ALICE Computing TDR

Federico Carminati

June 29, 2005

Layout

- Parameters
- Computing framework
- Distributed computing and Grid
- Project Organisation and planning
- Resources needed
- http://aliceinfo.cern.ch/NewAlicePortal/en/Collaboration/Documents/TDR/Computing.html

- Parameters
- Computing framework
- Distributed computing and Grid
- Project Organisation and planning
- Resources needed

Parameters

	Unit	pp	PbPb
T1	#	7	
T2	#	23	
Size raw	MB	0.2x5	12.5
Recording rate	Hz	100	100
ESD	MB	0.04	2.50
AOD	kB	4	250
Event Catalogue	kB	10	10
Running time	s	10^7	10^6
Events / y	#	10^9	10^8
Reconstruction passes (av)	#	3	
RAW duplication	#	2	
AOD/ESD duplication	#	2	
Scheduled analysis passes / rec ev / y (av)	#	3	
Chaotic analysis passes / rec ev / y (av)	#	20	

- Parameters
- Computing framework
- Distributed computing and Grid
- Project Organisation and planning
- Resources needed

Offline framework

- AliRoot in development since 1998
- Two main packages to install (ROOT and AliRoot)
- Ported on Linux (IA32/64 & AMD), Mac OS X (ppc & Intel), Digital True64, SunOS…
- Over 50 developers (30% CERN), one CVS
- Integration with DAQ (data recorder) and HLT (CVS)
- Abstract interfaces for modularity
- Subset of C++ used for maximum portability
- Very close integration of physicists and programmers

AliRoot layout

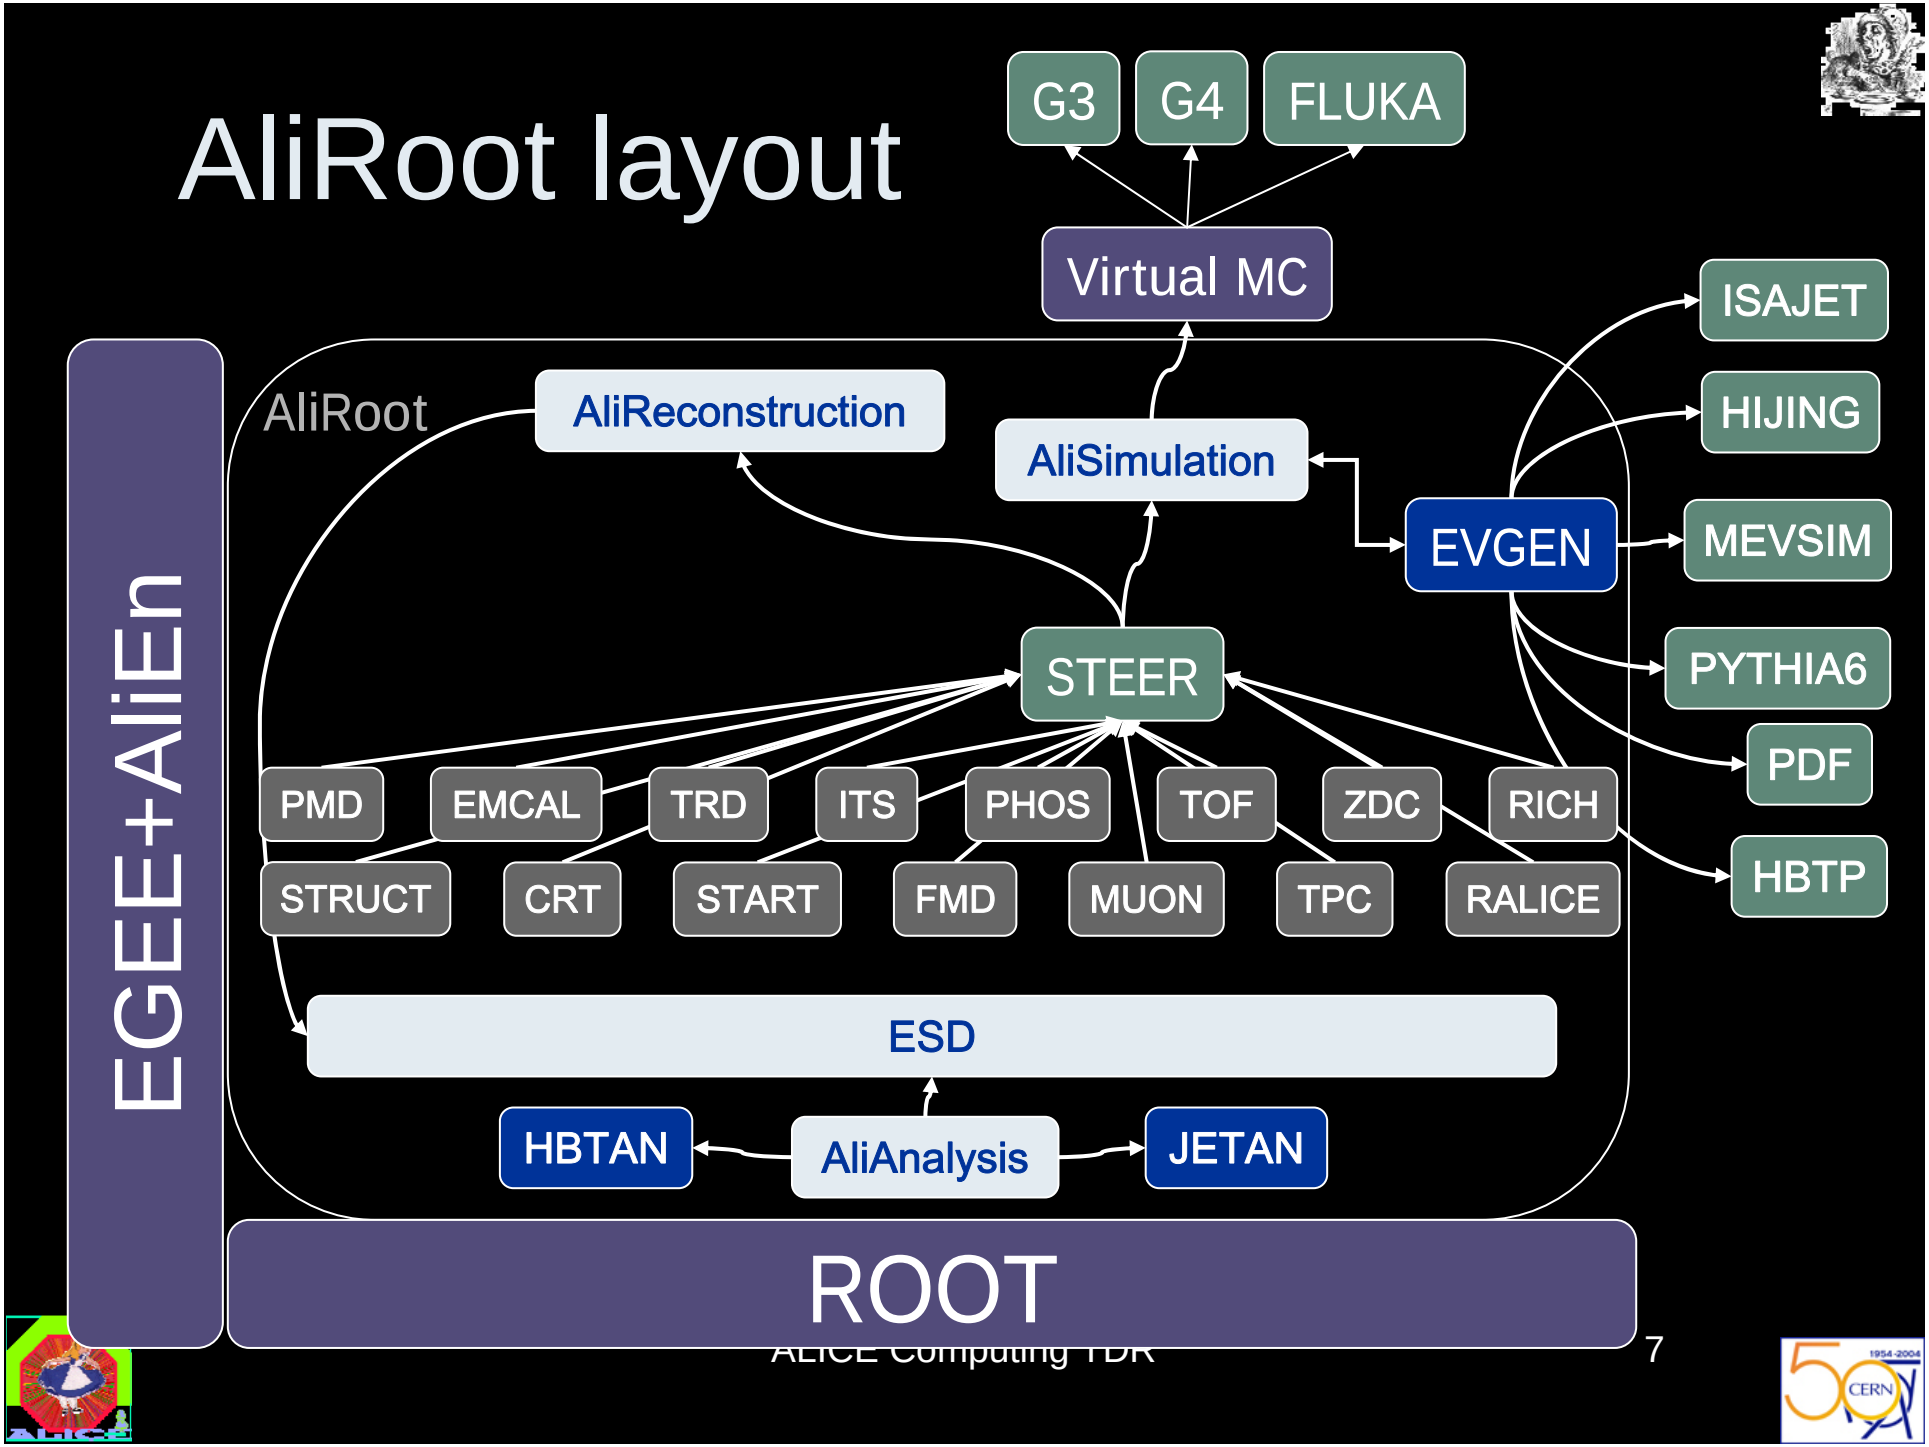

Software management

- Major release ~ six months, minor (tag) ~ monthly
- Emphasis on delivering production code
- Nightly produced UML diagrams code listing coding rule violations build and tests
  - No version management software
- Development of new coding tools (IRST)
  - Smell detection, aspect oriented programming, genetic testing

Detector Construction Database (DCDB)

- Detector construction in distributed environment:
  - Sub-detector groups work independently
  - Data collected in a central repository to facilitate movement of components between groups and during integration and operation

