Computing Infrastructure and Information Technologies for LHC Physics:

Grid Tools and the LHC Data Challenges

Lothar A.T. Bauerdick, Fermilab

LHC Physics Discovery

through Information Technology and Computing Infrastructure

+

LHC Computing Unprecedented in Scale and Complexity (and Costs)

—> Need an Advanced Coherent Global "**Info**rmation-Infra**structure**"

International and Interdisciplinary Partnerships

Empower the LHC Scientists at Universities and Labs to do Research on LHC Physics Data

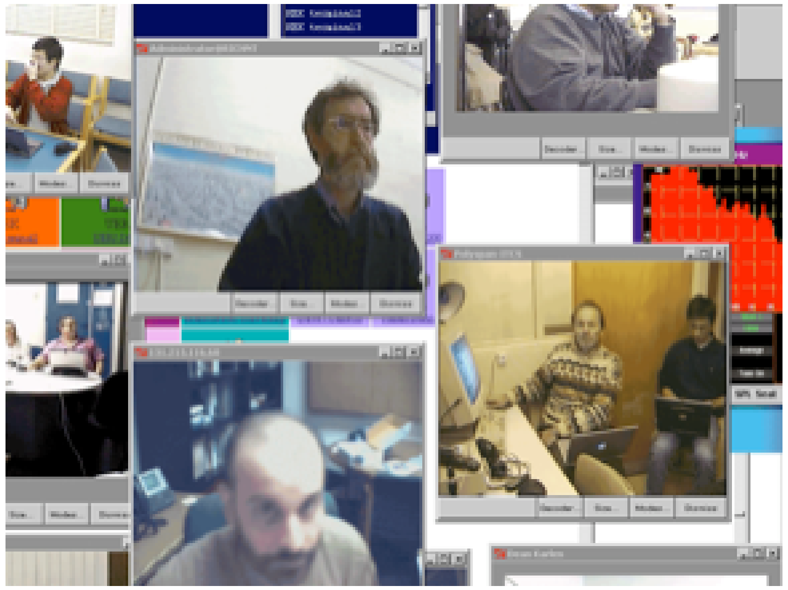

This is why we are pushing Grids and other Enabling Technology

We got to be ready on day 1!

LHC 2003

Commissioning a… Early Physics with …

Kerstin Hoepfner
RWTH Aachen, Institute of Physics IIIA
On Behalf of the CMS Collaboration

The Start-up LHC Scenario

PHASE 1: LHC commissioning (T_0 to T_0 + 4 months)

- LHC: Set-up machine. Start with one beam. Colliding beams and slowly increase # bunches and **L**. Collisions at L > 5 x 10^{32} at 25, 75 ns bunch spacing.
- CMS: Muon halo triggers, catalog detector problems, synchronization, debug data handling, record first collisions
  → This talk part I

+ 2 pileup events i.e. 10^{33} and 25ns bunch spacing

Phase2: Shutdown

PHASE 3: First physics run (T_0 + 7 mo.→ T_0 + 14 mo.)

- LHC: 25 ns and L ~ 10^{33} cm^{-2} s^{-1}
- CMS: Physics run, max. efficiency aiming for 5-10 fb^{-1}
- ~2 events per BX → This talk part II

+ 20 pileup events i.e. 10^{34} and 25ns bunch spacing

PHASE 4+n: High luminosity running

→ See talk by J.Rohlf on high luminosity

SUSY

Large cross-section for squark and gluino production

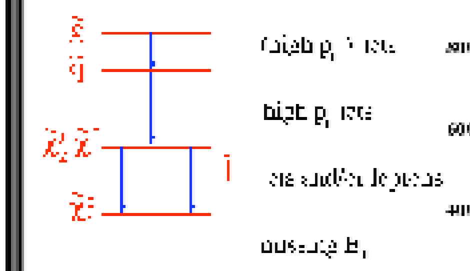

Decay chain leads to
- high p_T jets
- large missing E_T
- isolated leptons

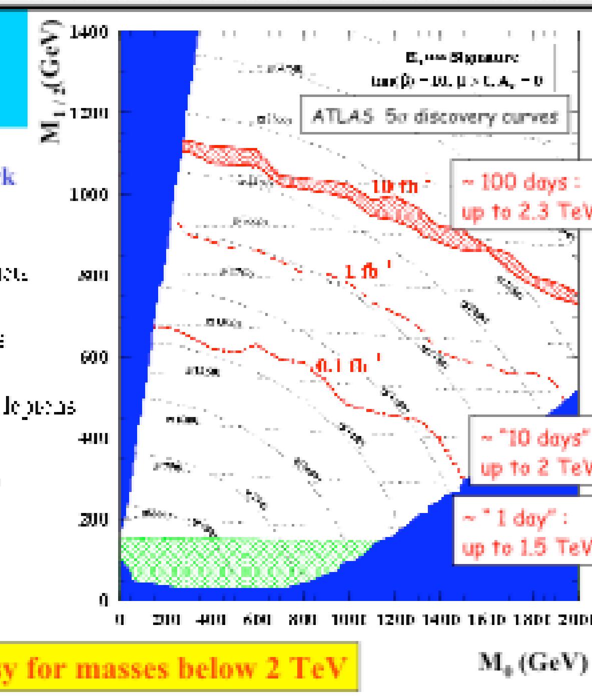

Discovery of SUSY is easy for masses below 2 TeV

May 2nd, 2003 Ivor Fleck ATLAS, Early physics reach

SM Higgs Discovery Potential with 10 fb^{-1}

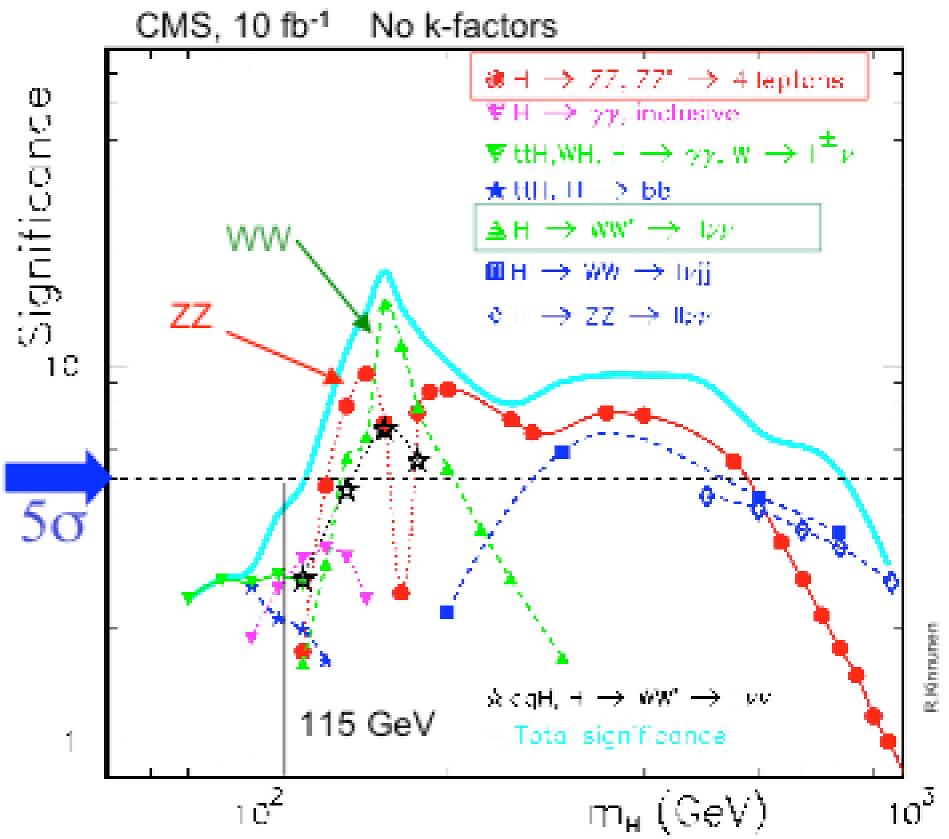

In the following:

H → 4 leptons,
m_H = 120….500 GeV

H → WW → lνlν
m_H = 110….200 GeV

K.Hoepfner, RWTH Aachen Early Physics with CMS, LHC Symposium May 2003 24

Conclusions

- **Commissioning of detector challenging**
- **procedures are being developed now**
- **within first days:**
  - **Alignment of central detector using muon tracks to < 2 mm**
  - **Calibration of EM using Z → ee to 0.6 %**
- **Impact of staging: Need ~ 10 - 15 % more integrated luminosity**

Physics results within first year:
- **Higgs boson may be discovered over full mass range (low mass region very challenging)**
- **MSSM Higgs likely to be seen**
- **LHC is factory for SUSY particles, discovery immediately first year reach is squark masses up to 2 TeV**

Looking forward to many interesting physics analyses

May 2nd, 2003 Ivor Fleck ATLAS, Early physics reach

Data Challenges Central to Being Ready!

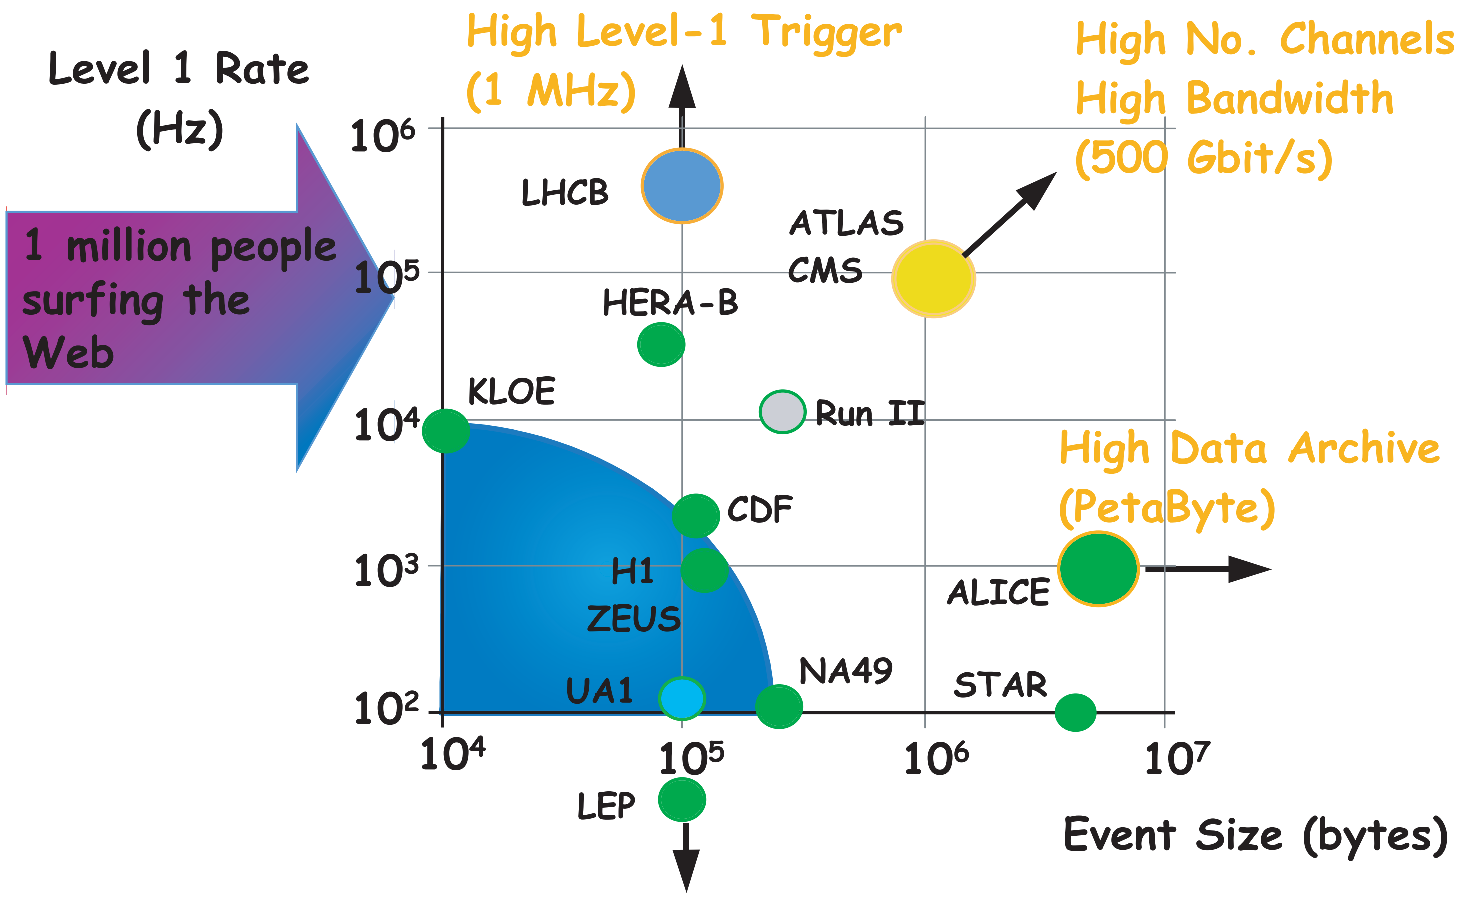

Experiment's Computing Models and Systems not yet fully defined

- Physics Model — Analysis Model — Data Model — Computing Model

Technical Design Reports are not yet written

- Functional implementations exist, but fundamental technologies are still beiong developed
  - e.g. Object store (POOL), many Grid Tools, data bases, etc

Approach: Series of Data Challenges; performance tests on international computing infrastructures

- example: CMS DC04, to be completed in April 2004
  Reconstruct 50 million events and cope with 25 Hz at 2_1033 cms-2 s-1 for 1 month
- Computing AND Physics Validation
- strong involvement of physicists at CERN and around the world

The Great Challenge of LHC and the Grid

"Virtualization": hiding useful functions behind an interface that conceals the details how they are implemented

- virtualization of computing and data services
- Grid of networked processing elements and data storage elements
- "middleware" provides the glue

Vision:
A Richly Structured, Global Dynamic System

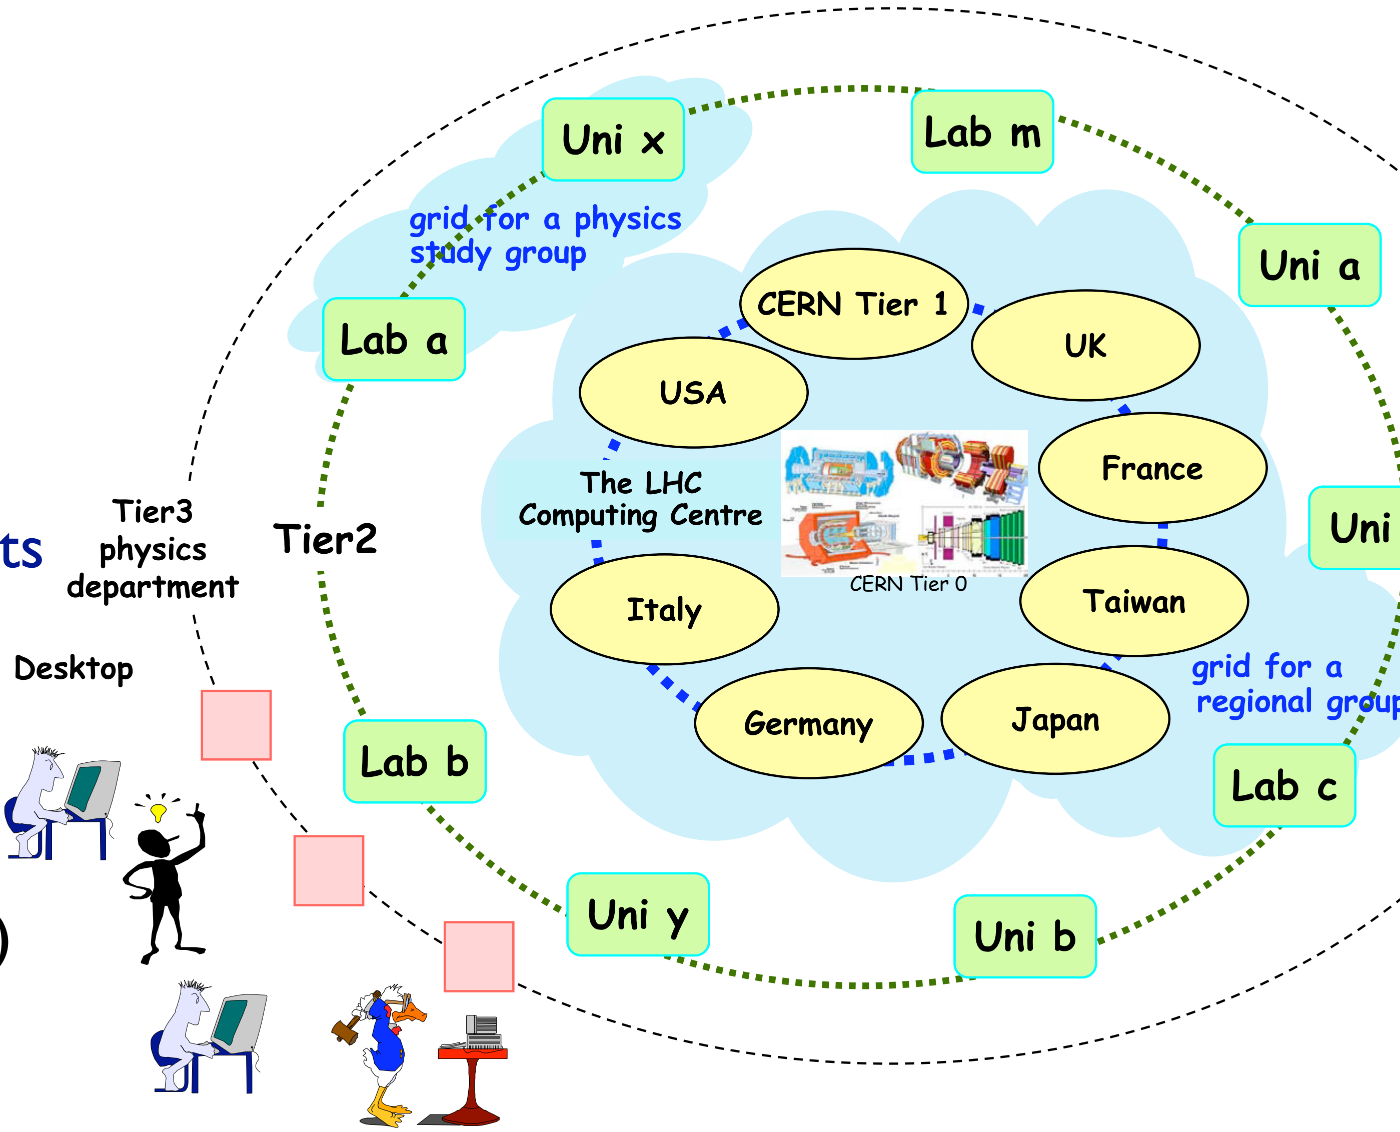

Virtual Computing Service for Experiments
Ubiquitous Responsive Environment for Physicists

- distributed yet coherent computing
- coordinated and efficient sharing of geographically distributed resources
- conditional sharing issues of trust, policy, negotiation, payment)
- optimization of the resources
- invisibility of the local architecture
- partnerships and collaboration

Grid Architecture — HEP Layers

HEP Grid Architecture: (H. Newman)

Layers Above the Collective Layer

Layered Grid Architecture

(I.Foster et al.)

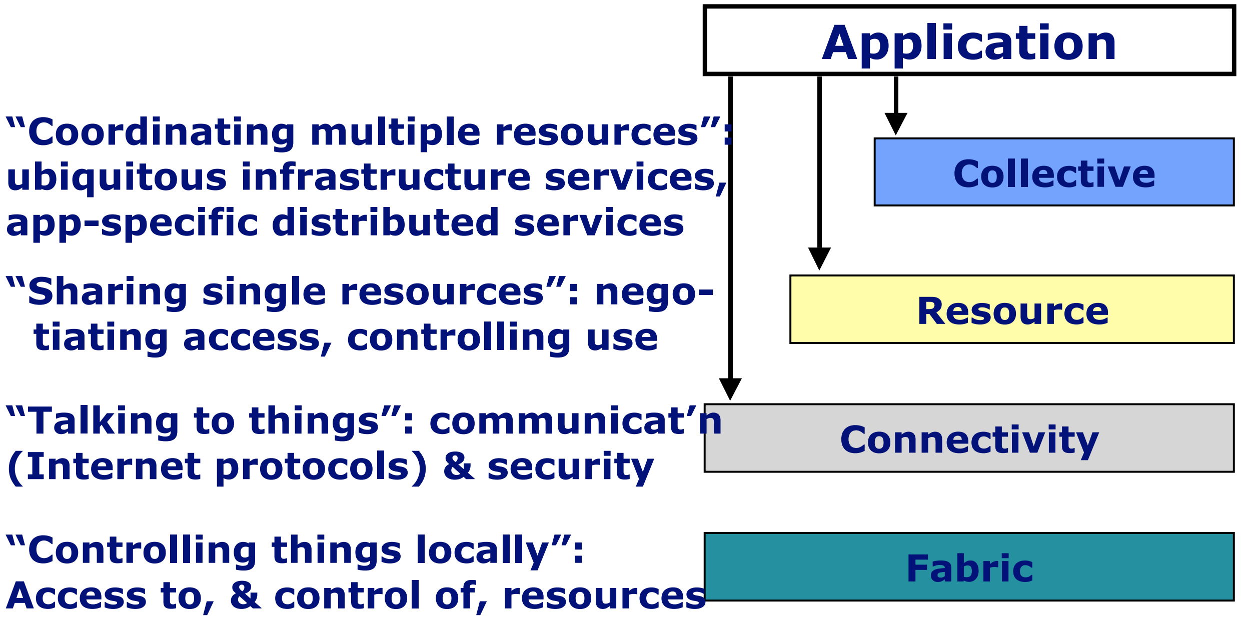

Physicist's Application Codes

- Reconstruction, Calibration, Analysis

Experiments' Software Framework Layer

- Modular and Grid-aware: Architecture able to interact effectively with the lower layers (above)

Grid Applications Layer
(Parameters and algorithms that govern system operations)

- Policy and priority metrics
- Workflow evaluation metrics
- Task-Site Coupling proximity metrics

Global End-to-End System Services Layer

- Workflow monitoring and evaluation mechanisms
- Error recovery and long-term redirection mechanisms
- System self-monitoring, steering, evaluation and optimization mechanisms
- Monitoring and Tracking Component performance

