLOC (generated) + macros: 77kLOC
  - FORTRAN: 13kLOC (ALICE) + 914kLOC (external packages)
  - Maintained on Linux (any version!), HP-UX, DEC Unix, Solaris, MacOSX, Intel Itanium
  - Works also with Intel icc compiler
- Two packages to install (ROOT+AliRoot) + MC's
  - Less that 1 second to link (thanks to shared libs)
  - 1-click-away install: download and make (non-recursive makefile)
- AliEn
  - 25kLOC of PERL5 (ALICE)
  - ~2MLOC mostly PERL5 (opens source components)
- Installed on almost 50 sites by physicists
  - >50 users develop AliRoot from detector groups
  - 70% of code developed outside, 30% by the core Offline team

# The Virtual MonteCarlo

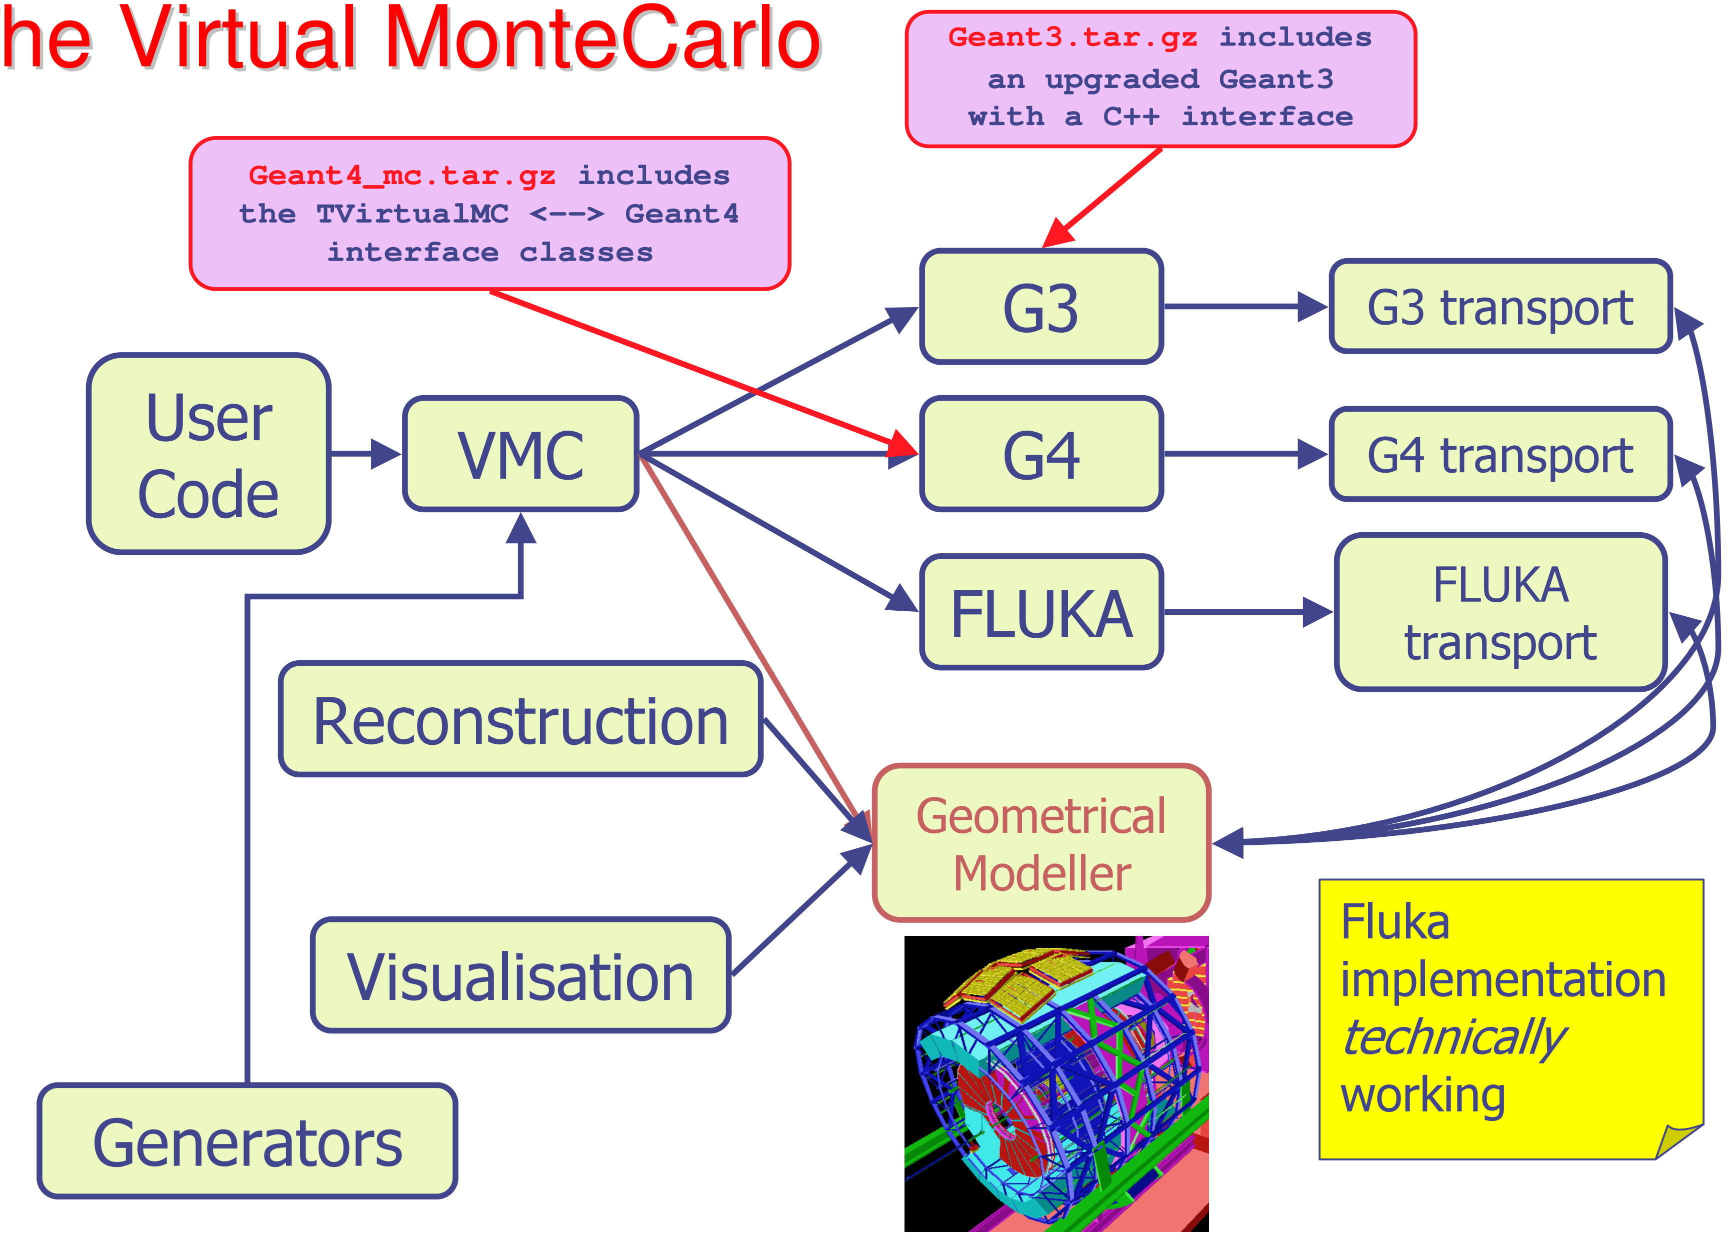

**Geant3.tar.gz** includes an upgraded Geant3 with a C++ interface

**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

Geometrical Modeller

Fluka implementation *technically* working

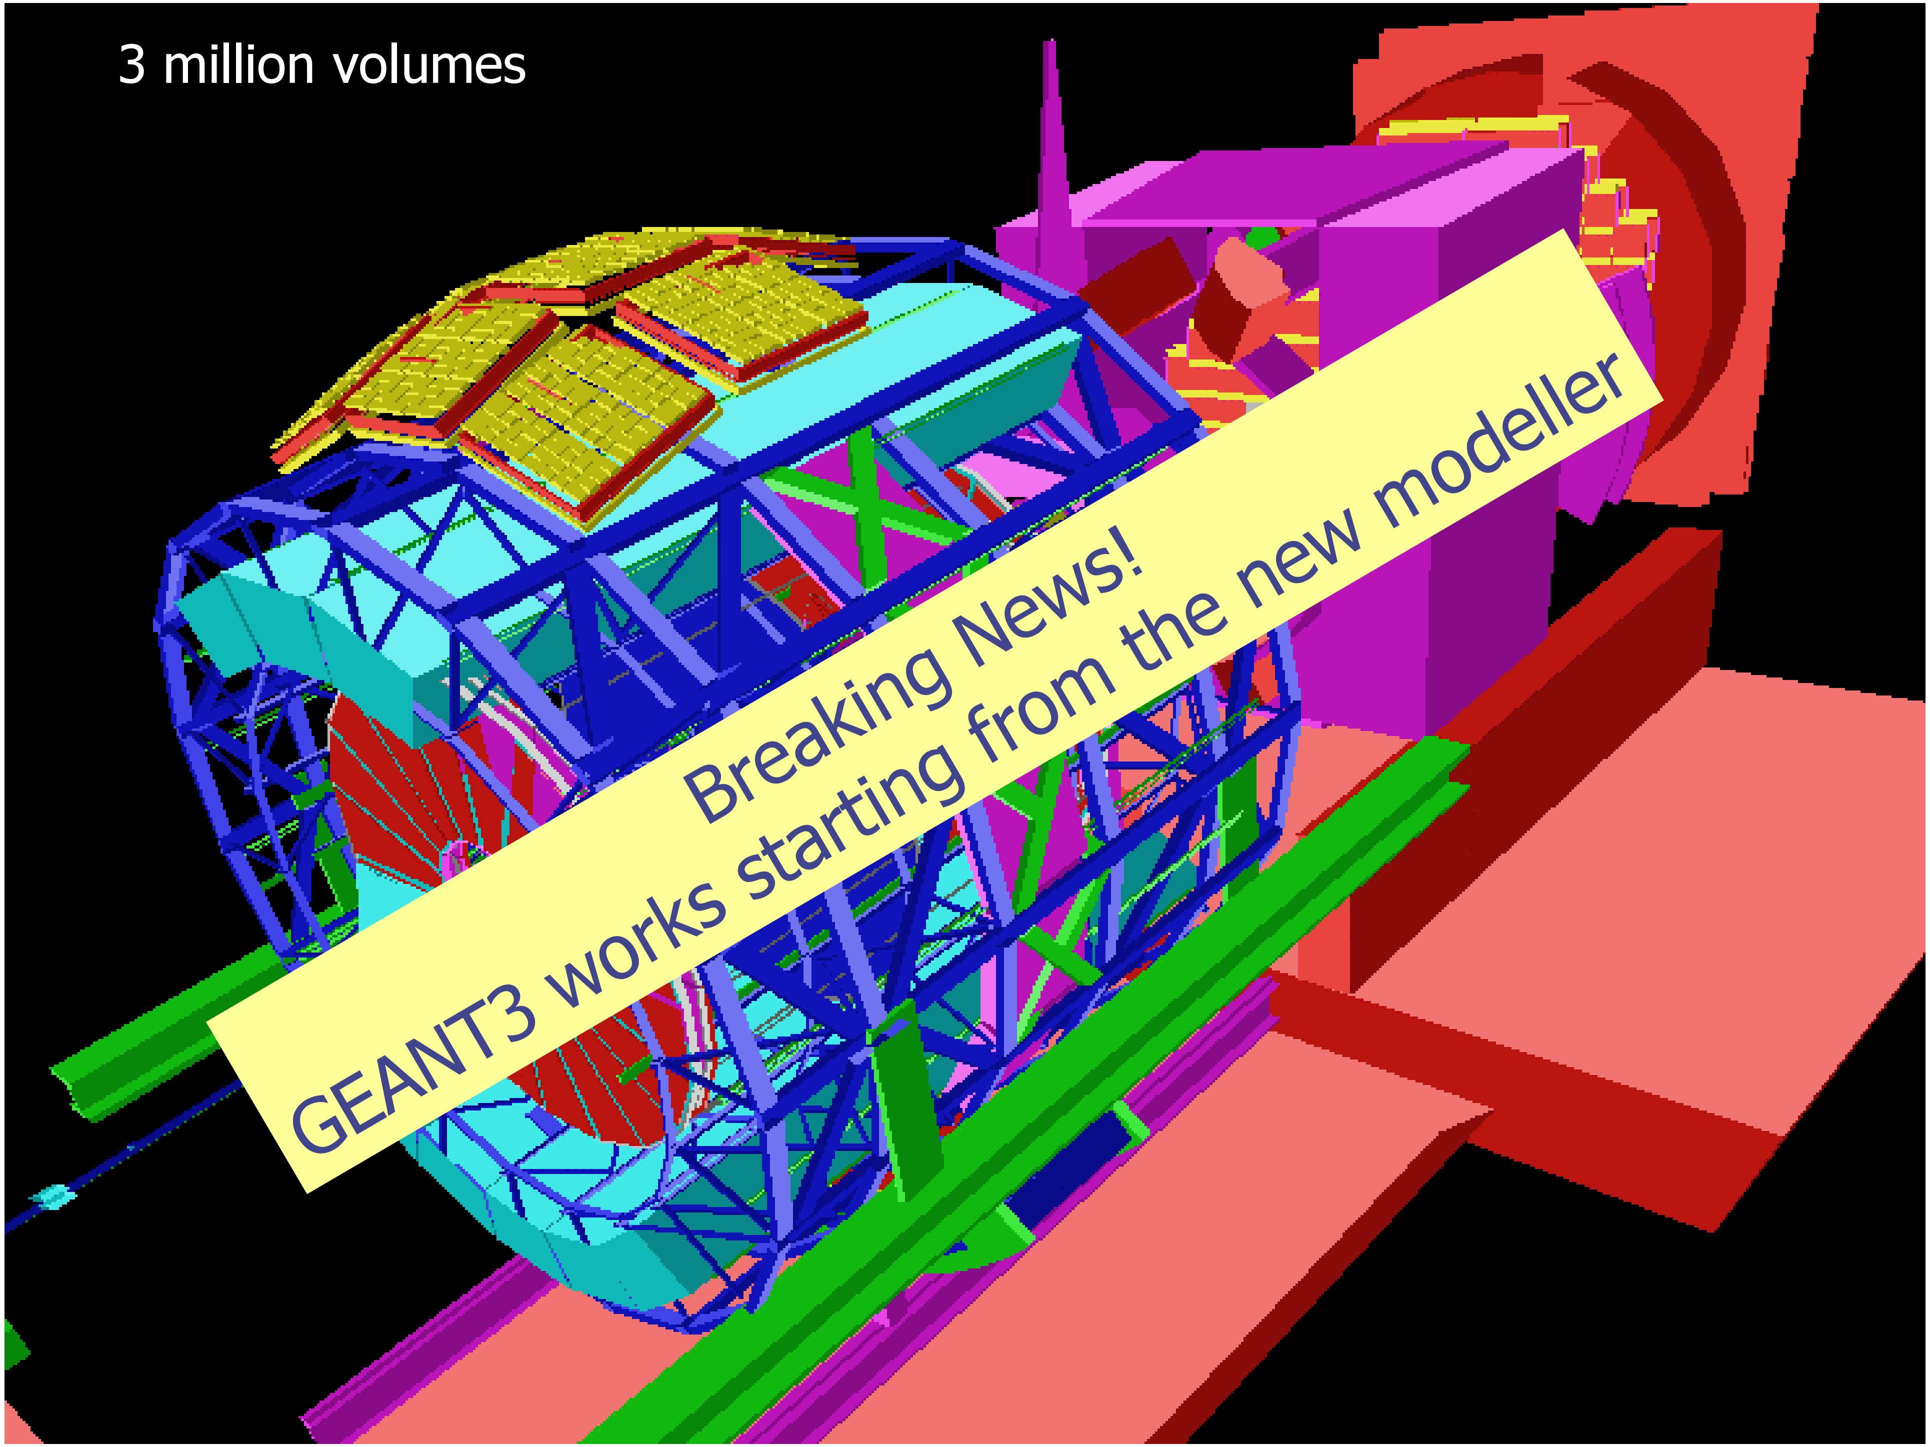

3 million volumes

Breaking News!
GEANT3 works starting from the new modeller

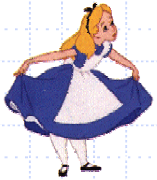

# NewIO

- The New I/O will be used for the first time in production in PDC 3
- I/O managed by a set of dedicated classes (AliLoader, AliRunLoader, AliDataLoader…)
- Class dependency *reduced* by the use of root folders
  - Better modularity
  - Possibility to reprocess single detectors
  - Possibility to access and update single detector information without updating files
  - Different version of derived files available