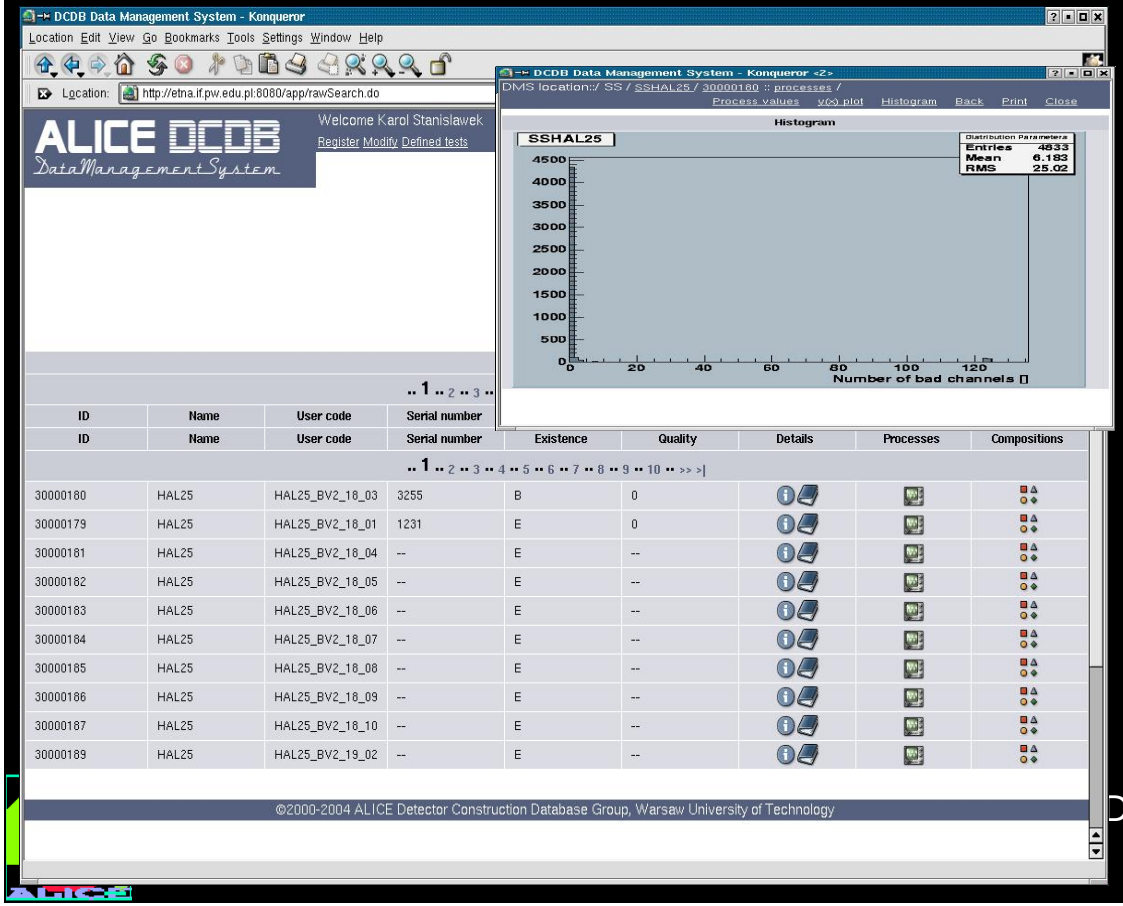

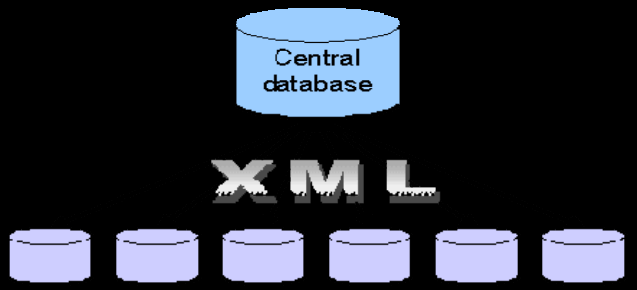

- Different user interfaces
  - WEB portal
  - LabView, XML
  - ROOT for visualisation
- In production since 2002
- Important spin-offs
  - Cable Database
  - Calibration Database

Simulation

- Simulation performed with G3 till now
- Will move to FLUKA and G4
- VMC insulates users from the transport MC
- New geometrical modeller in production
- Physics Data Challenge ‘05 with FLUKA
- Interface with G4 designed (4Q’05?)
- Test-beam validation activity ongoing

The Virtual MC

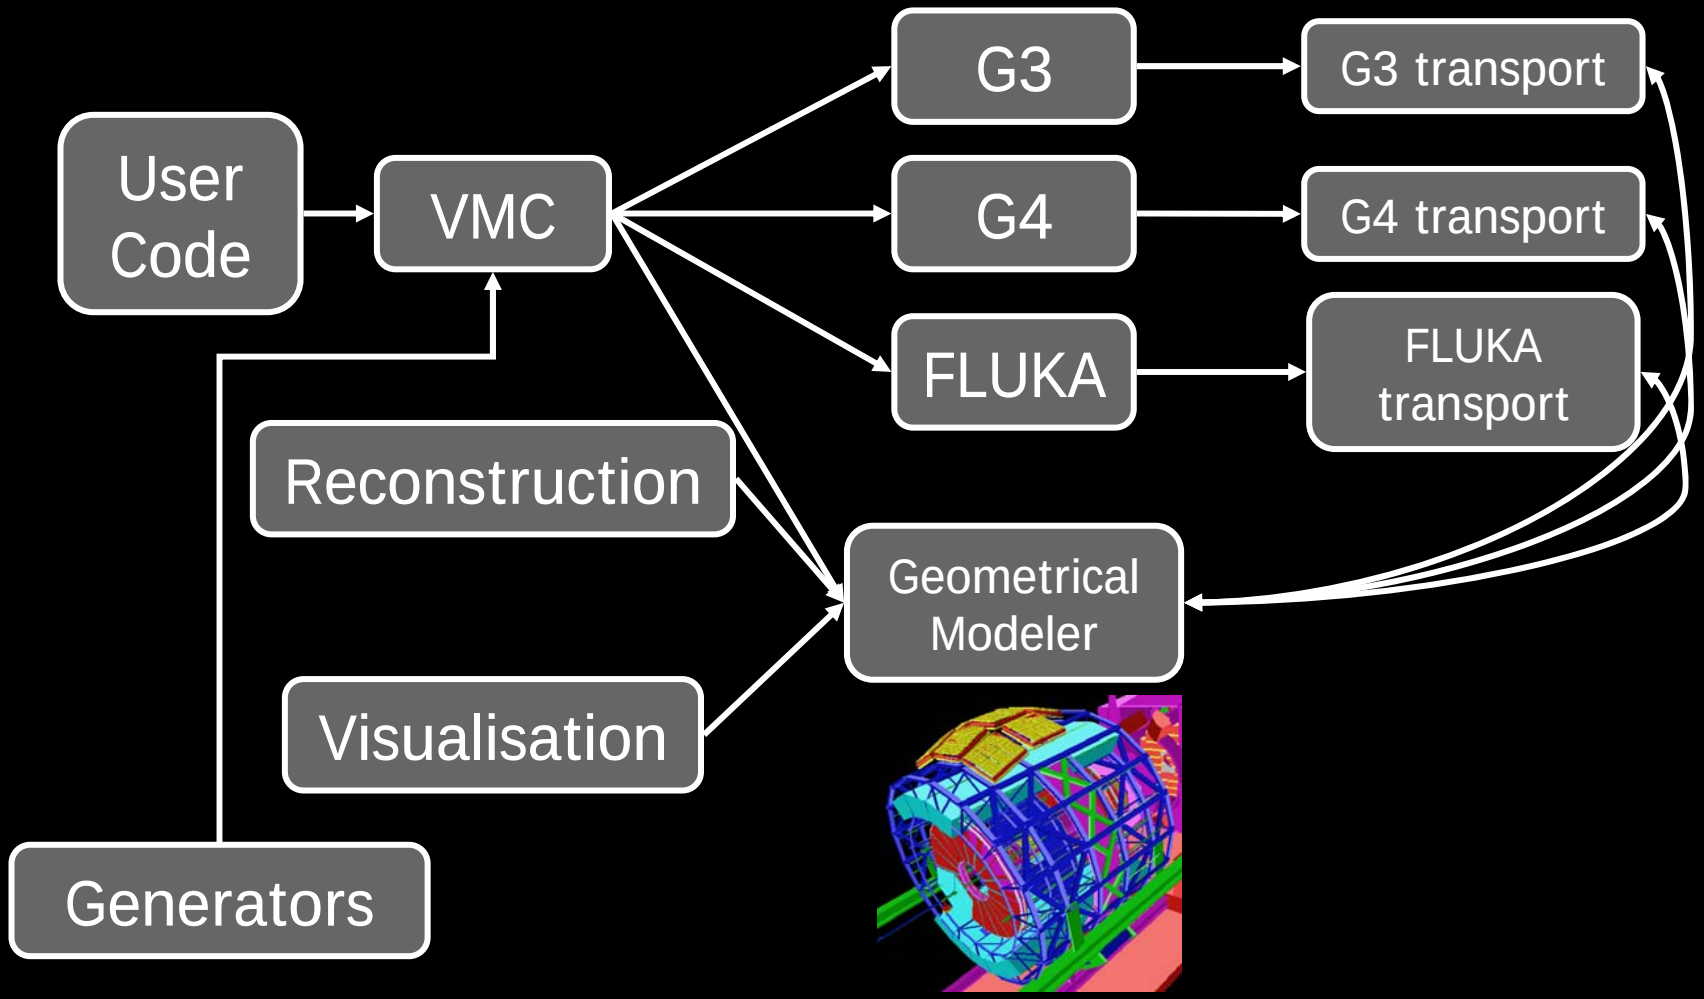

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

ime™ and a
sed) decompressor

are needed to see this pic

me™ and a
sed) decompressor
are needed to see this picture.