LCG Prototype: LCG-1 Architecture of Middleware Layers

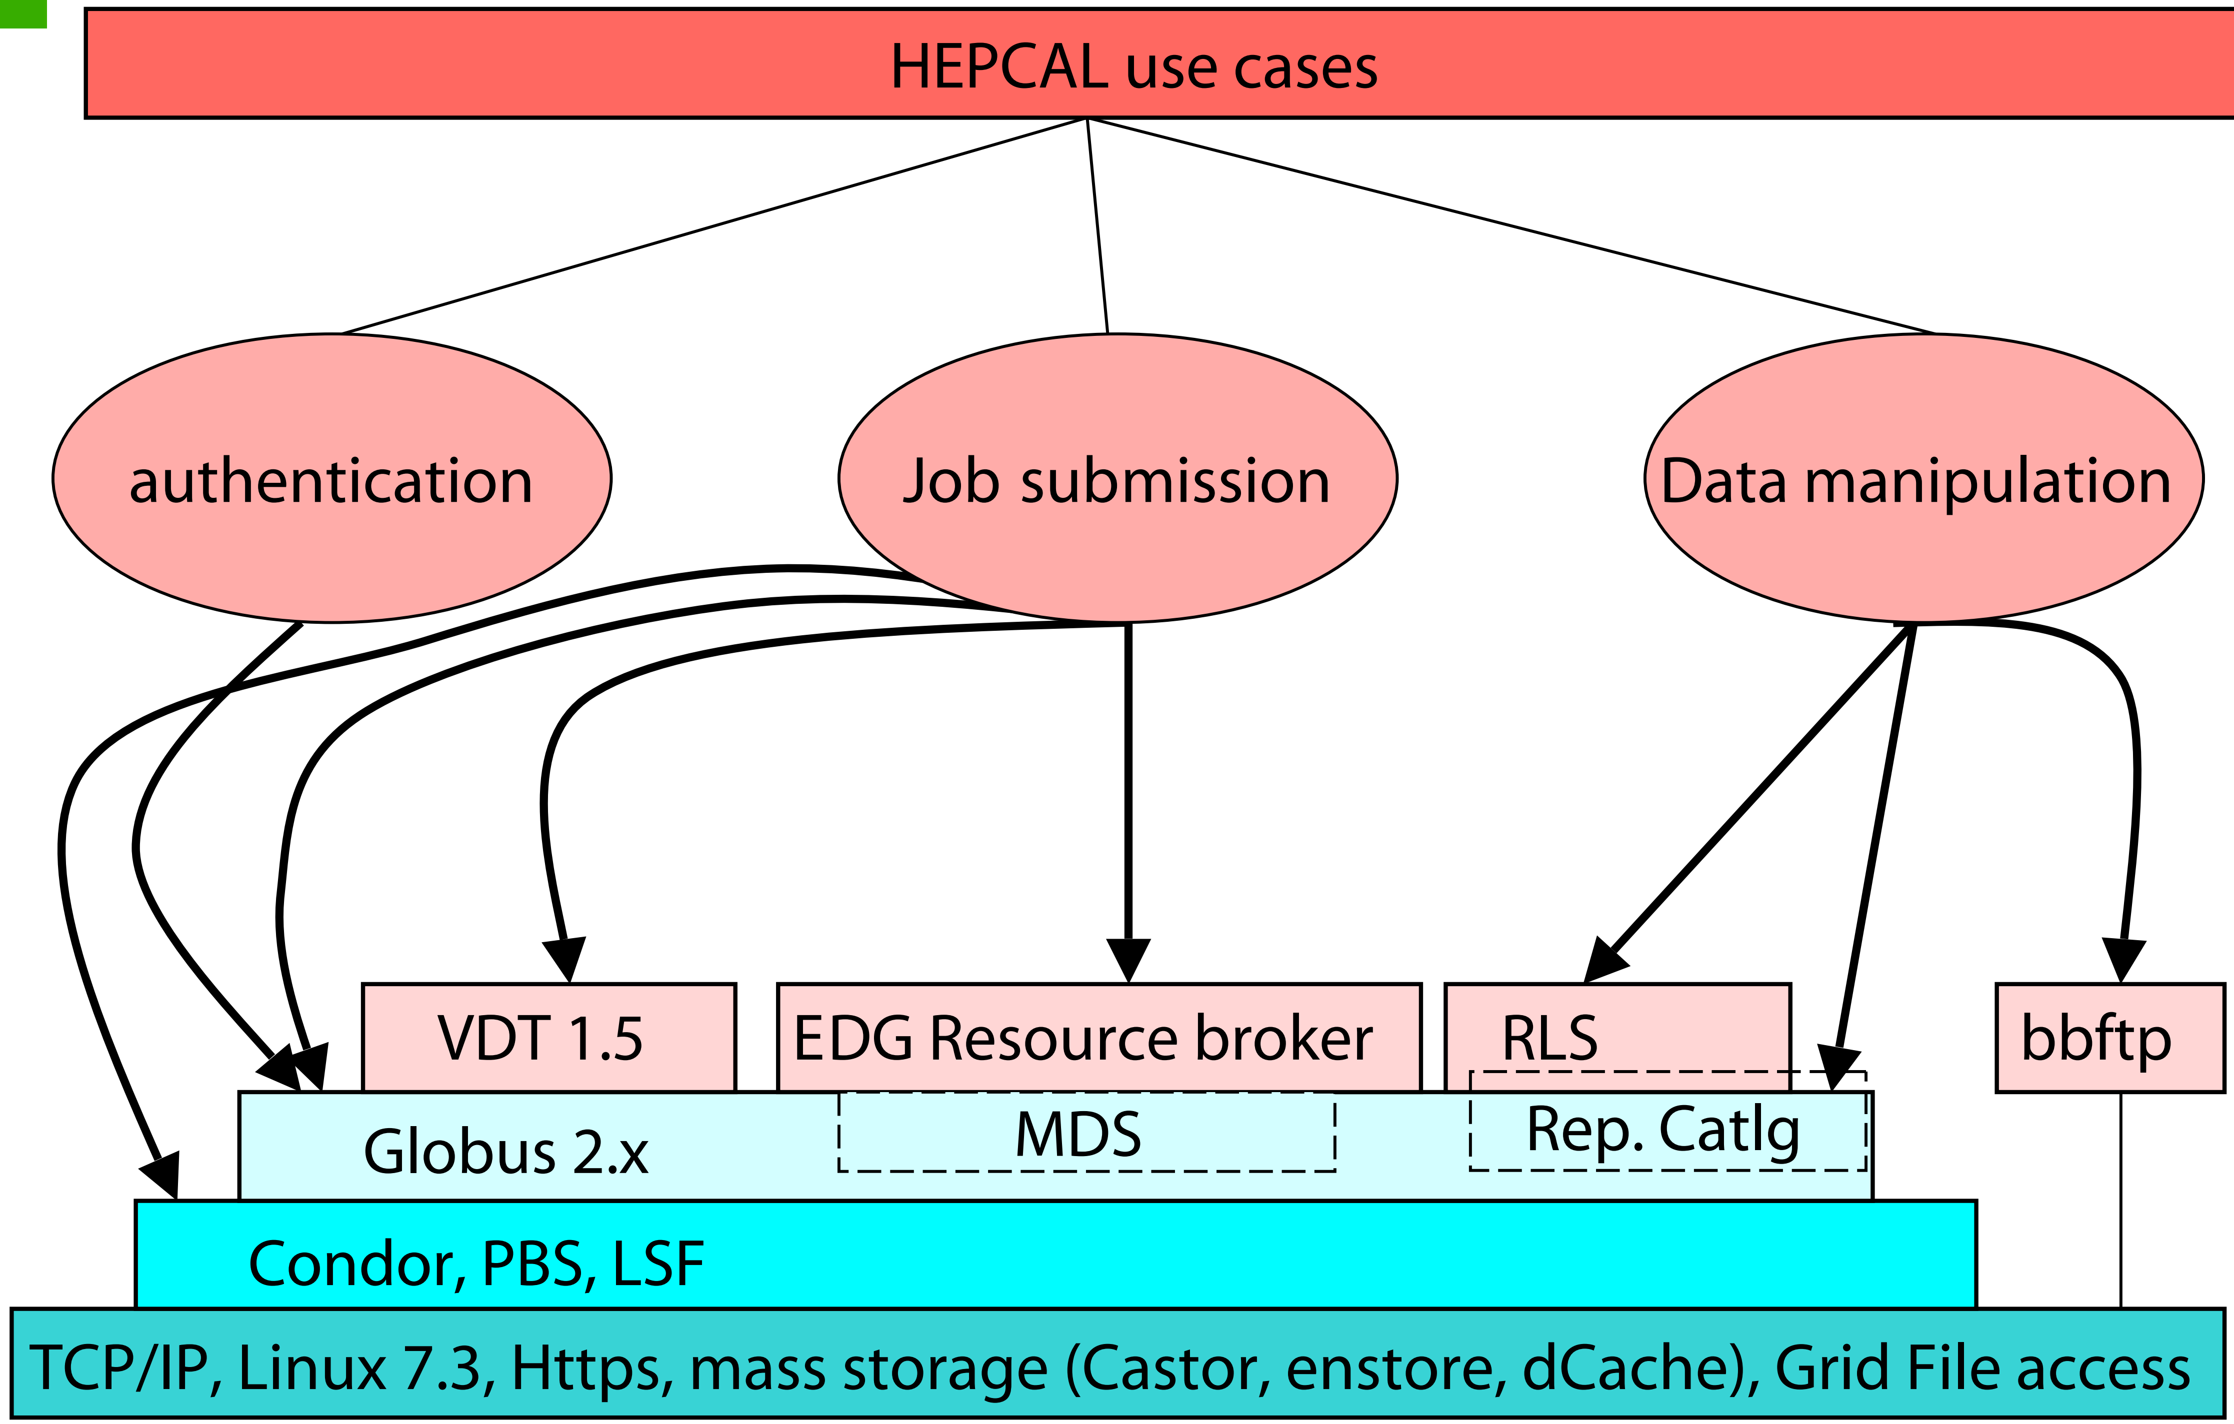

Transition to Production-Quality Grid

Centers for LHC Grid 2003 Around the World ➔ Around the Clock!

Preparing for next round of Physics and Computing Data Challenges

Experiment's Data Challenges

Data, e.g. Atlas, CMS

- Every event will consist of 1-1.5 MB (all detectors together)
- After on-line selection, events will be written to permanent storage at a rate of 100-200 Hz
- Total amount of "raw" data: ~1 PB/year
- To reconstruct and analyze this data:
  Complex "Worldwide Computing Model" and "Event Data Model"
  - Raw Data @ CERN (and Tier-1?)
  - Reconstructed data "distributed"
  - All members of the collaboration must have access to "ALL" public copies of the data (at a reasonable access speed)

Data Challenges

- The emerging world wide computing model "is an answer" to the LHC computing challenge
- In this model the Grid must take care of:
  - data replicas and catalogues
  - condition data base replicas, updates and synchronization
  - access authorizations for individual users, working groups, production managers
  - access priorities and job queues
- Validation of the new Grid computing paradigm in the period before the LHC requires Data Challenges of increasing scope and complexity
- Build up and involve the physics community to be ready at day 1!

e.g. LHCb Data Challenge Plans

Feb-May 2003: physics data challenge: Centers and Data Grid integration

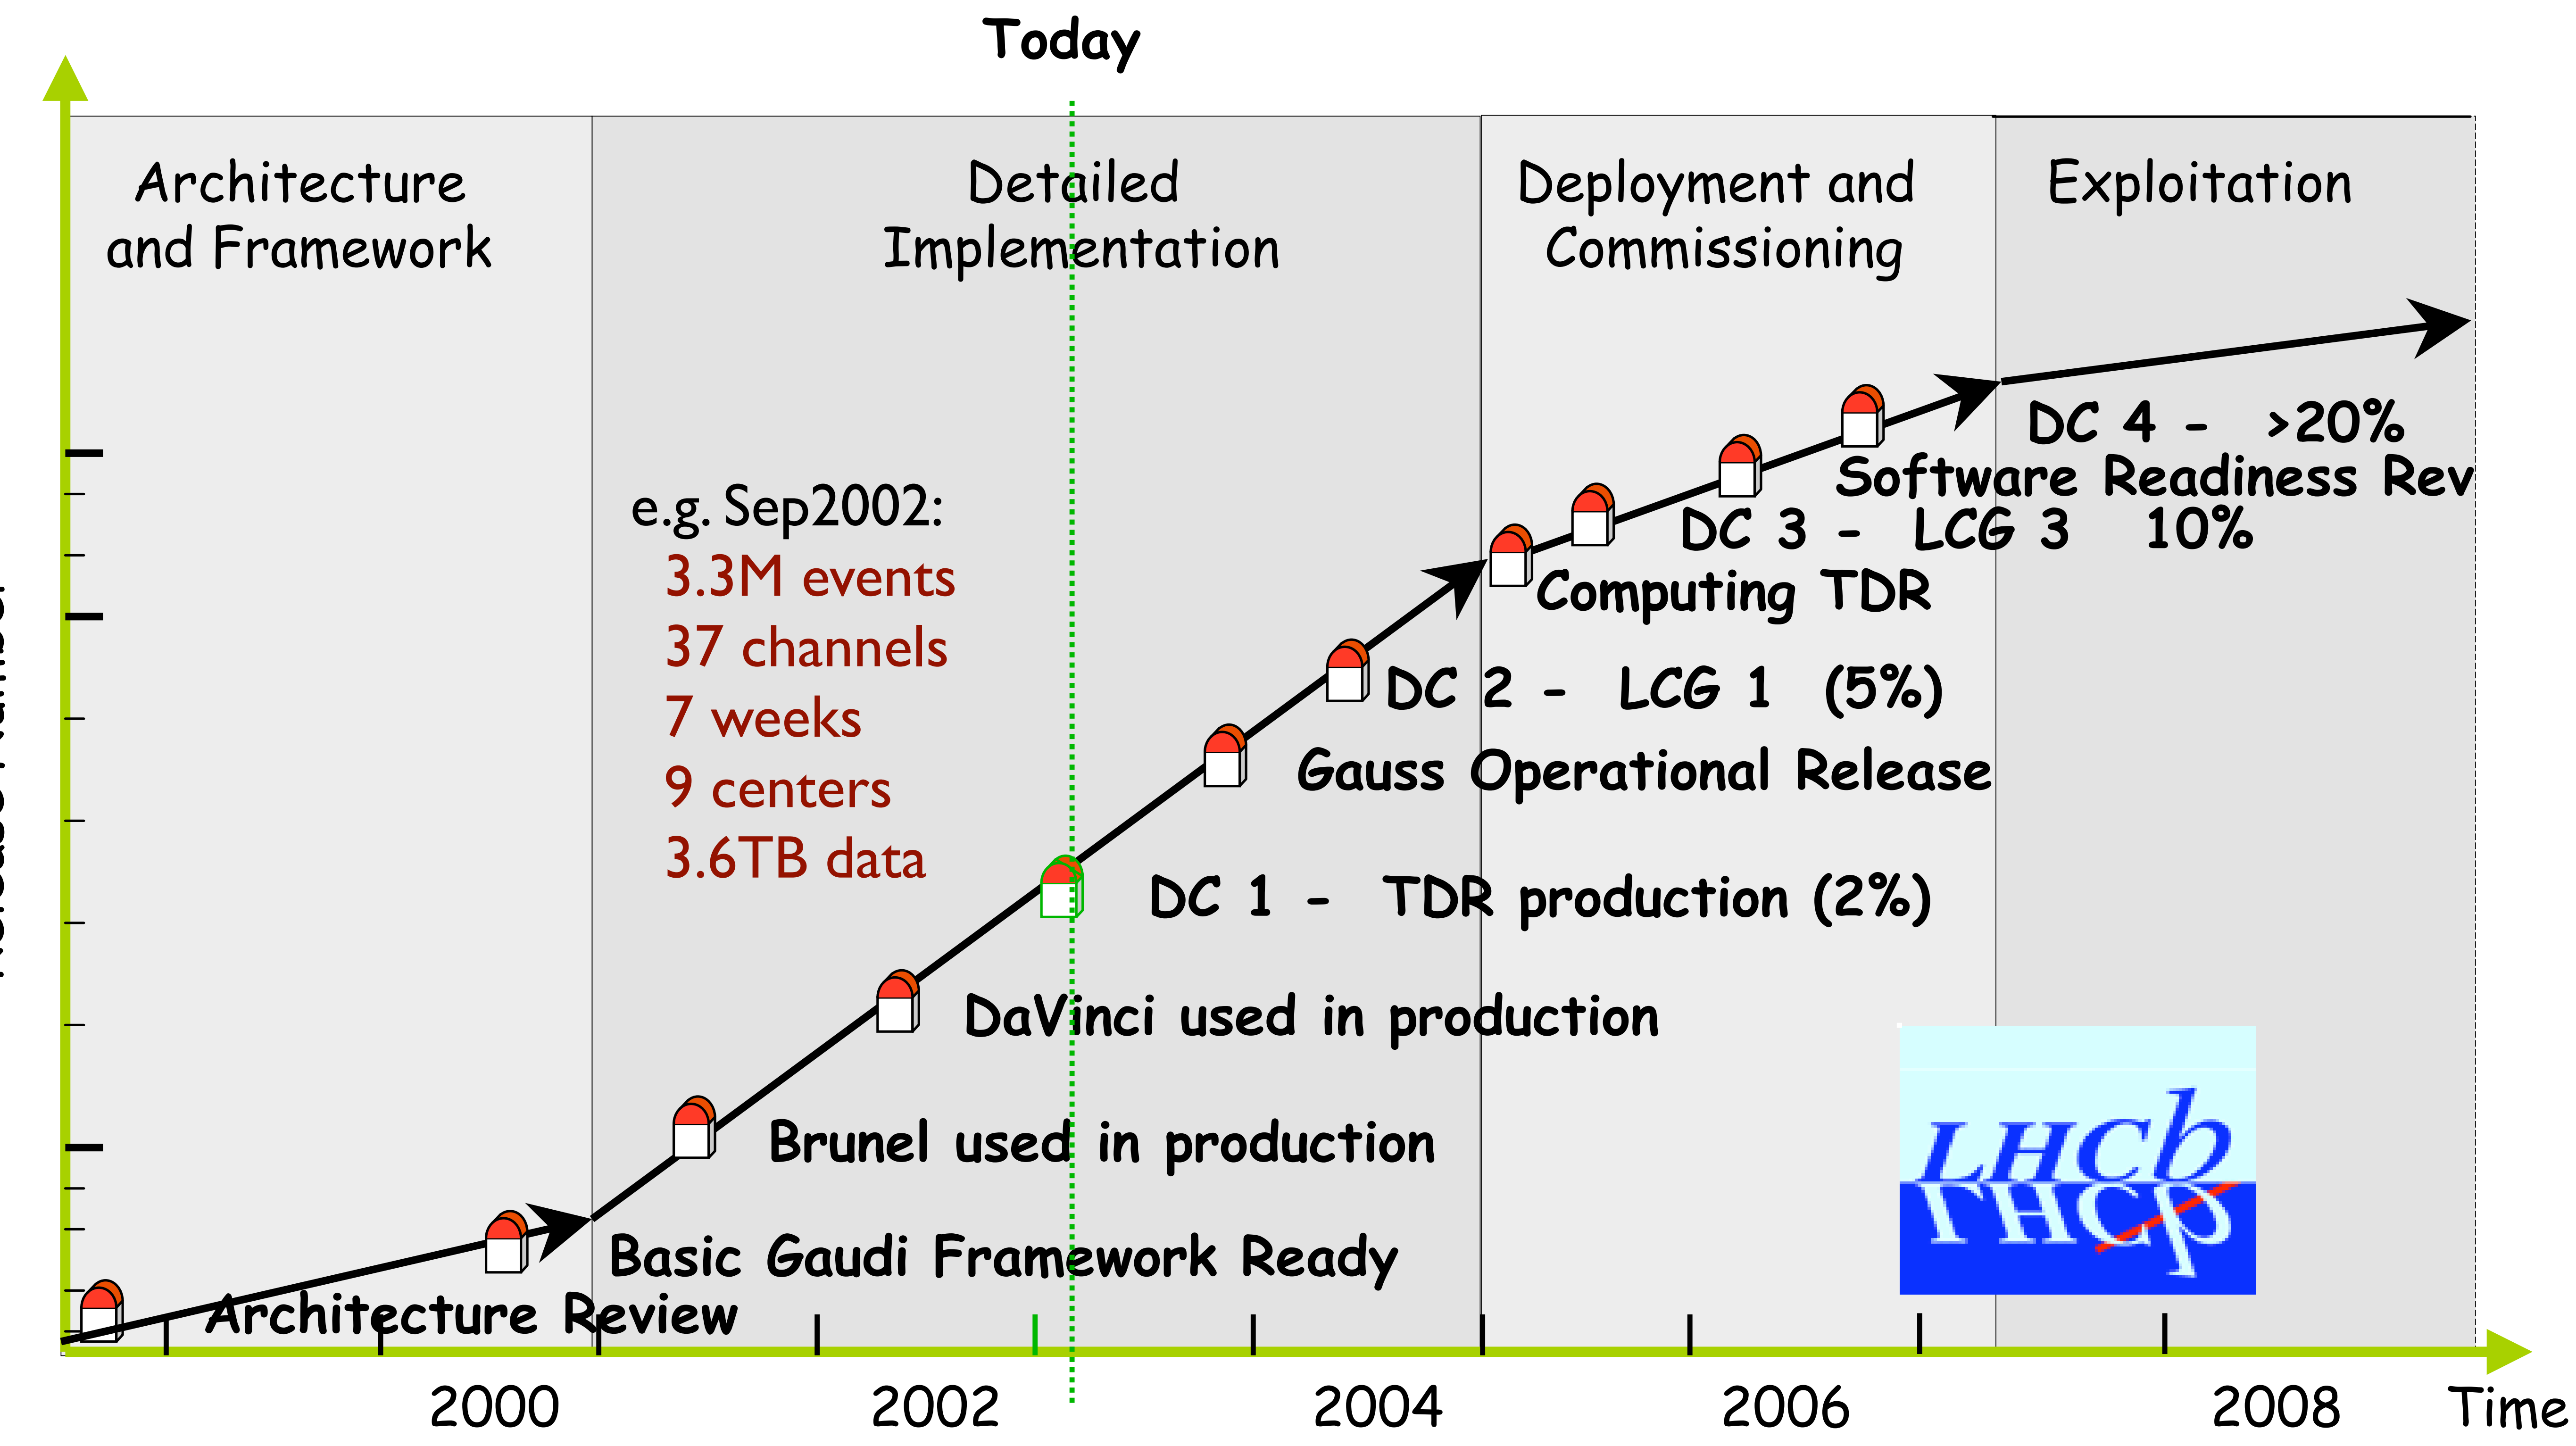

Alice: running physics data challenges with AliEn

AliEn @GRID

- AliEn integrating distributed computing and mass storage resources into a coherent production environment, and can be interfaced with any Grid environment
- Have completed several rounds of production
- First users are beginning to use AliEn to access simulated data and run their analysis

PHYSICS DATA CHALLENGES SCHEDULE & OBJECTIVES

DC1 (1%) 06/01-12/01 pp studies, reconstruction of TPC and ITS

DC2 (5%) 06/02-12/02 First test of the complete chain from simulation to reconstruction for the PPR. Simple analysis tools. Digits in ROOT format.

DC3 (10%) 01/04-06/04 Complete chain used for trigger studies. Prototype of the analysis tools. Comparison with parameterised MonteCarlo. Simulated raw data.

DC4 (20%) 01/06-06/06 Test of the final system for reconstruction and analysis.