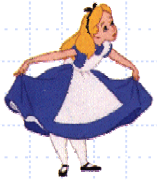

# Whiteboard data communication

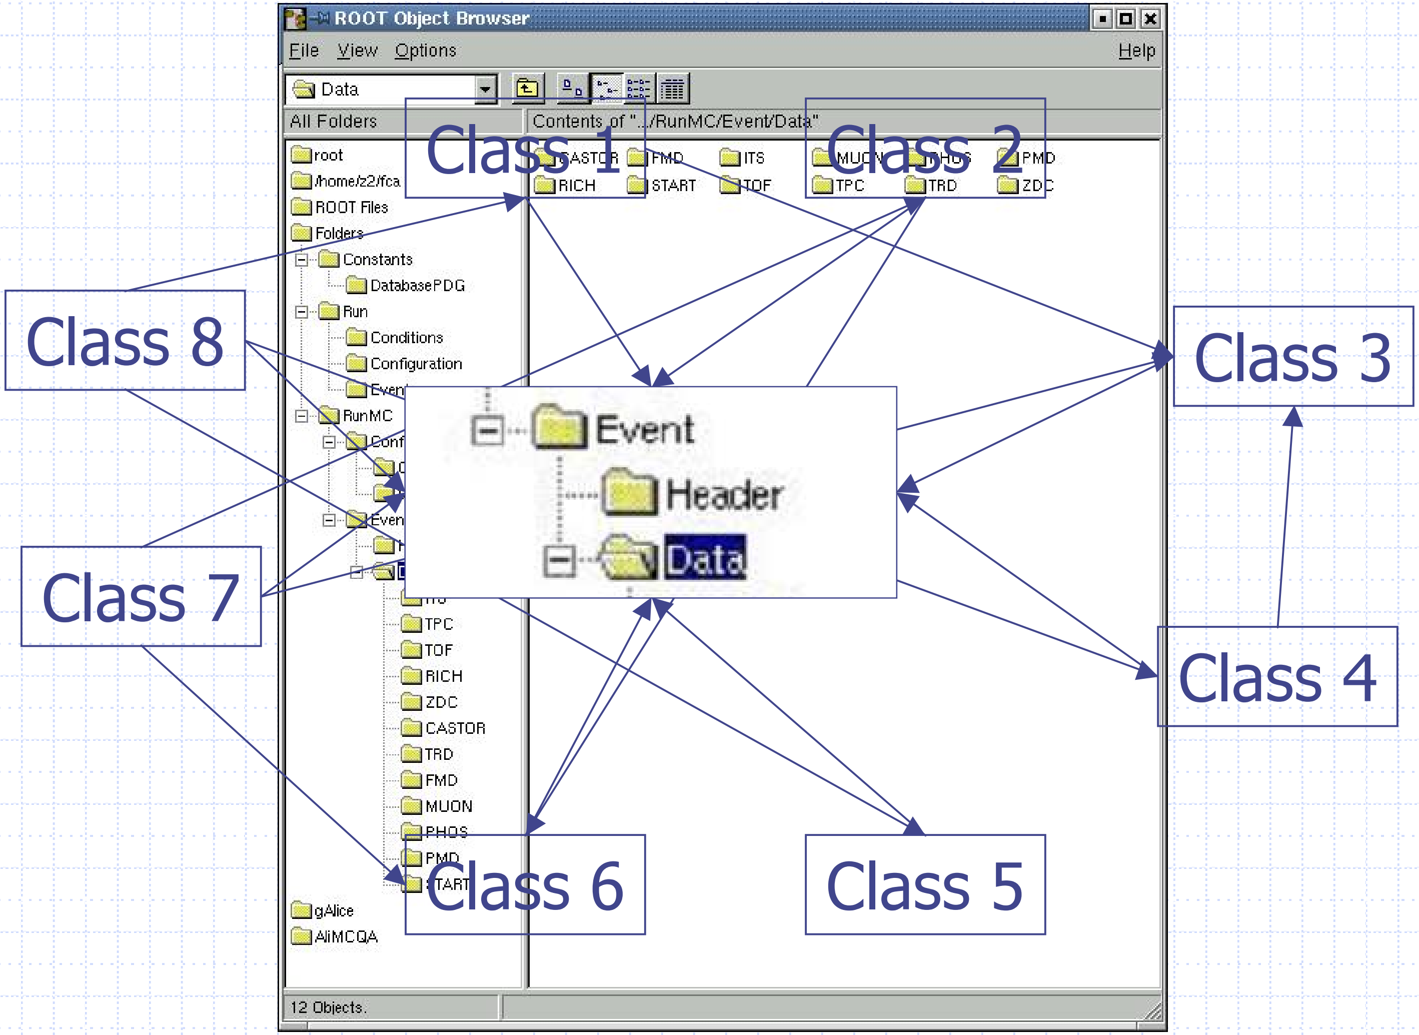
ROOT Object Browser
File View Options Help
Data
All Folders
Contents of ".../RunMC/Event/Data"
root
/home/z2/fca
ROOT Files
Folders
Constants
DatabasePDG
Run
Conditions
Configuration
RunMC
TPC
TOF
RICH
ZDC
CASTOR
TRD
FMD
MUON
PHOS
PMD
START
gAlice
AliMCQA
CASTOR FMD ITS MUON PHOS PMD
RICH START TOF TPC TRD ZDC
Event
Header
Data
12 Objects.
Class 1
Class 2
Class 3
Class 4
Class 5
Class 6
Class 7
Class 8

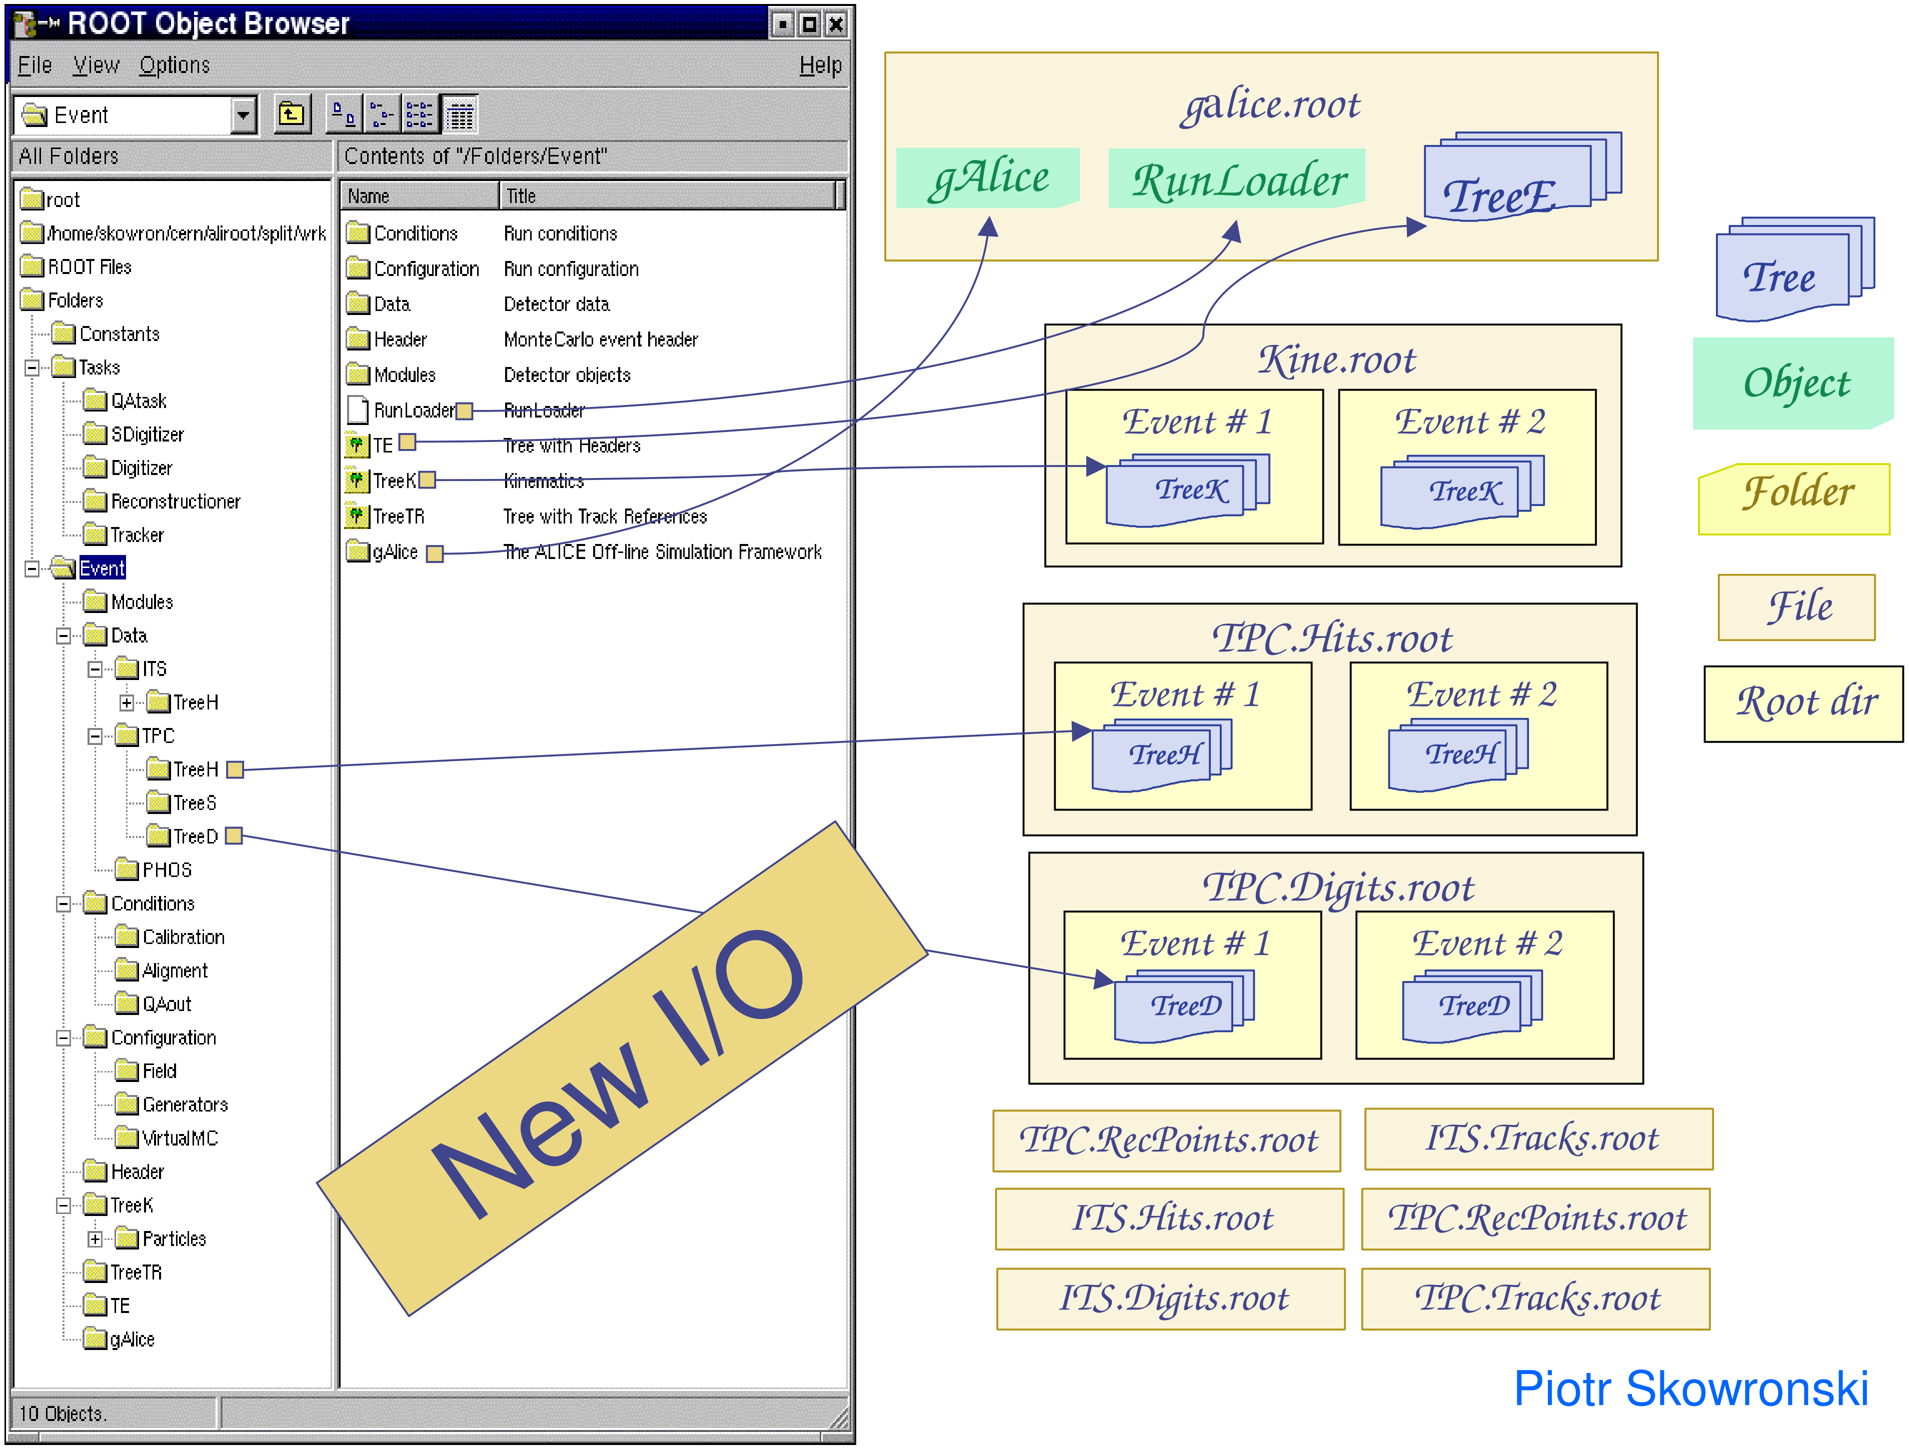

- root
- /home/skowron/cern/aliroot/split/wrk
- ROOT Files
- Folders
  - Constants
  - Tasks
    - QAtask
    - SDigitizer
    - Digitizer
    - Reconstructioner
    - Tracker
  - Event
    - Modules
    - Data
      - ITS
        - TreeH
      - TPC
        - TreeH
        - TreeS
        - TreeD
      - 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.

New I/O

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

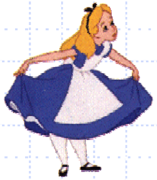

# Architecture

- The *loaders* manage files and directories and move data from and to (T)Folders
  - Old I/O: dynamic_cast<TTree*> (fEventFolder->FindObject(fgkHeaderContainerName));
  - New I/O: loader->TreeE();

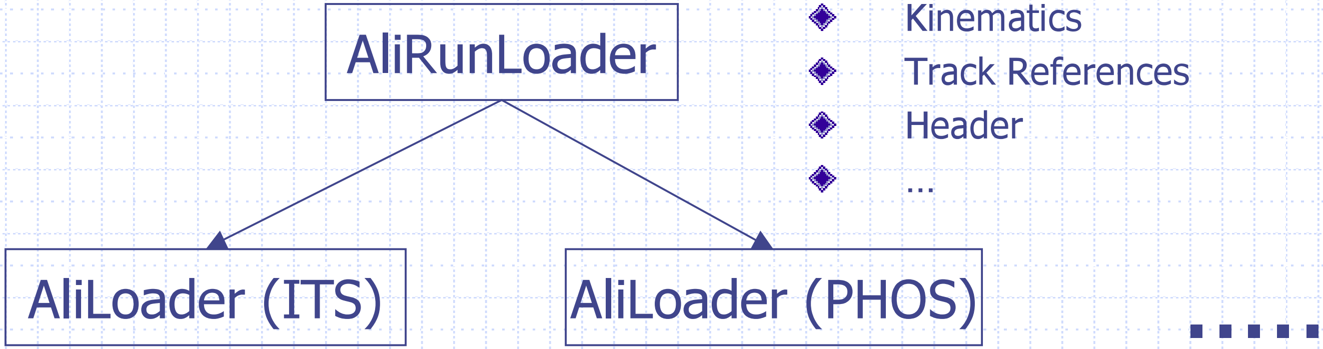

AliRunLoader
AliLoader (ITS)
AliLoader (PHOS)
.....

Takes care of

- Kinematics
- Track References
- Header
- …

Take care of all detector data

- Hits
- Summable Digits
- Digits
- …

# Architecture

- Better scalability
- Easier addition of new types

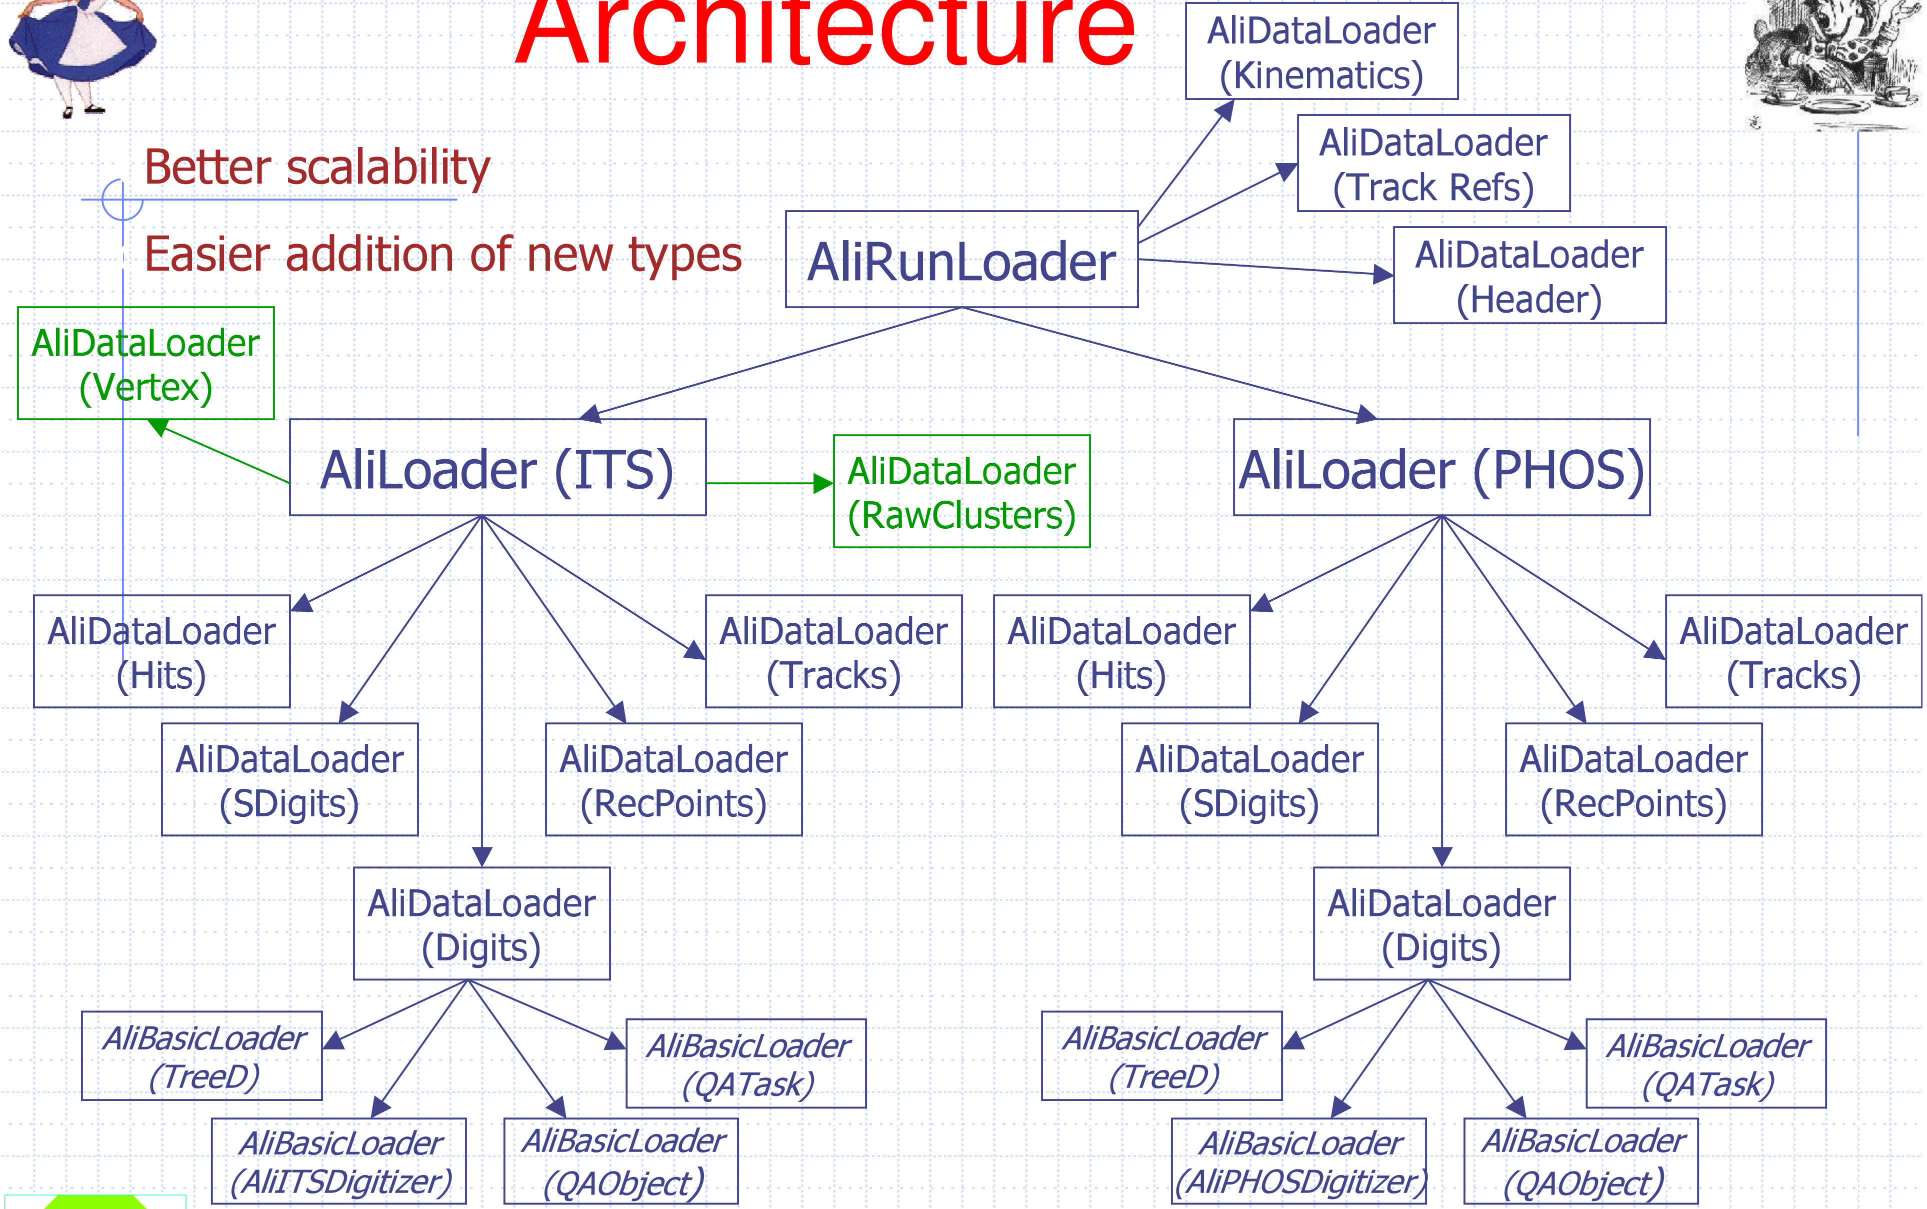

AliRunLoader

AliDataLoader (Kinematics)