HMPID
5 GeV Pions

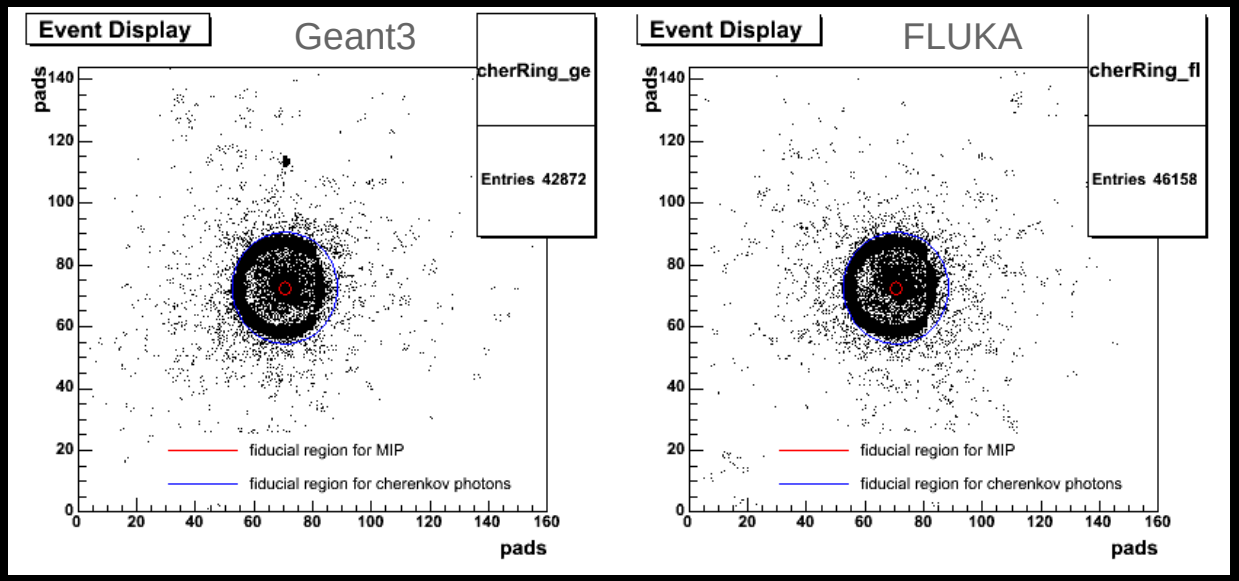

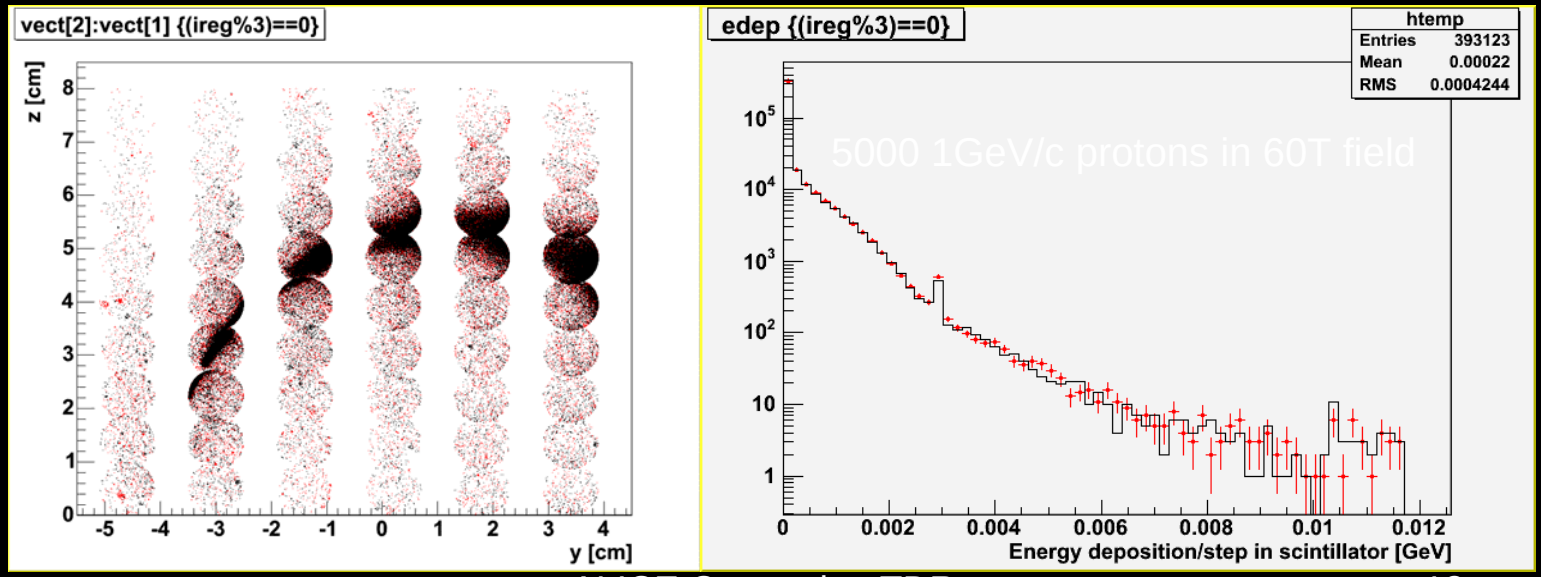

Reconstruction strategy

- Very high flux, TPC occupancy up to 40%
- Maximum information approach
- Optimization of access and use of information
  - Localize relevant information
  - Keep this information until it is needed

Tracking & PID

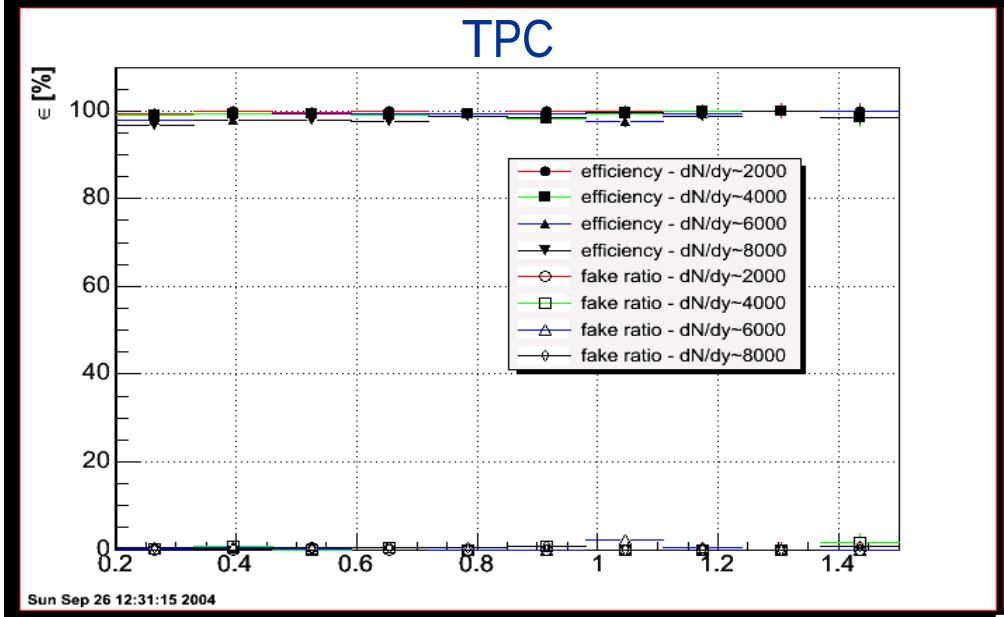

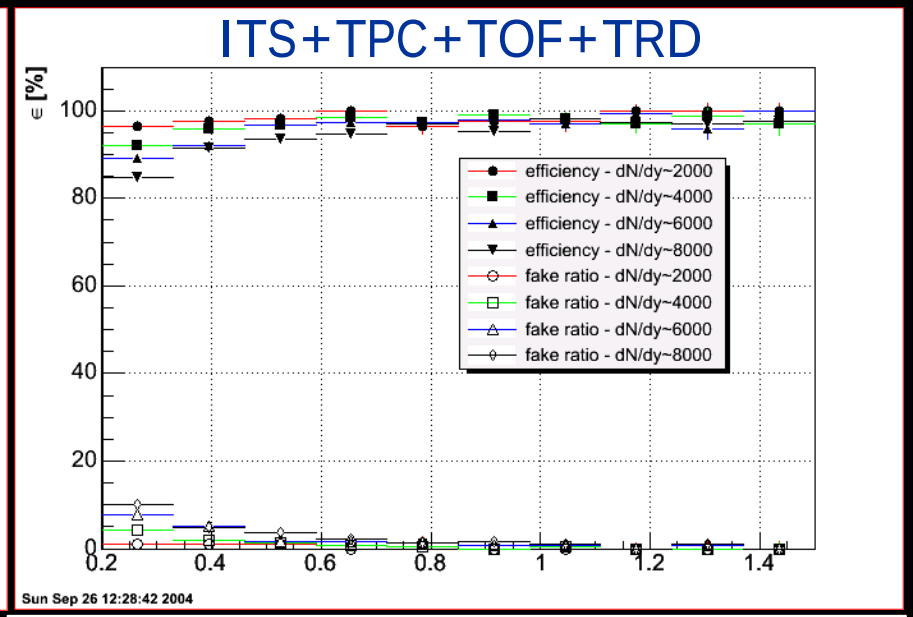

- PIV 3GHz – (dN/dy – 6000)
  - TPC tracking - ~ 40s
  - TPC kink finder ~ 10 s
  - ITS tracking ~ 40 s
  - TRD tracking ~ 200 s

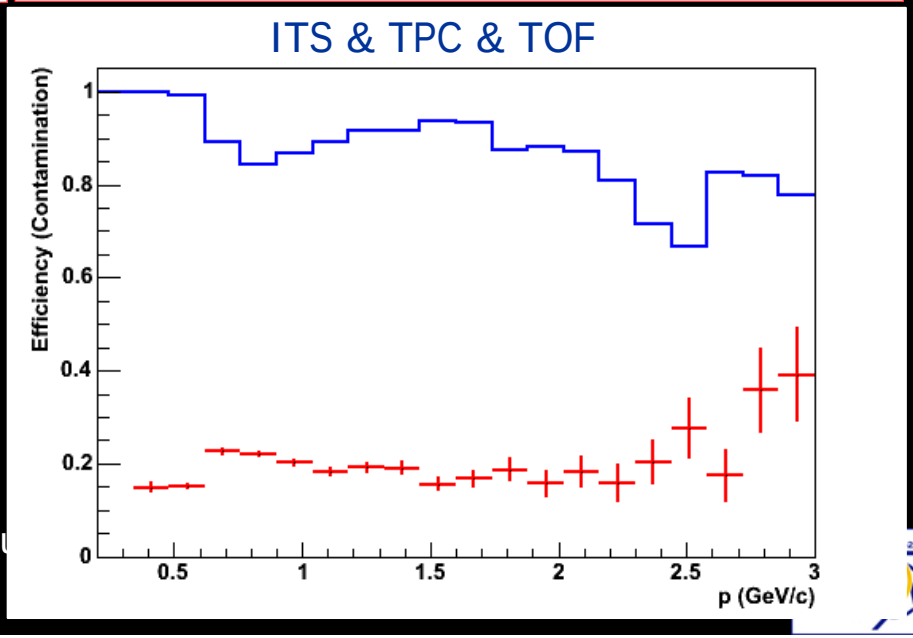

Condition and alignment

- Heterogeneous sources periodically polled
- ROOT files with condition information created
- Published on the Grid and distributed as needed by the Grid DMS
- Files contain validity information and are identified via DMS metadata
- No need for a distributed DBMS
- Reuse of the existing Grid services