Total jobs per site

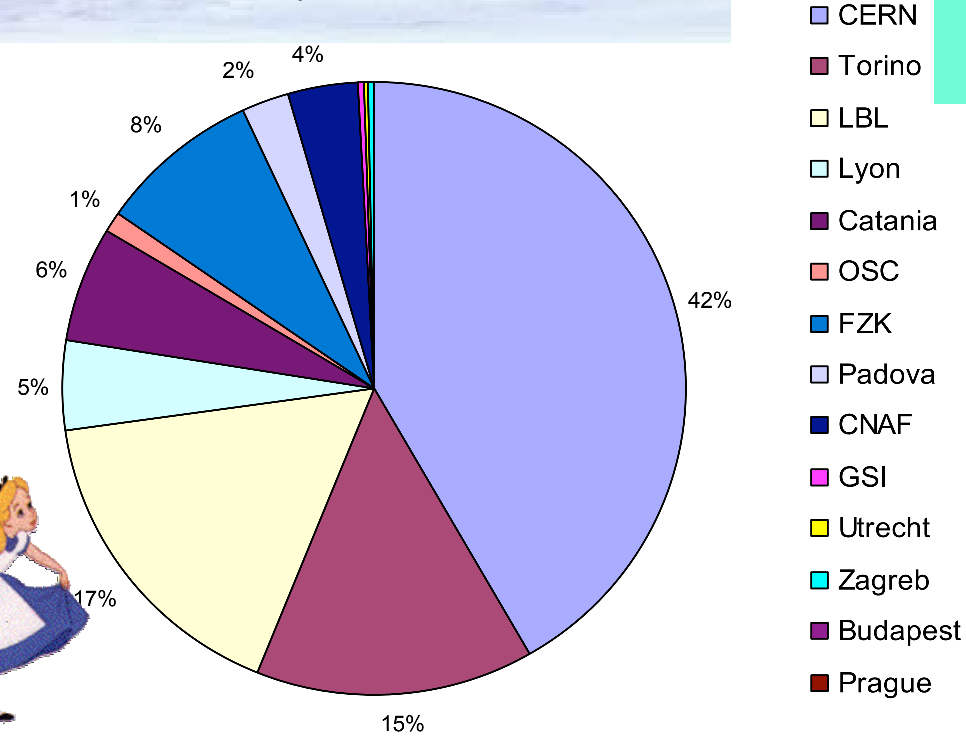

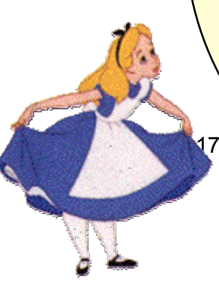

22713 jobs, ~12CPUh/job, ~1GB output/job
up to 450 concurrently running jobs

35 sites configured, at present ~14 contributing with CPU cycles

Fully distributed production controlled from one point

Alice DC IV performance

5-days running flat out: event building and data recording to tape

- ➡ successful integrated test high throughput into tape store

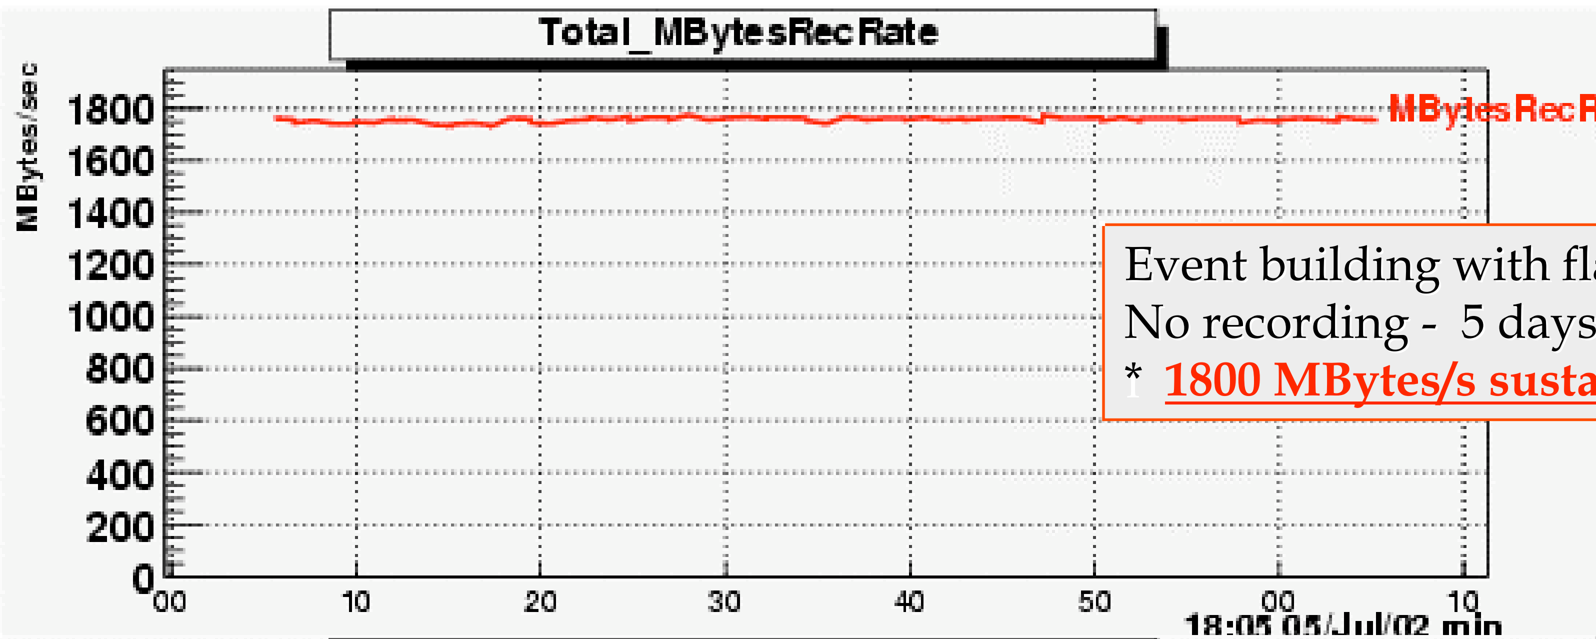

Event building with flat data traffic
No recording - 5 days non-stop
* <u>1800 MBytes/s sustained (milestone: 1000 Mbytes/s)</u>

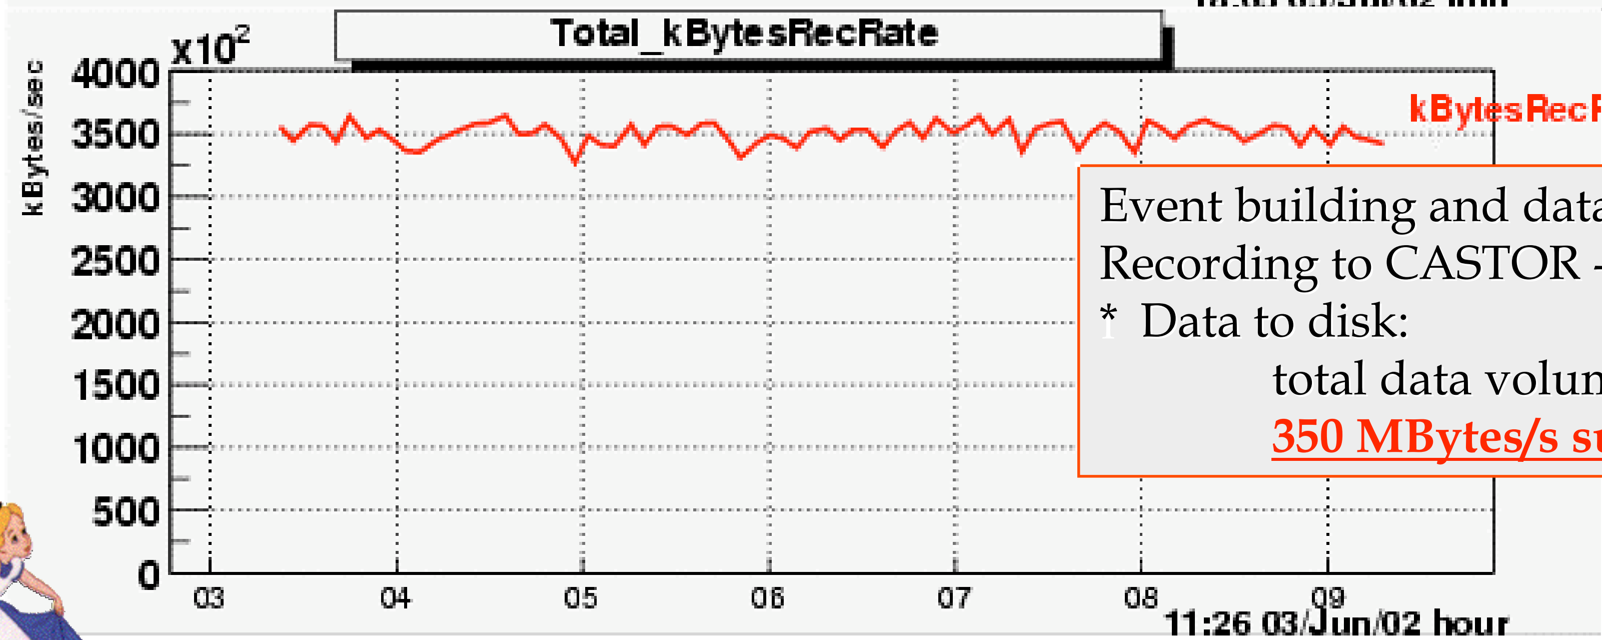

Event building and data recording with ALICE-like data traffic
Recording to CASTOR - 4.5 days non-stop
* Data to disk:
  total data volume : ~ 140 Tbytes
  <u>350 MBytes/s sustained (milestone 300 MBytes/s)</u>

Atlas DC1 Phase 1 : July-August 2002

$5*10^7$ events generated
$1*10^7$ events simulated
$3*10^7$ single particles
30 TBytes, 35 000 files

3200 CPU's
110 kSI95
71000 CPU days

39 Institutes in
18 Countries
grid tools
used at 11 sites

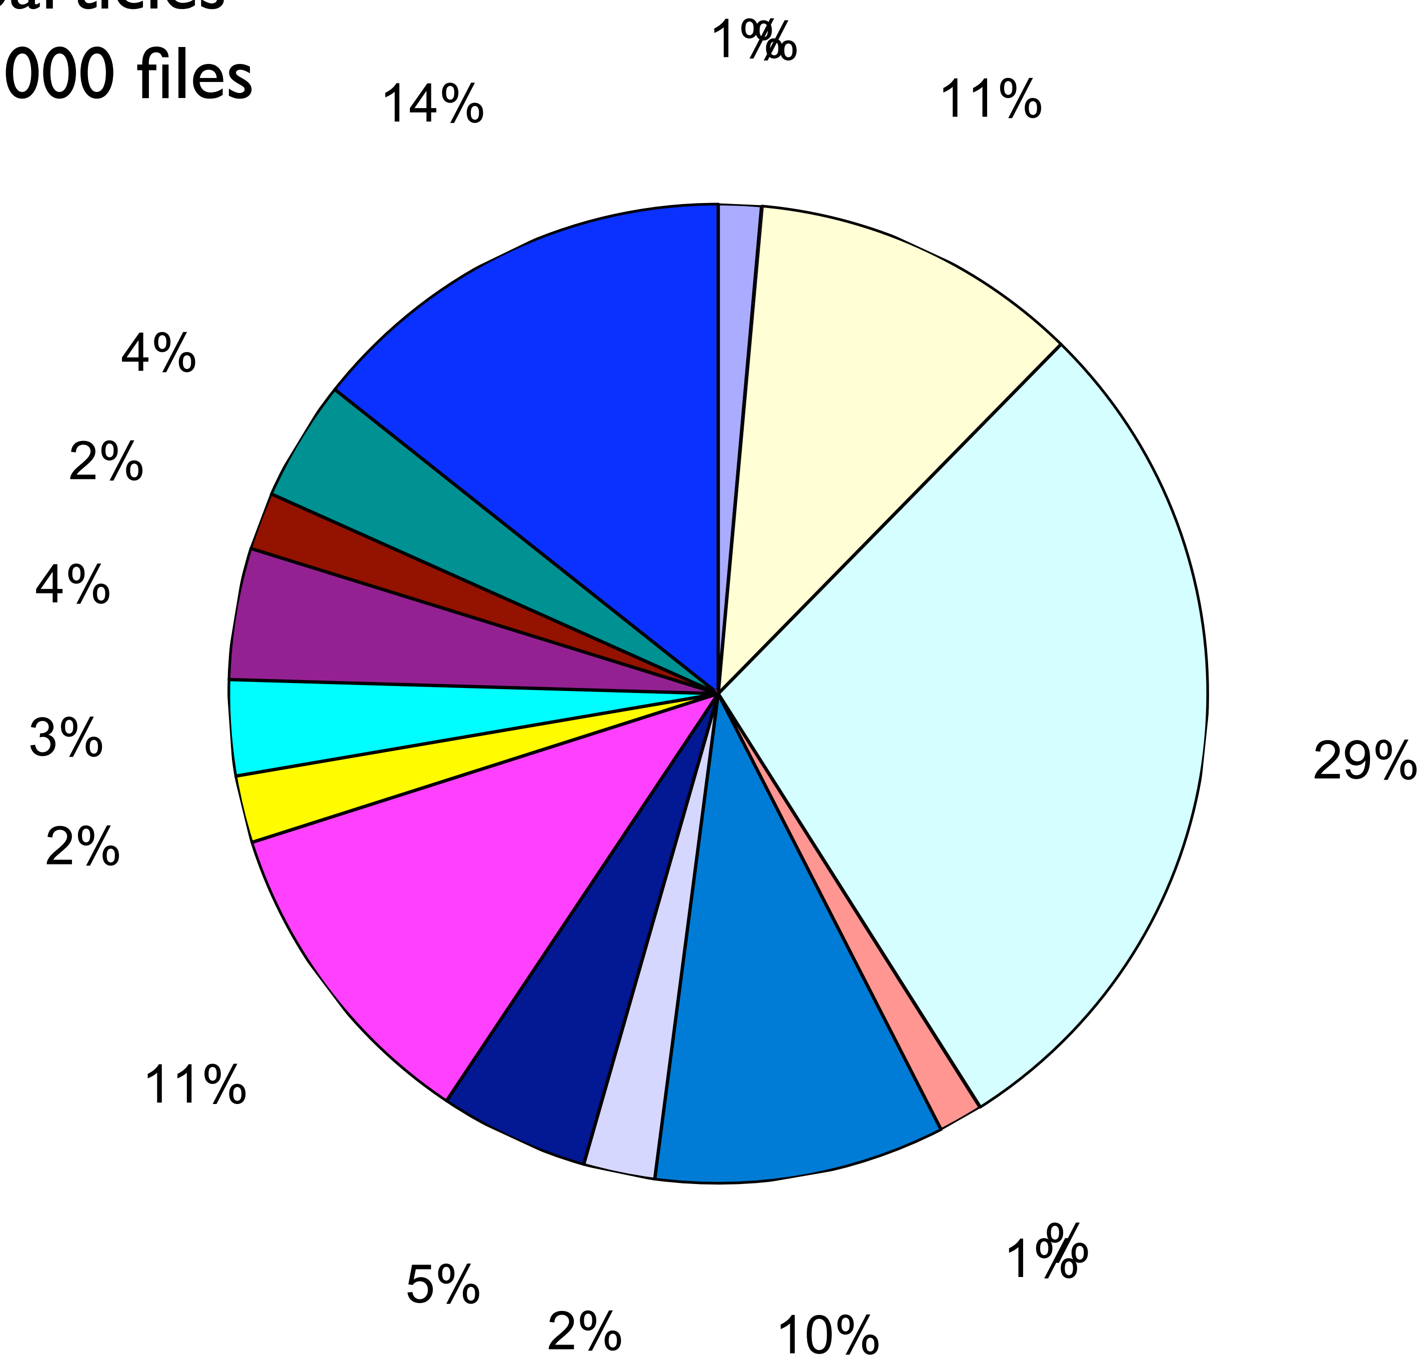

1. Australia
2. Austria
3. Canada
4. CERN
5. Czech Republic
6. France
7. Germany
8. Israel
9. Italy
10. Japan
11. Nordic
12. Russia
13. Spain
14. Taiwan
15. UK
16. USA

ATLAS

Grids in Atlas DC1

US-ATLAS

- part of Phase 1 production
- full Phase 2 production

EDG Testbed Prod

- reproduce part of phase 1 data
- several tests

NorduGrid

- full phase 1 & 2 production

ATLAS

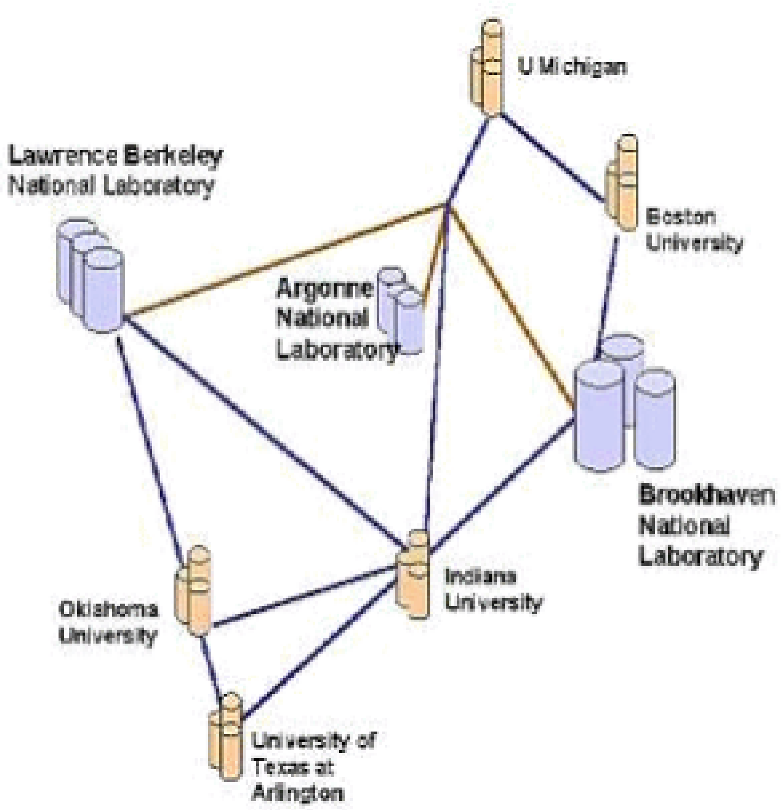

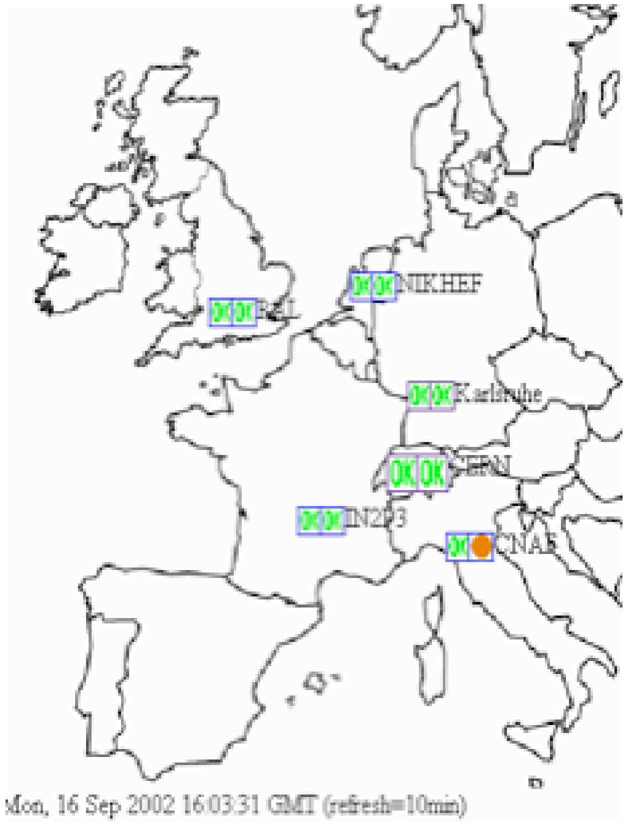

CMS First Data Challenge Results

Previous Round of CMS Production Physics Data Challenge

- ended Summer 2002
- full simulation and pileup
- events used in HLT studies mostly using "traditional" analysis tools
- full simulation, including pileup important to produced dependable event selection and trigger tables
- next: replacement of Objectivity Geant3—> Geant4
- preparation for 5% data challange DC04 —> Computing TDR

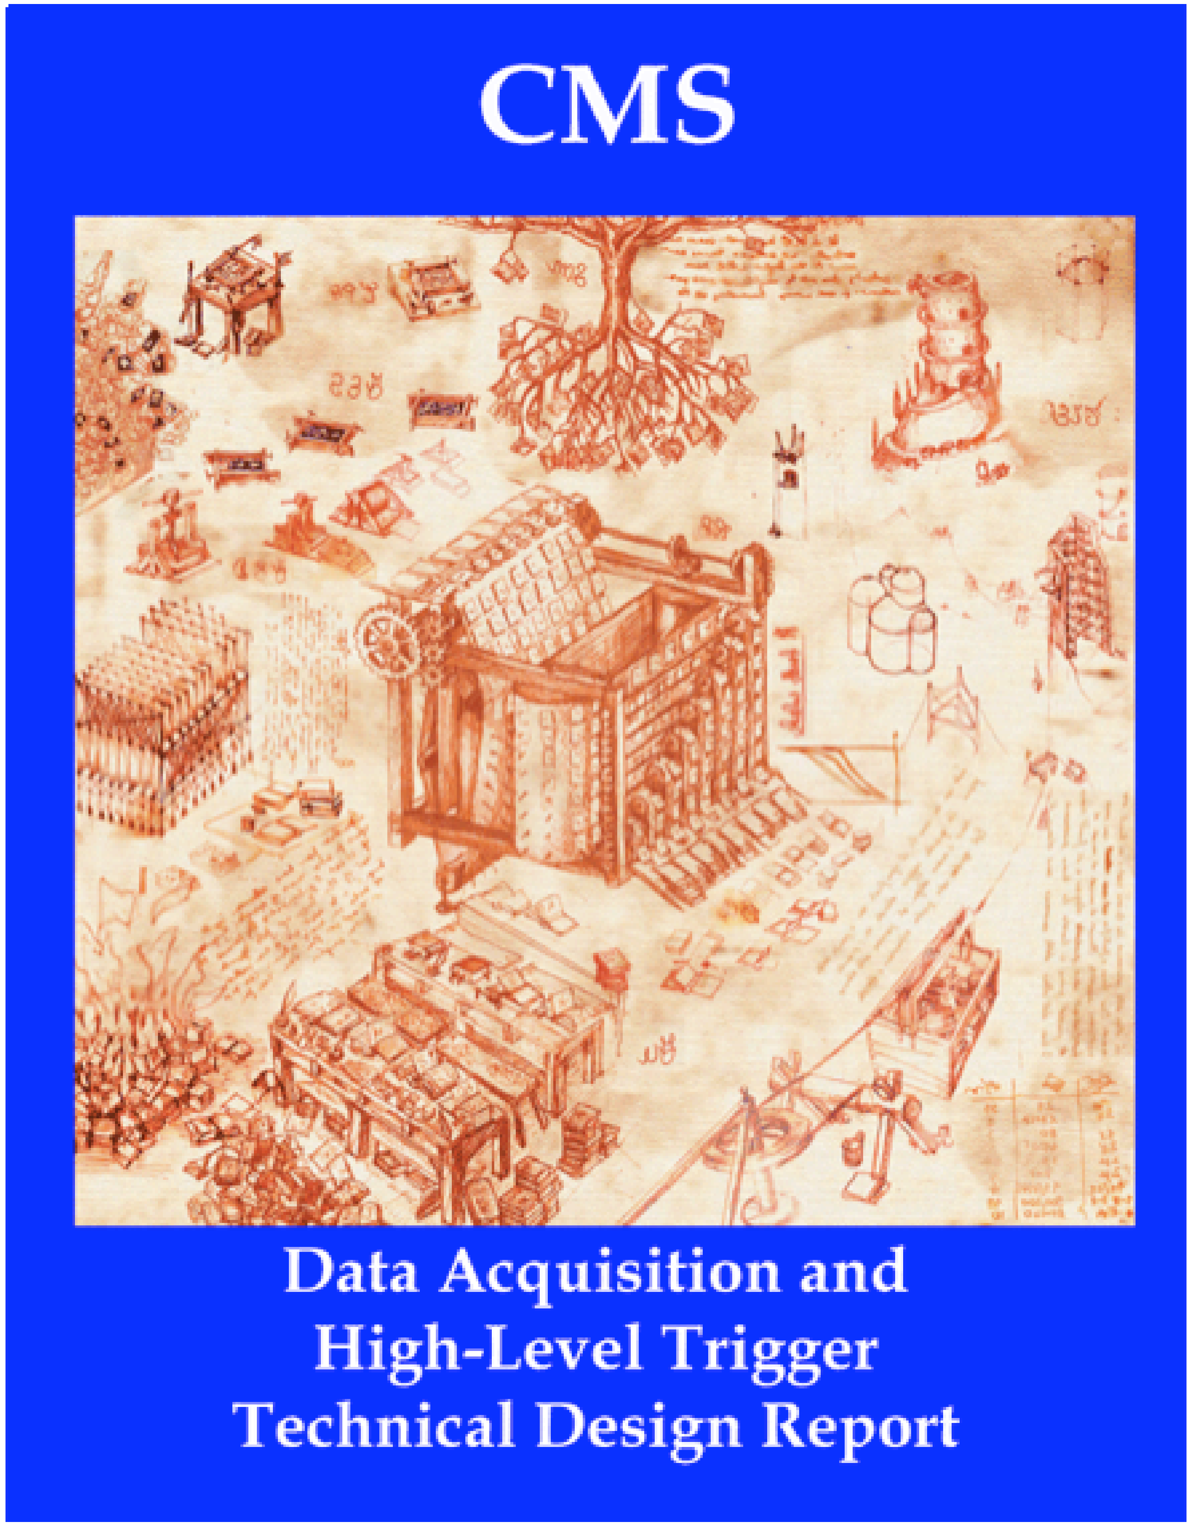

Goals and Status of DCs

LHCb

- first physics DC ended: simulation and reconstruction of large event samples
  - 45M events produced —> CERN tapes; next: HLT studies
- Spring 2004: Computing DC —> TDR

Alice

- Physics DC2 ended end 2002
  - testing complete chain from simulation to reconstruction to simple analysis tools
- Computing DC
  - high throughput into CERN tape system

ATLAS

- DC1 - Phase 2 with pile-up simulation, reconstruction ending now
- next: moving to higher level Grid tools

CMS

- (physics) DC 2002 finished in summer —> HLT and DAQ TDR
- next: Computing DC —> Computing TDR

Plans: e.g. Atlas DC 2-4

DC0 and DC1: finished

DC2: Probably Q4/2003 – Q2/2004

- **Goals**
  - **Full deployment of EDM & Detector Description**
  - **Geant4 replacing Geant3**
  - **Test the calibration and alignment procedures**
  - **Use LCG common software (POOL, …)**
  - **Use widely GRID middleware**
  - **Perform large scale physics analysis**
  - **Further tests of the computing model (Analysis)**
  - **Run on LCG-1**
- **Scale**
  - **As for DC1: ~ 10**7 fully simulated events**