AliDataLoader (Track Refs)

AliDataLoader (Header)

AliLoader (ITS)

AliDataLoader (Vertex)

AliDataLoader (RawClusters)

AliDataLoader (Hits)

AliDataLoader (SDigits)

AliDataLoader (Digits)

AliDataLoader (RecPoints)

AliDataLoader (Tracks)

*AliBasicLoader (TreeD)*

*AliBasicLoader (AliITSDigitizer)*

*AliBasicLoader (QAObject)*

*AliBasicLoader (QATask)*

AliLoader (PHOS)

AliDataLoader (Hits)

AliDataLoader (SDigits)

AliDataLoader (Digits)

AliDataLoader (RecPoints)

AliDataLoader (Tracks)

*AliBasicLoader (TreeD)*

*AliBasicLoader (AliPHOSDigitizer)*

*AliBasicLoader (QAObject)*

*AliBasicLoader (QATask)*

# An update on AliEn @GRID

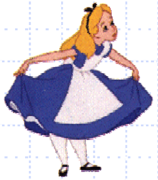

Time for analysis!!

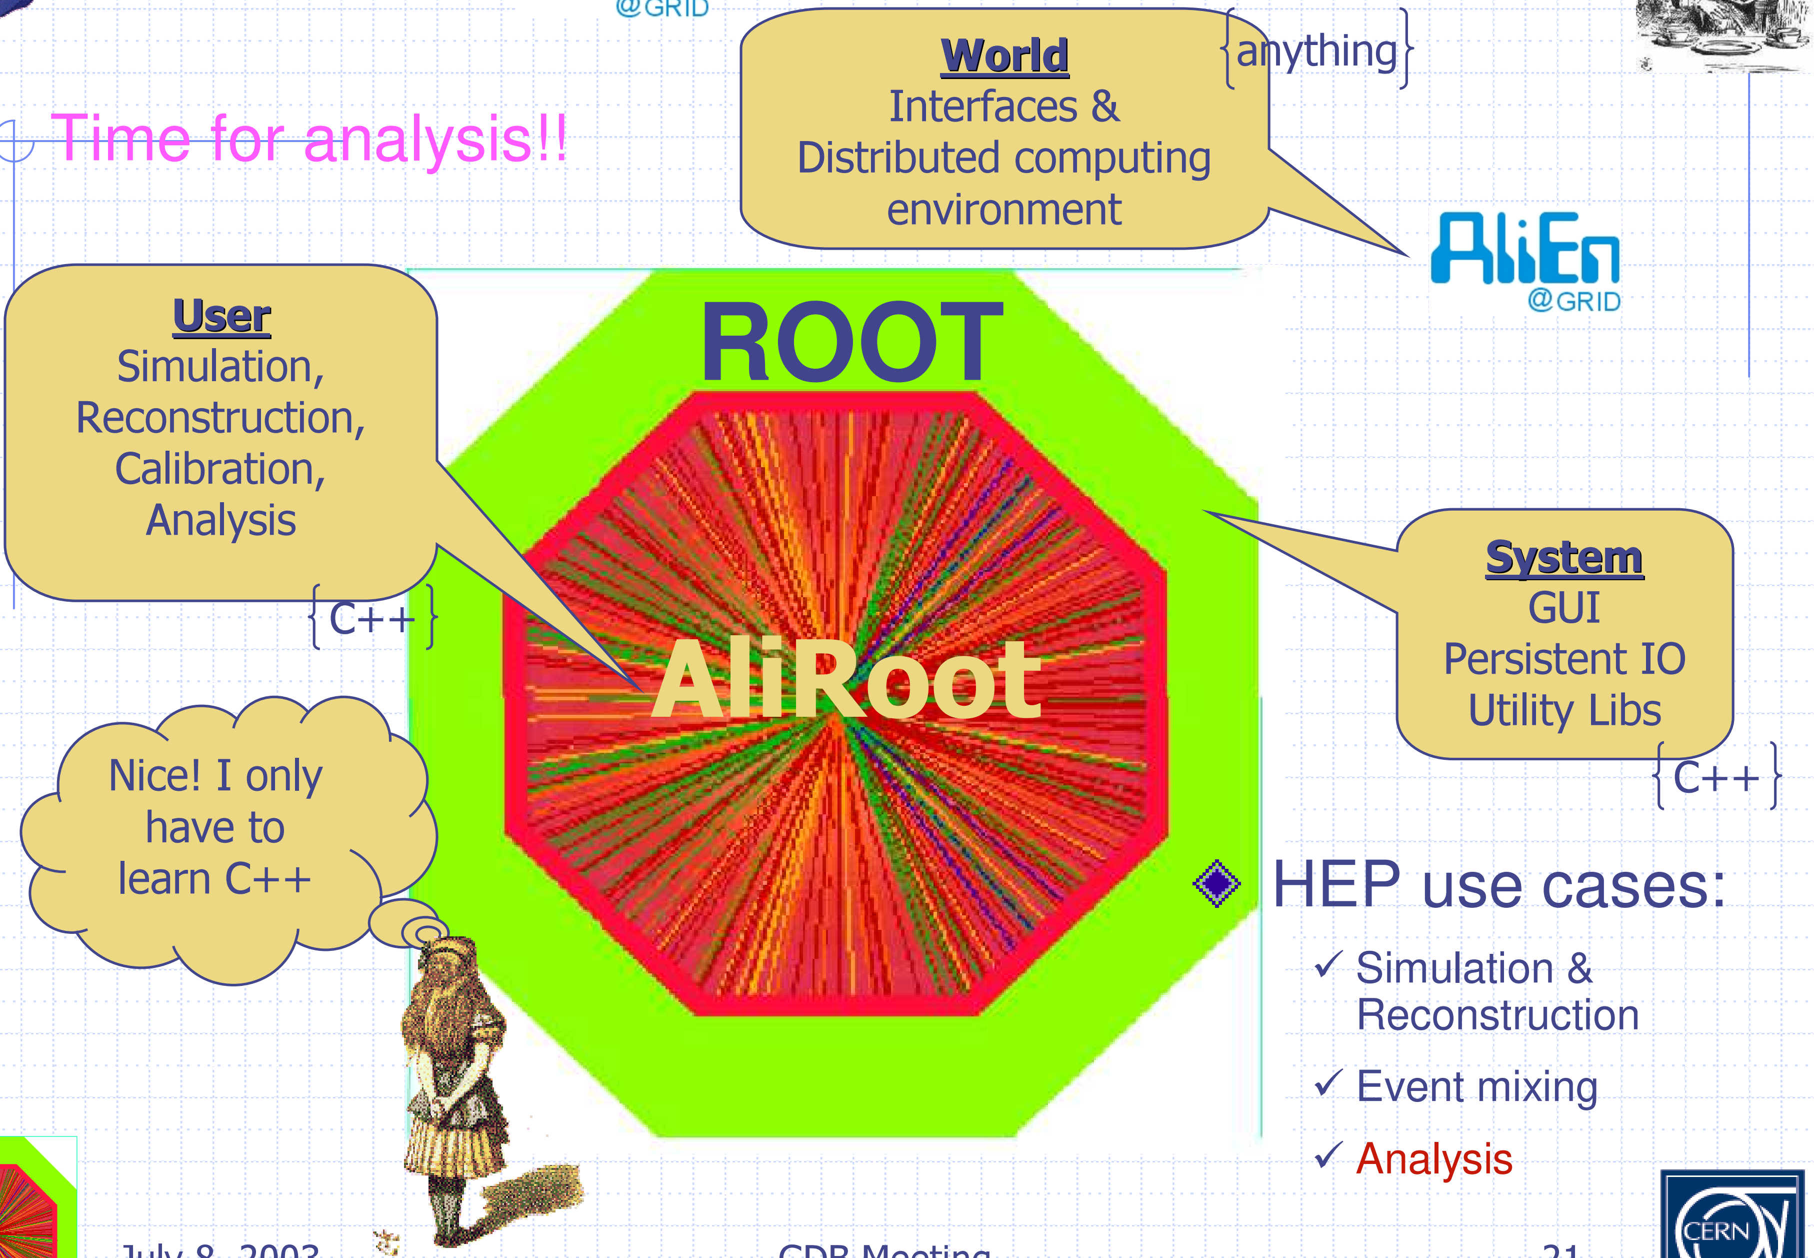

World
Interfaces & Distributed computing environment
{anything}
AliEn @GRID
ROOT
User
Simulation, Reconstruction, Calibration, Analysis
{C++}
AliRoot
System
GUI
Persistent IO
Utility Libs
{C++}
Nice! I only have to learn C++

- HEP use cases:
  - ✓ Simulation & Reconstruction
  - ✓ Event mixing
  - ✓ Analysis

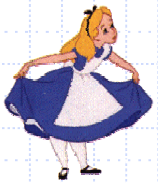

# PPR production summary

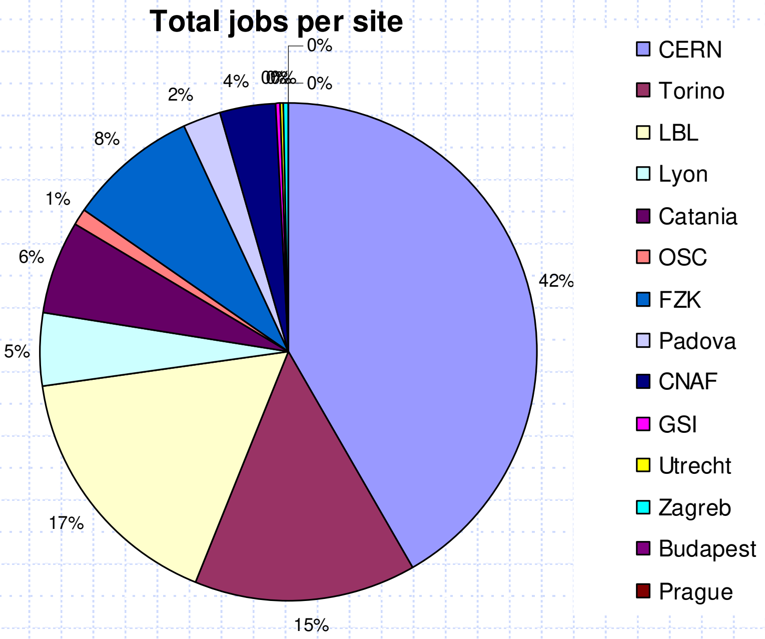

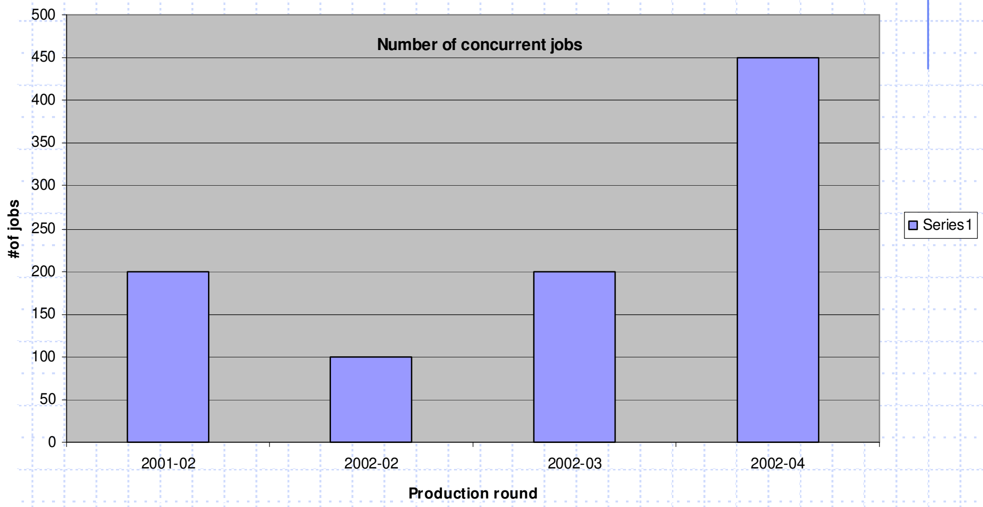

- ◈ 32 (was 28) sites configured
- ◈ 5 (was 4) sites providing mass storage capability
- ◈ 12 production rounds
- ◈ 22773 jobs validated, 2428 failed (10%)
- ◈ Up to 450 concurrent jobs
- ◈ 0.5 operators

# ALICE and GRID

http://www.to.infn.it/activities/experiments/alice-grid

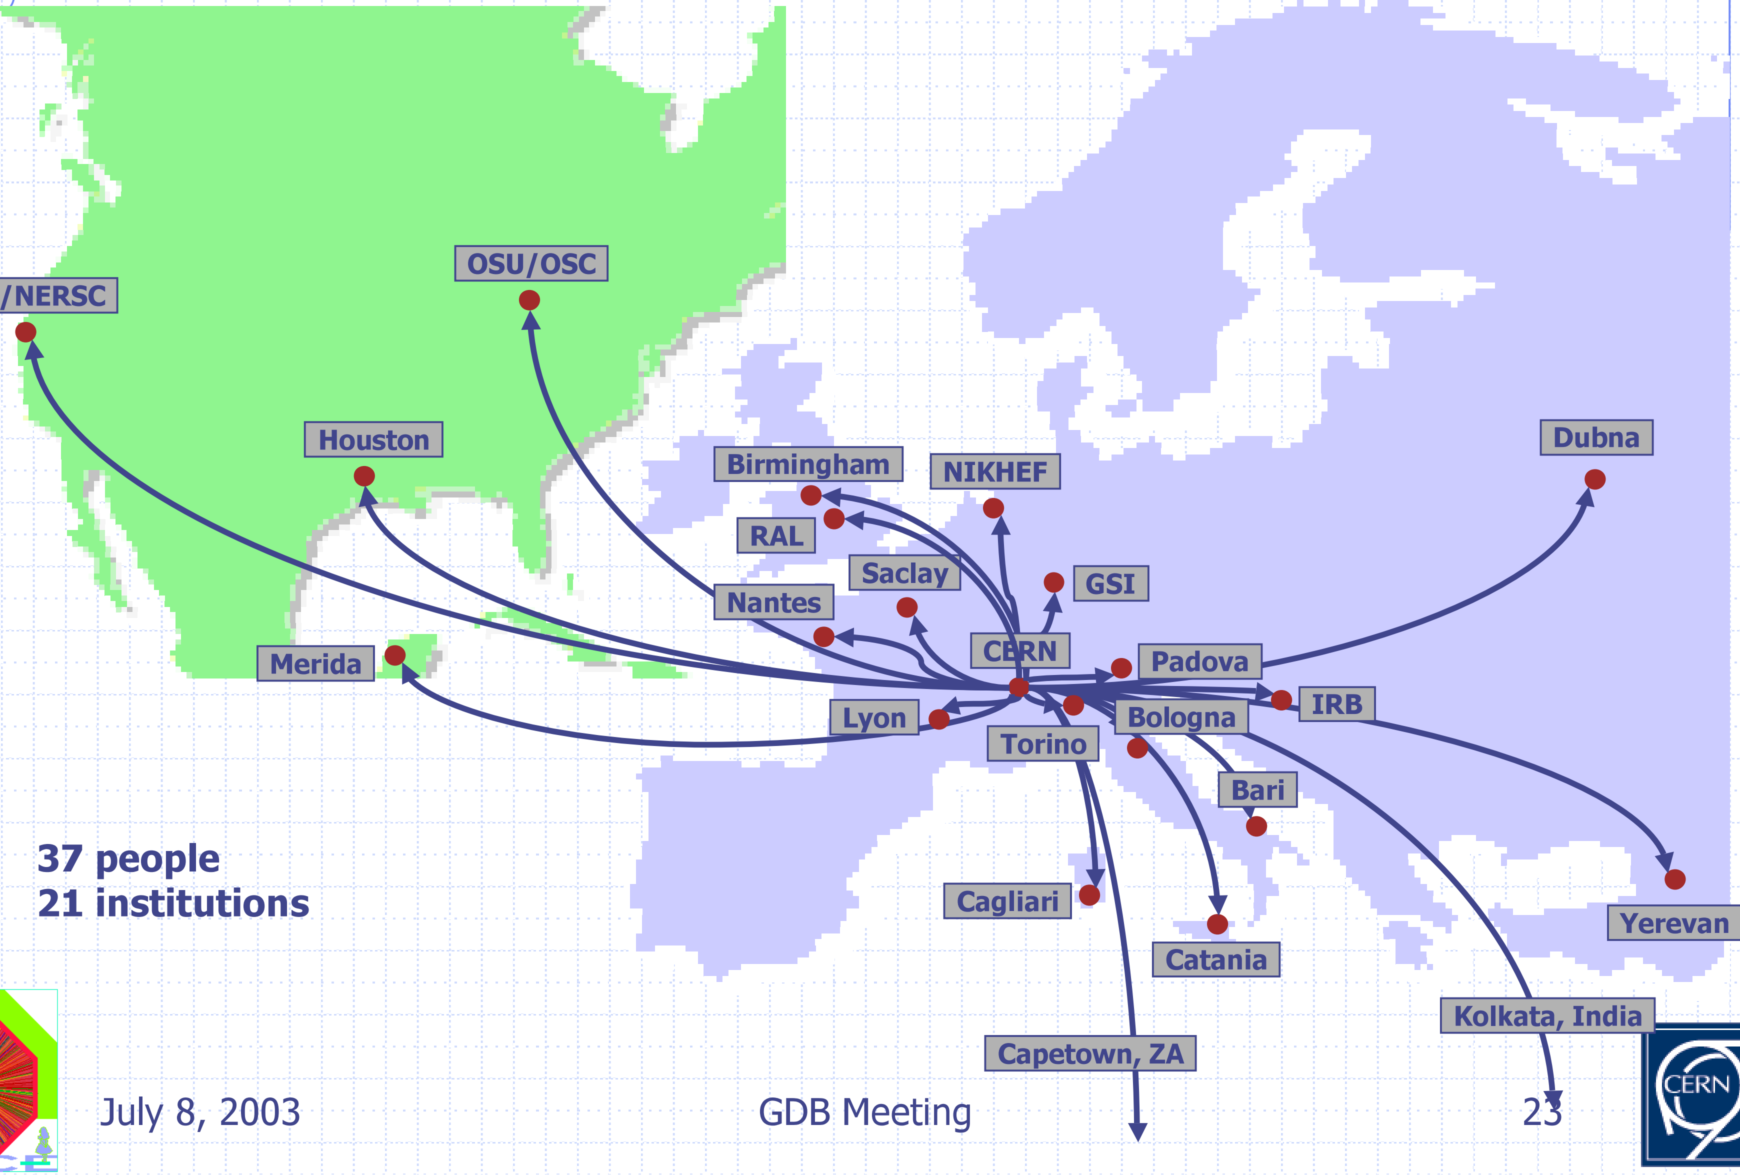

LBL/NERSC
OSU/OSC
Houston
Merida
Birmingham
NIKHEF
RAL
Saclay
Nantes
GSI
CERN
Padova
Dubna
Lyon
Torino
Bologna
IRB
Bari
Cagliari
Catania
Yerevan
Kolkata, India
Capetown, ZA