External relations and DB connectivity

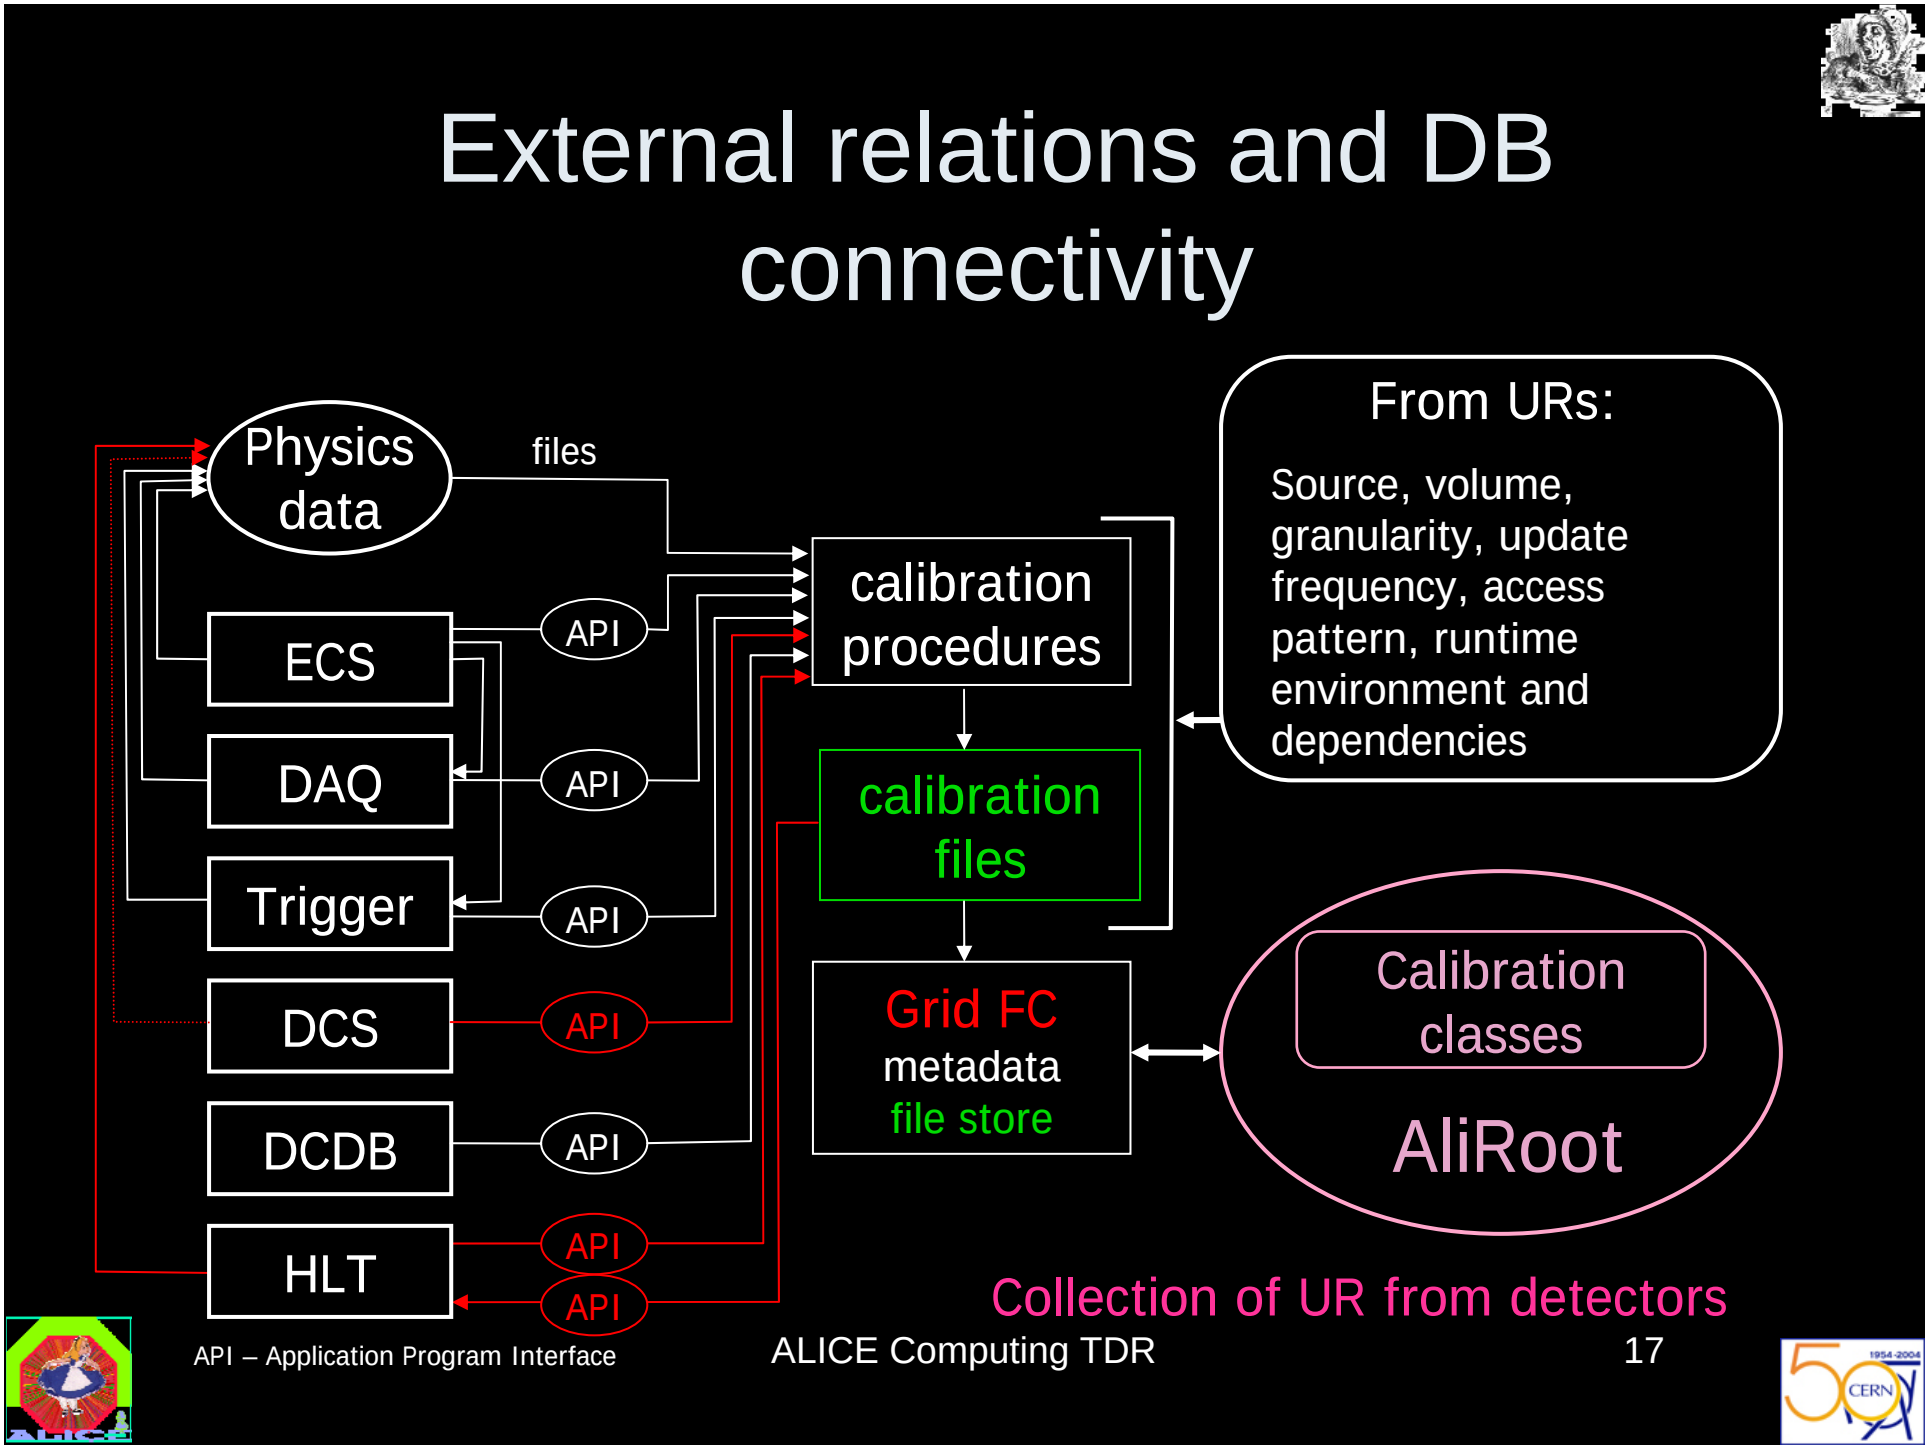

API – Application Program Interface

Metadata

- Essential for the selection of events
- Grid file catalogue for file-level MD
- Need event-level MetaData
- Collaboration with STAR
- Prototype in preparation for PDC’05-III

ALICE CDC's

Date	MBytes/s	Tbytes to MSS	Offline milestone
10/2002		200	C and ITS
9/2003			
5/2004–4/2005			w digits for
5/2005	750 (or higher)	750 (or higher)	(Raw digits for all detectors)
5/2006	750 (or higher)	750 (or higher)	Final test (Final system)

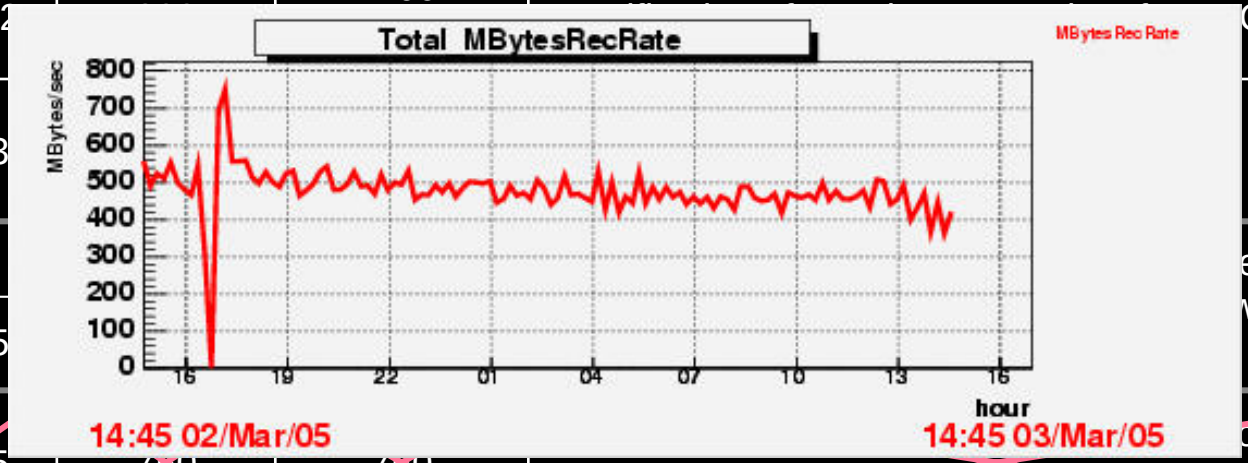

- Parameters
- Computing framework
- Distributed computing and Grid
- Project Organisation and planning
- Resources needed

ALICE Physics 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
06/05-12/05	15%	Test of condition infrastructure and FLUKA To be combined with SDC 3 Speed test of distributing data from CERN
01/06-06/06?	20%	Test of the final system for reconstruction and analysis

PDC04 schema

Production of RAW

Shipment of RAW to CERN

Reconstruction of RAW in all T1's

Analysis

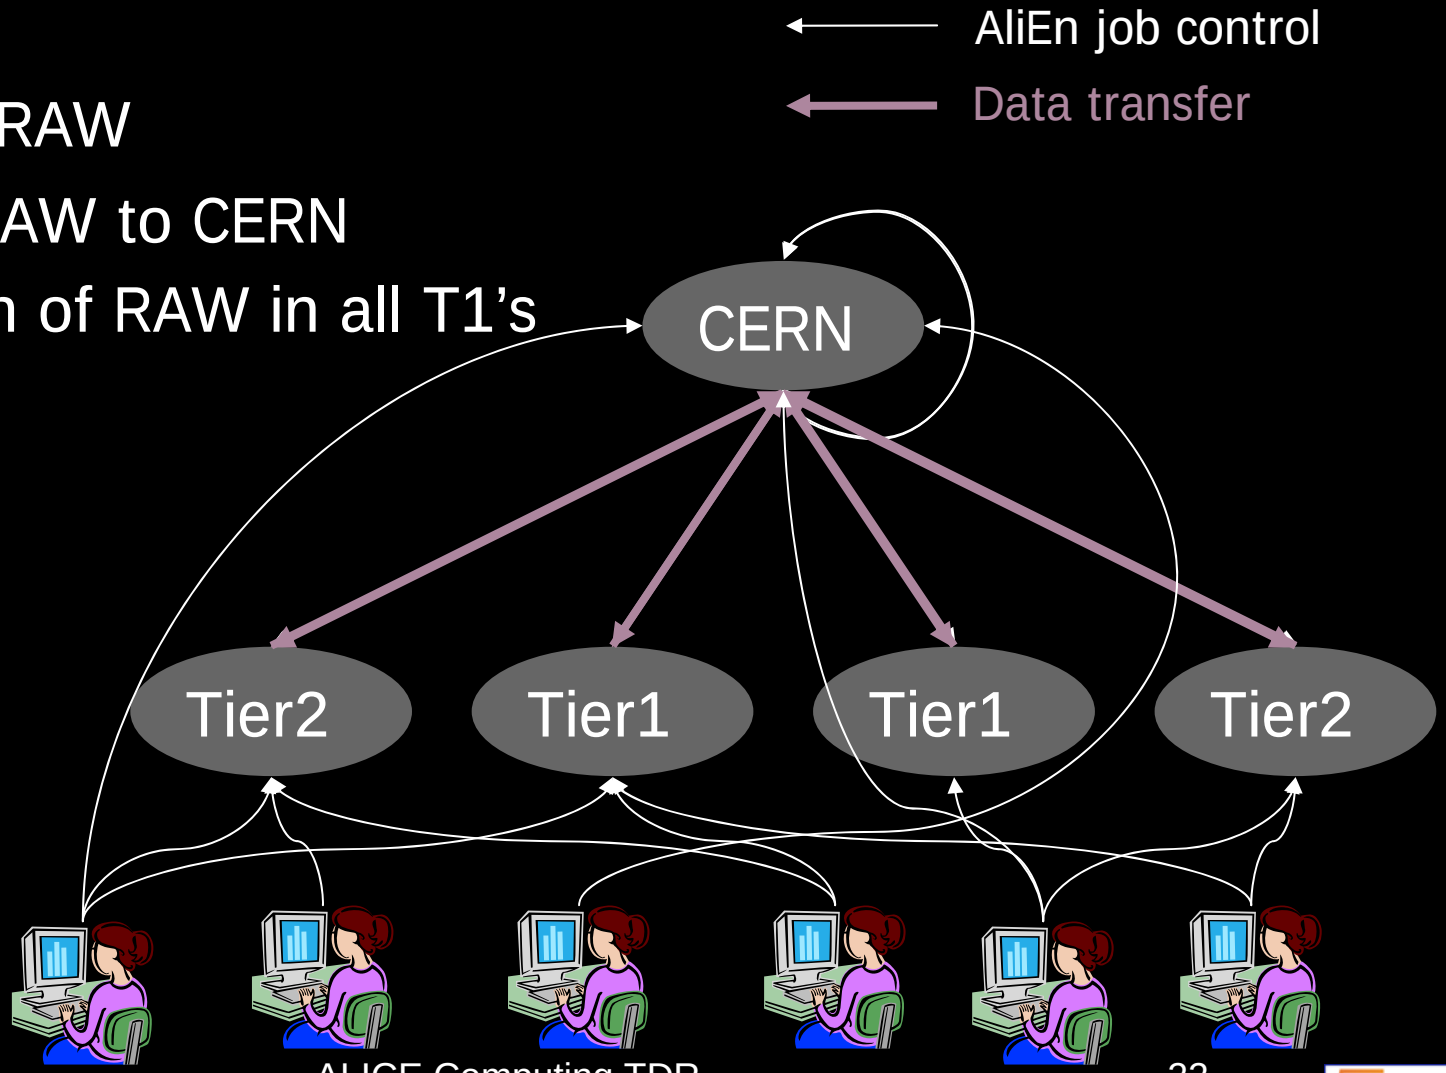

Goals, structure and tasks

- Validate the computing model with ~10% of the SDTY data
- Use the offline chain, the Grid, PROOF and the ALICE ARDA prototype