DC3: Q3/2004 – Q2/2005

- **Goals to be defined; Scale: 5 x DC2**

DC4: Q3/2005 – Q2/2006

- **Goals to be defined; Scale: 2 X DC3**

ATLAS

e.g. CMS Data Challenge Planning

CMS

❖ Computing TDR, Fall 2004

- Technical specifications of the computing and core software systems
  - for DC06 Data Challenge and subsequent real data taking
- Includes results from DC04 Data Challenge
  - successfully copes with a sustained data-taking rate equivalent to 25Hz at 2x1033 for a period of 1 month

❖ Physics TDR, Dec 2005

❖ CMS Physics, Summer 2007

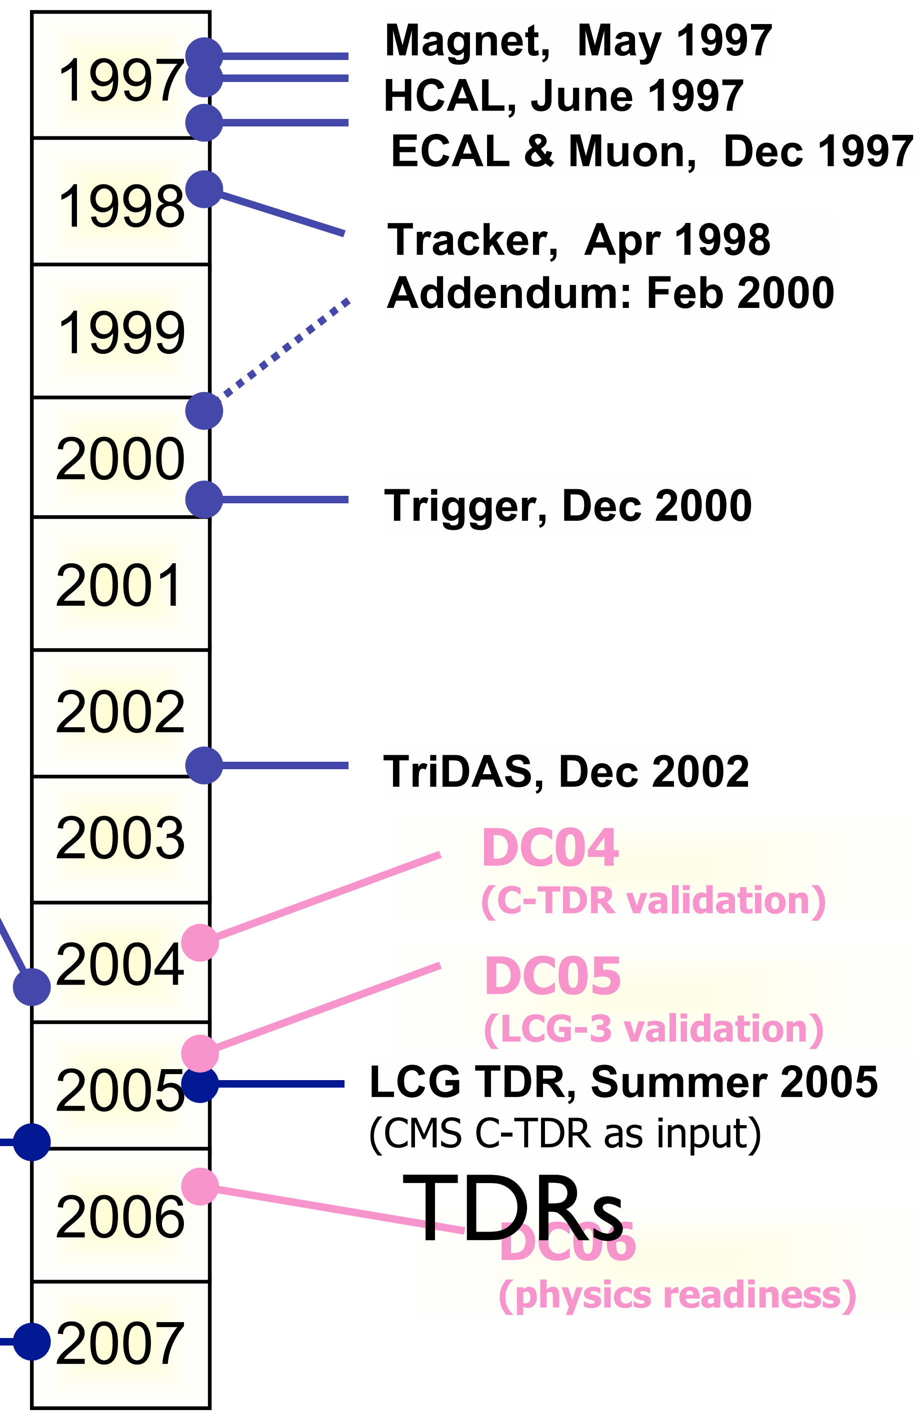

CMS Strategy for Computing TDR

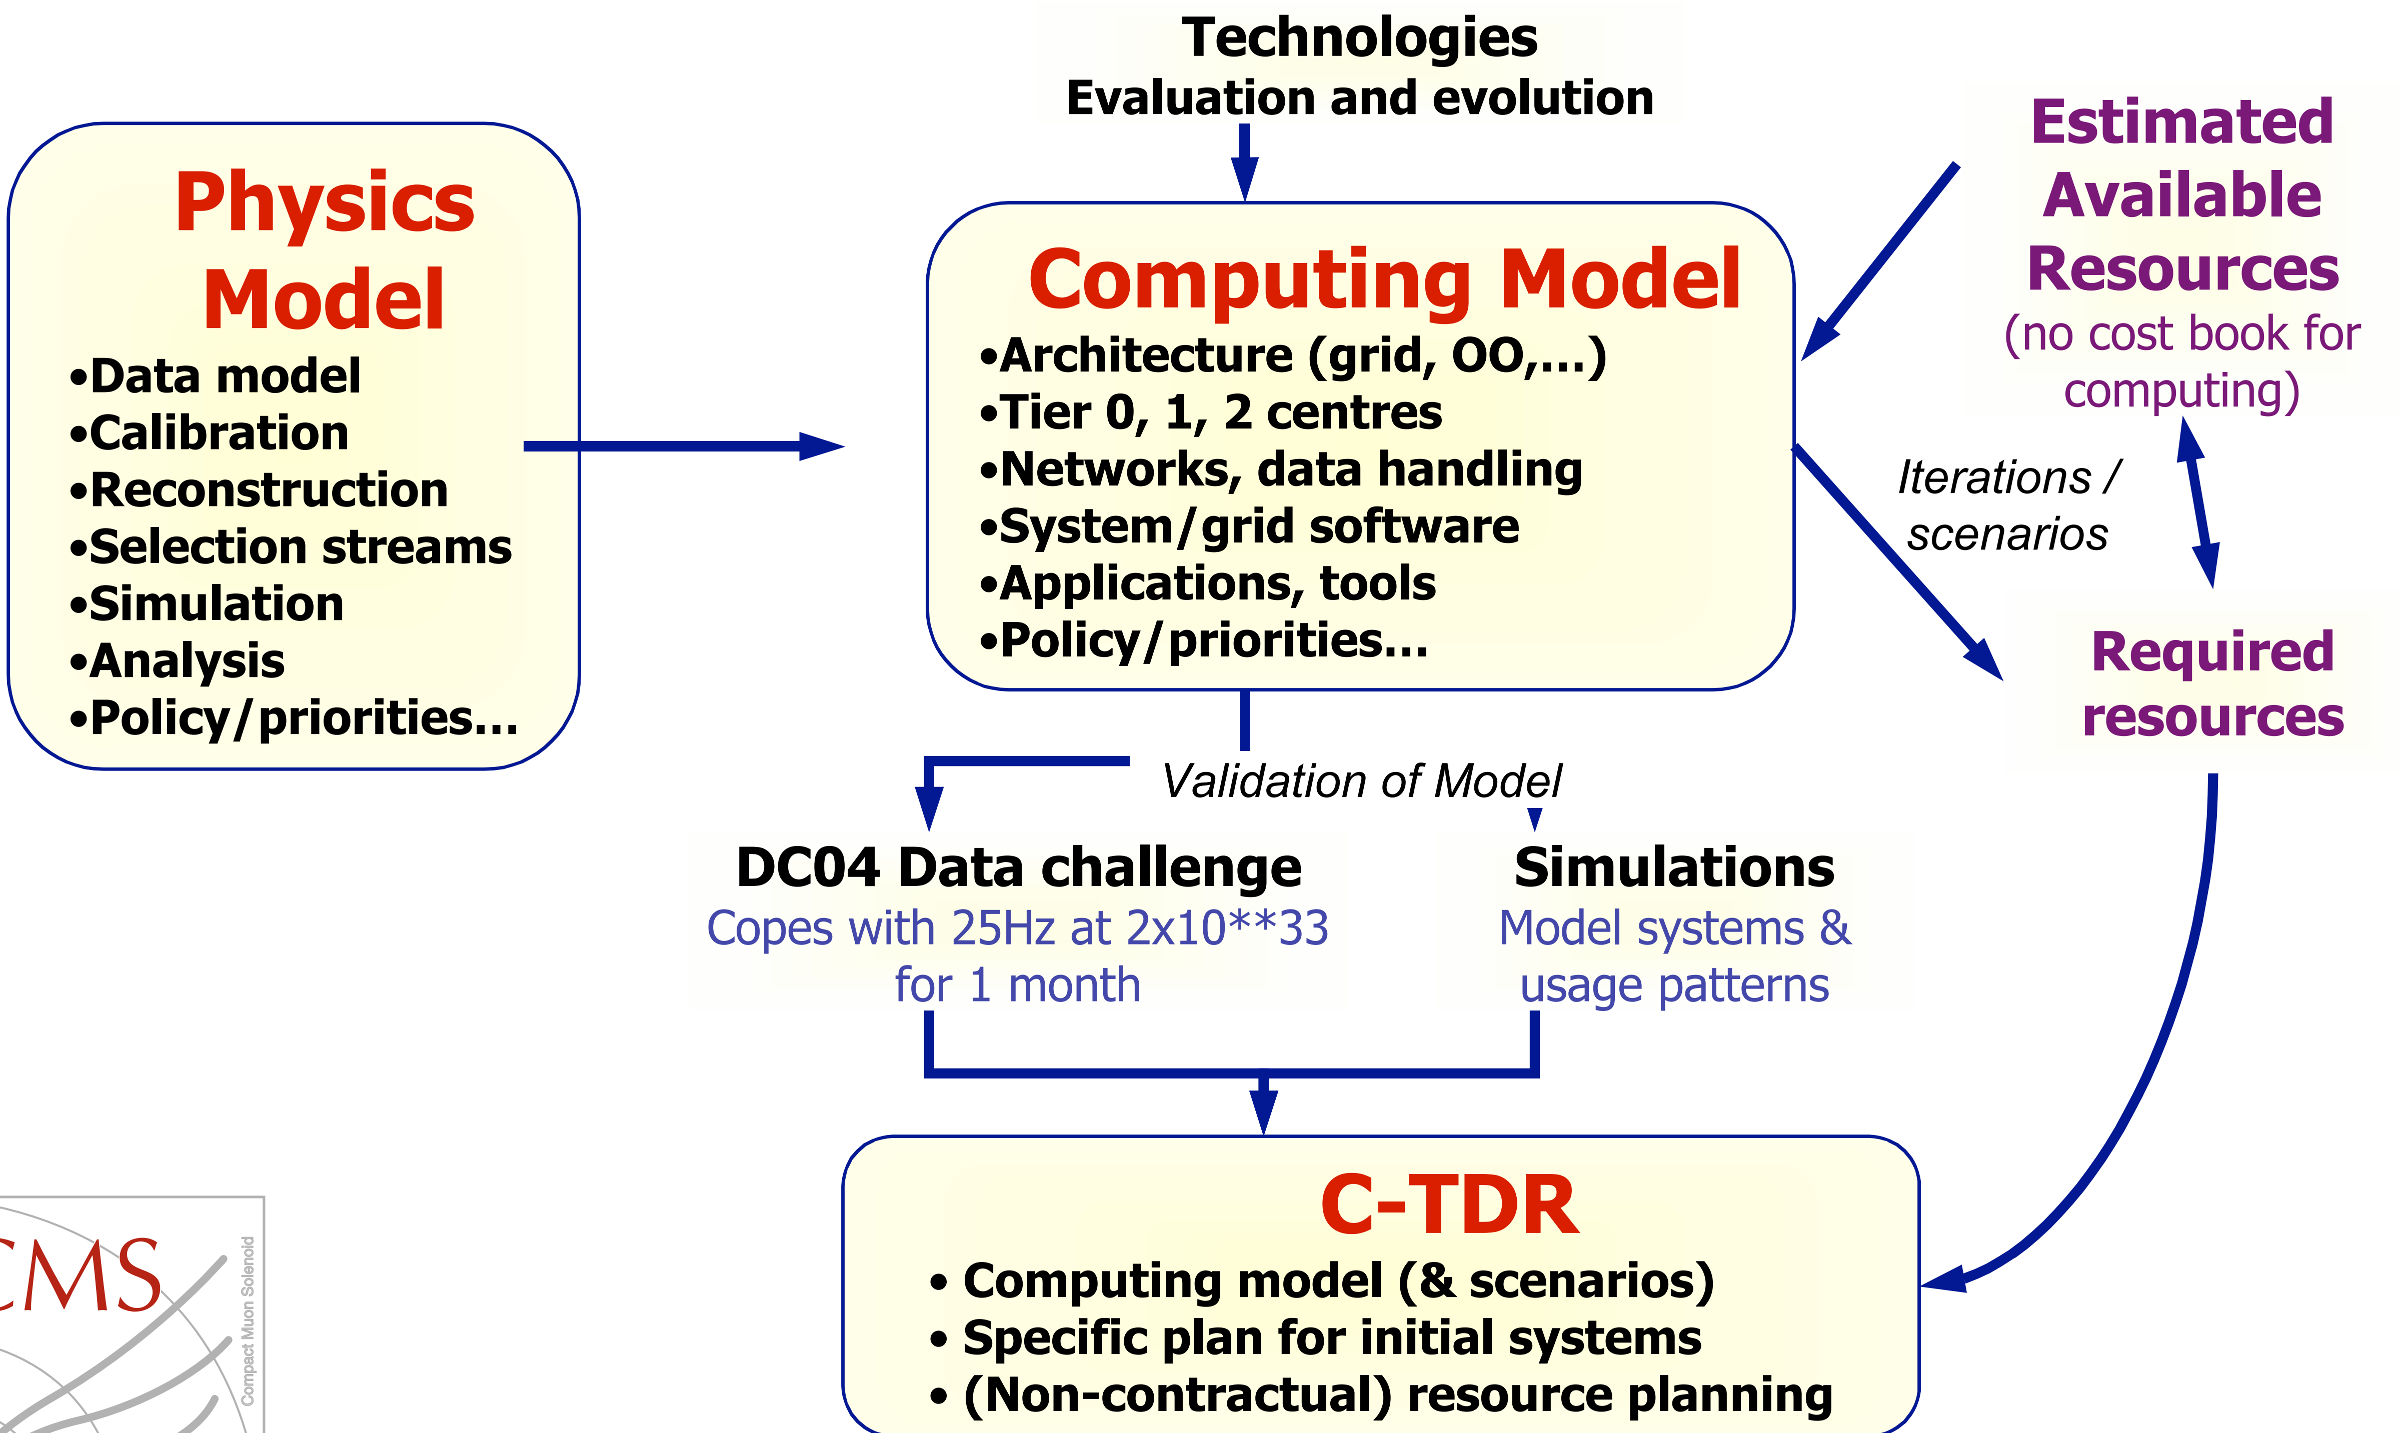

CMS Computing TDR Schedule

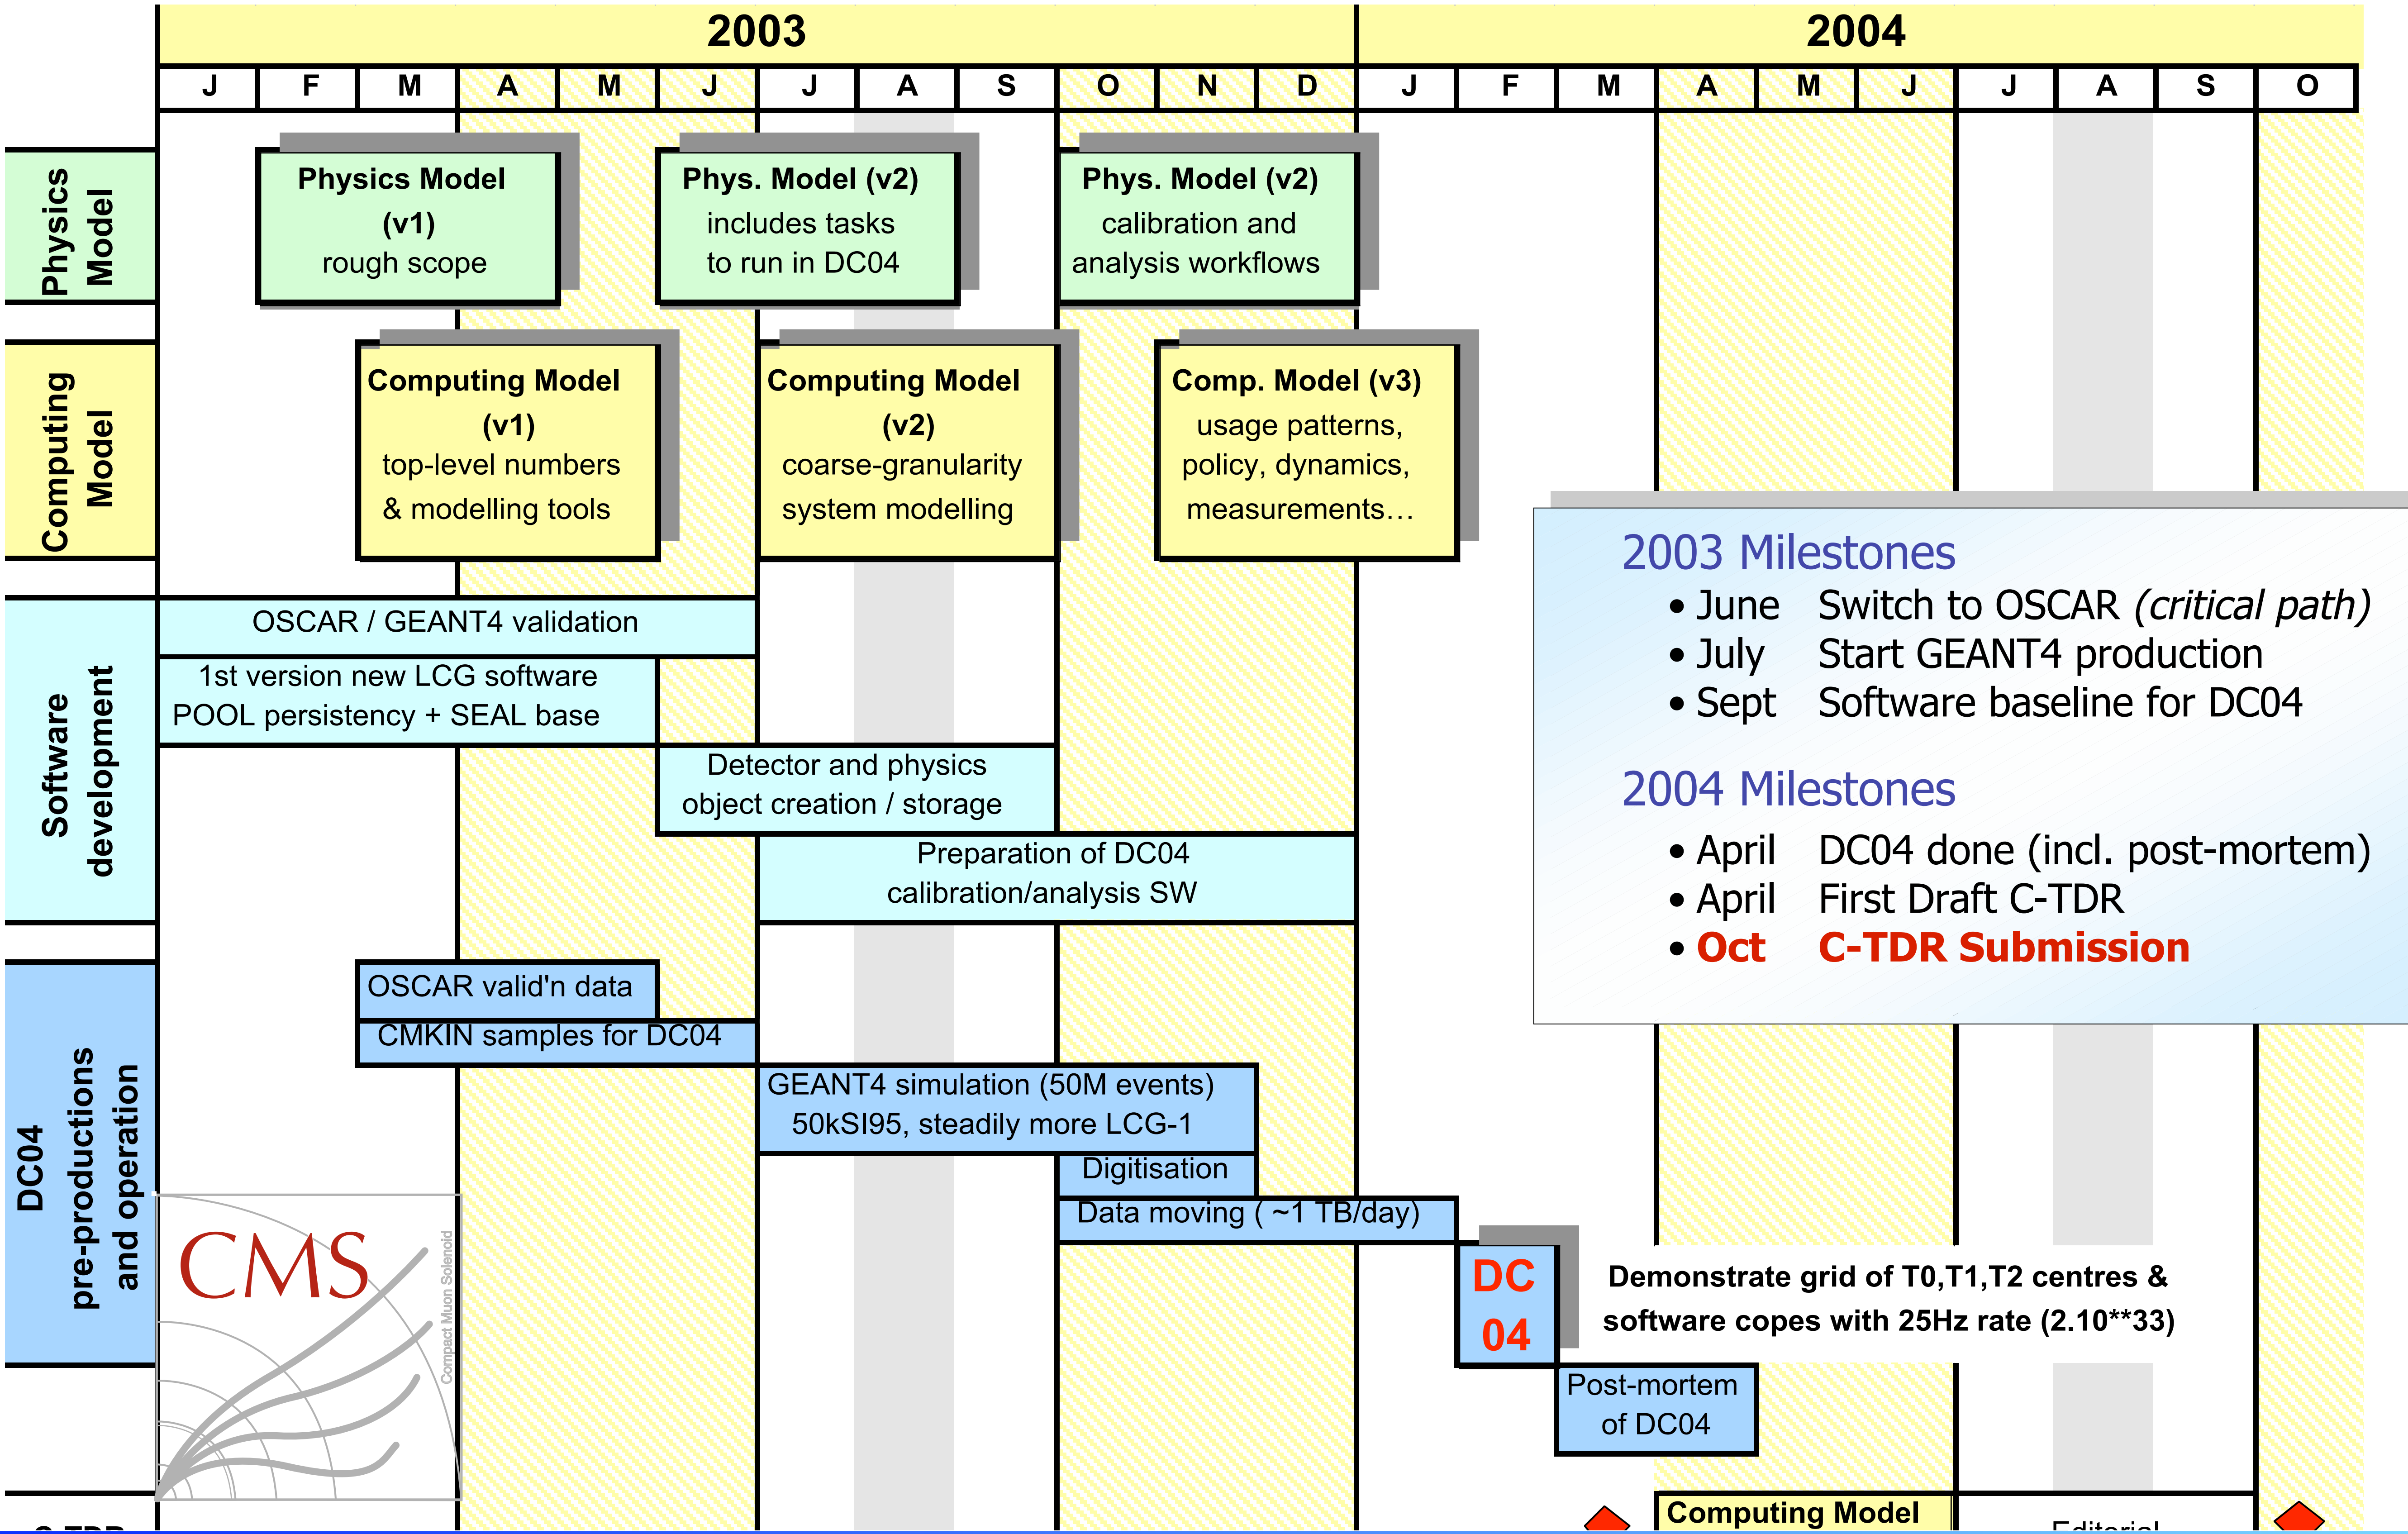

Plan for CMS DC04

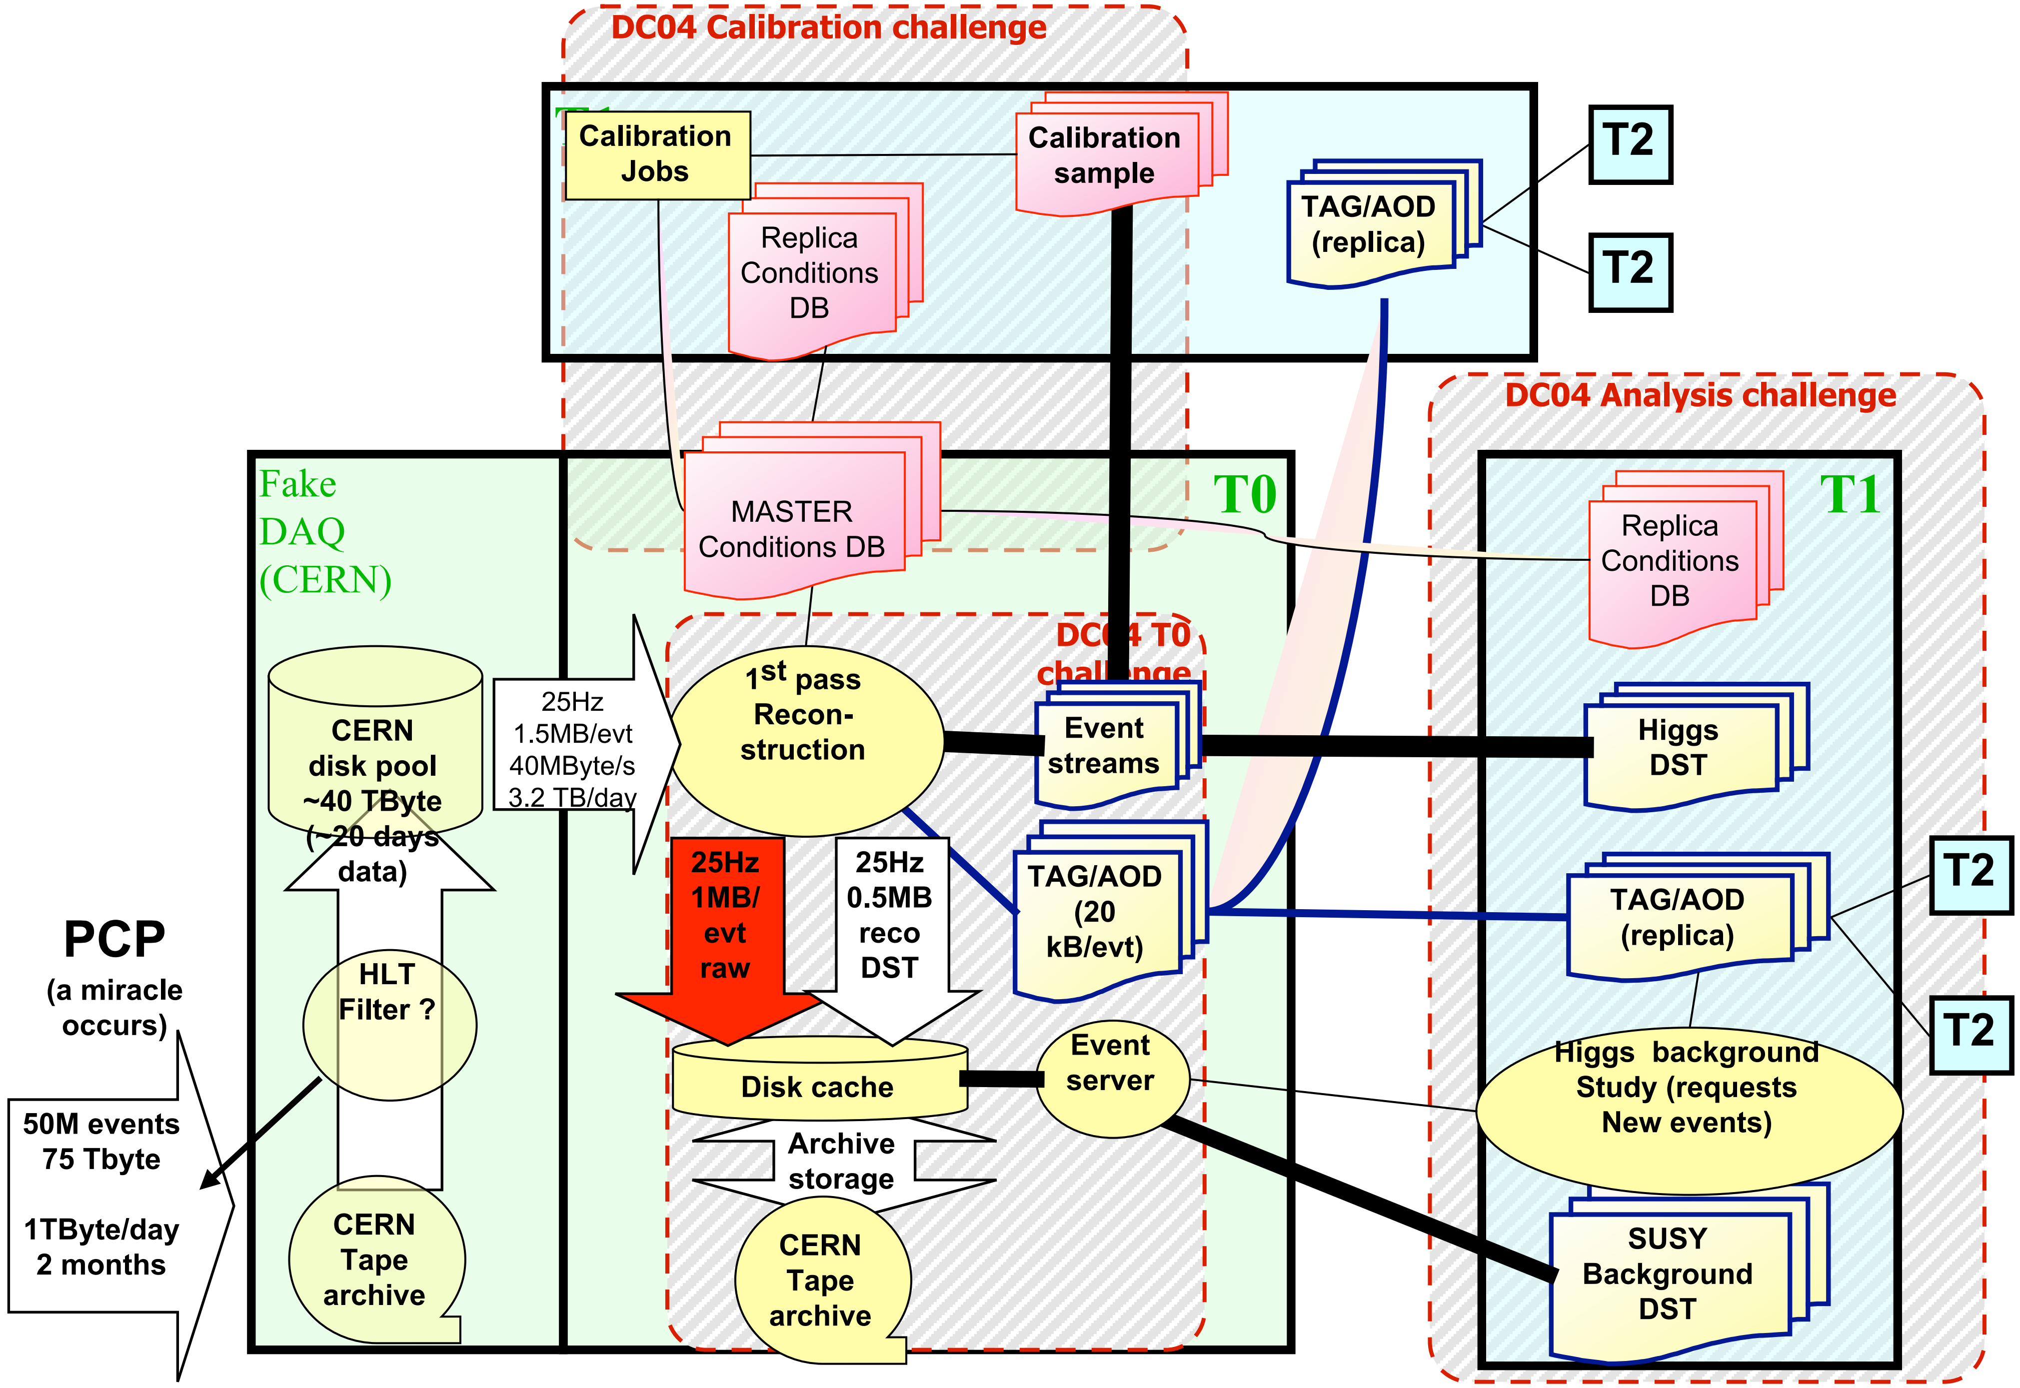

Getting CMS DC04 underway

- Generation already started,
  - Generator level pre-selections in place to enrich simulated background samples
  - Goal is 50 million "useful" events Simulated and Reconstructed
    - To fit scale of DC04
    - Also fits scale of first round of Physics TDR work
- Simulation will start in July
  - Probably some G4 and some CMSIM
  - Some samples simulated with both
  - Expect G4 version to go through a few versions during production
  - Mix and choice of G4/CMSIM not yet determined
  - Production rate will be about 1 TB/day
- Simulated Data kept at Tier 1 centers
  - Reconstructed data sent to CERN for DC04 proper (Spring 2004)

LHCb — DIRAC Architecture

Production Service keeps list of "desired" simulations,
SW agents discover and pull assignments