**37 people**
**21 institutions**

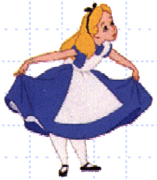

# AliEn as a meta-GRID

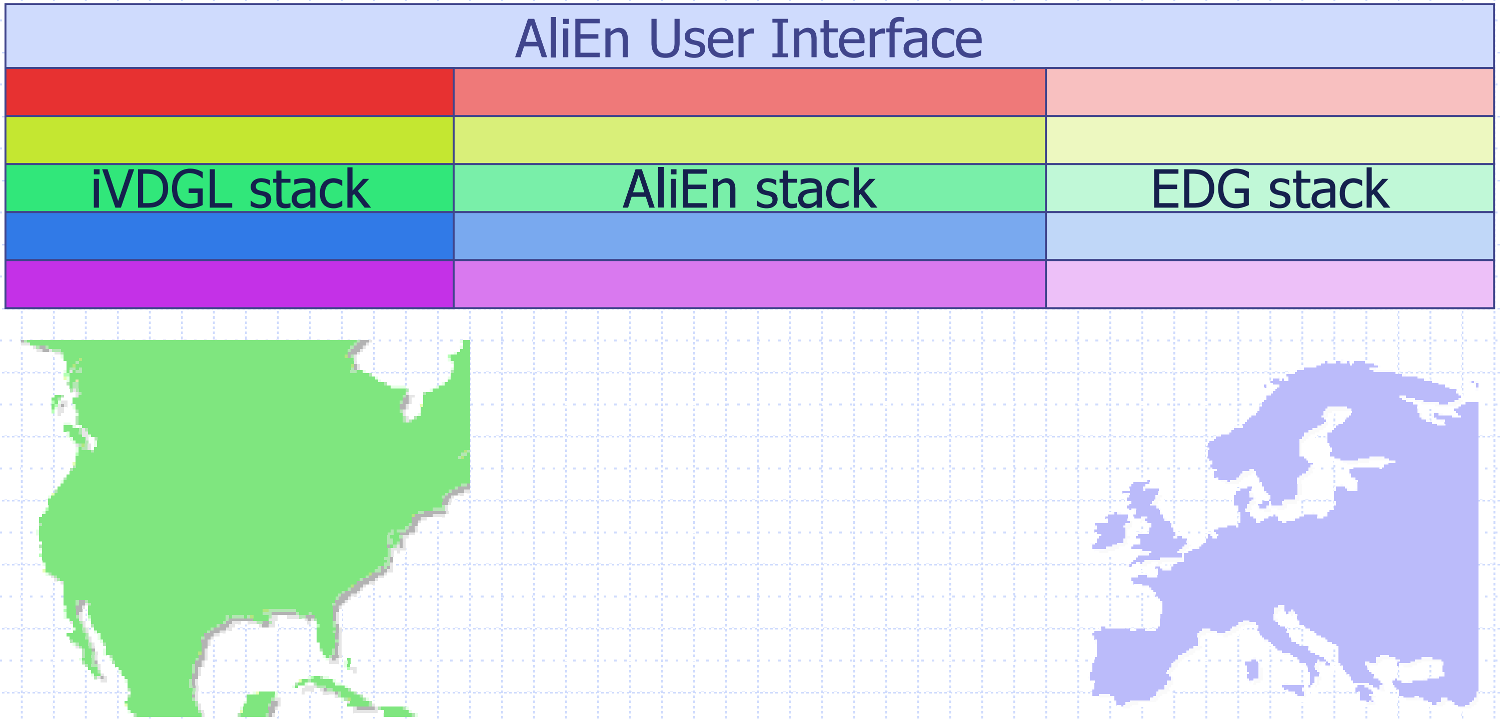

AliEn User Interface

iVDGL stack

AliEn stack

EDG stack

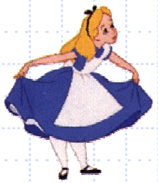

# Communicating with EDG

The first events submitted from AliEn to the EDG 1.4 testbed run on 11/03/2003

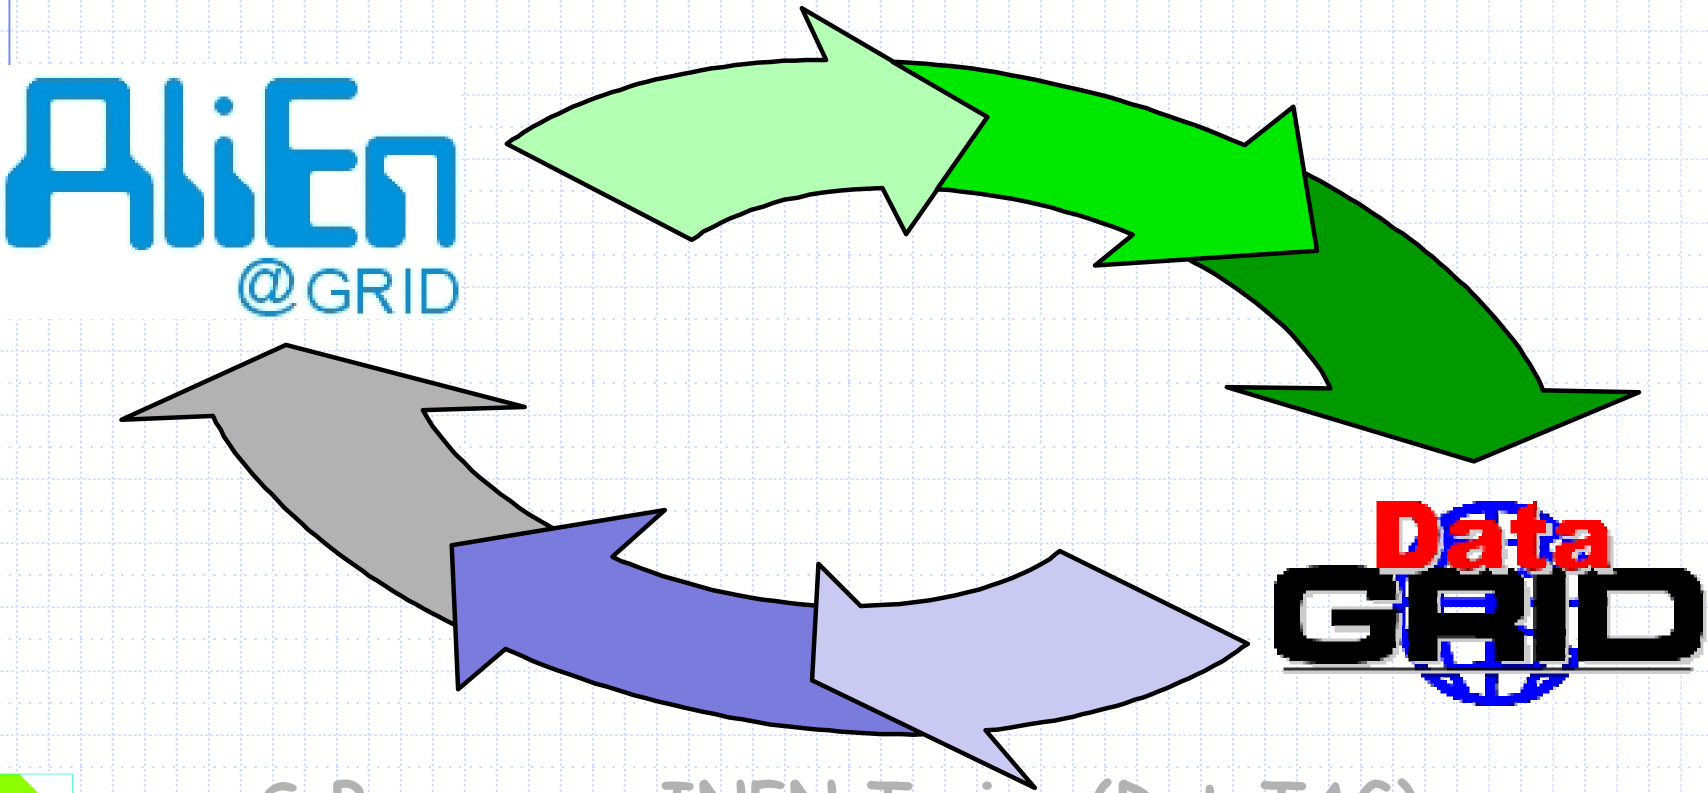
AliEn
@GRID
DataGRID

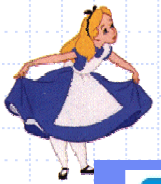

# Communicating with EDG

**The EDG CE's are seen as single AliEn Computing Element and the EDG SE's as single AliEn SE.**

- The interface site runs the EDG UI and the AliEn client suite - ClusterMonitor, CE e SE
- The interface client takes ("pulls") a job from the server, generates the GLUE-compliant JDL and submits to the EDG resource broker (RB)
- On the WN the job activates AliEn that communicates with the server

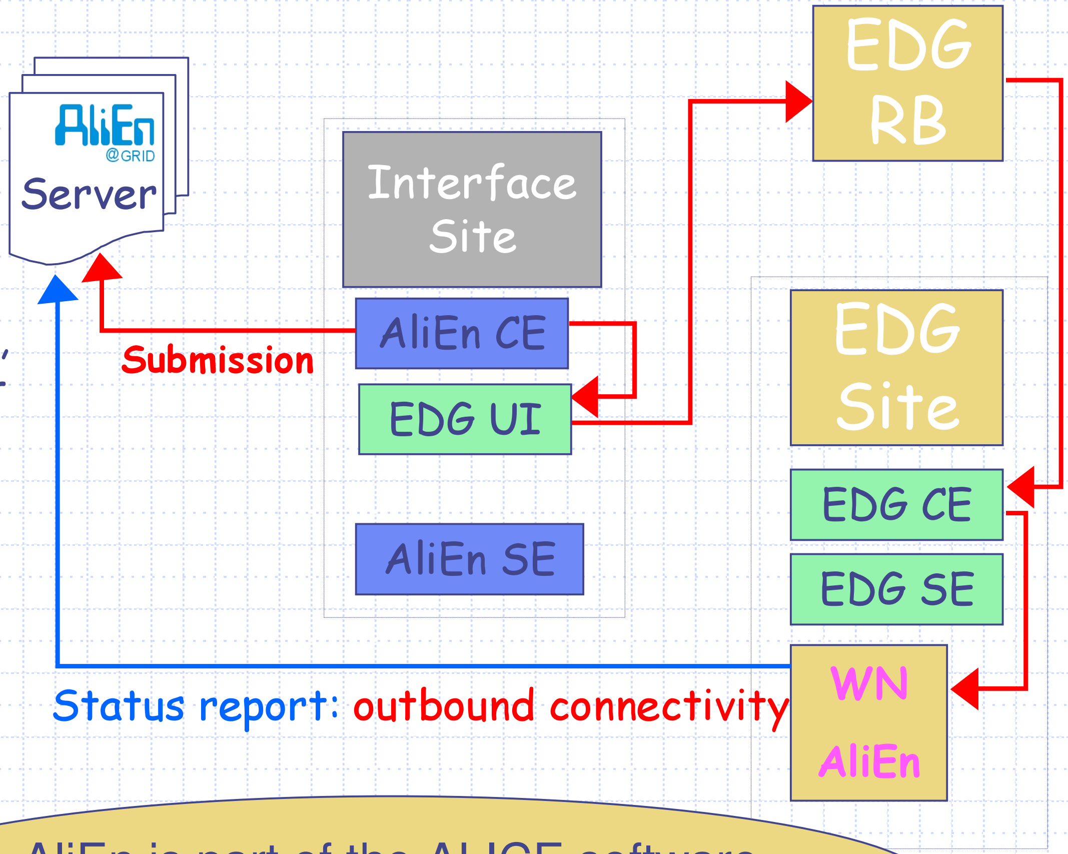
AliEn @GRID
Server
Interface Site
AliEn CE
EDG UI
AliEn SE
Submission
EDG RB
EDG Site
EDG CE
EDG SE
WN
AliEn
Status report: outbound connectivity

AliEn is part of the ALICE software
No daemons on the WNs

# Communicating with EDG

- Output generated by the AliEn job on an EDG CE is saved on an EDG SE. Information is passed to AliEn via .Brokerinfo file
- Files are registered in the EDG RC
- Files are also registered in the AliEn Data Catalogue using the EDG LFN as AliEn PFN
- File access from AliEn becomes a two steps process via AliEn RC and EDG RC
- The interface node can handle this and stage the PFN via the two systems

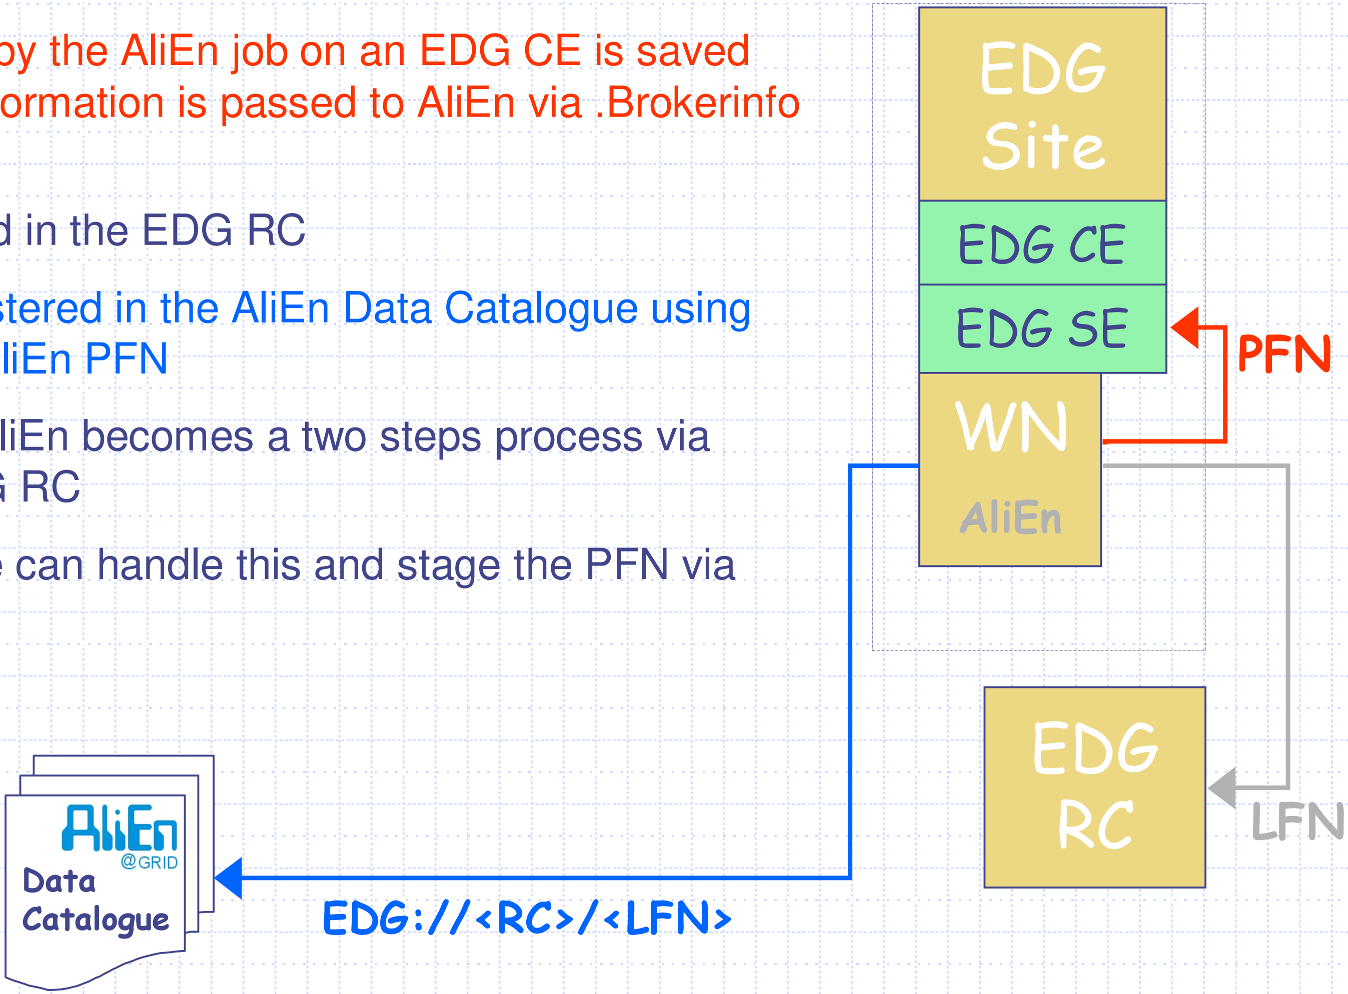
EDG Site
EDG CE
EDG SE
PFN
WN
AliEn
EDG RC
LFN
AliEn @GRID
Data Catalogue
EDG://<RC>/<LFN>

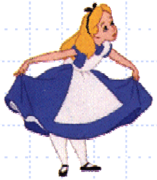

# How does it work?

- Failure rate for file transfer is still 90%
  - EDG instability (service disruption)
  - Known EDG problems with some sites
  - Not enough space on the interface staging area (it was nobody's responsibility, should be solved now)
  - Occasional crashes of the File Transfer Daemons due to slow file transfer performance
- With EDG 1.4 problem tracking is very difficult
- No fundamental problems, should be better with EDG 2.0