1. Production of underlying Pb+Pb and p+p events
   - Completed June 2004
2. Mixing of different signal events with underlying Pb+Pb events (up to 50 times)
   - Completed September 2004
3. Distributed analysis
   - Delayed

Phase 2 job structure

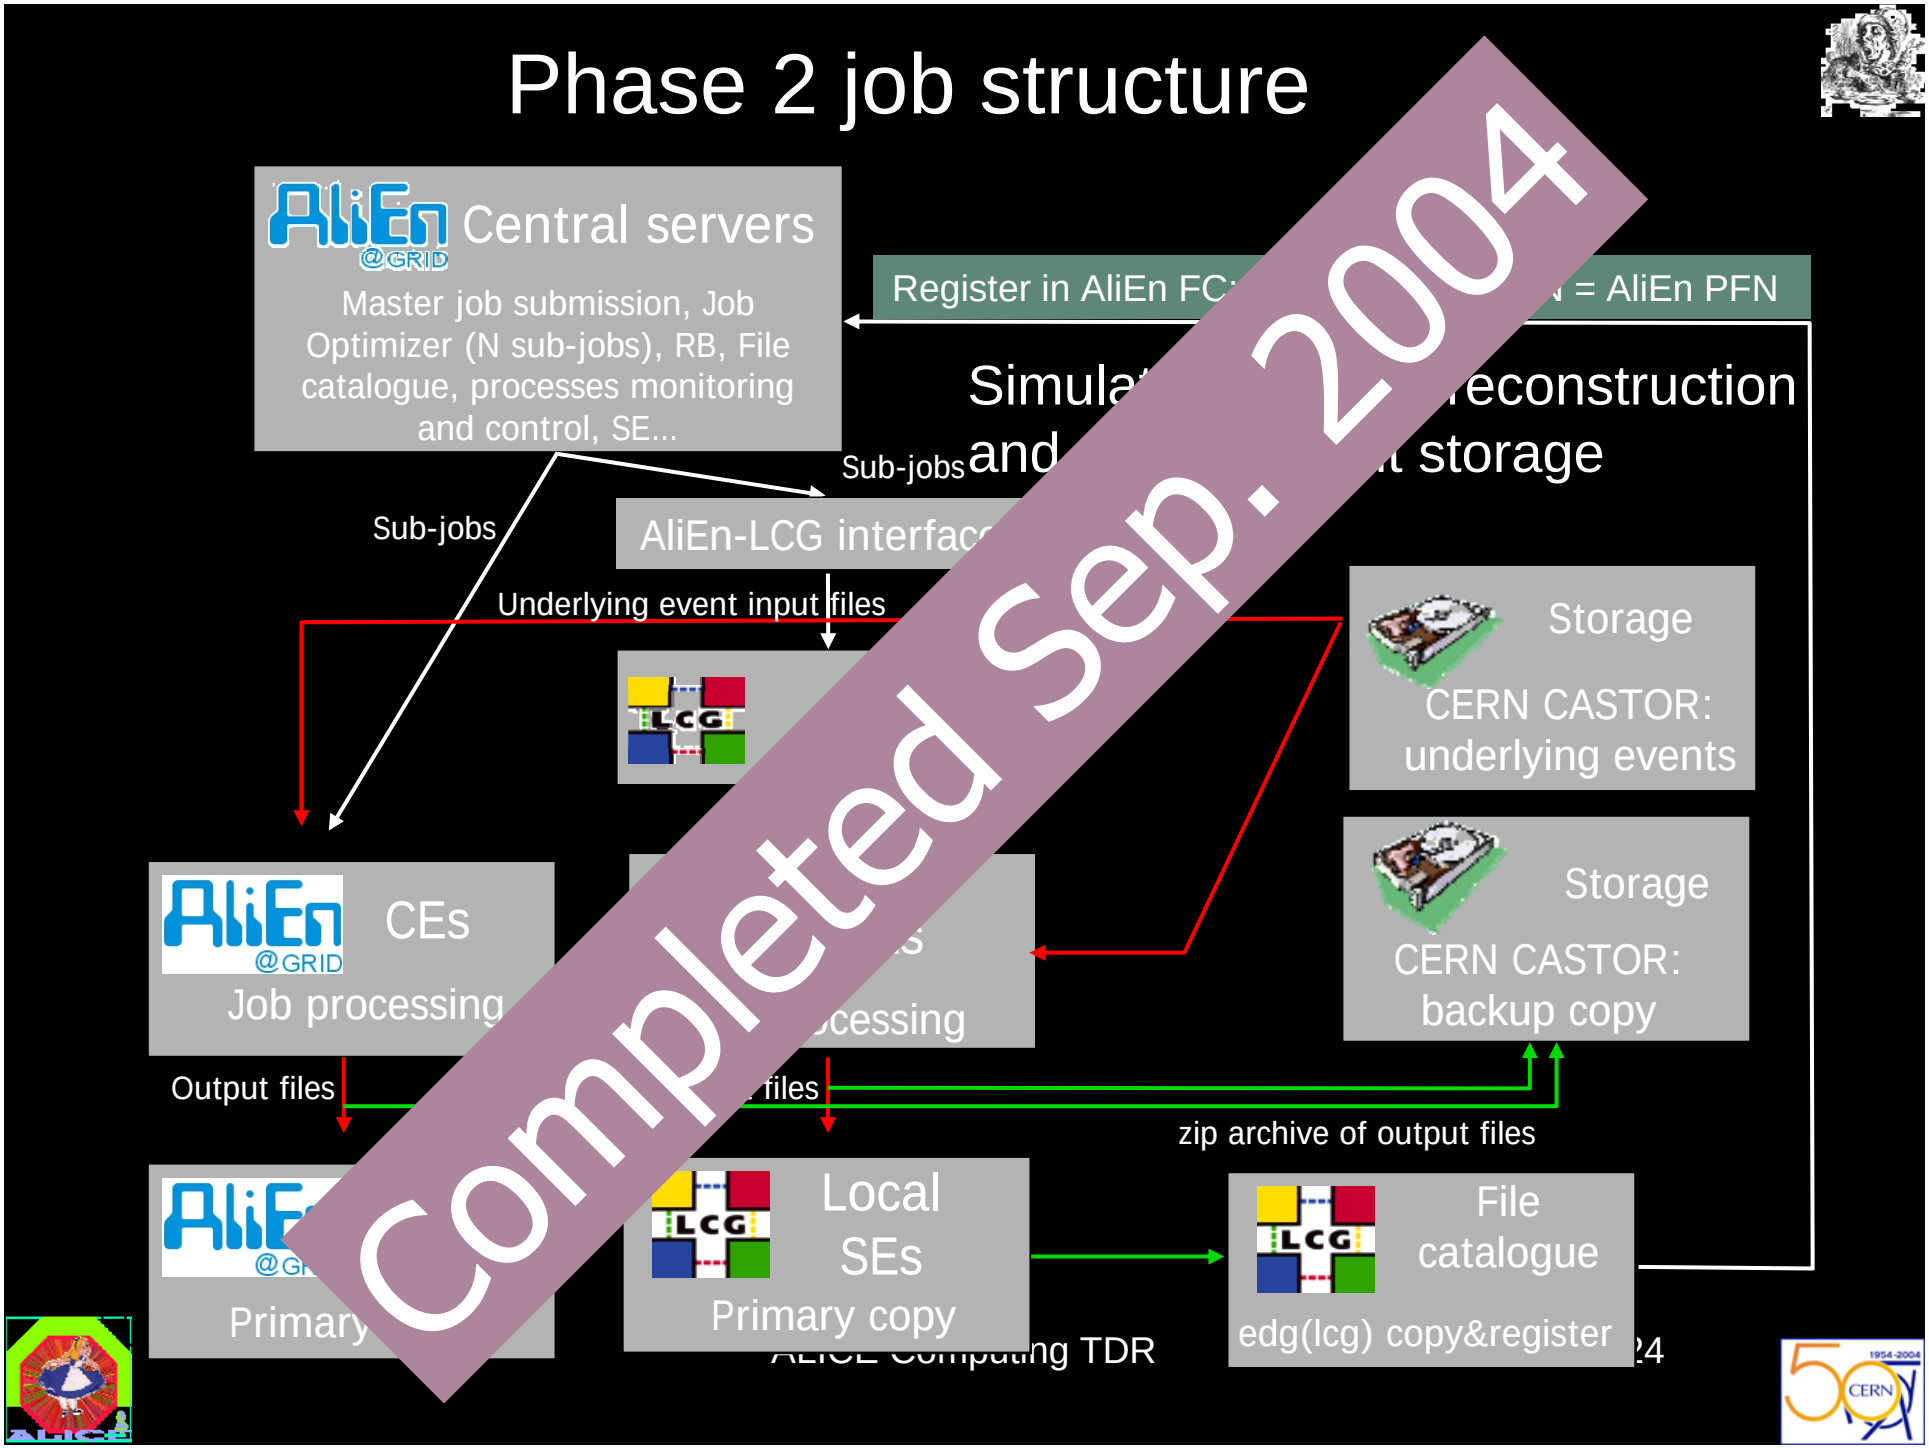

Summary

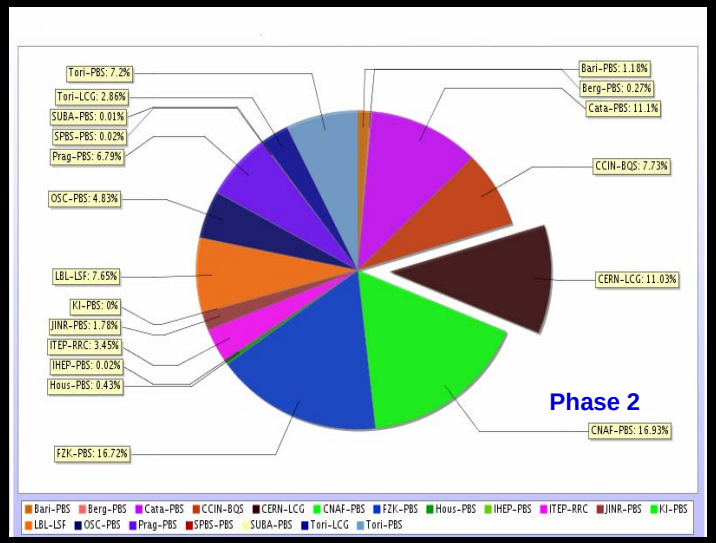

33 sites, 3 grids
AliEn/LCG: P1 75/25%, P2 89/11%

400 000 jobs, 6 hours/job, 750 MSi2K hours
9M entries in the AliEn file catalogue
4M physical files at 20 AliEn SEs world-wide
30 TB@CERN CASTOR
10 TB@remote SEs + 10 TB backup@CERN
200 TB network transfer CERN –> remote centres