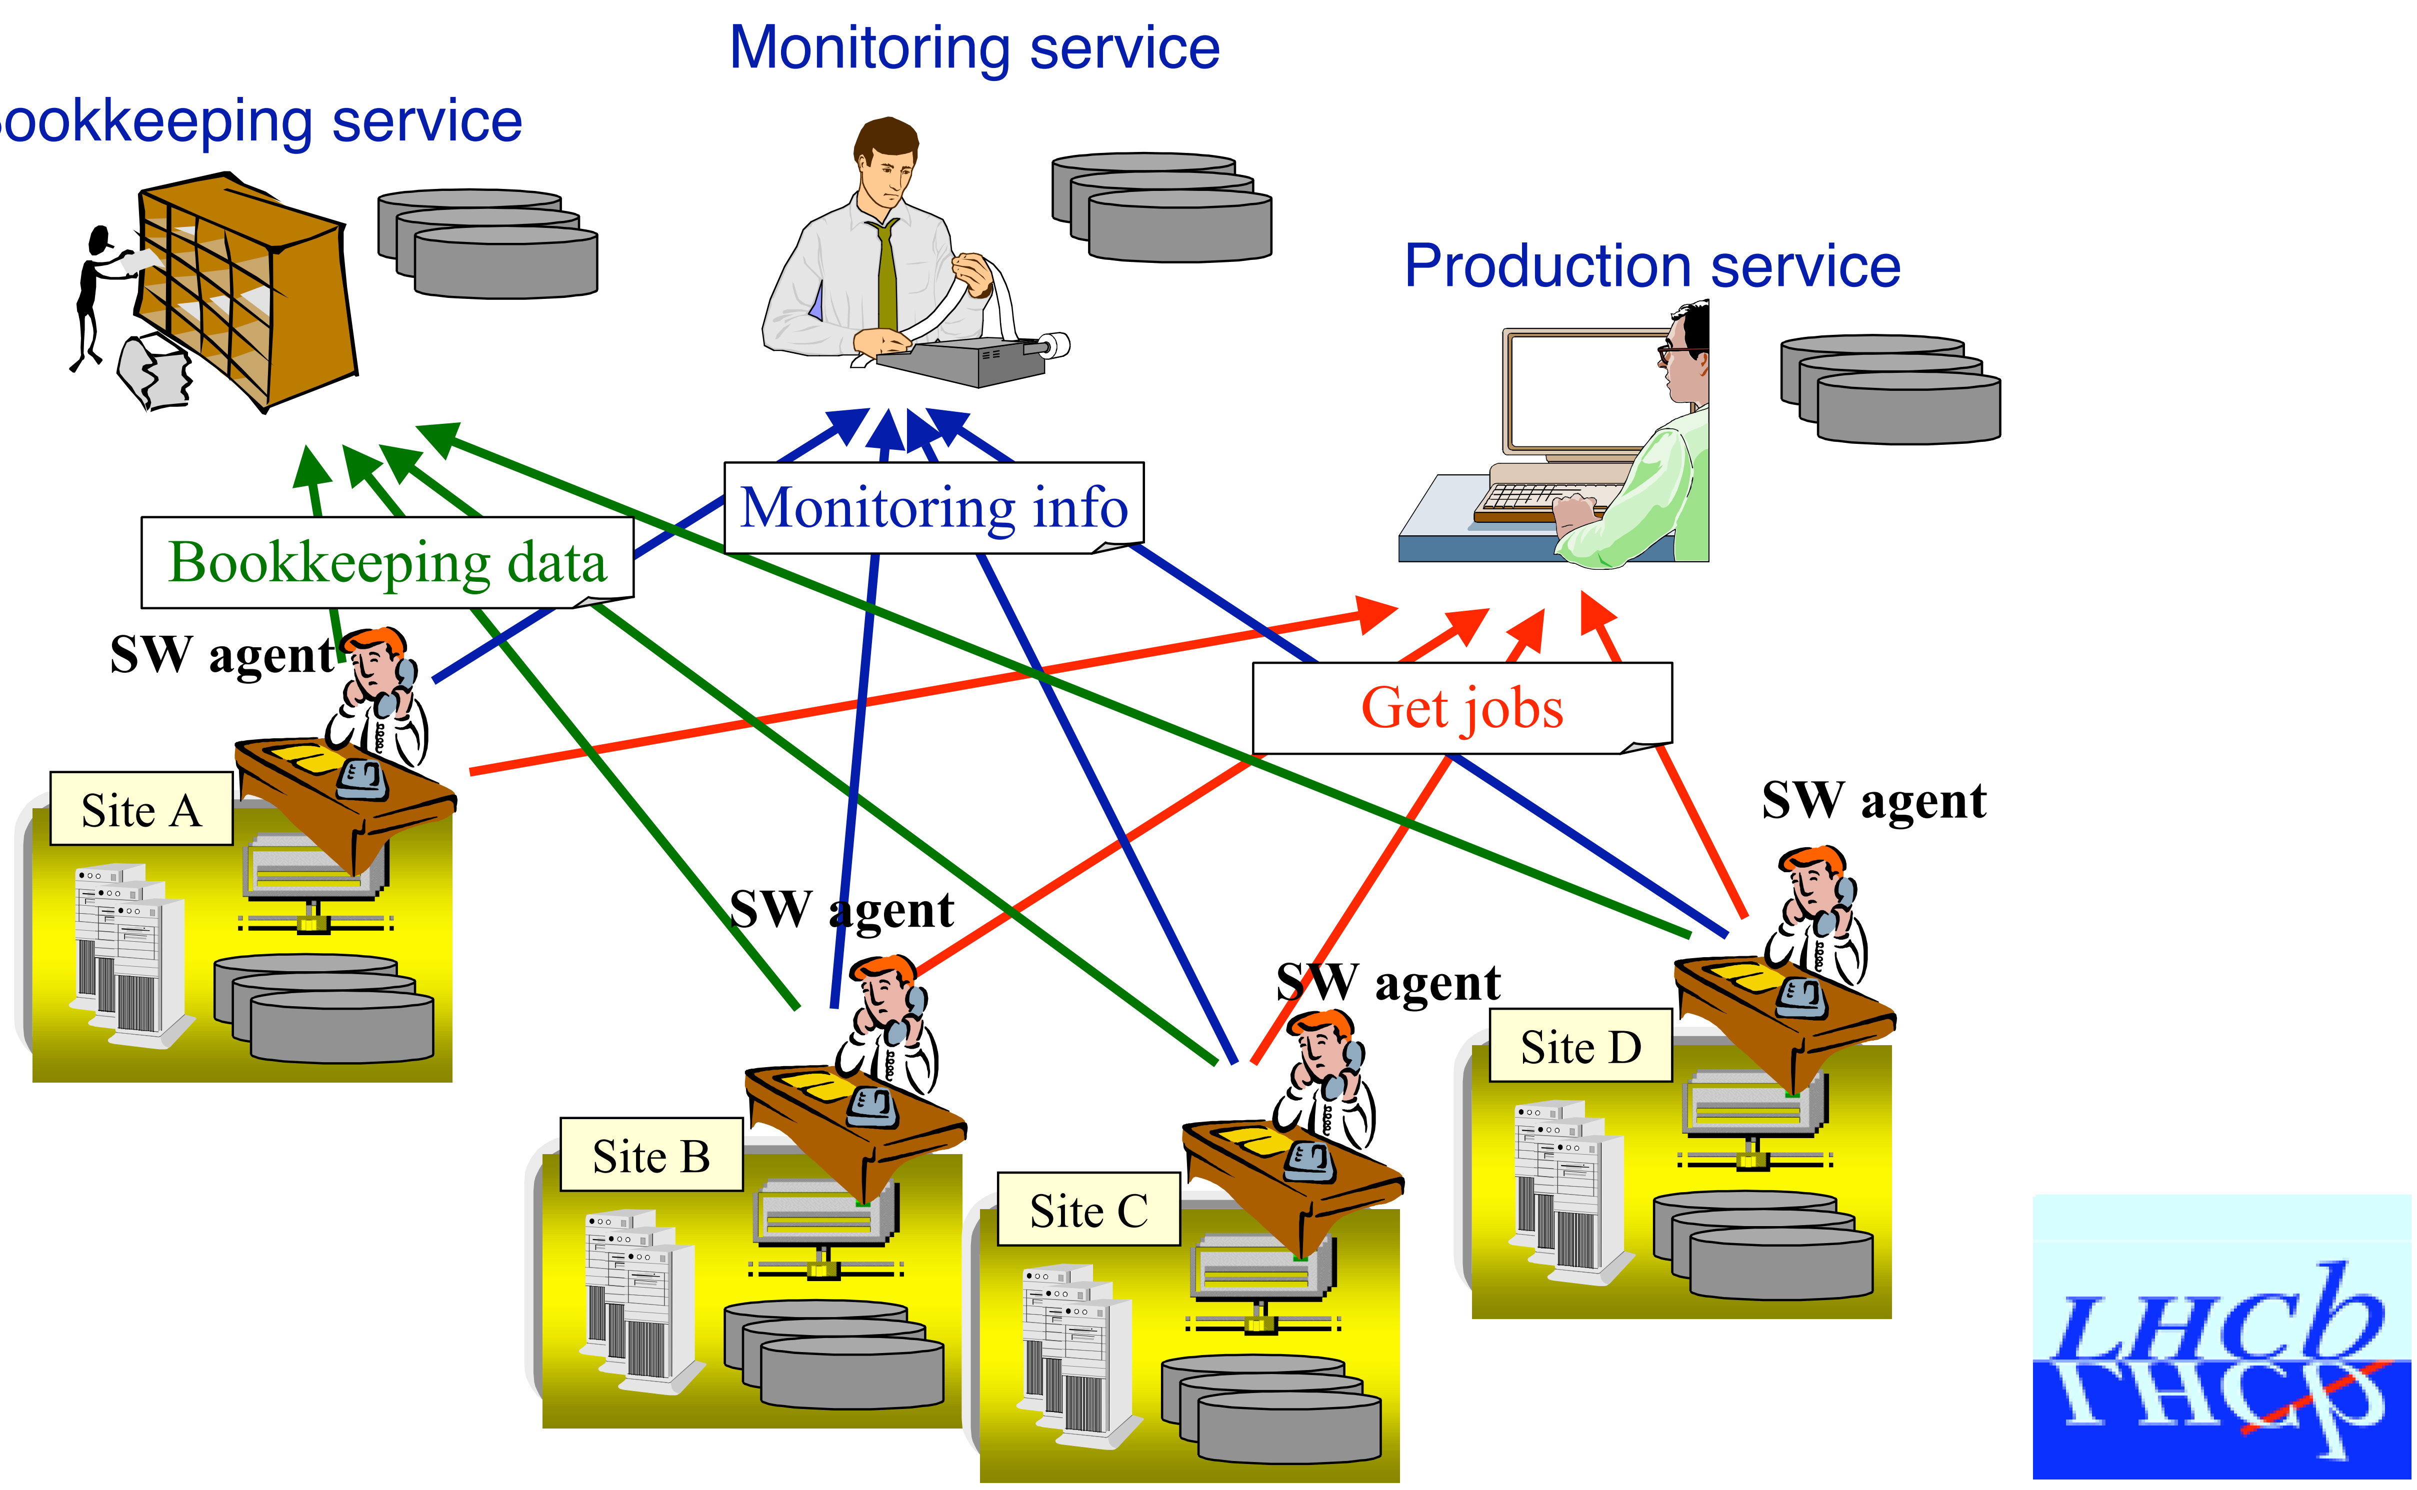

LHCb — DIRAC on the DataGRID

moved the "SW Agents" to Worker Nodes on the EU DataGRID Testbed

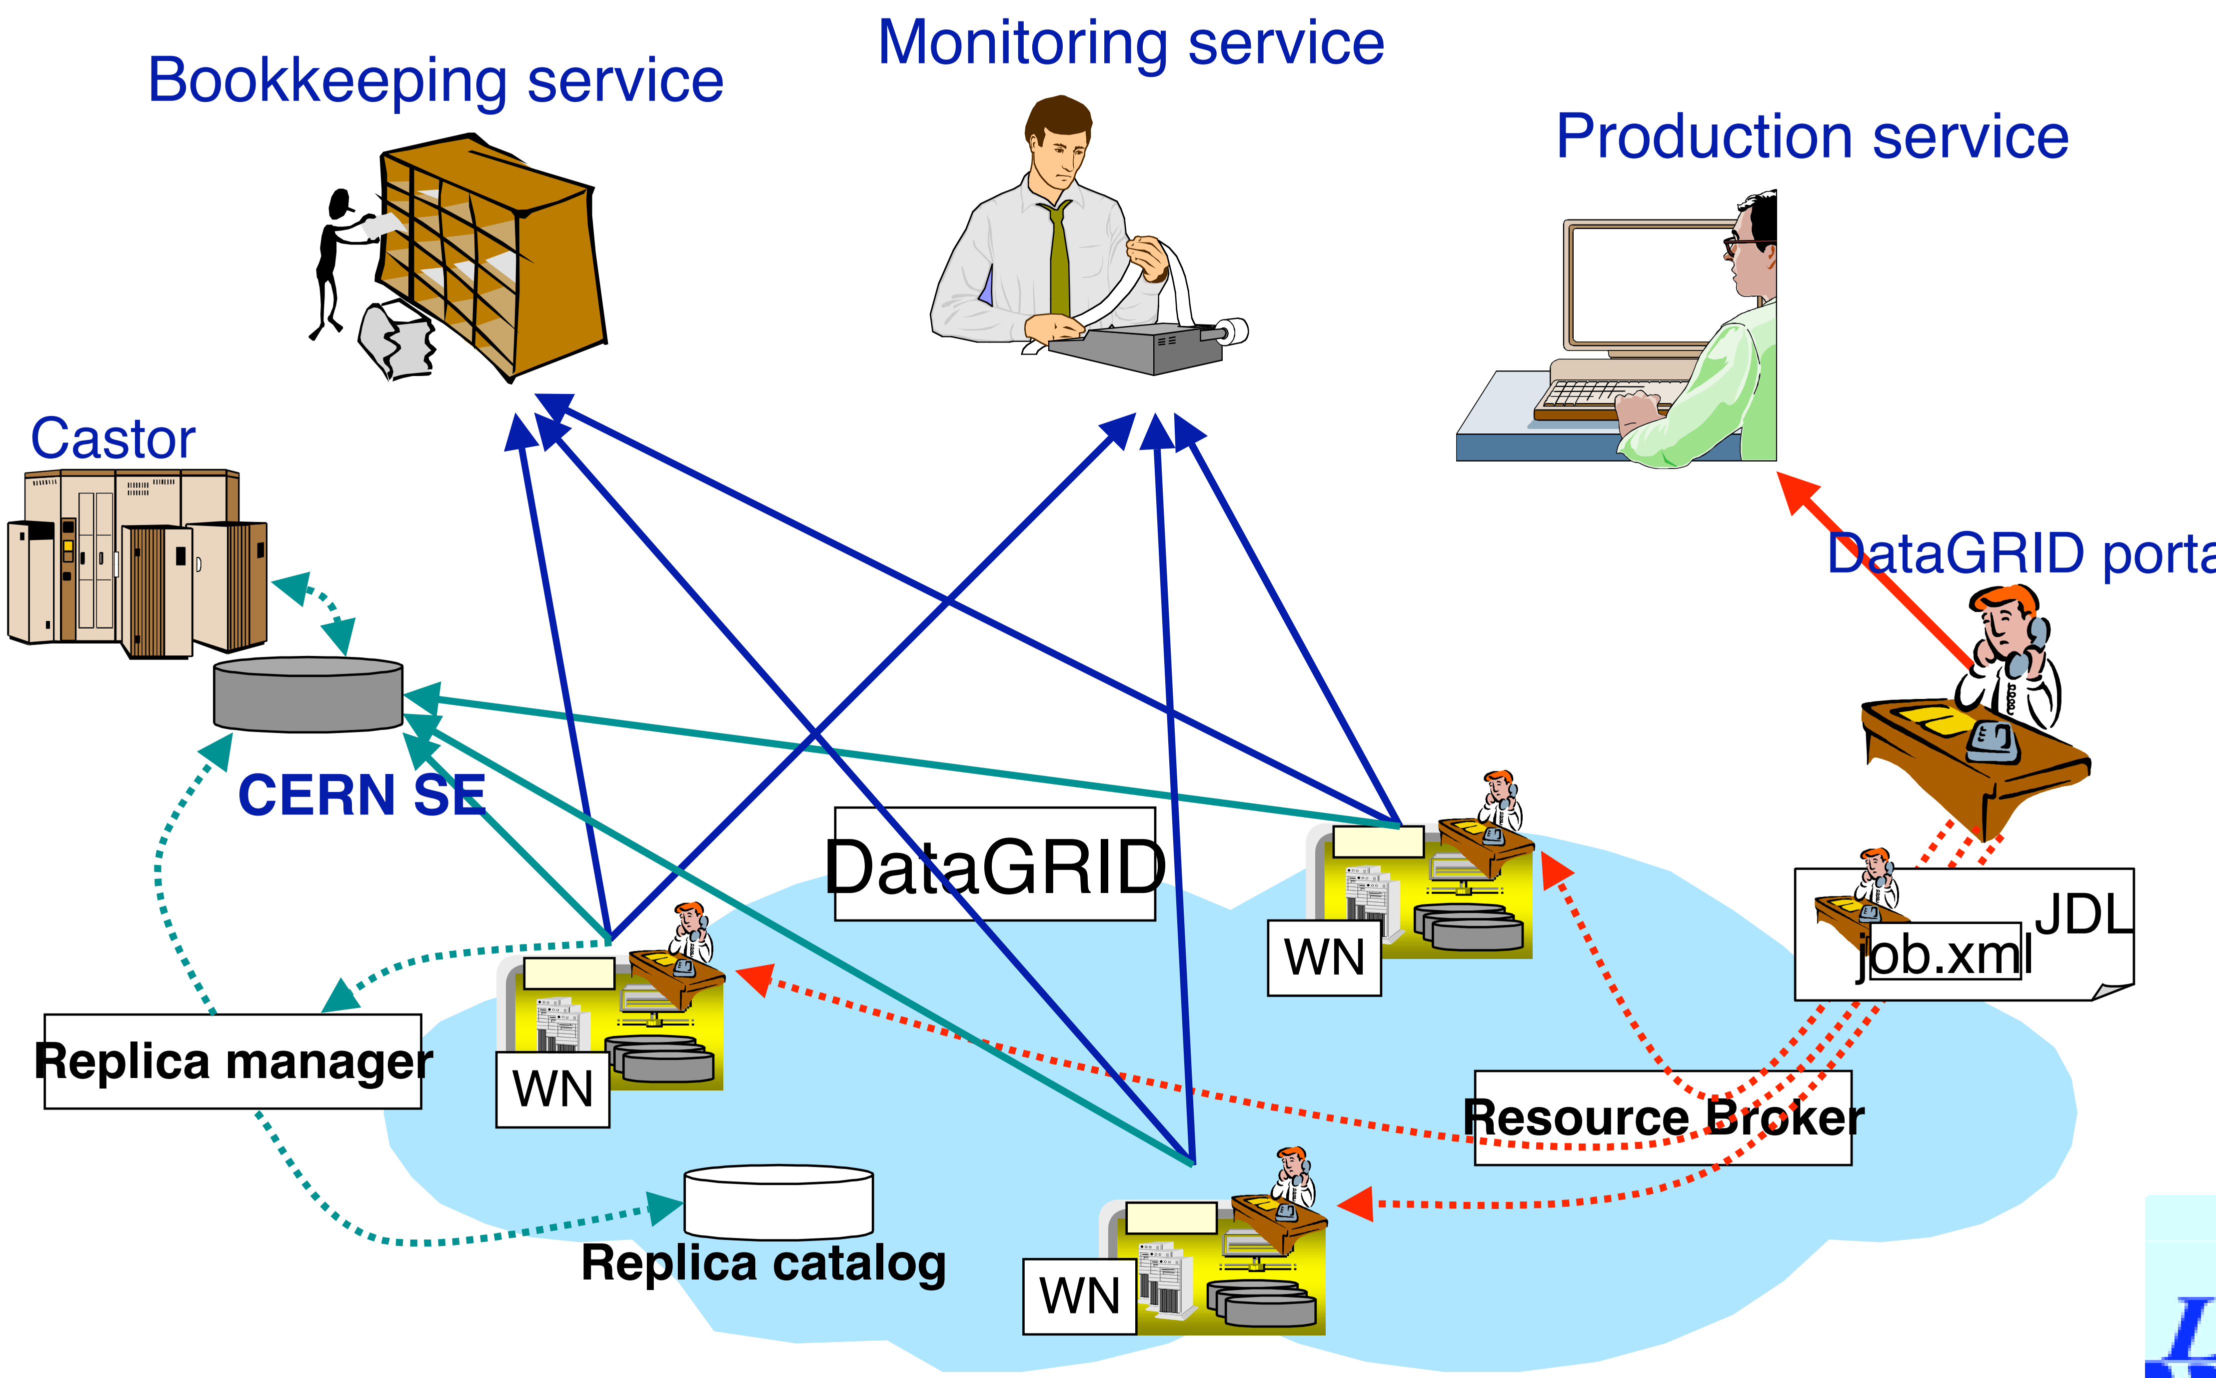

Grid Testbeds and Technology Cycles

"Rolling Prototypes": evolution of the facility and data systems

Prototyping and early roll out to production quality services

Participation in Computing and Physics **Data Challenges**

Emphasis on Quality, Documentation, Dissemination,

Tracking of external "practices"

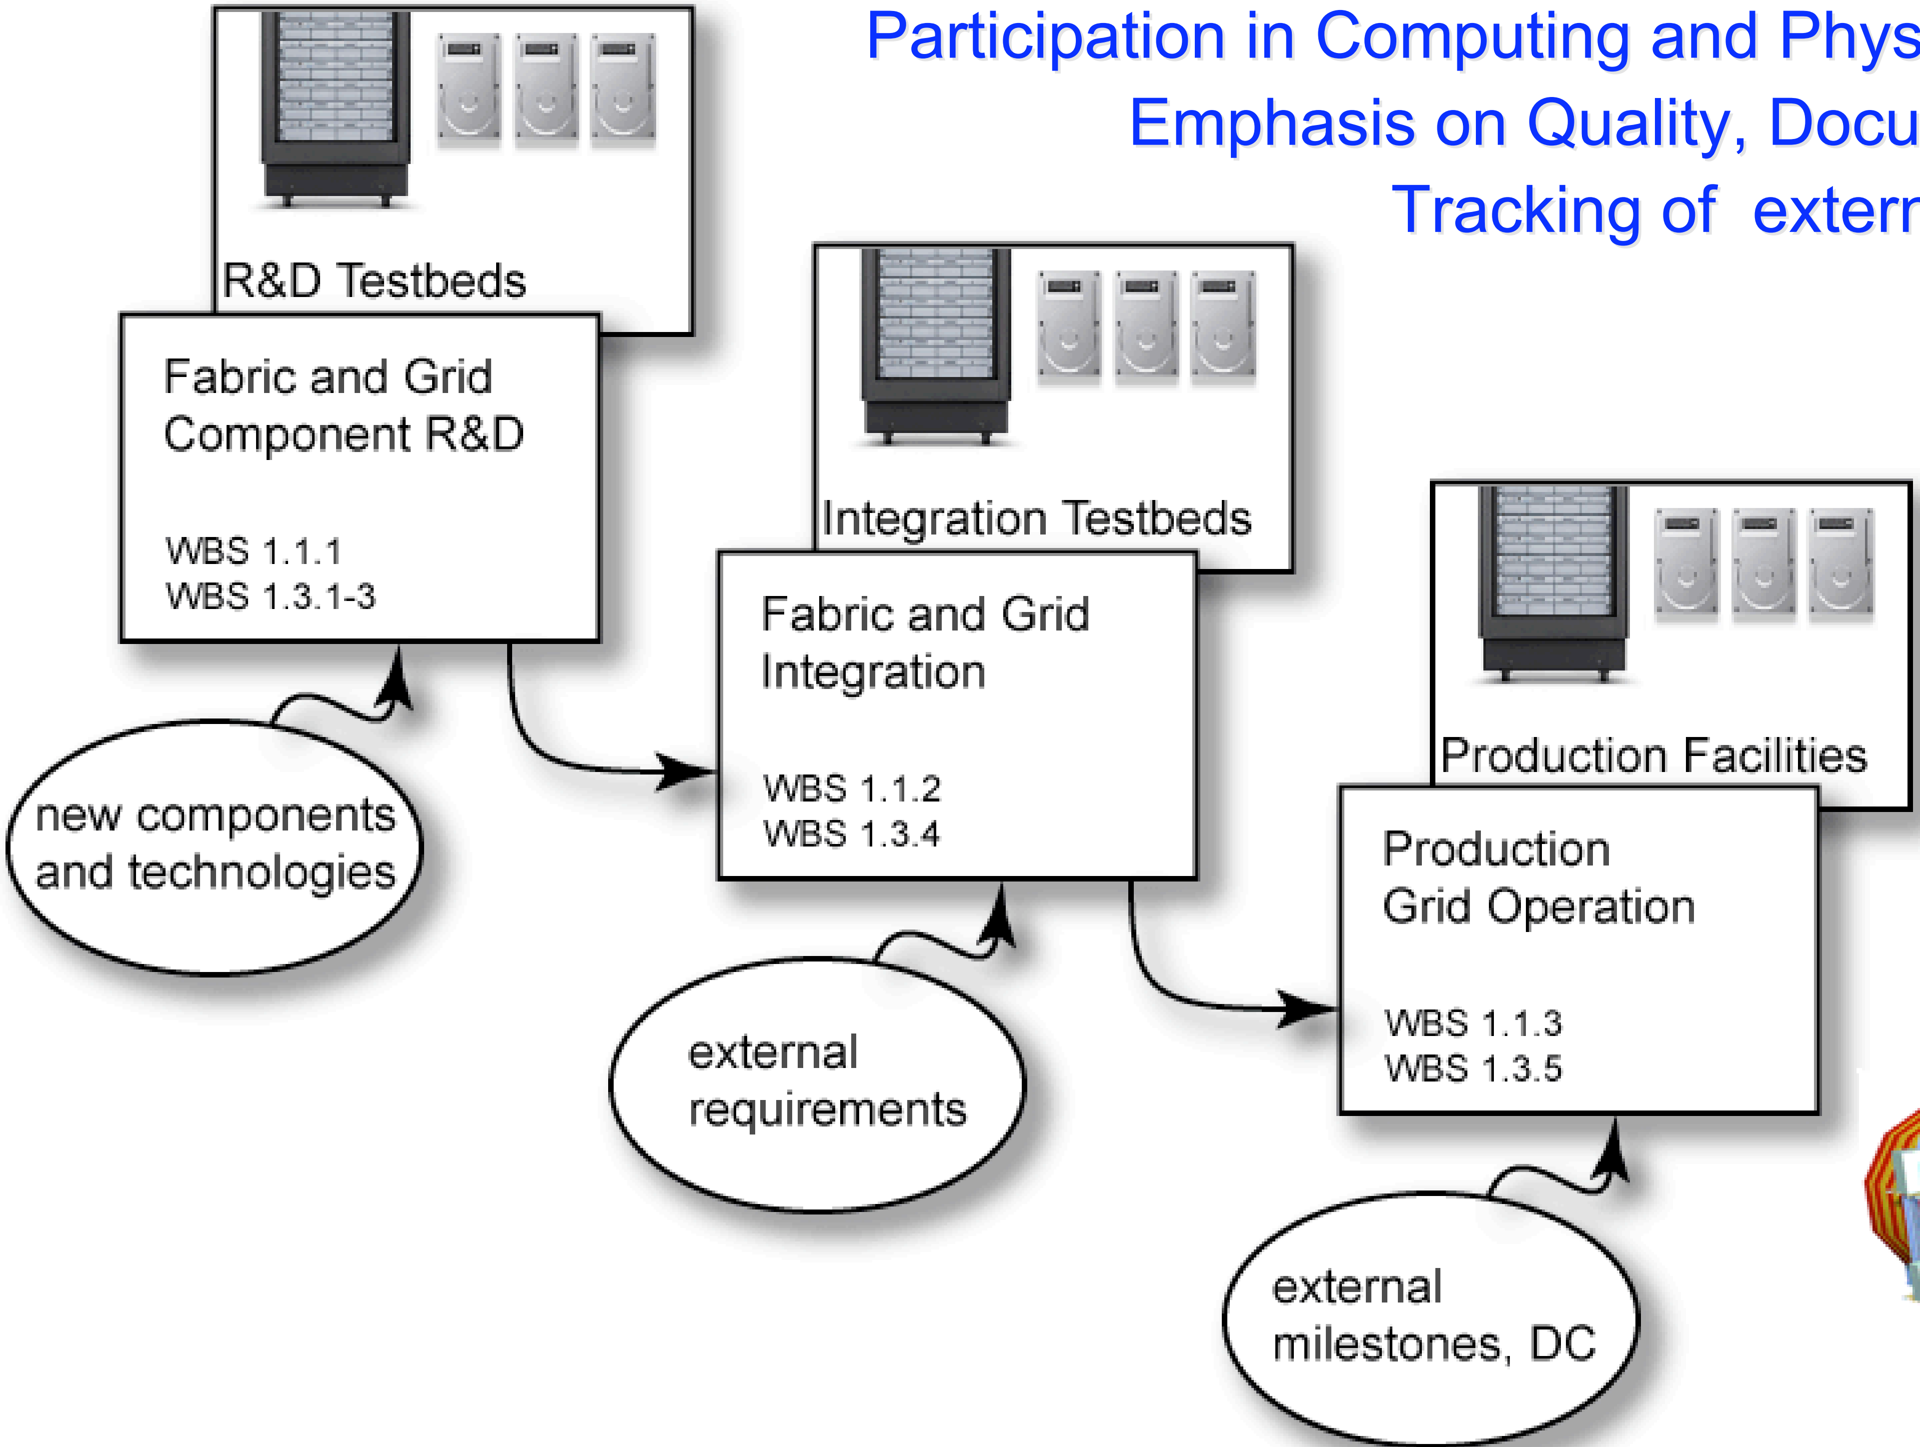

Grid Testbeds to develop Production Grids

Grid Testbeds for Research, Development and Dissemination!

- USCMS Testbeds real-life large Grid installations, becoming production quality
- Strong Partnership between Labs, Universities, with Grid (iVDGL, GriPhyN, PPDG) and Middleware Projects (Condor, Globus)
- Strong dissemination component, together with Grid Projects
- Caltech, UCSD, U.Florida, UW Madison, Fermilab, CERN