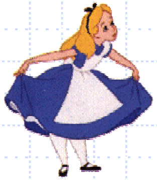

# AliEn Timeline

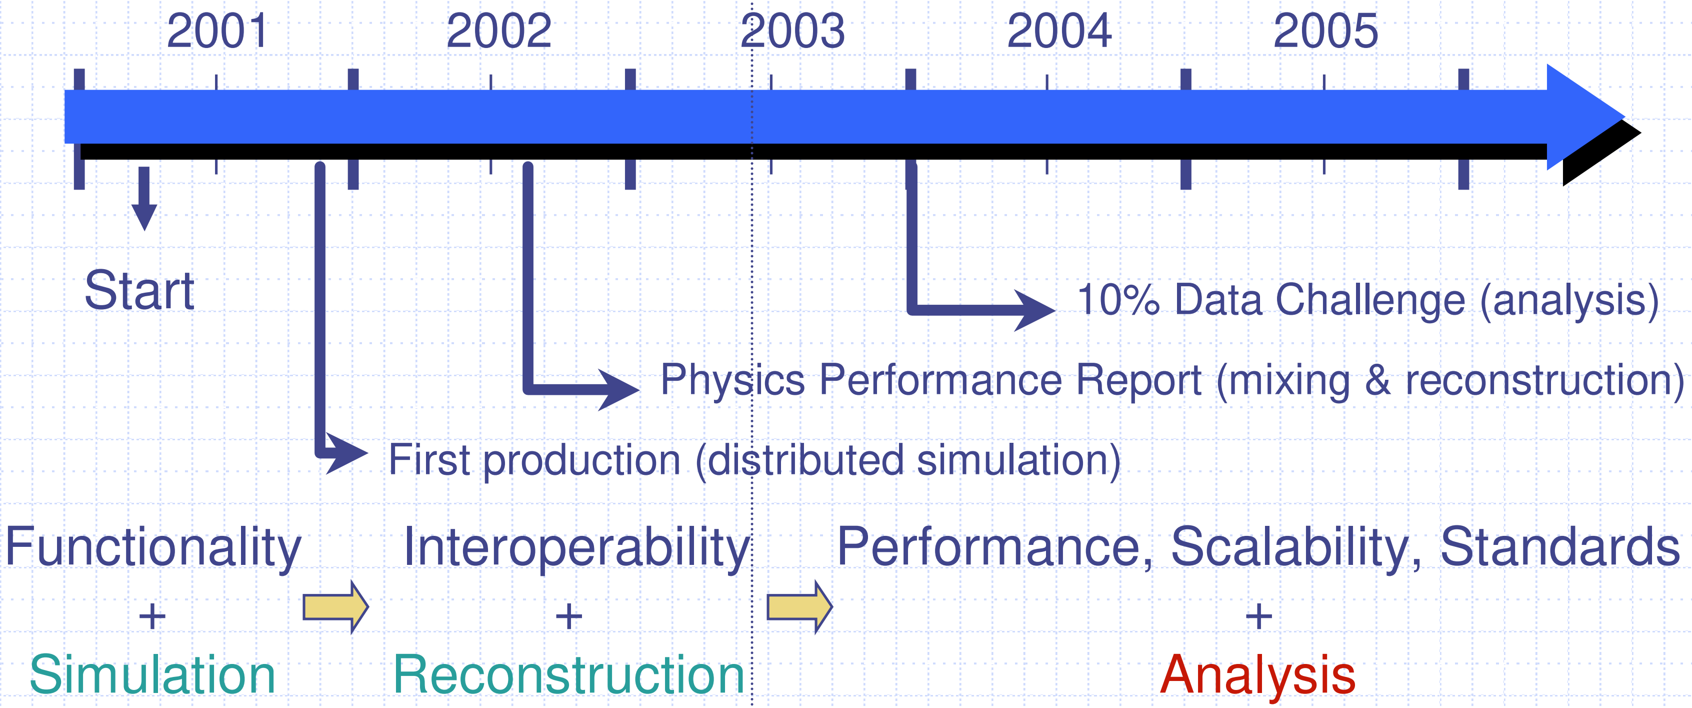

2001
2002
2003
2004
2005
Start
10% Data Challenge (analysis)
Physics Performance Report (mixing & reconstruction)
First production (distributed simulation)

Functionality
+
Simulation

Interoperability
+
Reconstruction

Performance, Scalability, Standards
+
Analysis

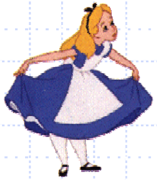

# Distributed Analysis @ GRID

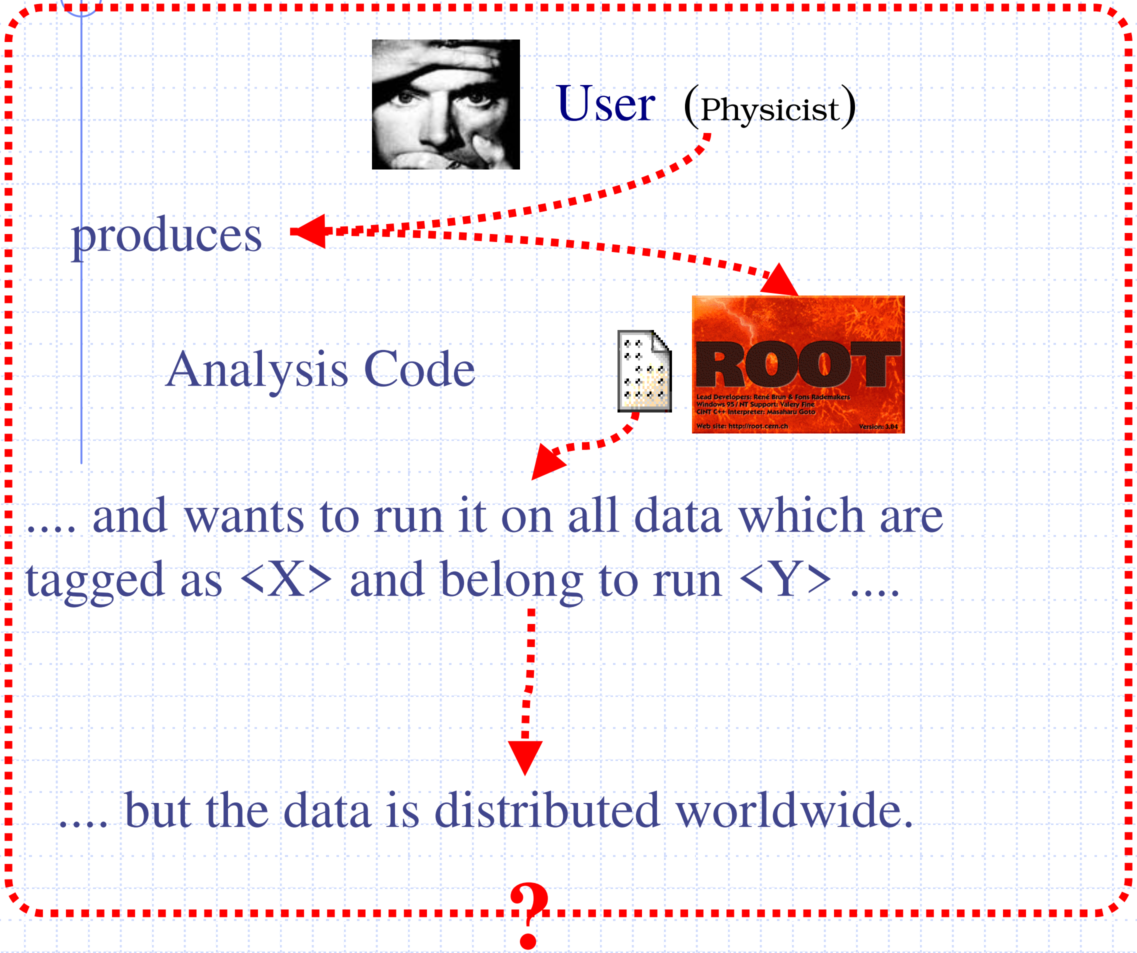

User (Physicist)

produces

Analysis Code

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

.... but the data is distributed worldwide.

?

A. Asynchronous
- Pseudo-interactive
- Job splitting + batch processing (transparent to end user)
- Can be done using existing tools (AliEn + ROOT)

B. True Interactive
- Highly synchronous system
- Instantaneous analysis results
- Needs
  - New functionality (AliEn + PROOF)
  - High system availability

# Going to Analysis, AliEnfs

- AliEn catalogue can be *mounted* as a Linux fs
- It looks and smells like a normal Linux fs
- All the storage in the world as a single (large!) file system
- HOWEVER: access for files on MSS depends on local access time (no miracle!)

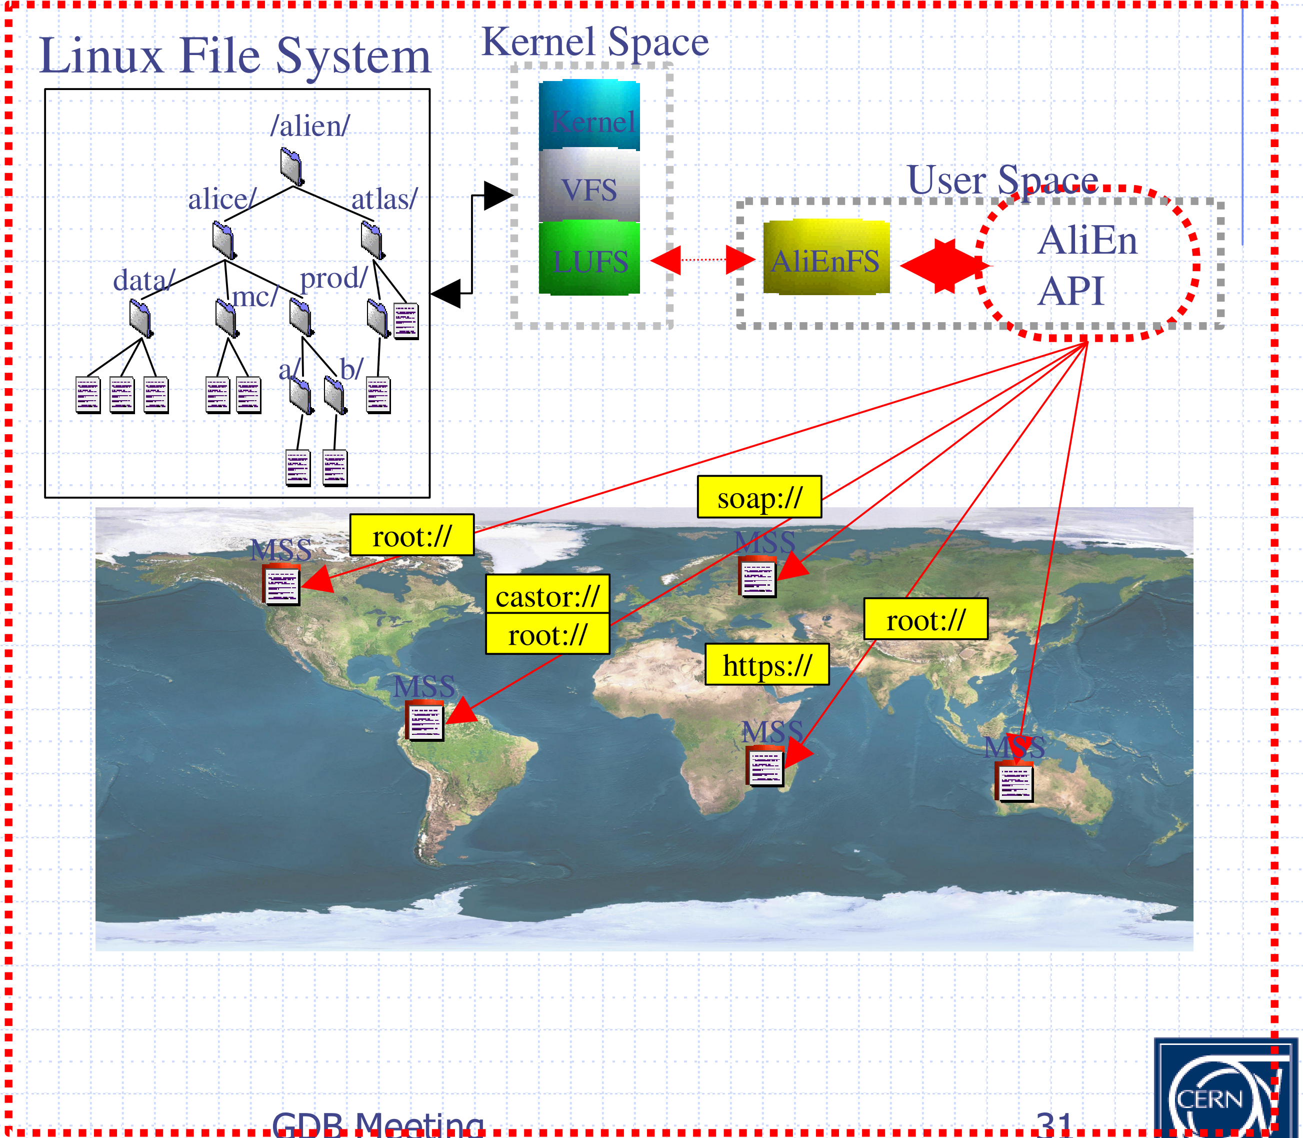
Linux File System
/alien/
alice/
atlas/
data/
mc/
prod/
a/
b/
Kernel Space
Kernel
VFS
LUFS
User Space
AliEnFS
AliEn API
root://
soap://
castor://
root://
root://
https://
MSS
MSS
MSS
MSS
MSS

# Crosslink-Cache

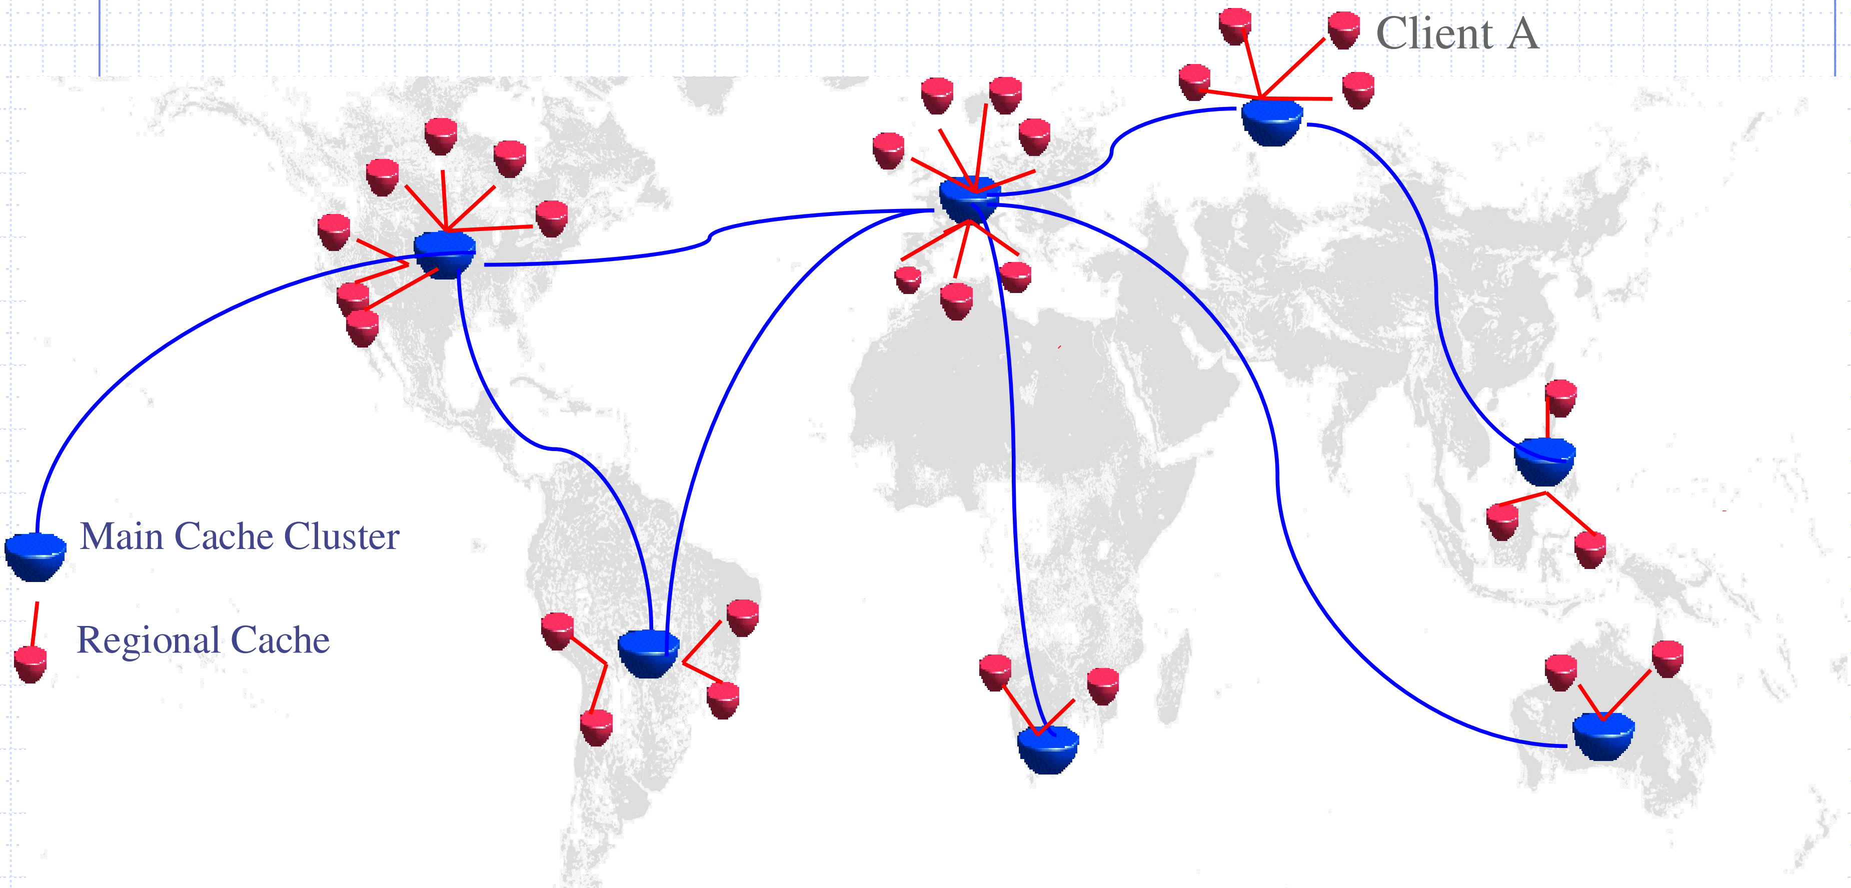
Client A
Main Cache Cluster
Regional Cache