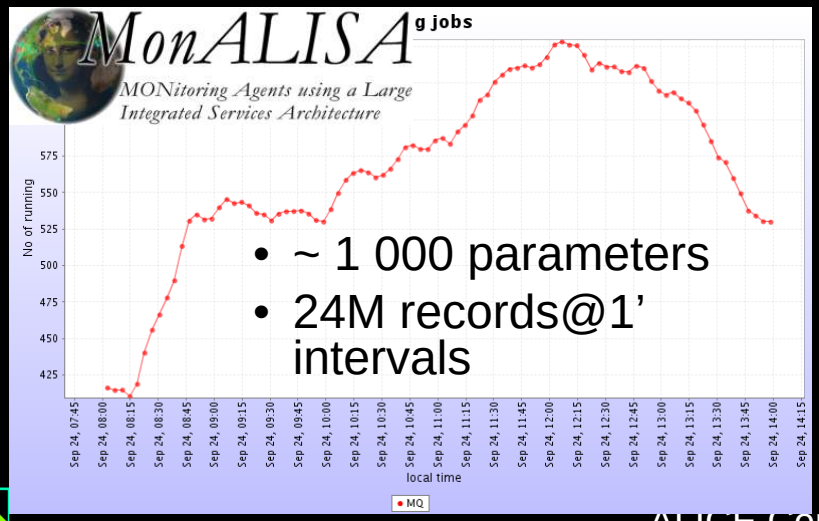

Summary of PDC'04

- Computer Centres
  - Tuning of environment
  - Availability of resources
- Middleware
  - Phase 1&2 successful
  - AliEn fully functional
  - LCG not yet ready
  - No AliEn development for phase 3, LCG not ready
- Computing model validation
  - AliRoot worked well
  - Data Analysis partially tested on local CN, distributed analysis prototype demonstrated

Development of Analysis

- ROOT & a small library on AOD's and ESD's
- Work on distributed analysis from ARDA
- Batch prototype tested at the end 2004
- Interactive prototype demonstrated end 2004
- Physics Working Groups providing requirements
- Planning for fast analysis facility at CERN

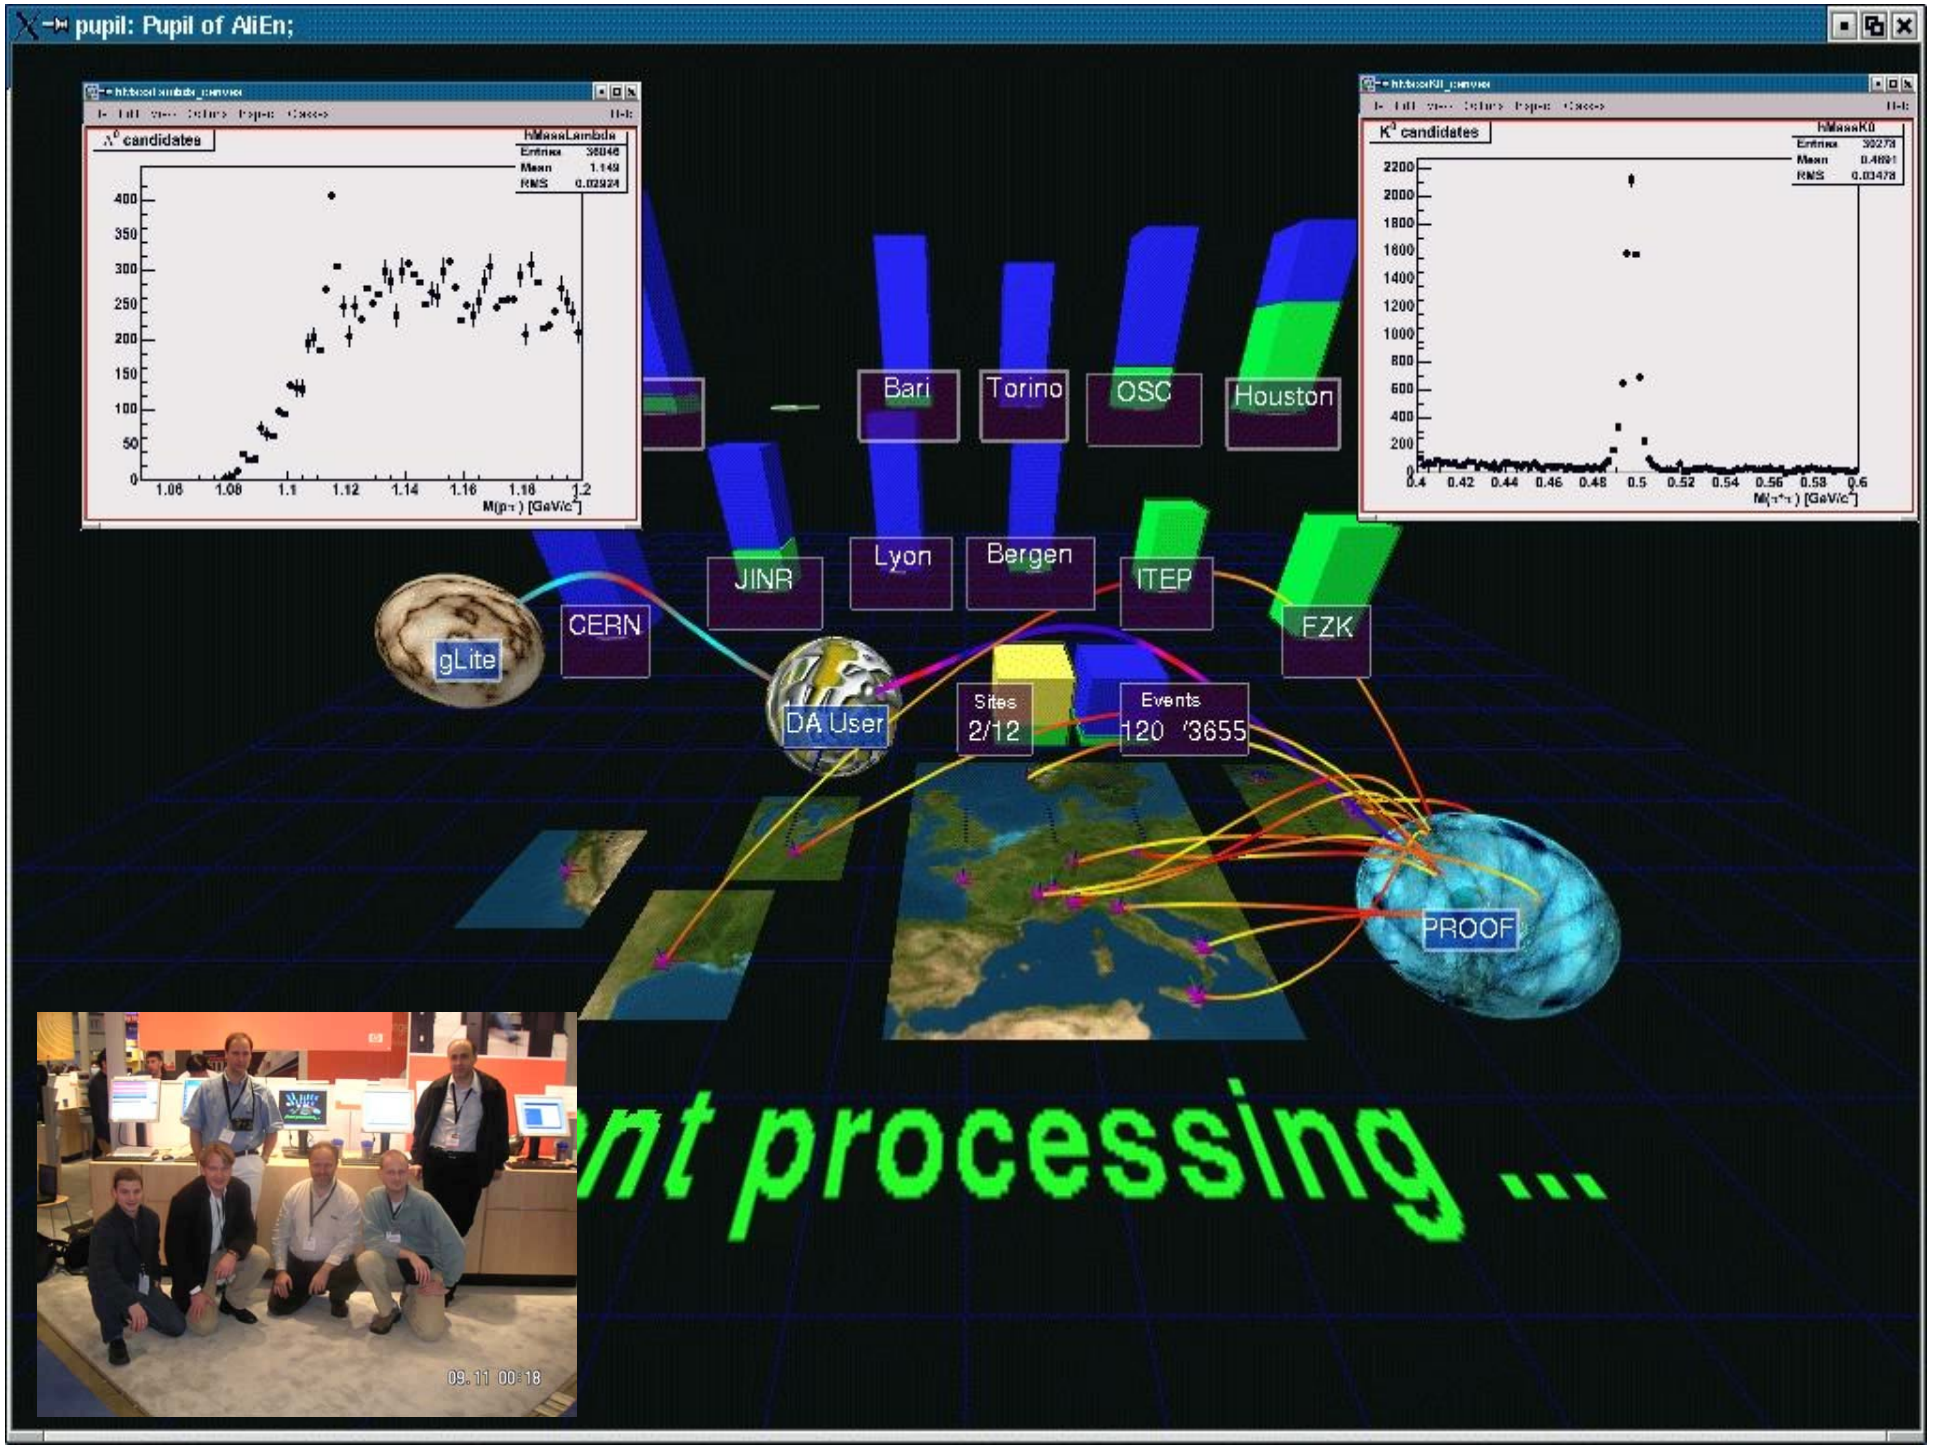
pupil: Pupil of AliEn;
Bari
Torino
OSC
Houston
JINR
Lyon
Bergen
ITEP
FZK
CERN
gLite
DA User
Sites
2/12
Events
120 '3655
PROOF
nt processing ...