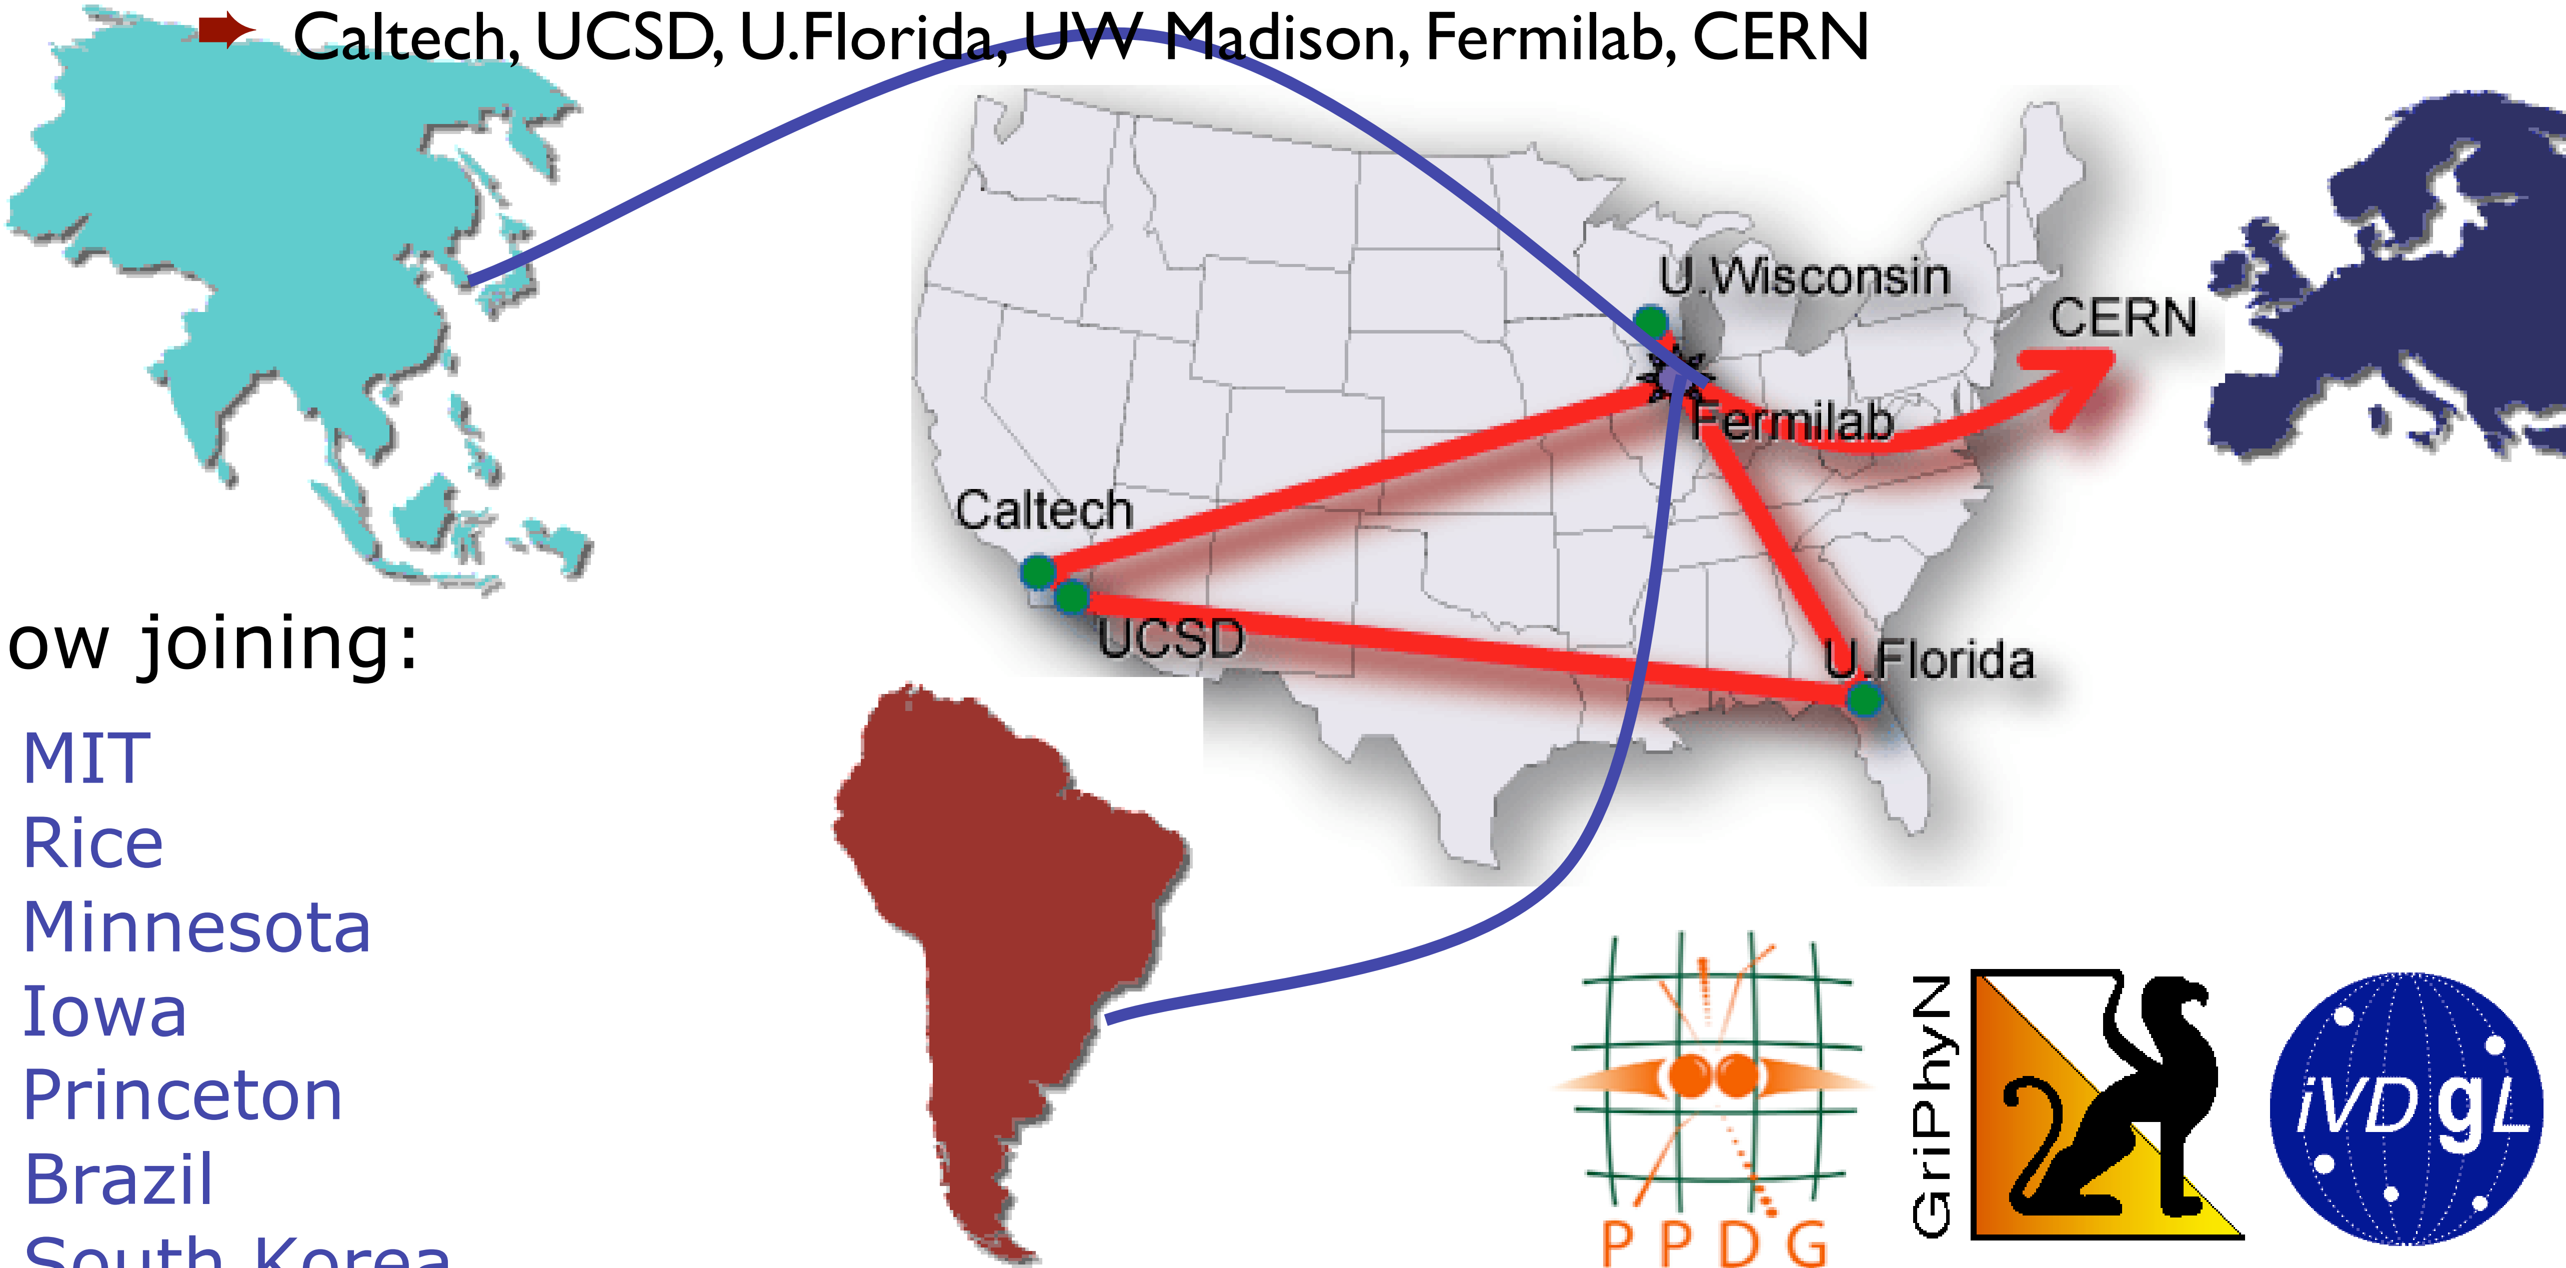

- Now joining:
  - MIT
  - Rice
  - Minnesota
  - Iowa
  - Princeton
  - Brazil
  - South Korea

Technology Example: Grid Monitoring and Information Services

- MonALISA Monitoring System (Caltech) deployed in U.S. & CERN Grid Testbed
- Dynamic information services and Grid resource discovery mechanisms using "intelligent agents"
  - Use and further develop novel Grid Technologies and Grid Interfaces
- "Grid Control Room" For LHC Grid
- Technology driver for other projects