Cache Levels: ***main – regional – local***

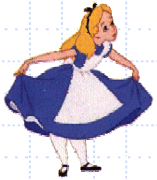

# Distributed analysis: job splitting

- ROOT/AliEn protype 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

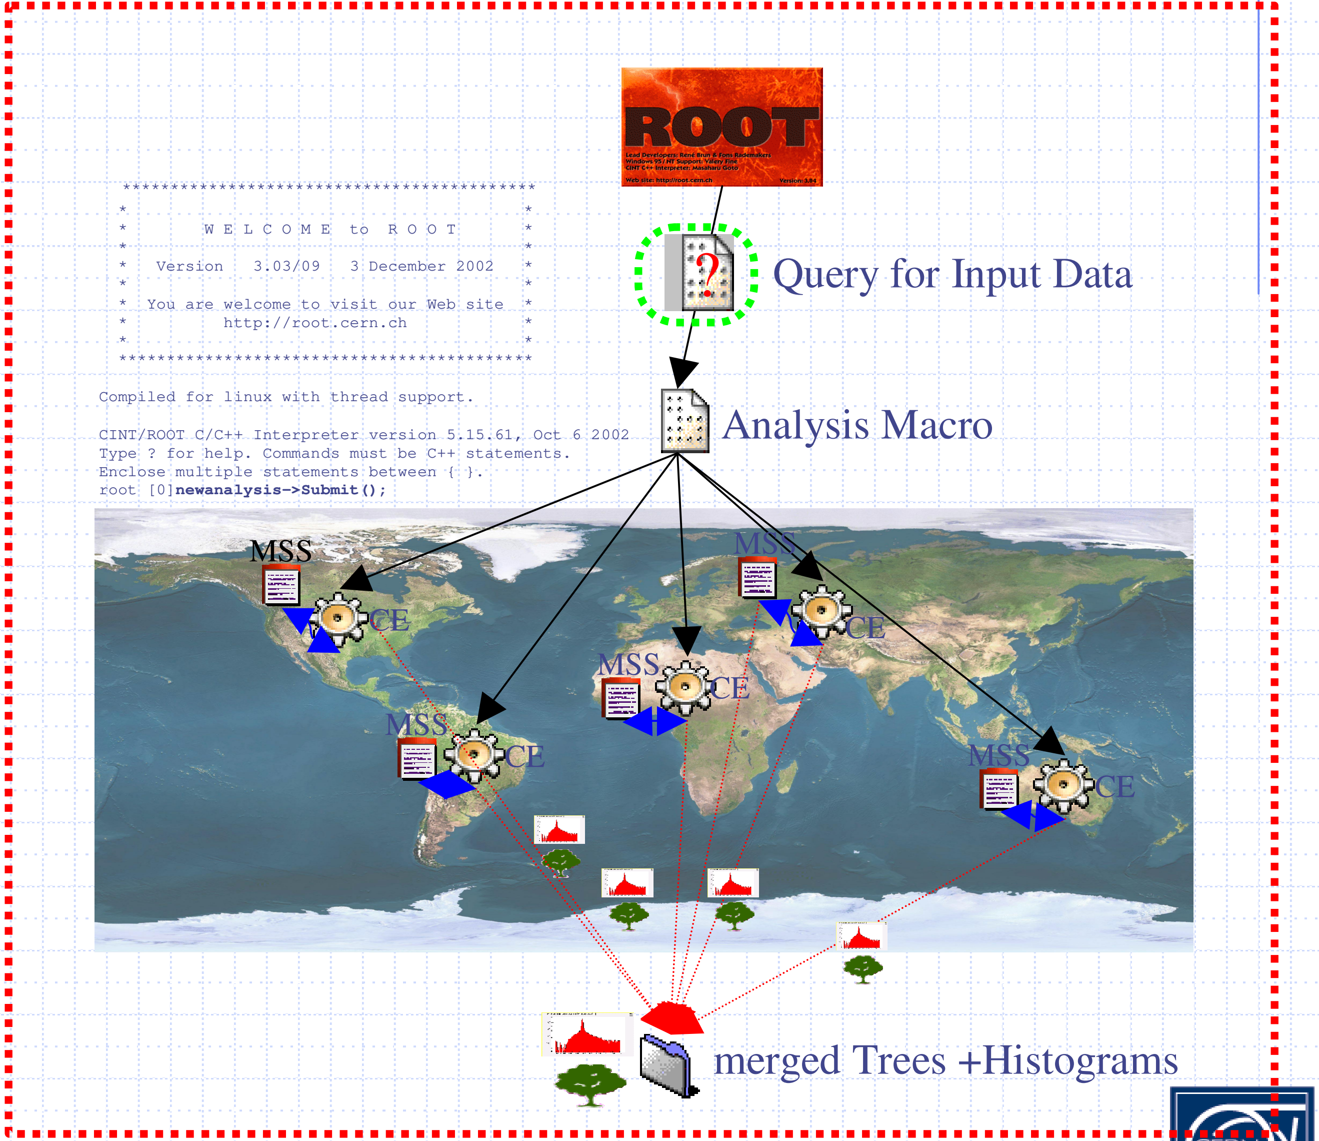

```
 ******************************************
*                                          *
*        W E L C O M E  to  R O O T        *
*                                          *
*   Version   3.03/09   3 December 2002    *
*                                          *
*  You are welcome to visit 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();
```

ROOT
Query for Input Data
Analysis Macro
MSS
CE
MSS
CE
MSS
CE
MSS
CE
MSS
CE
merged Trees +Histograms

# Classic PROOF

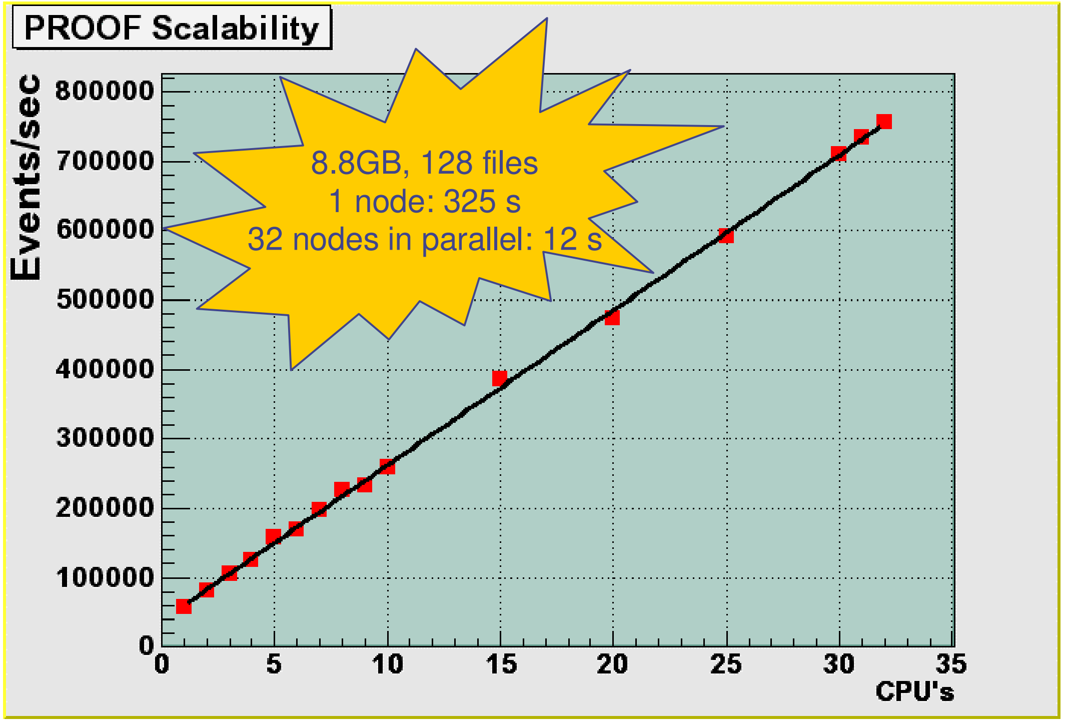
PROOF Scalability

8.8GB, 128 files
1 node: 325 s
32 nodes in parallel: 12 s

32 nodes: dual Itanium II 1 GHz CPU's,
2 GB RAM, 2x75 GB 15K SCSI disk,
1 Fast Eth, 1 GB Eth nic (not used)

Each node has one copy of the data set
(4 files, total of 277 MB), 32 nodes:
8.8 Gbyte in 128 files, 9 million events

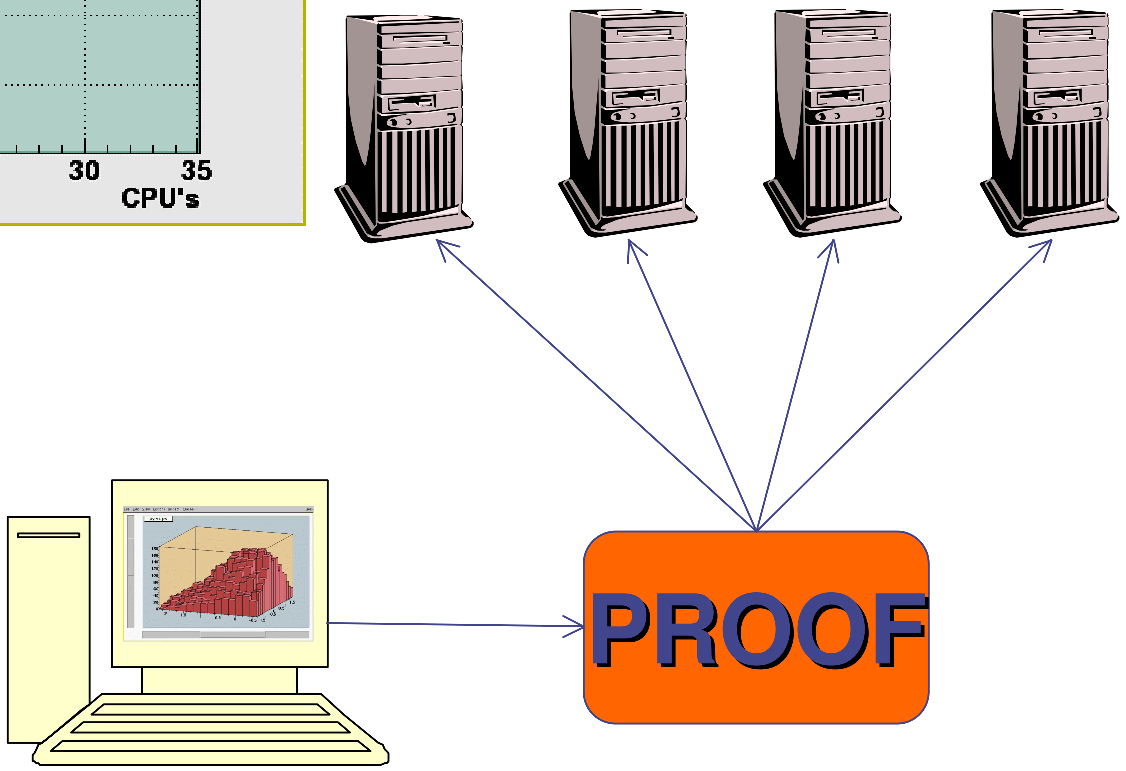
**PROOF SLAVE SERVERS**

**PROOF**

**USER SESSION**

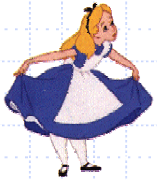

# GRID (Super) PROOF

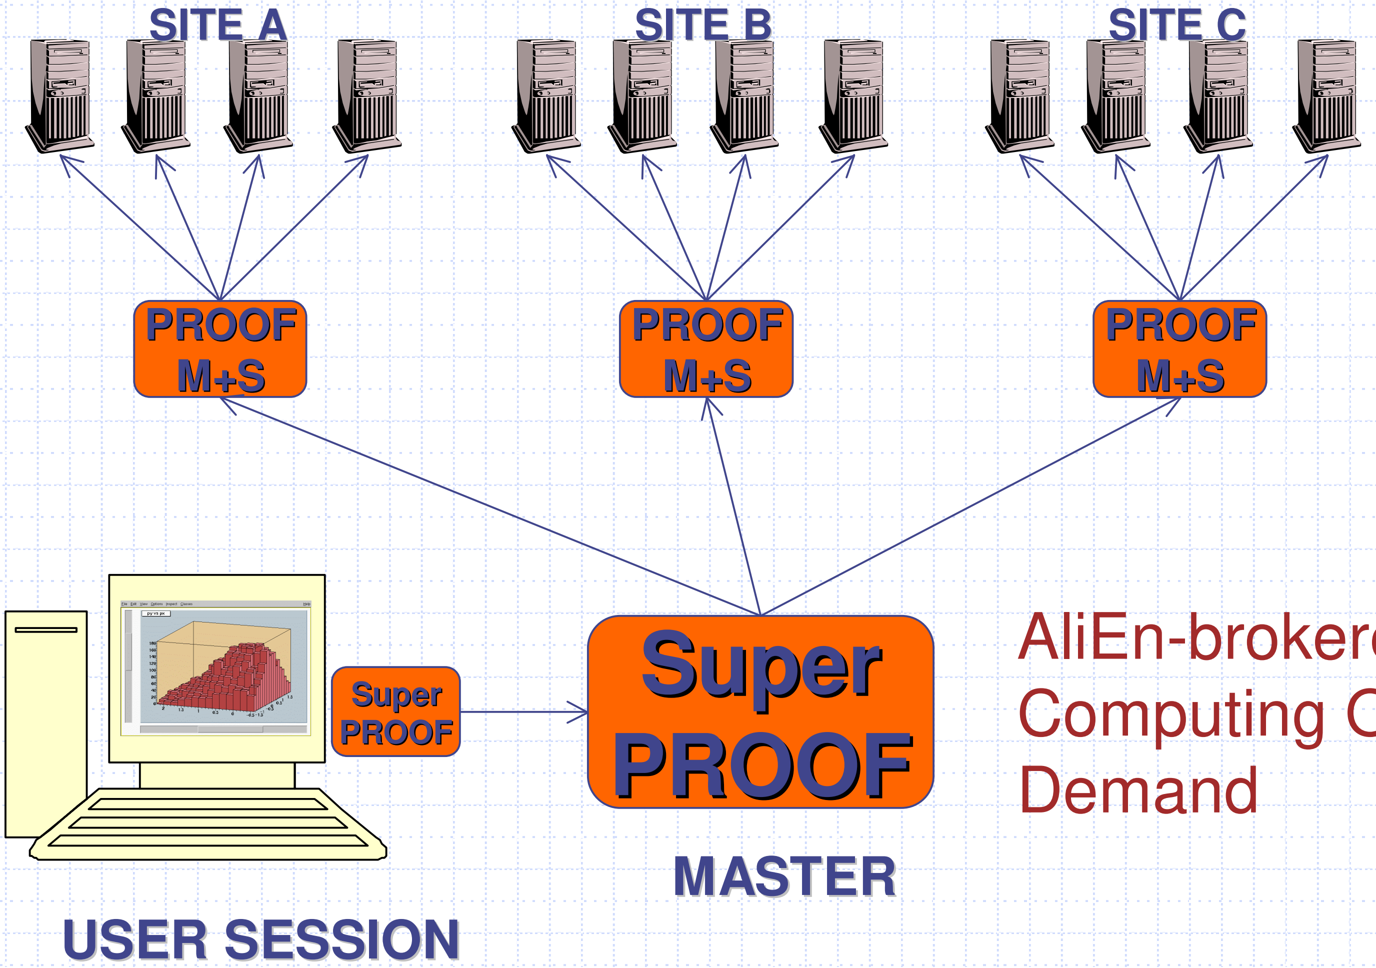

SITE A

SITE B

SITE C

PROOF M+S

PROOF M+S

PROOF M+S

Super PROOF

Super PROOF

MASTER

USER SESSION

AliEn-brokered
Computing On
Demand

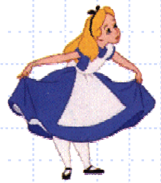

# AliEn + ROOT

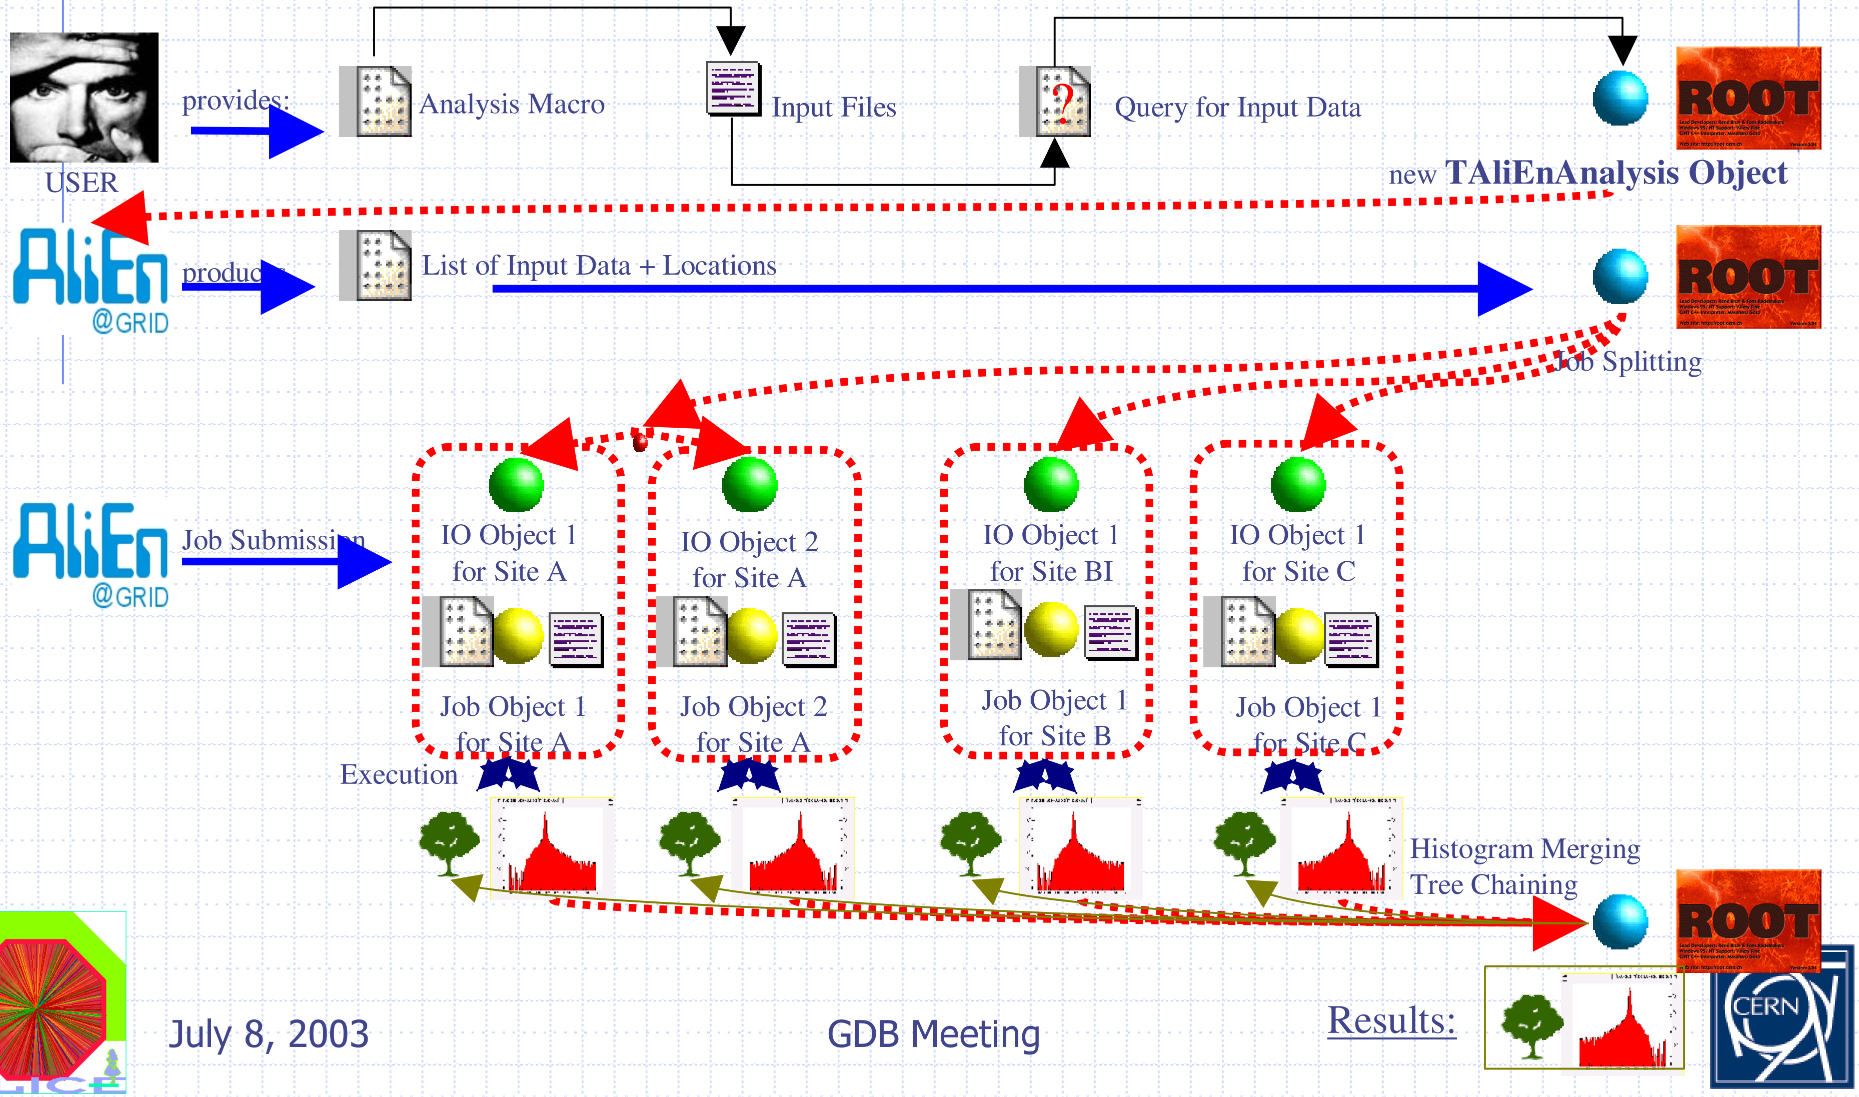

USER provides:

Analysis Macro

Input Files

Query for Input Data

ROOT

new **TAliEnAnalysis Object**

AliEn @GRID produces

List of Input Data + Locations

ROOT

Job Splitting

AliEn @GRID Job Submission

IO Object 1 for Site A

Job Object 1 for Site A

IO Object 2 for Site A

Job Object 2 for Site A

IO Object 1 for Site BI

Job Object 1 for Site B

IO Object 1 for Site C

Job Object 1 for Site C

Execution

Histogram Merging

Tree Chaining

ROOT

Results:

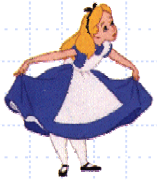

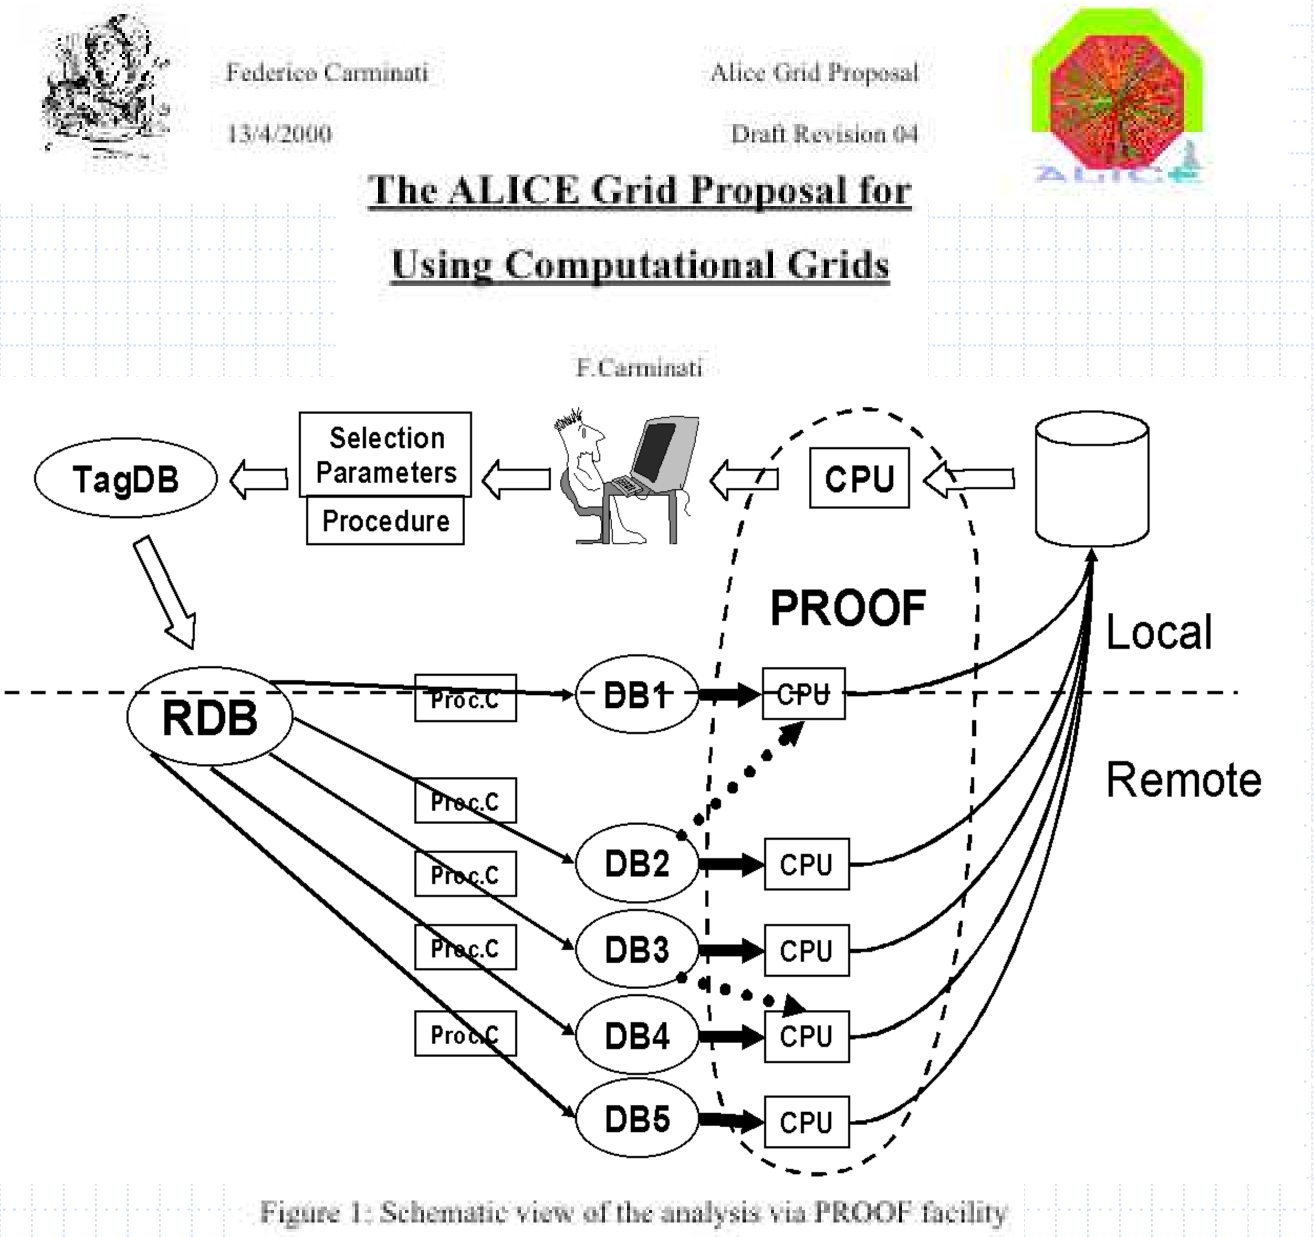

Federico Carminati

13/4/2000

Alice Grid Proposal

Draft Revision 04

# The ALICE Grid Proposal for Using Computational Grids

F.Carminati

TagDB
Selection Parameters
Procedure
CPU
PROOF
Local
Remote
RDB
Proc.C
Proc.C
Proc.C
Proc.C
Proc.C
DB1
DB2
DB3
DB4
DB5
CPU
CPU
CPU
CPU
CPU

Figure 1: Schematic view of the analysis via PROOF facility

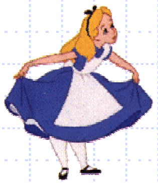

# The Caltech High Energy Physics Department
# Grid Enabled Analysis Workshop
# June 23-25, 2003

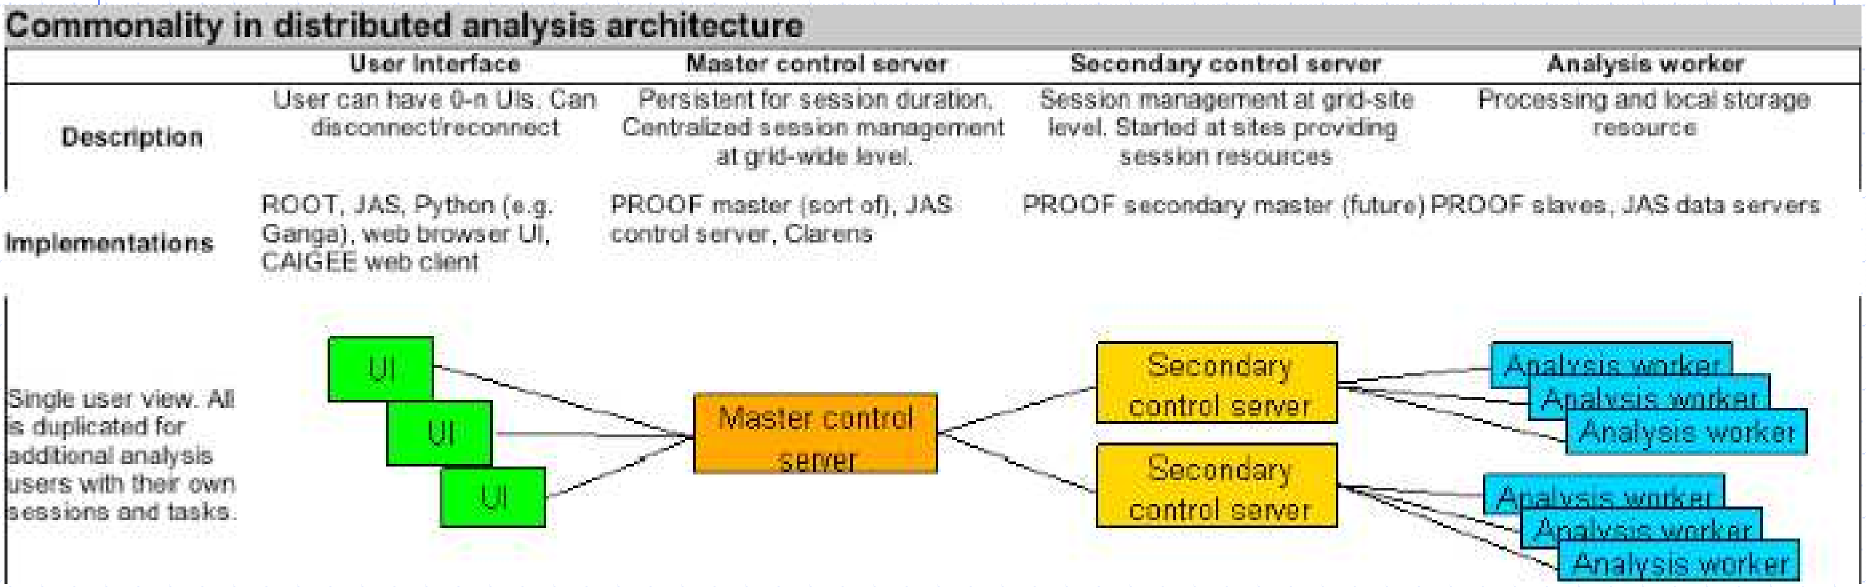

**Commonality in distributed analysis architecture**

| | **User Interface** | **Master control server** | **Secondary control server** | **Analysis worker** |
|---|---|---|---|---|
| **Description** | User can have 0-n UIs. Can disconnect/reconnect | Persistent for session duration. Centralized session management at grid-wide level. | Session management at grid-site level. Started at sites providing session resources | Processing and local storage resource |
| **Implementations** | ROOT, JAS, Python (e.g. Ganga), web browser UI, CAIGEE web client | PROOF master (sort of), JAS control server, Clarens | PROOF secondary master (future) | PROOF slaves, JAS data servers |

Single user view. All is duplicated for additional analysis users with their own sessions and tasks.

UI
UI
UI
Master control server
Secondary control server
Secondary control server
Analysis worker
Analysis worker
Analysis worker
Analysis worker
Analysis worker
Analysis worker

Status:
Missing & needed
Work needed
Essentials exist
Low/future priority

**Bold=priority work**

All this is basically there with AliEn/Proof/Root!

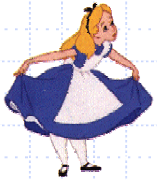

# ALICE’s experience with EDG

- Produce and analyse 5000 events Pb-Pb (HBT) on EDG 1.4 testbed
    - 1 job/event, 12-24 CPU hours, 1.8 GB
    - Total of 9TB, ~90,000 CPU h
- Start March 15-03 -- Stop May 15-03
- No *task force* provided by EDG as it was the case for ATLAS and CMS
    - But a very active support
- 450 events produced (6/day)
    - Enough to study pion correlations
    - Output registered in the EDG ALICE RC and stored
        - EDG disk SE's (300TB) and MSS SE's (150TB)
        - CASTOR at CNAF and CERN
        - All registered in the AliEn Data Catalogue
- No analysis done

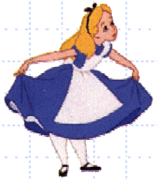

# ALICE’s experience with EDG

- RB overloaded
  - Service crash, jobs get lost even if executing, and all results are gone
- WN disk space full
  - Disk space on the WN is filled up and the job crashes
- The "Lyon" problem
  - WN's publish their **total** memory in the IS
  - The JDL memory requirement is compared to the published values, but more than one job is allowed on the WN (workaround by F. Hernandez)
- Behaviours not understood
  - Jobs go to "OutputReady" status after 6-8 days
  - MSS jobs fail more frequently (and job information only available for CNAF jobs)

| MSS jobs | |
|---|---|
| OK | 74 |
| LDAP failure | 23 |
| RC failure | 35 |
| Disk full | 16 |
| Lost | 32 |
| Wrapper | 39 |
| Running | 36 |
| Submit | 15 |
| Total | 270 |

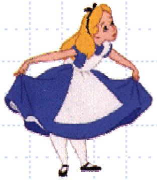

# ALICE’s experience with EDG

- 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
    - Its use is very frustrating: instability, limited functionality, low efficiency
    - At the experienced rate, it would take 18 months  to complete the production :-(
  - We will try another time when EDG 2.0 will be available

# Pre PDC 3 network "Stress test"

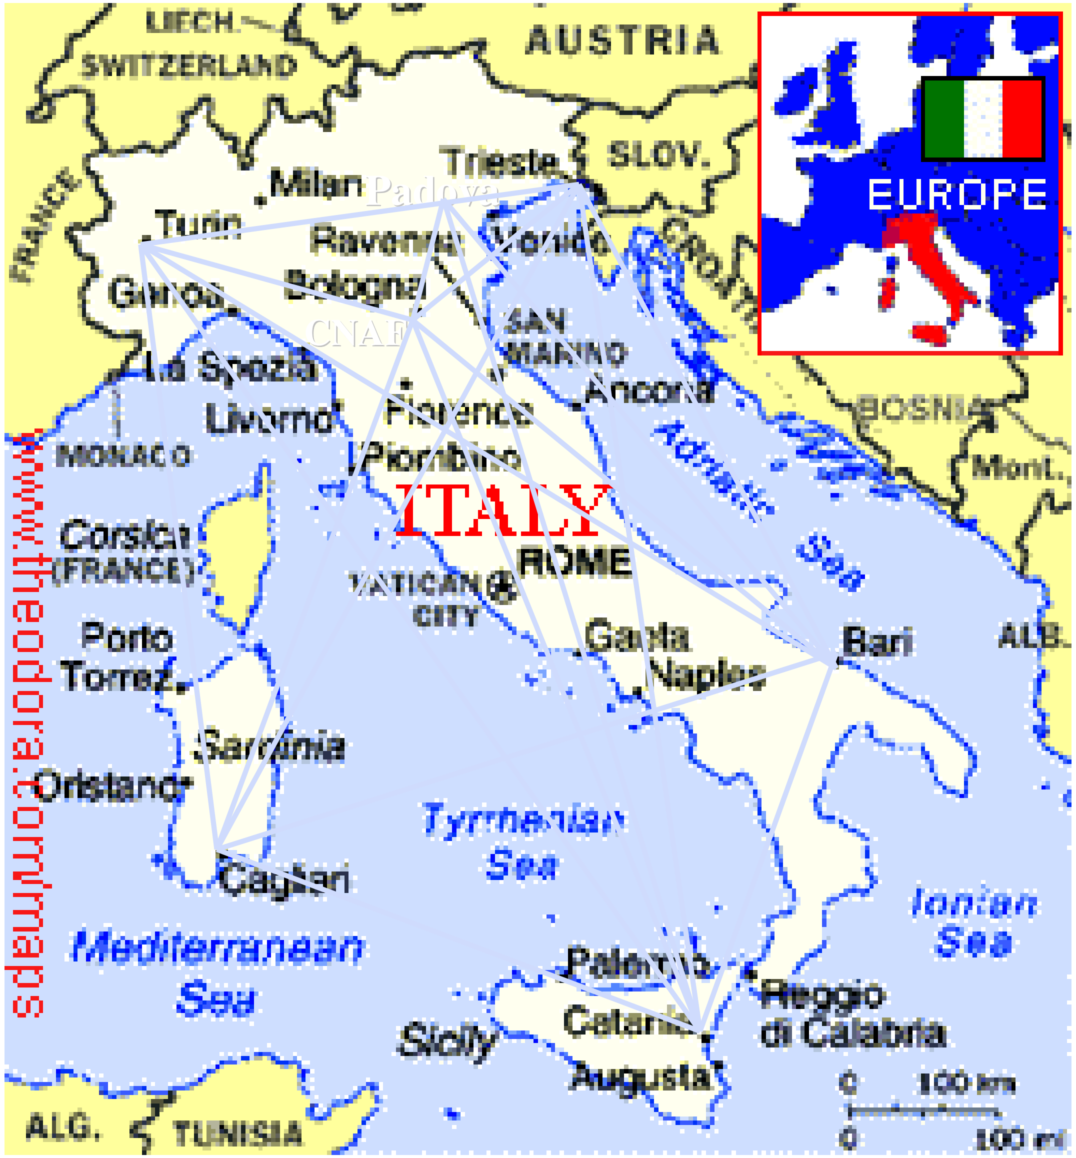
LIECH.
SWITZERLAND
AUSTRIA
EUROPE
FRANCE
SLOV.
CROATIA
Trieste
Milan
Padova
Turin
Ravenna
Venice
Genoa
Bologna
CNAF
SAN MARINO
La Spezia
Ancona
BOSNIA
Livorno
Firenze
Piombino
MONACO
Adriatic Sea
Mont.
ITALY
Corsica (FRANCE)
ROME
VATICAN CITY
Porto Torrez
Gaeta
Naples
Bari
ALB.
Sardinia
Oristano
Tyrrhenian Sea
Cagliari
Ionian Sea
Mediterranean Sea
Palermo
Reggio di Calabria
Sicily
Catania
Augusta
0 100 km
0 100 mi
ALG.
TUNISIA
www.theodora.com/maps

- ALICE has made a stress test on the Italian network
- Objectives
  - Where are the bottlenecks
  - Are bandwidths enough?
- Results
  - Bandwidth saturated in some sites at PDC nominal conditions
  - Farm architecture: clear nfs inadequacy for large files and multiple threads
- An international version of the test is in preparation