ALICE computing model

- pp
  - Quasi-online data distribution and first reconstruction at T0
  - Further reconstructions at T1's
- AA
  - Calibration, alignment and pilot reconstructions during data taking
  - Data distribution and first reconstruction at T0 during four months after AA
  - Further reconstructions at T1's
- One copy of RAW at T0 and one distributed at T1's

ALICE computing model

- T0
  - First pass reconstruction, storage of one copy of RAW, calibration data and first-pass ESD's
- T1
  - Reconstructions and scheduled analysis, storage of the second collective copy of RAW and one copy of all data to be kept, disk replicas of ESD's and AOD's
- T2
  - Simulation and end-user analysis, disk replicas of ESD's and AOD's
- Difficult to estimate network load

ALICE T1's & T2's

- T1 for ALICE: CERN, CCIN2P3, CNAF, GridKa, NIKHEF, NGDF, RAL, USA (under discussion)
- T2 for ALICE: CERN, CCIN2P3, Nantes, Clermont-Ferrand, Paris, Bari, Catania, Legnaro, Torino, GSI, RUSSIA, Prague, Korea, Kolkata, Wuhan, Cape Town, USA, UK Grid, Athens

ALICE MW requirements

- Baseline Services available on LCG (in three flavours?)
- An agreed standard procedure to deploy and operate VO-specific services
- The tests of the integration of the components have started

- Parameters
- Computing framework
- Distributed computing and Grid
- Project Organisation and planning
- Resources needed

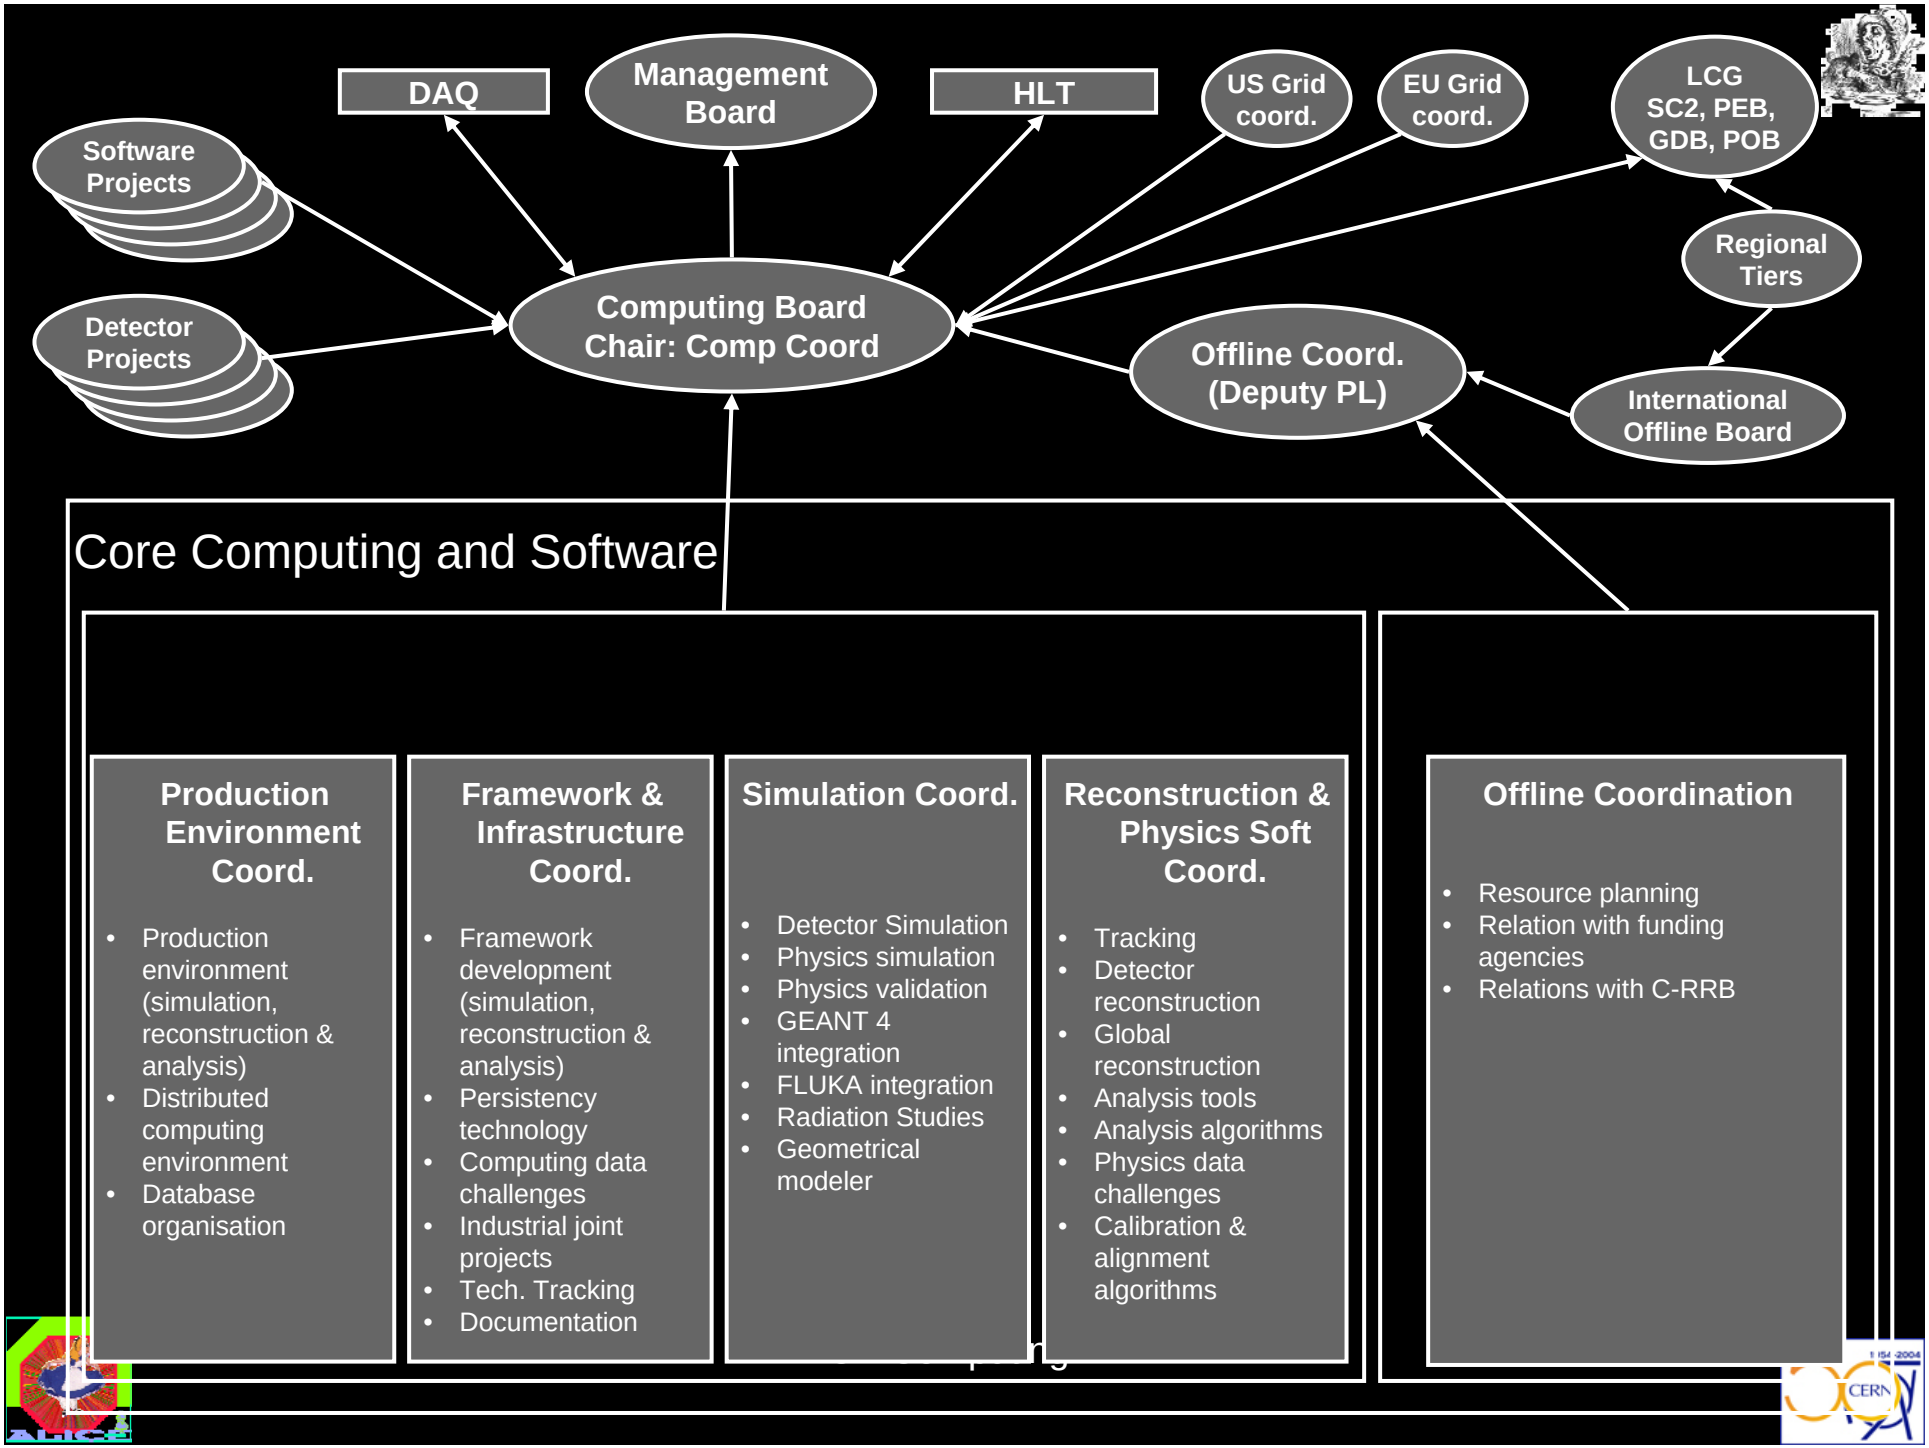
DAQ
Management Board
HLT
US Grid coord.
EU Grid coord.
LCG SC2, PEB, GDB, POB
Software Projects
Detector Projects
Computing Board Chair: Comp Coord
Regional Tiers
Offline Coord. (Deputy PL)
International Offline Board
Core Computing and Software
Production Environment Coord.
• Production environment (simulation, reconstruction & analysis)
• Distributed computing environment
• Database organisation
Framework & Infrastructure Coord.
• Framework development (simulation, reconstruction & analysis)
• Persistency technology
• Computing data challenges
• Industrial joint projects
• Tech. Tracking
• Documentation
Simulation Coord.
• Detector Simulation
• Physics simulation
• Physics validation
• GEANT 4 integration
• FLUKA integration
• Radiation Studies
• Geometrical modeler
Reconstruction & Physics Soft Coord.
• Tracking
• Detector reconstruction
• Global reconstruction
• Analysis tools
• Analysis algorithms
• Physics data challenges
• Calibration & alignment algorithms
Offline Coordination
• Resource planning
• Relation with funding agencies
• Relations with C-RRB

Core Computing Staffing

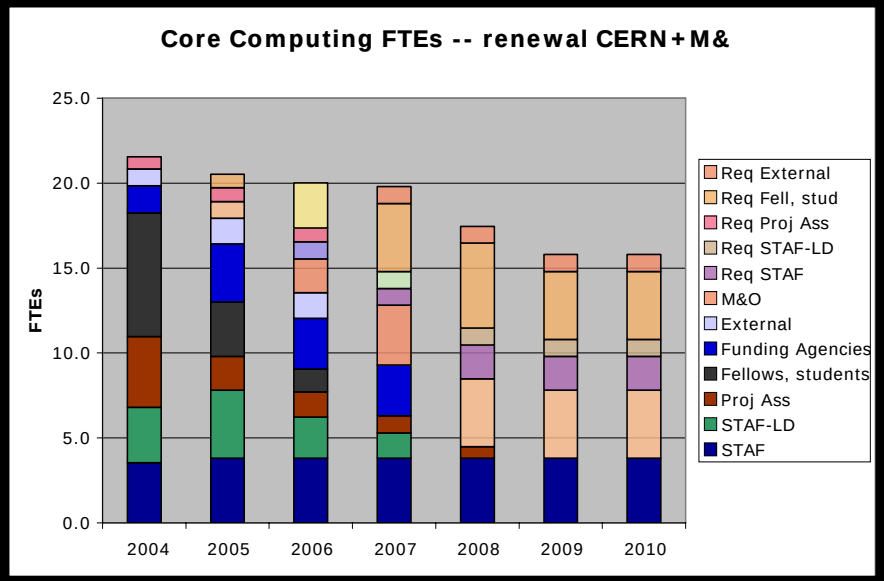

- The call to the collaboration has been successful
- Four people have been sent to CERN to help with CORE Computing
- Securing few long-term position is still a very high priority