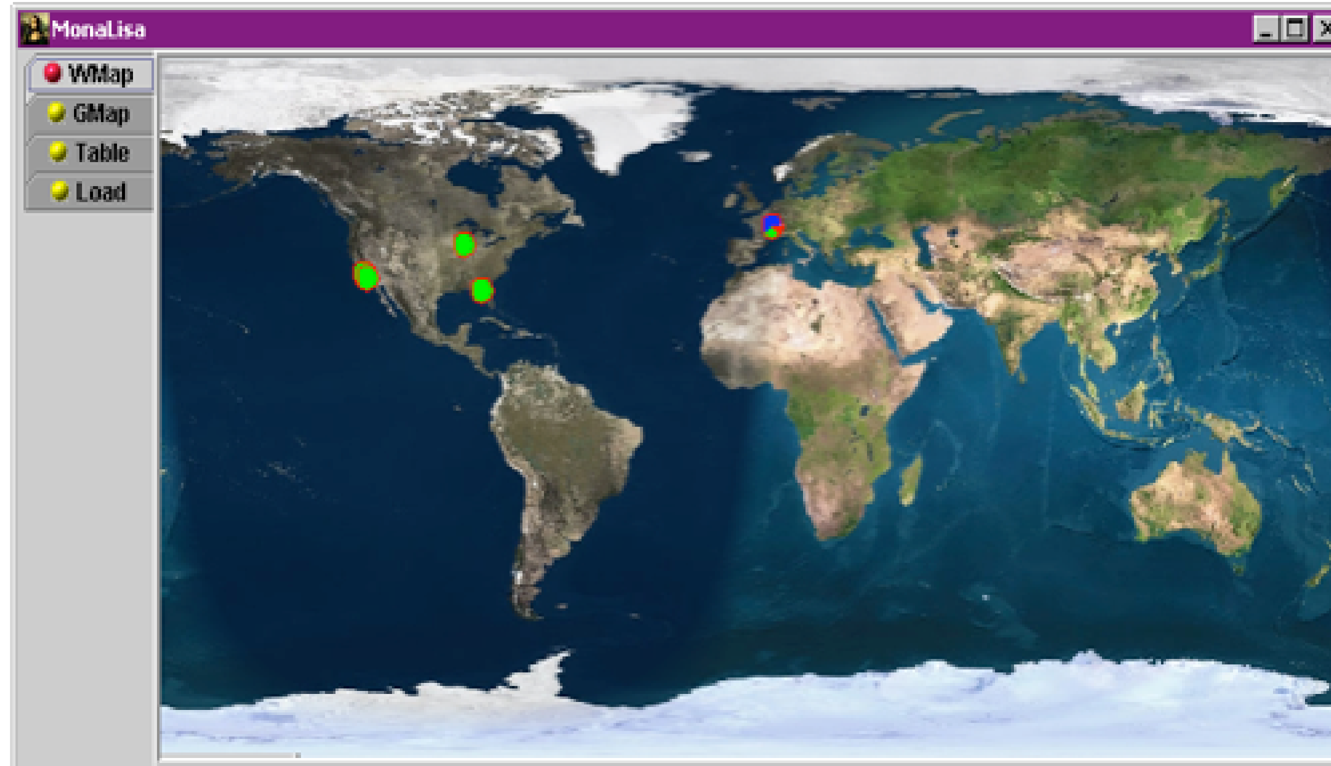

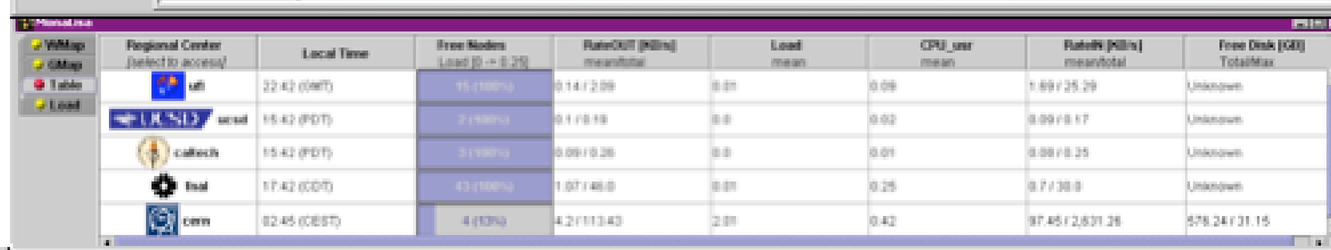

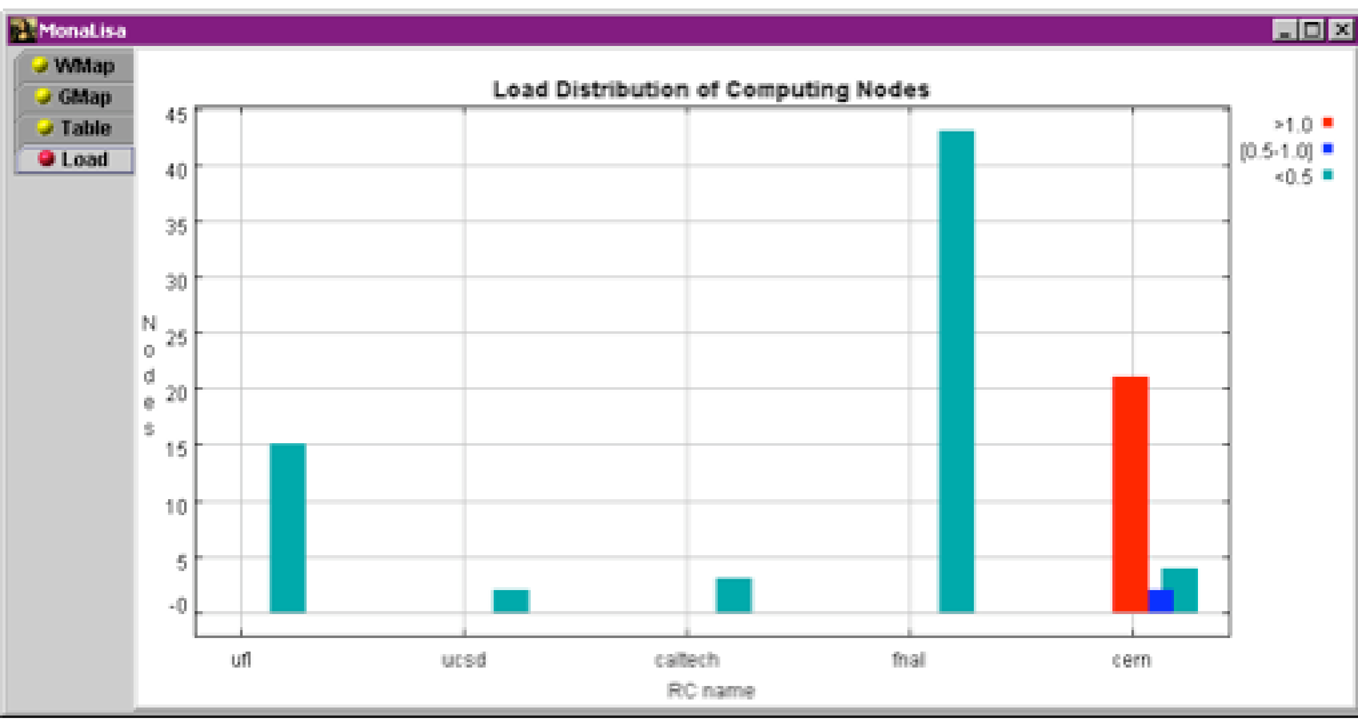

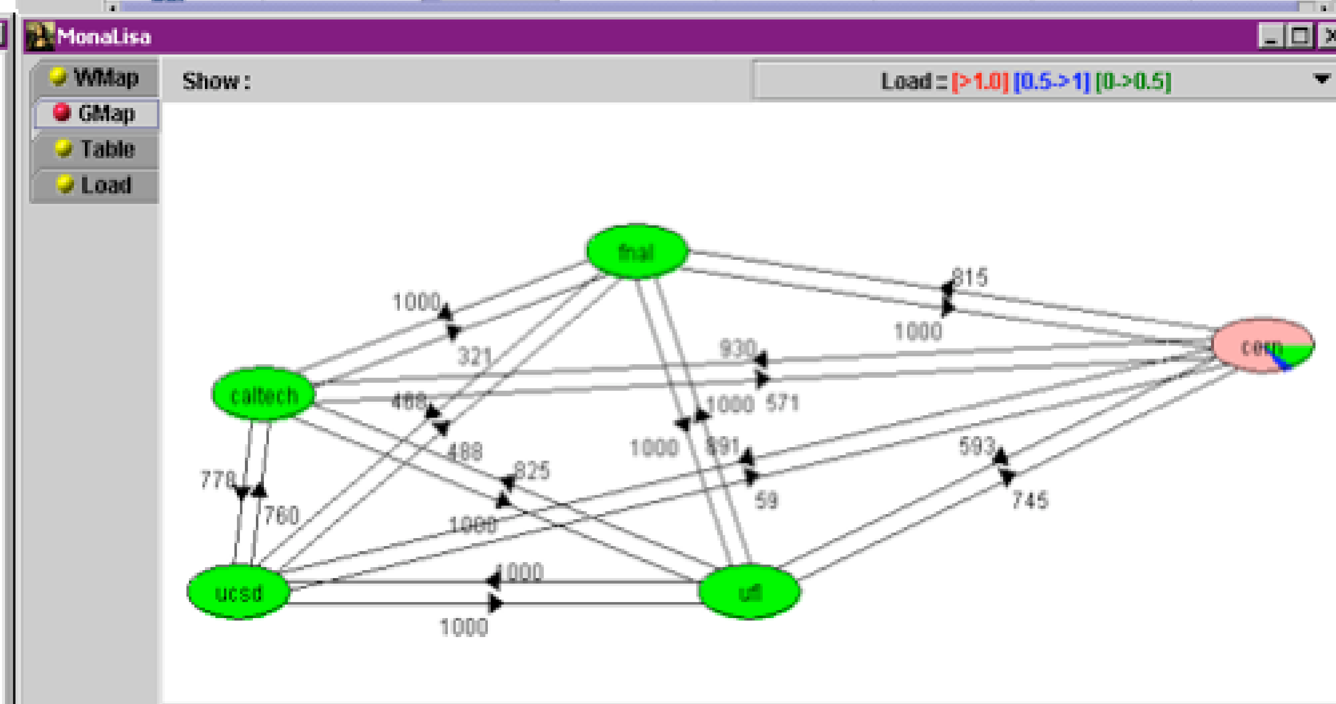

LHC Multi-Tier Structured Computing Resources

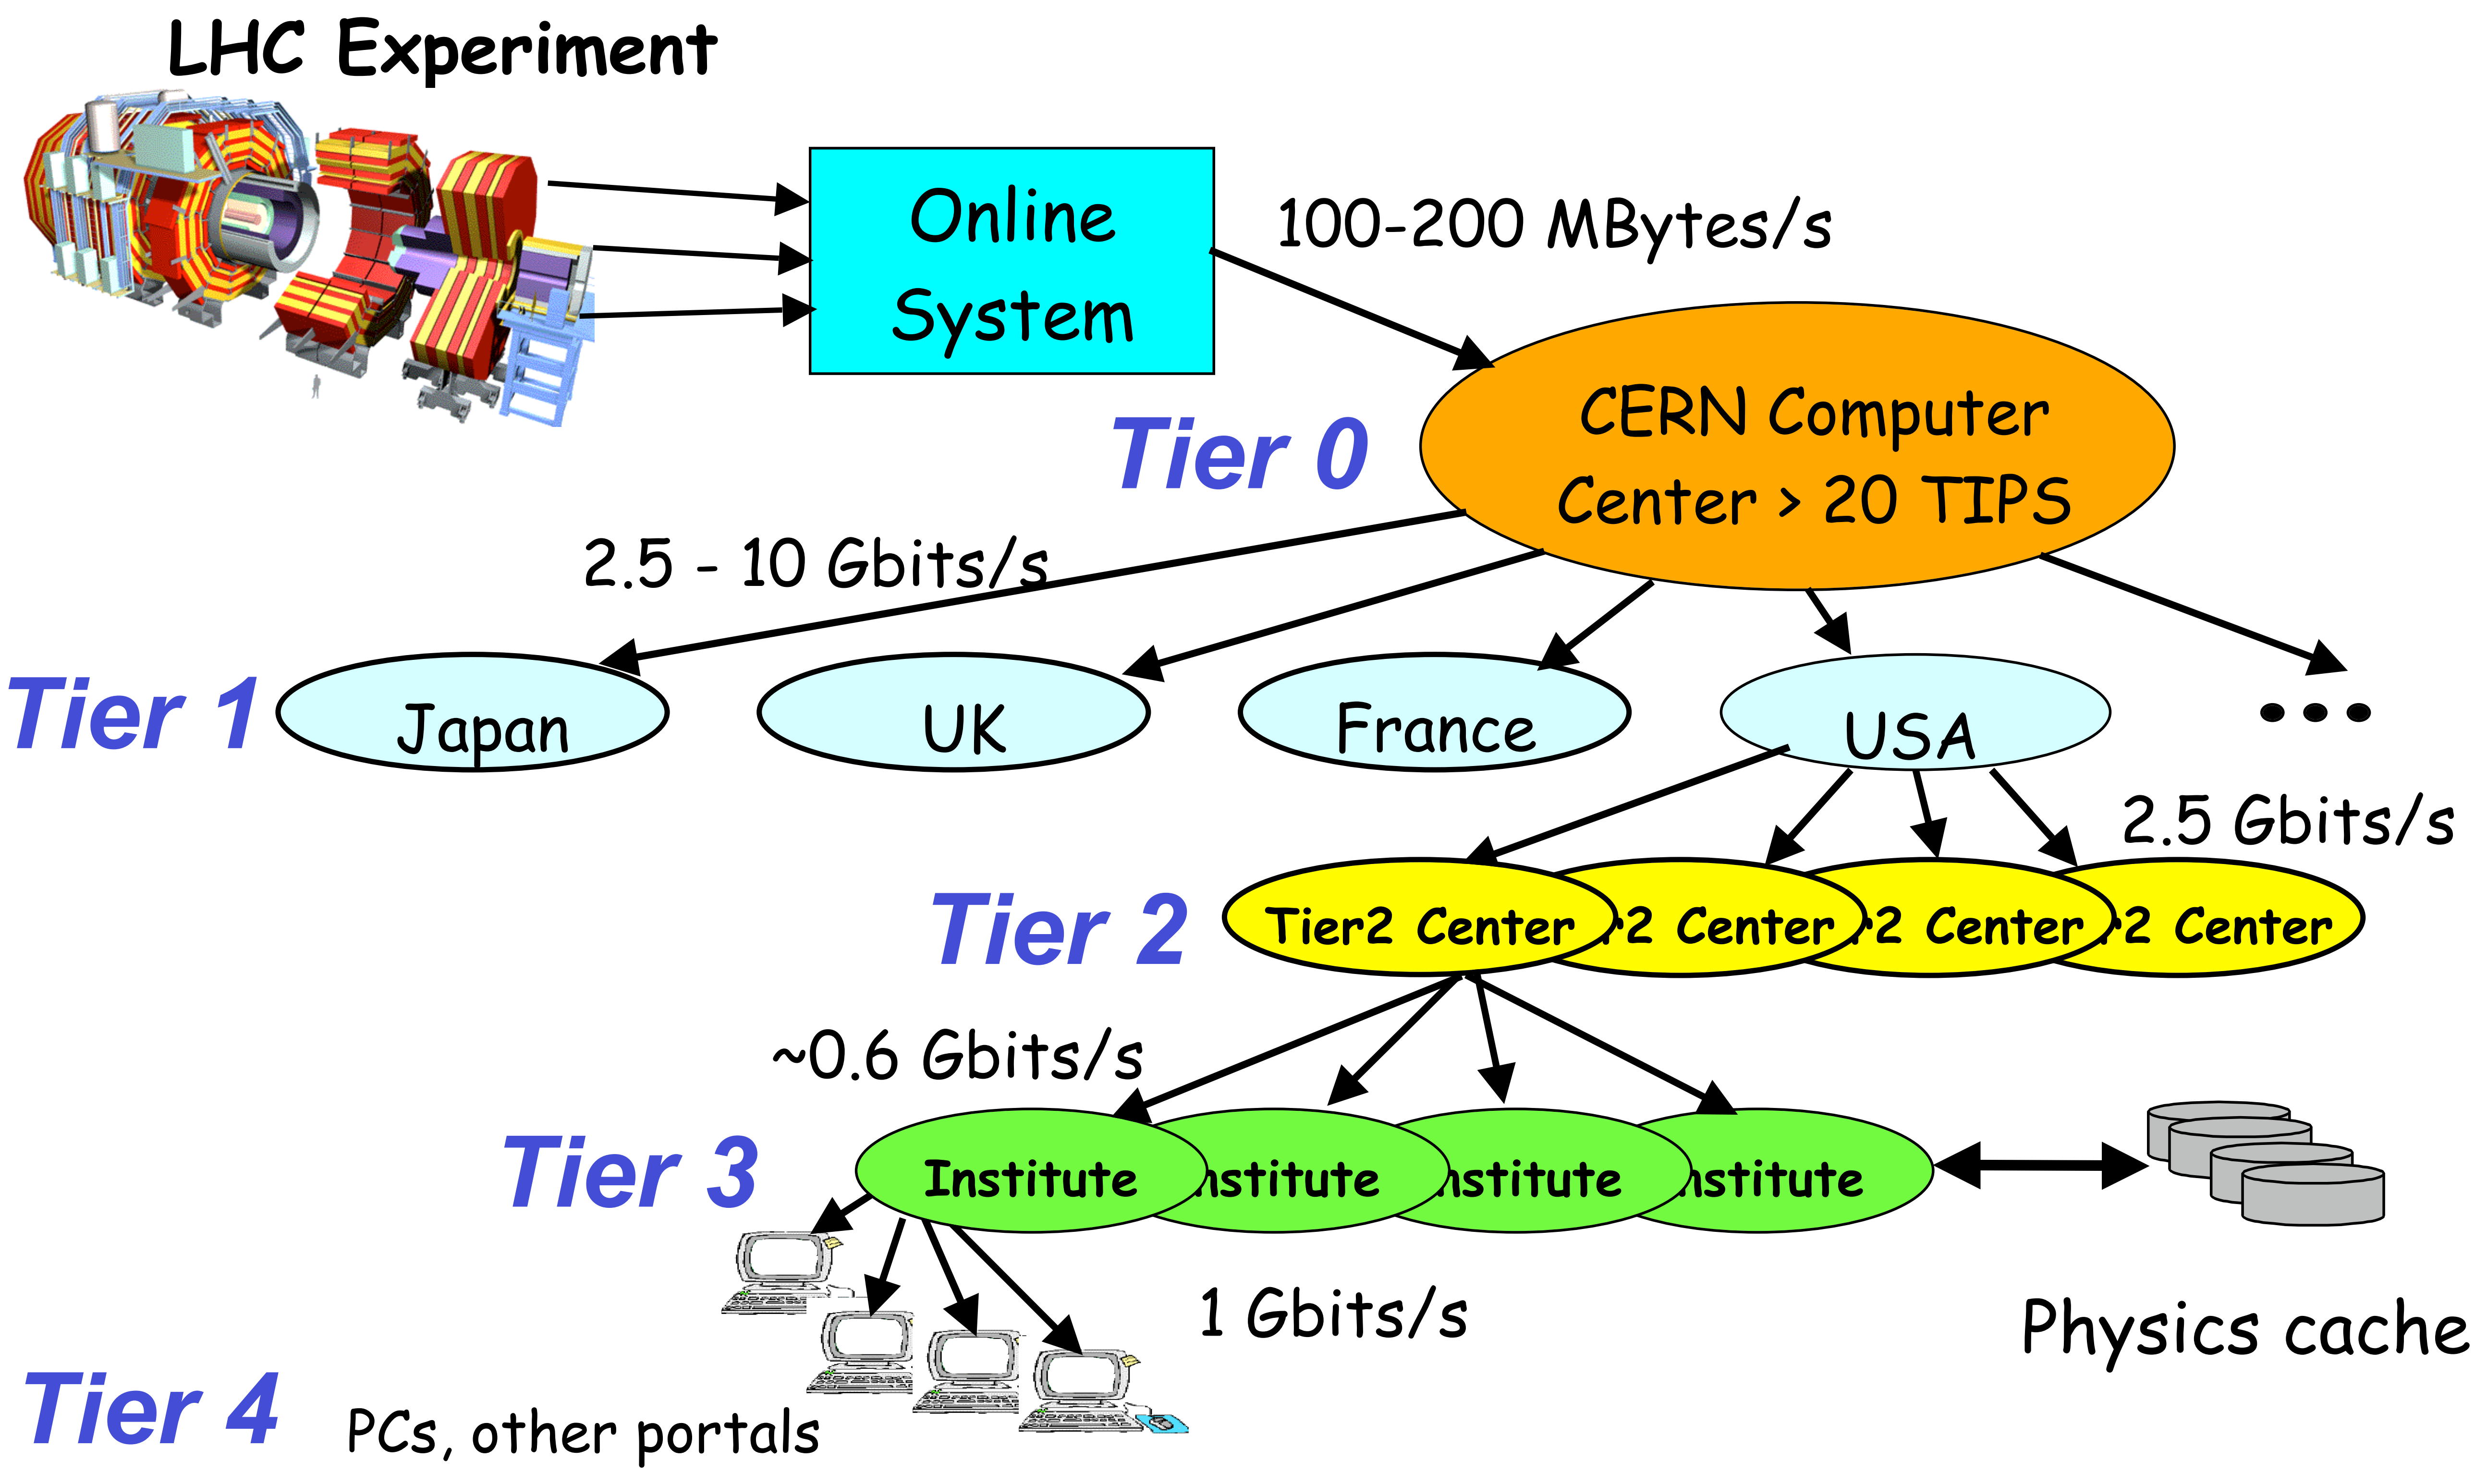

LHC Multi-Tier Structured Computing Resources

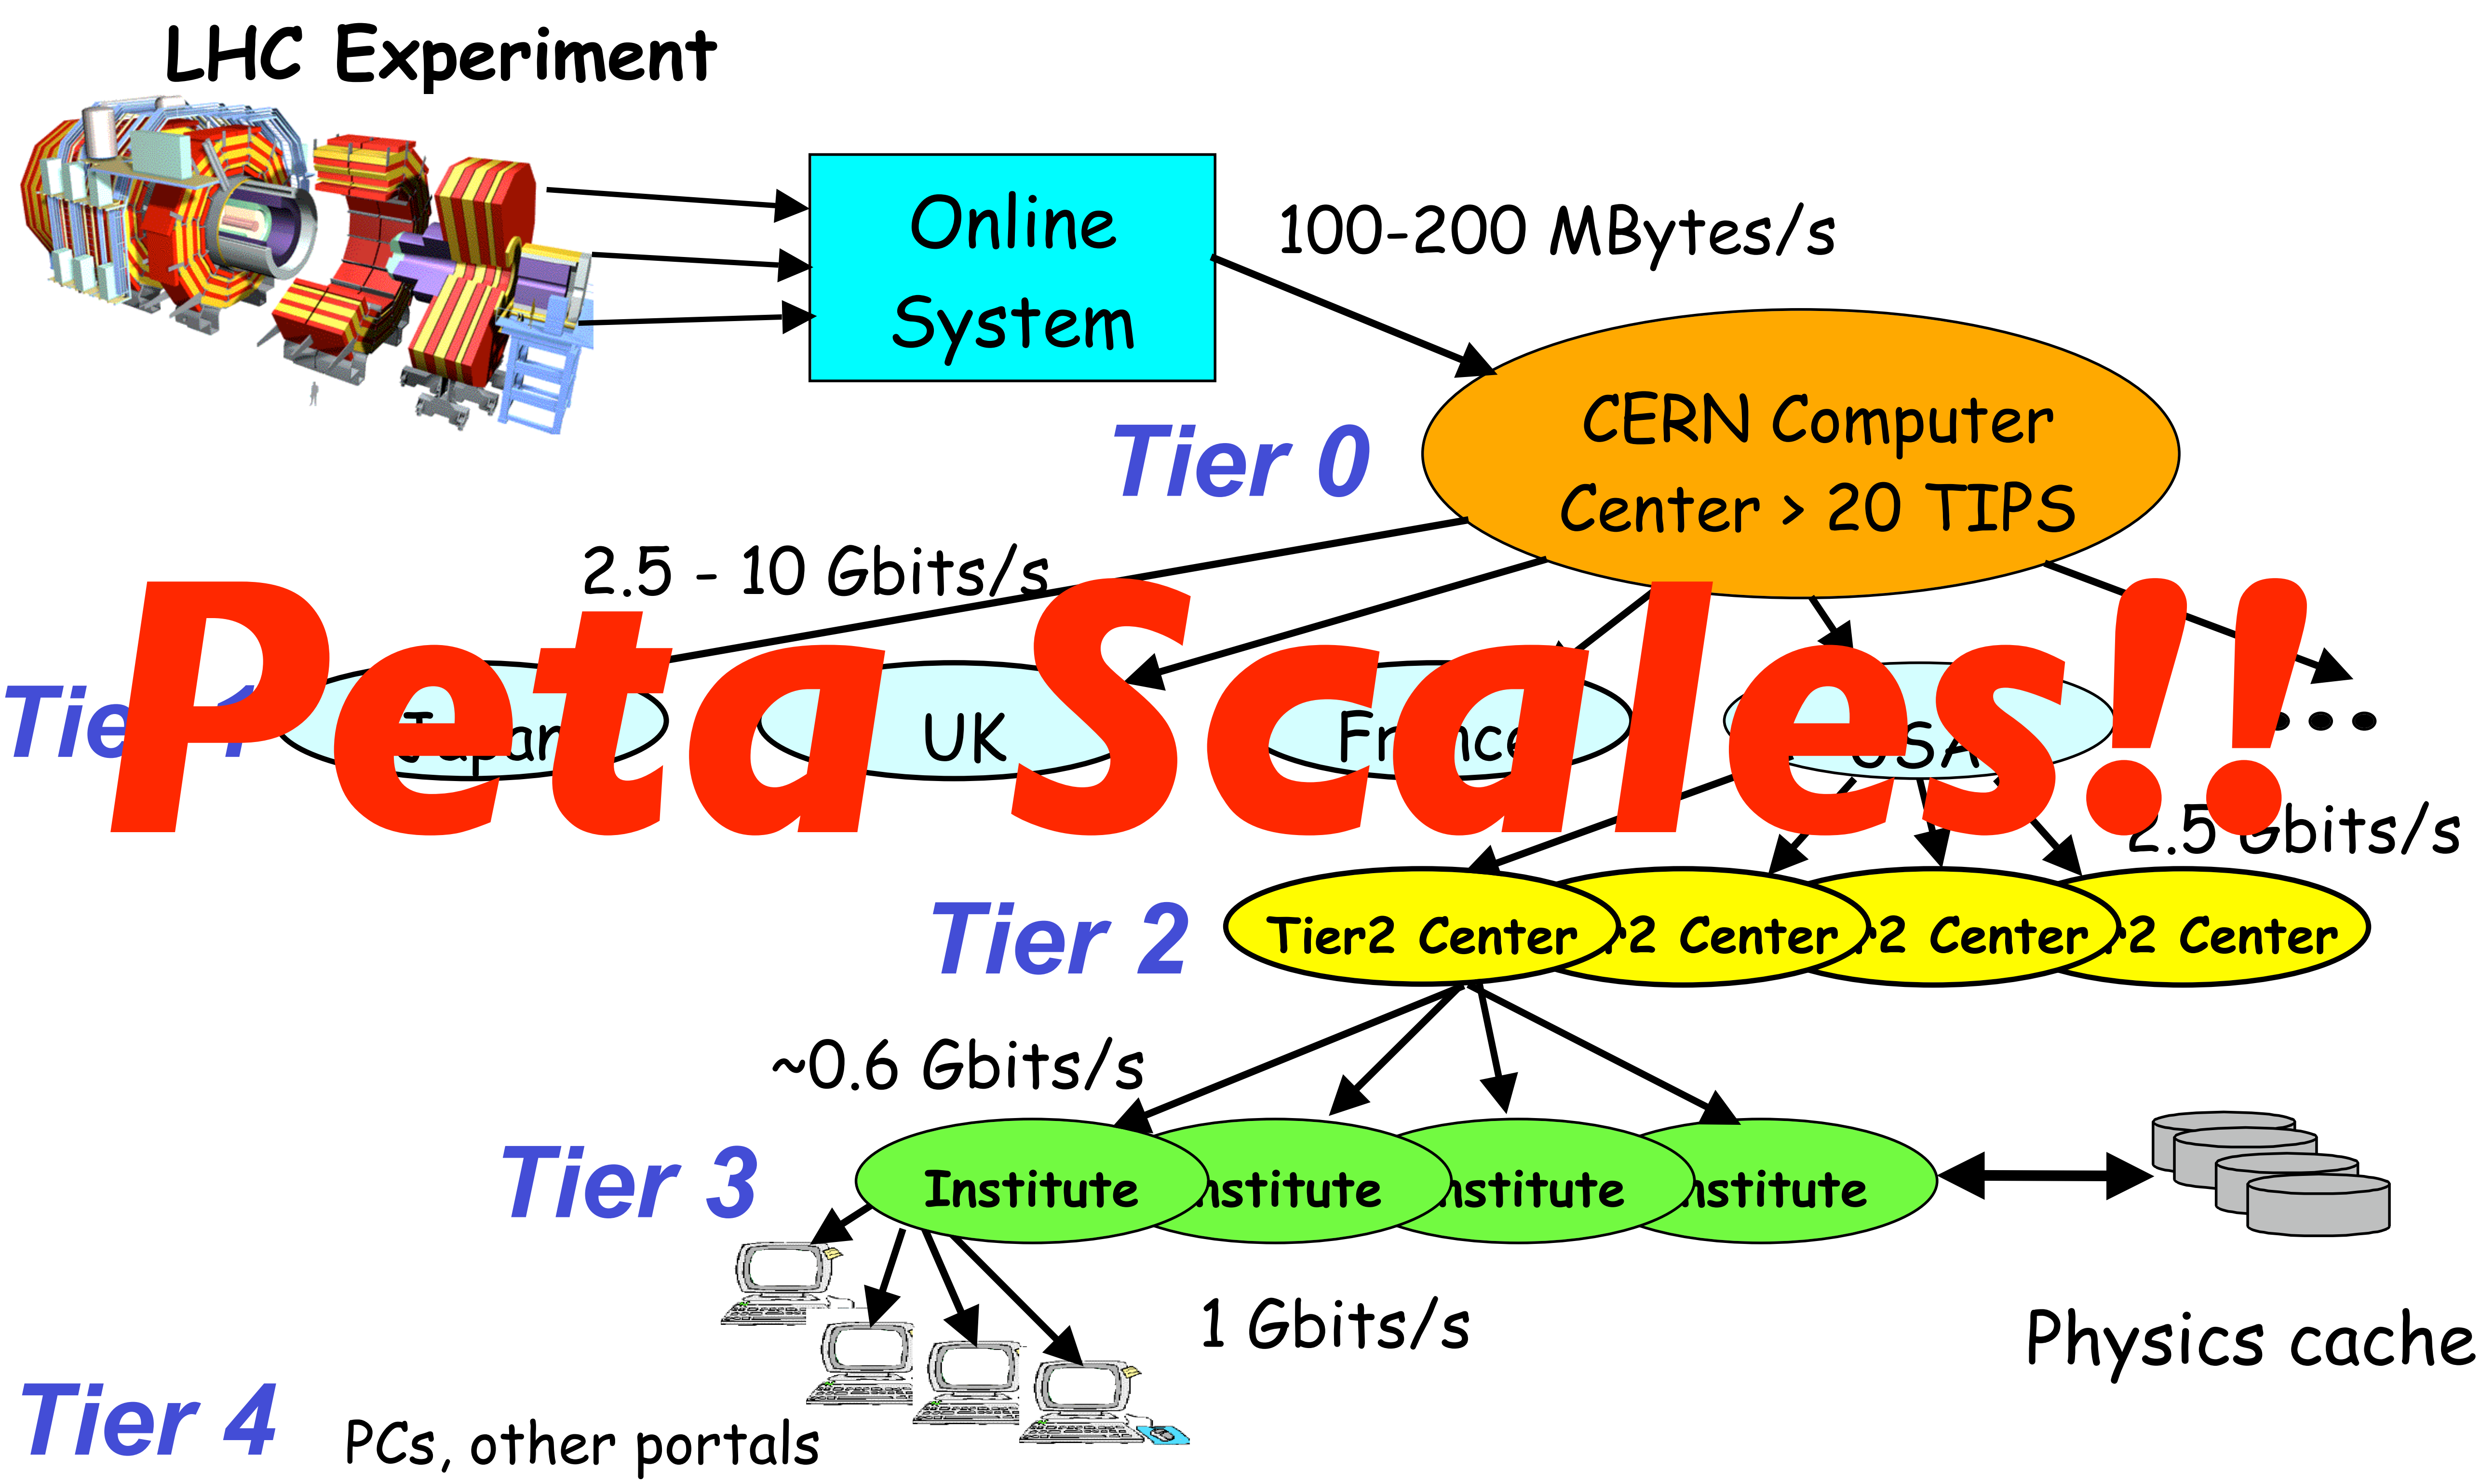

Scientists within Dynamic Workspaces!