**BA: 3 servers (2 ALICE, 1 CMS)**
**BO: 6 servers**
**CA: 2 servers**
**CNAF: 2 servers**
**CT: 2 servers**
**PD: 6 servers**
**TO: 2 servers**
**TS: 1 server**

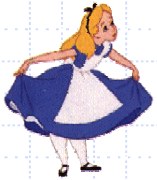

# PDC 3 schema

- Production of RAW
- Shipment of RAW to CERN
- Reconstruction of RAW
- Analysis

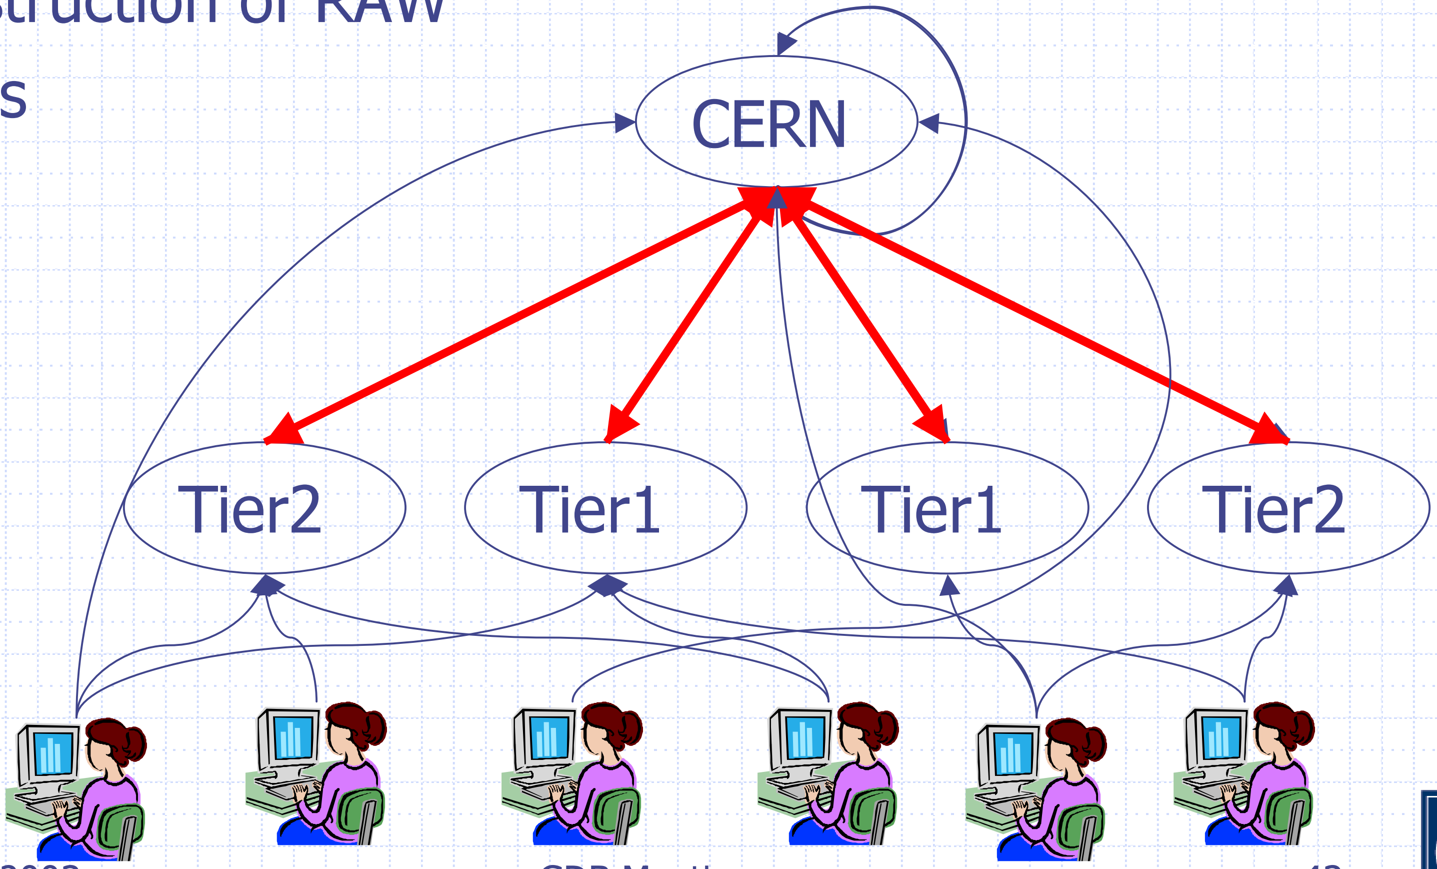

← AliEn job control

← Data transfer

CERN
Tier2
Tier1
Tier1
Tier2

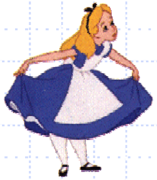

# Resource estimate

- Simulation
  - $10^5$ Pb-Pb + $10^7$ p-p
  - Distributed production, (partial) data replication at CERN
- Reconstruction and analysis
  - $5\times10^6$ Pb-Pb + $10^7$ p-p
  - Reconstruction at CERN and outside depending on resource availability
  - Data source is CERN
- Resources (CPU and Storage)
  - 2004 Q1: 1354 KSI2K and 165 TB
  - 2004 Q2: 1400 KSI2K and 301 TB
  - Defined in the Data Challenge document presented at the PEB, reviewed in the light of the reduced mixing factor
- Bandwidth
  - Simulation in 2004 Q1
  - ~90 TB will be shipped to CERN in about 2 months → ~10 days using 10% of the CERN bandwidth.
- Adjustments done at the June Offline week and being finalised now

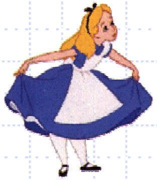

# Foreseen resource profile

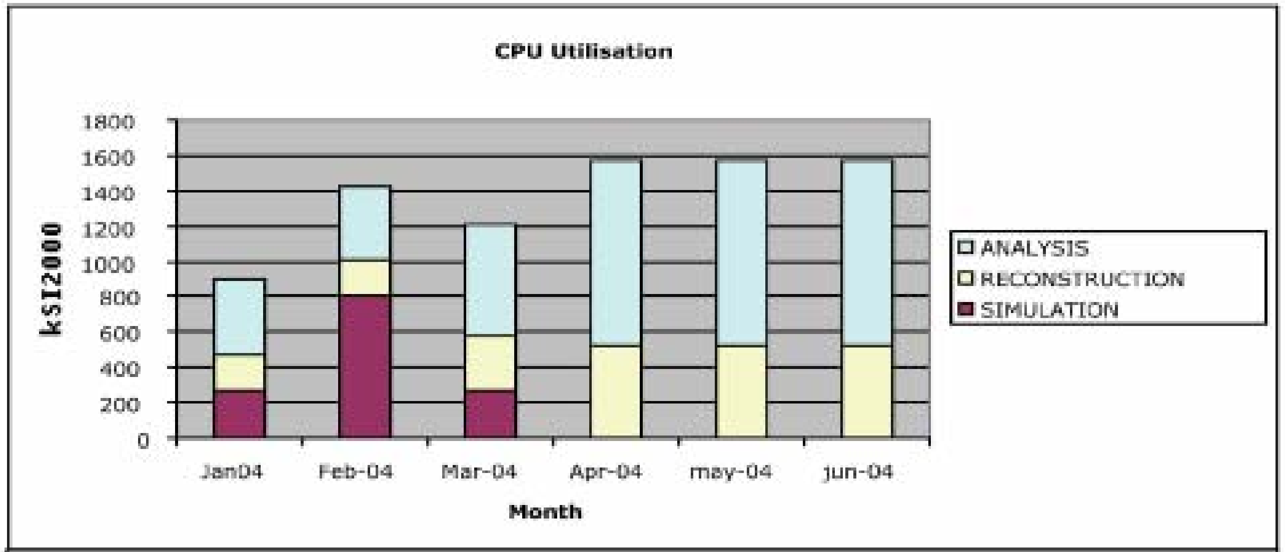

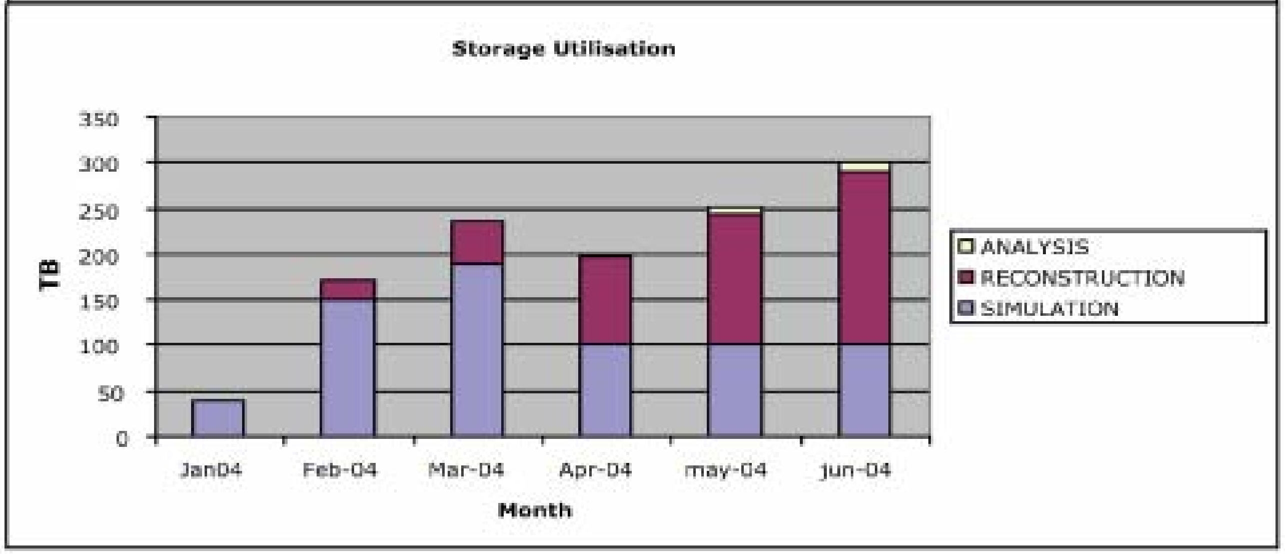

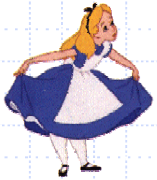

# PDC 3 resources

| | | O4Q1 | O4Q2 |
|---|---|---|---|
| **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 |

USA quota to be confirmed

Details in the “ALICE Data Challenges” paper taking into account

- Results of previous PDC
- Estimation of simulations in a standard year (2009)
- Storage: 200TB must be kept beyond the PDC end!!
- The numbers indicating the LCG resources for ALICE assume simultaneous use of the resources by all the experiments!
- A dynamic resource allocation would easily solve the deficit

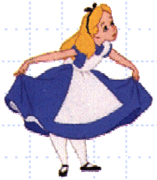

# LGC resources pledged for ALICE in 2003

| Sum LCG Values | | Resource | | | |
|---|---|---|---|---|---|
| Year | Country | CPU capacicty (kSI2K) | Disk Capacity (TB) | LCG Support (FTE) | Tape Capacity (TB) |
| 2004 | CERN | 175 | 40 | 3 | 250 |
| | Czech Rep. | 18 | 2 | 1 | 2 |
| | France | 120 | 23 | 3 | 154 |
| | Germany | 69 | 14 | 3 | 22 |
| | Holland | 41 | 1 | 1 | 4 |
| | Italy | 145 | 16 | 4 | 25 |
| | Poland | 0 | 0 | 0 | 0 |
| | Russia | 30 | 8 | 3 | 10 |
| | UK | 69 | 24 | 2 | 41 |
| | USA | 150 | 40 | 1 | 50 |
| | Sweden | 90 | 20 | 1 | 20 |
| | Hungary | 35 | 6 | 1 | 0 |
| Total 2004 | | 941 | 192 | 23 | 578 |

**USA quota to be confirmed**

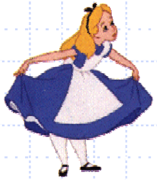

# ALICE PDC-3 & LCG

- 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 need to take any *decision*: the performance of the system will decide for us
- The figure of merit will be

$$Q^{LCG} = \frac{CPU^{LCG}_{good\ jobs}}{CPU^{LCG}_{available}}; \qquad Q^{AliEn} = \frac{CPU^{AliEn}_{good\ jobs}}{CPU^{AliEn}_{available}}$$

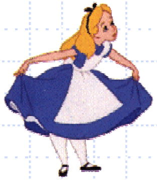

# Backup solutions

- If LCG *really* underperforms we would lack resources and the PDC 3 could be jeopardised
    - Evaluation of non-LCG resources is difficult and in any case these are not enough to run the PDC 3
- One possibility would be to *access* the LCG resources directly via AliEn
    - Provided the problem lies with EDG
    - Anyway AliEn is already installed on the system as part of the ALICE software
- Conversion of WN nodes is very simple
- Conversion of SE and CE may be more tricky
- Conversion back to LCG could require mirroring of MSS data

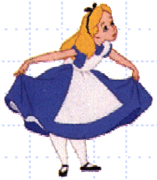

# Planning for PDC 3

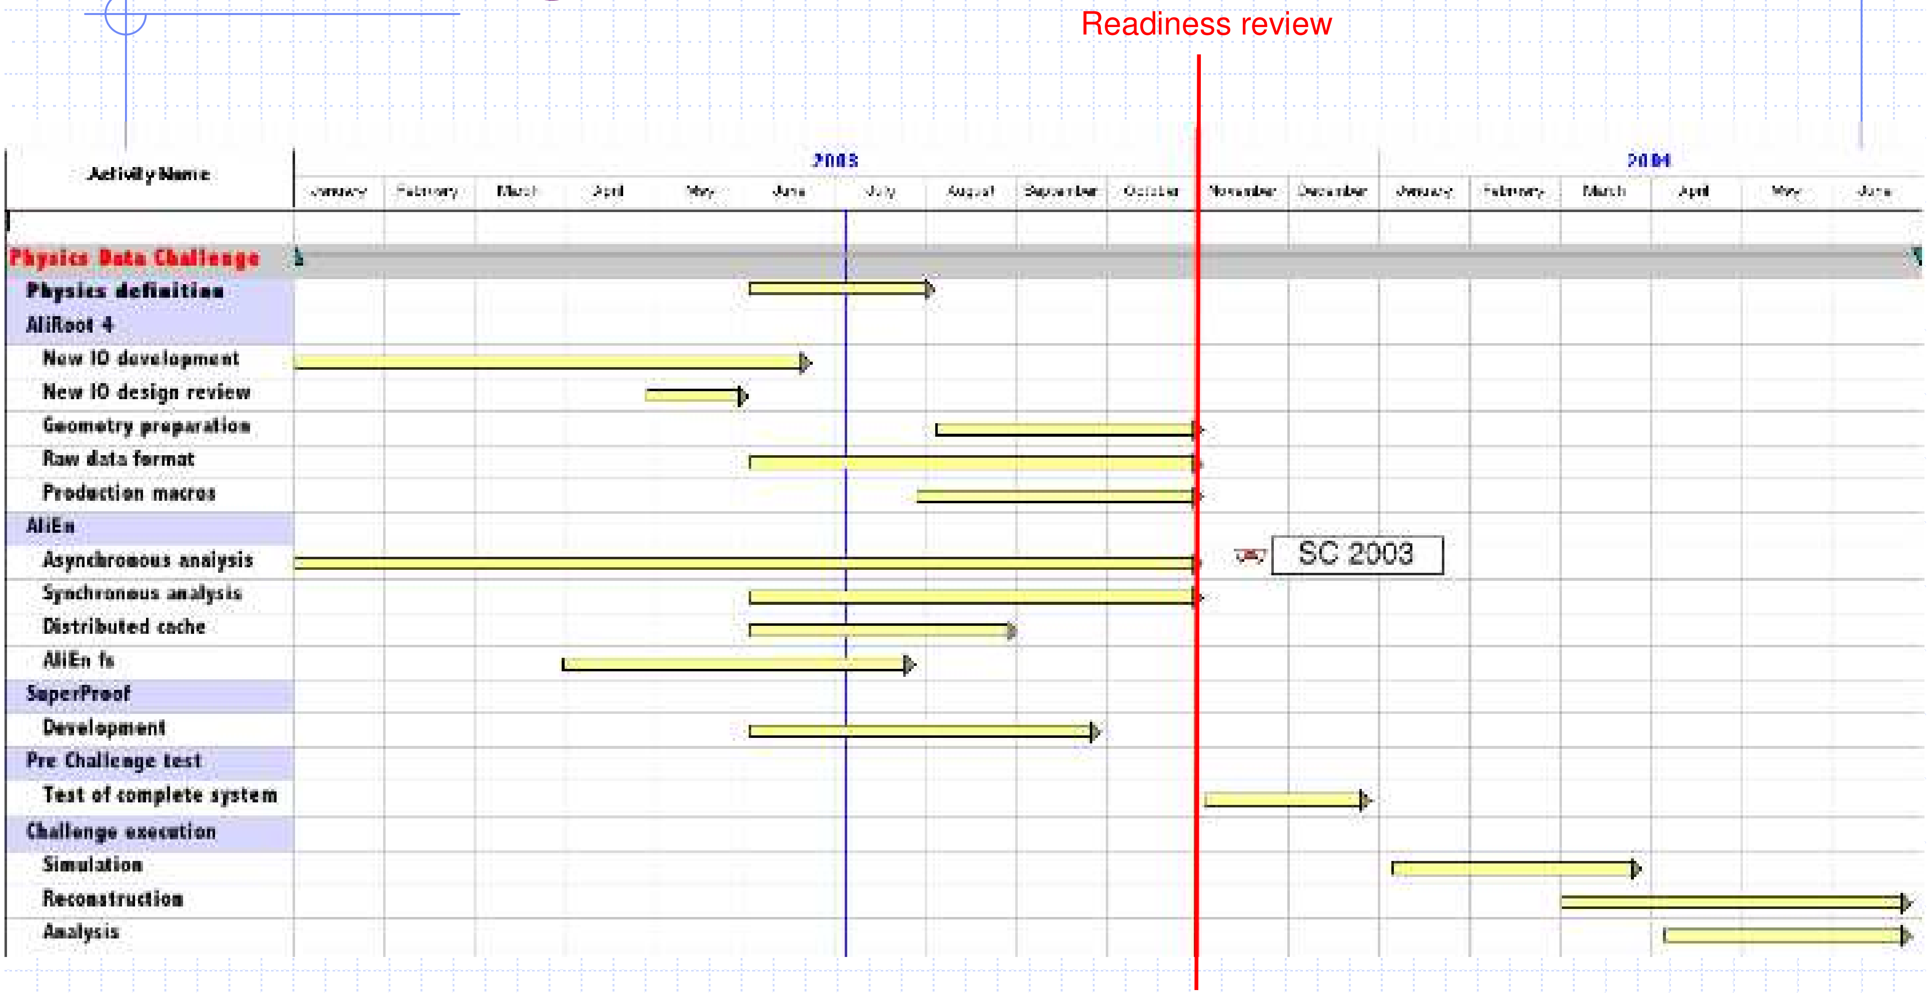

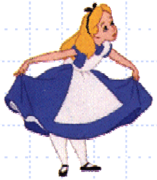

# Conclusion

- Preparation for PDC 3 is continuing on schedule
- PDC 3 will serve as a test of a complicated system with many variables
  - LCG-1, AliEn, AliRoot, Proof etc…
- We will certainly learn a lot… however we also need the results for ALICE physics
- So

wish us (all) good luck…