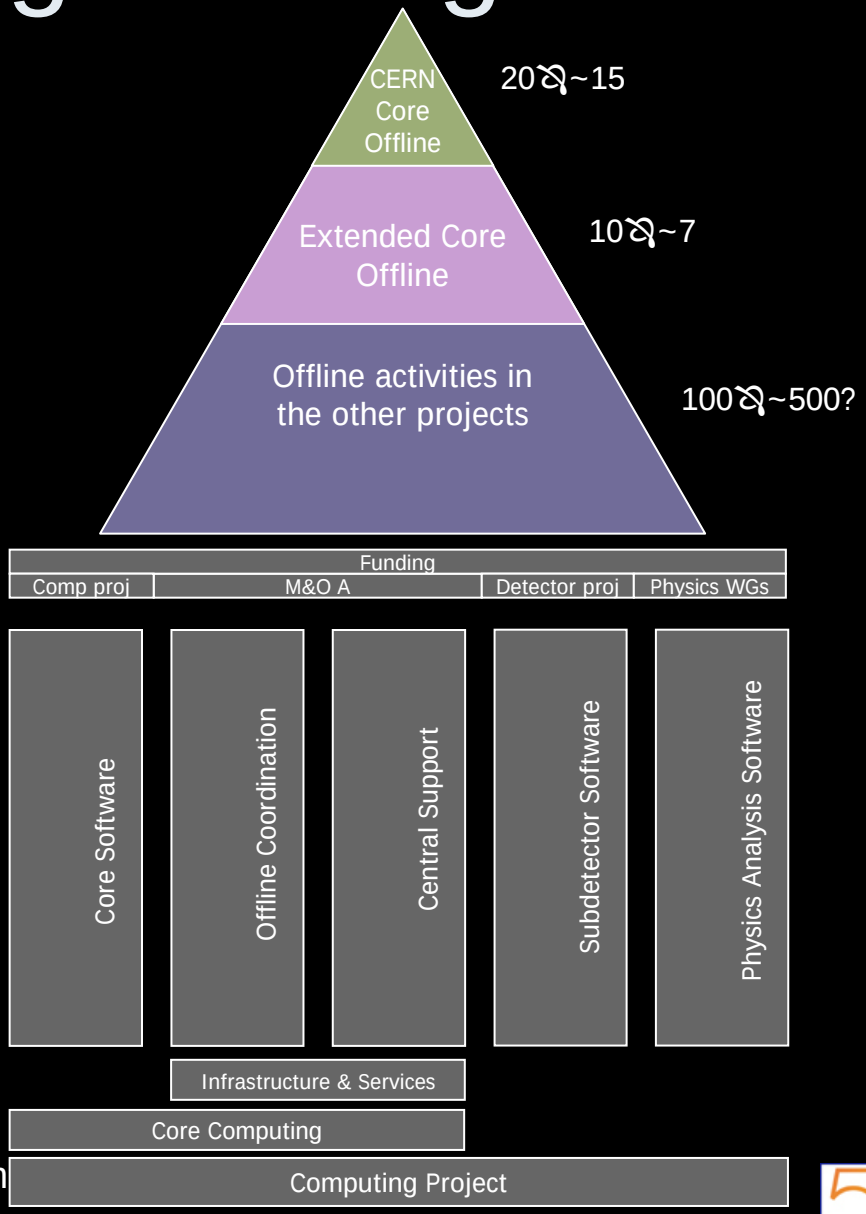

Milestones

- Jun 05: prototype of the condition infrastructure
- Jun 05: FLUKA MC in production with new geometrical modeller
- Jun 05: Computing TDR completed
- Aug 05: PDC05 simulation with Service Data Challenge 3
- Sep 05: MetaData prototype infrastructure ready
- Nov 05: Analysis of PDC05 data
- Dec 05: condition infrastructure deployed
- Dec 05: alignment and calibration algorithms for all detectors
- Jun 06: PDC 06 successfully executed (depends on SDC4)
- Jun 06: alignment and calibration final algorithms
- Dec 06: AliRoot ready for data taking

- Parameters
- Computing framework
- Distributed computing and Grid
- Project Organisation and planning
- Resources needed

Summary of needs

	Tier0	Tier1	Tier1ex	Tier2	Tier2ex	Total	CERN
CPU (MSI2k)	8,3	18,7	12,3	21,4	14,4	35,0	8,3
			35%		41%	100%	24%
Disk (PB)	0,2	8,6	7,4	5,3	5,1	14,1	1,7
	2%	60%	52%	38%	36%	100%	12%
MS (PB/y)	2,5	8,1	6,9			10,6	3,6
	23%	77%	66%			100%	34%
Network in (Gb/s)	8,00	2,00		0,01			
Network out (Gb/s)	6,00	1,50		0,27			

Capacity profile

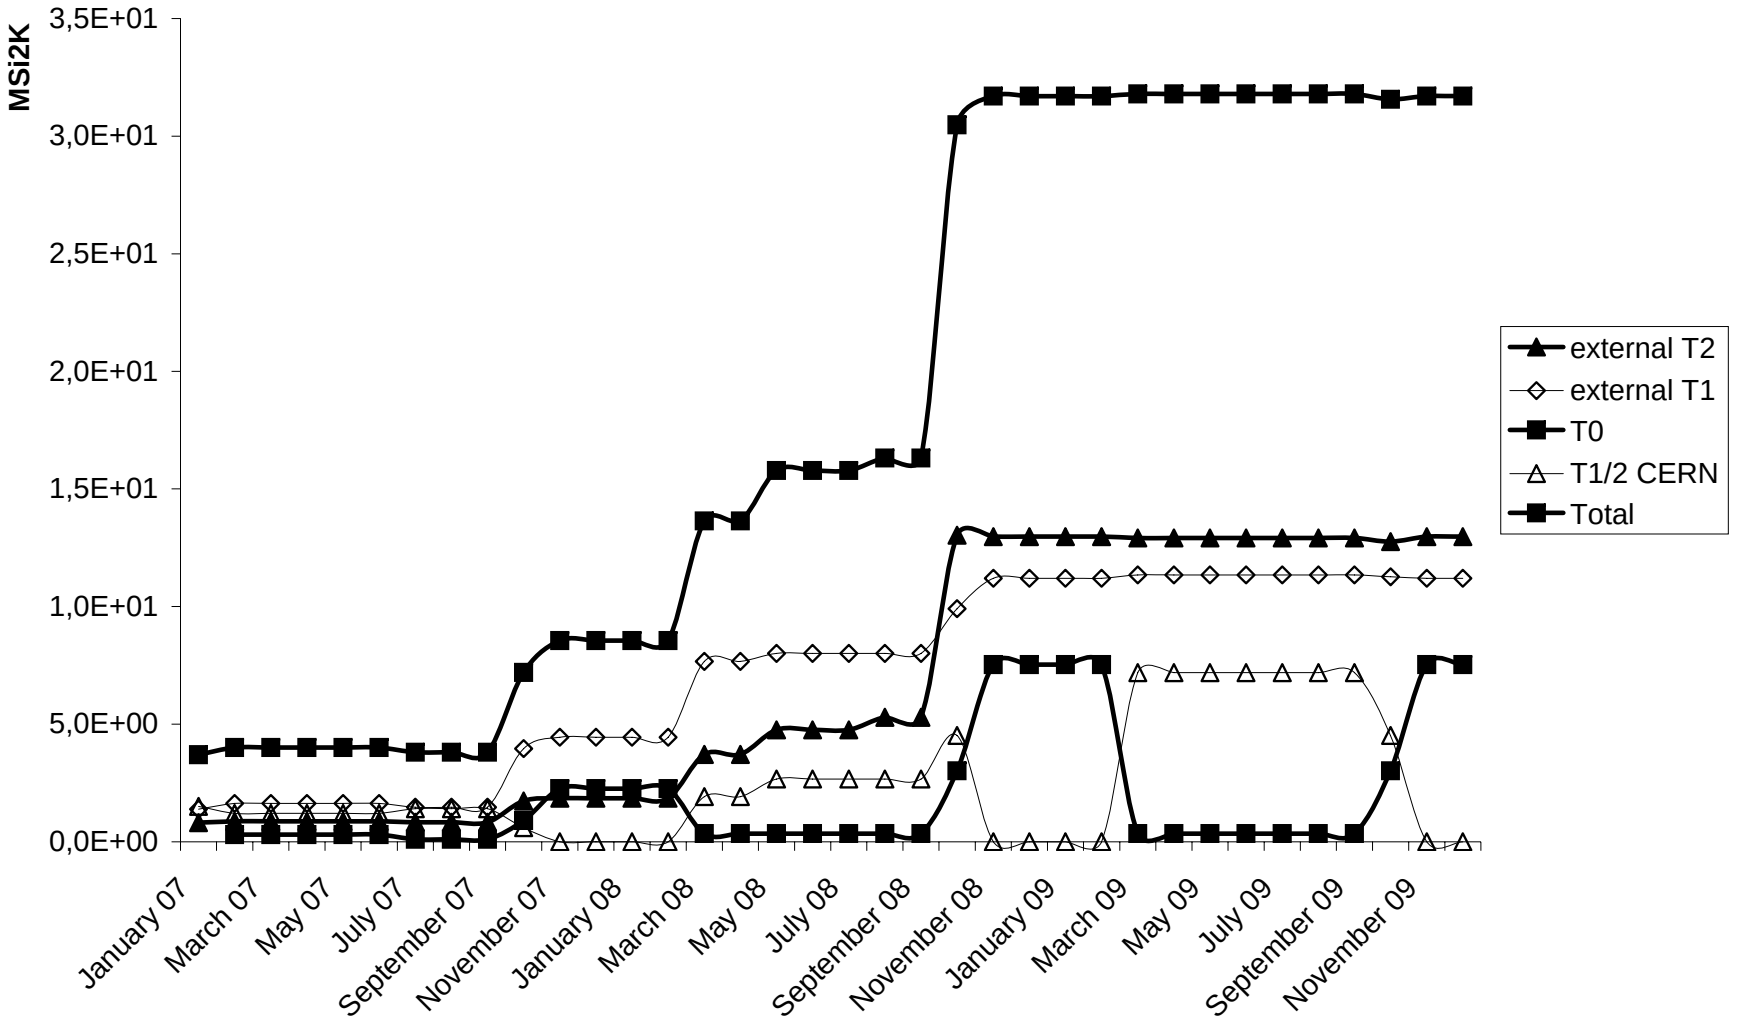

Summary pledged resources

	2005	2006	2007	2008	2009	2010
Tier1						
CPU (MSI2k)	0,48	1,03	2,91	8,94	14,88	14,81
Disk (PB)	0,09	0,50	1,53	3,48	5,70	5,88
MS (PB)	0,11	0,85	2,53	5,86	9,93	8,59
Tier2						
CPU (MSI2k)	1,82	2,92	4,81	6,18	8,34	9,01
Disk (PB)	0,28	0,61	1,07	1,68	2,58	3,41
Pledged versus required						
CPU (%)			61%	48%	56%	44%
Disk (%)			47%	37%	46%	40%
MS (%)			78%	72%	94%	63%

Conclusions

- ALICE computing choices have been validated by experience
  - The Offline development is on schedule, although contingency is scarce
- Collaboration between physicists and computer scientists is excellent
- Integration with ROOT allows fast prototyping and development cycle
- Early availability of Baseline Grid services and their integration with the ALICE-specific services will be crucial
  - This is a major “risk factor” for ALICE readiness
  - The development of the analysis infrastructure is particularly late
- The manpower situation for the core team has been stabilised thanks to the help from the collaboration
- The availability of few CERN long-term positions is of the highest priority