Communities of Scientists Working Locally within a Global Context

Infrastructure for sharing, consistency of physics and calibration data, software

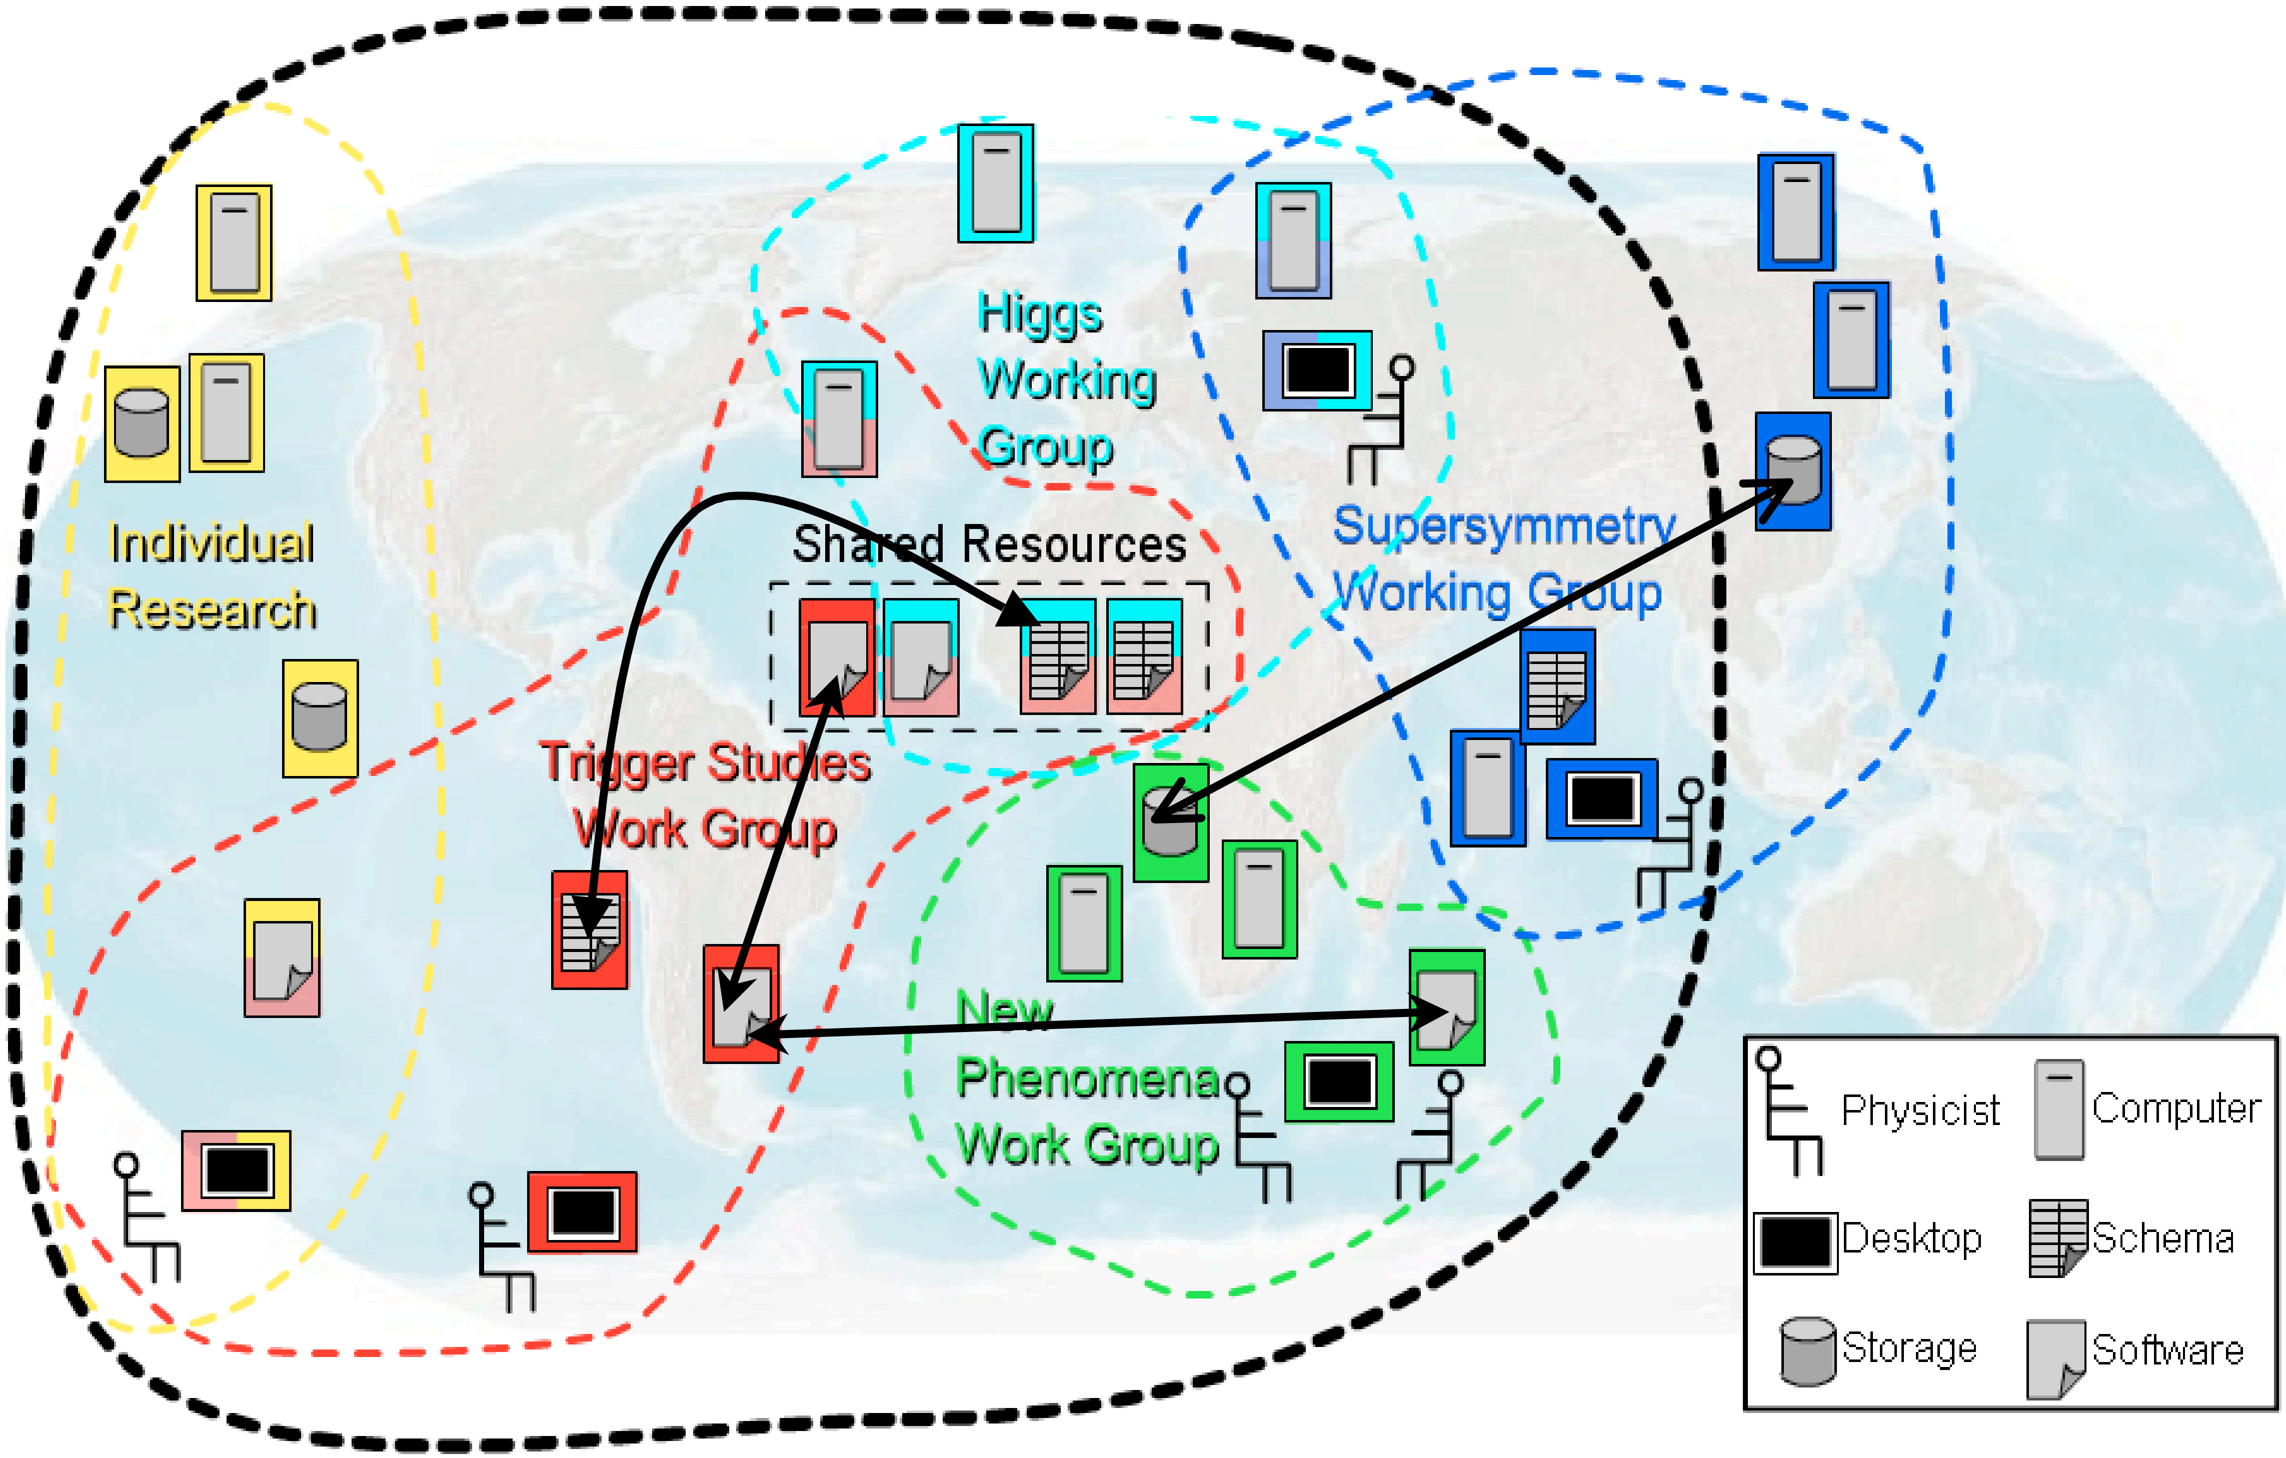

Scientists within Dynamic Workspaces!

Communities of Scientists Working Locally within a Global Context

Infrastructure for sharing, consistency of physics and calibration data, software

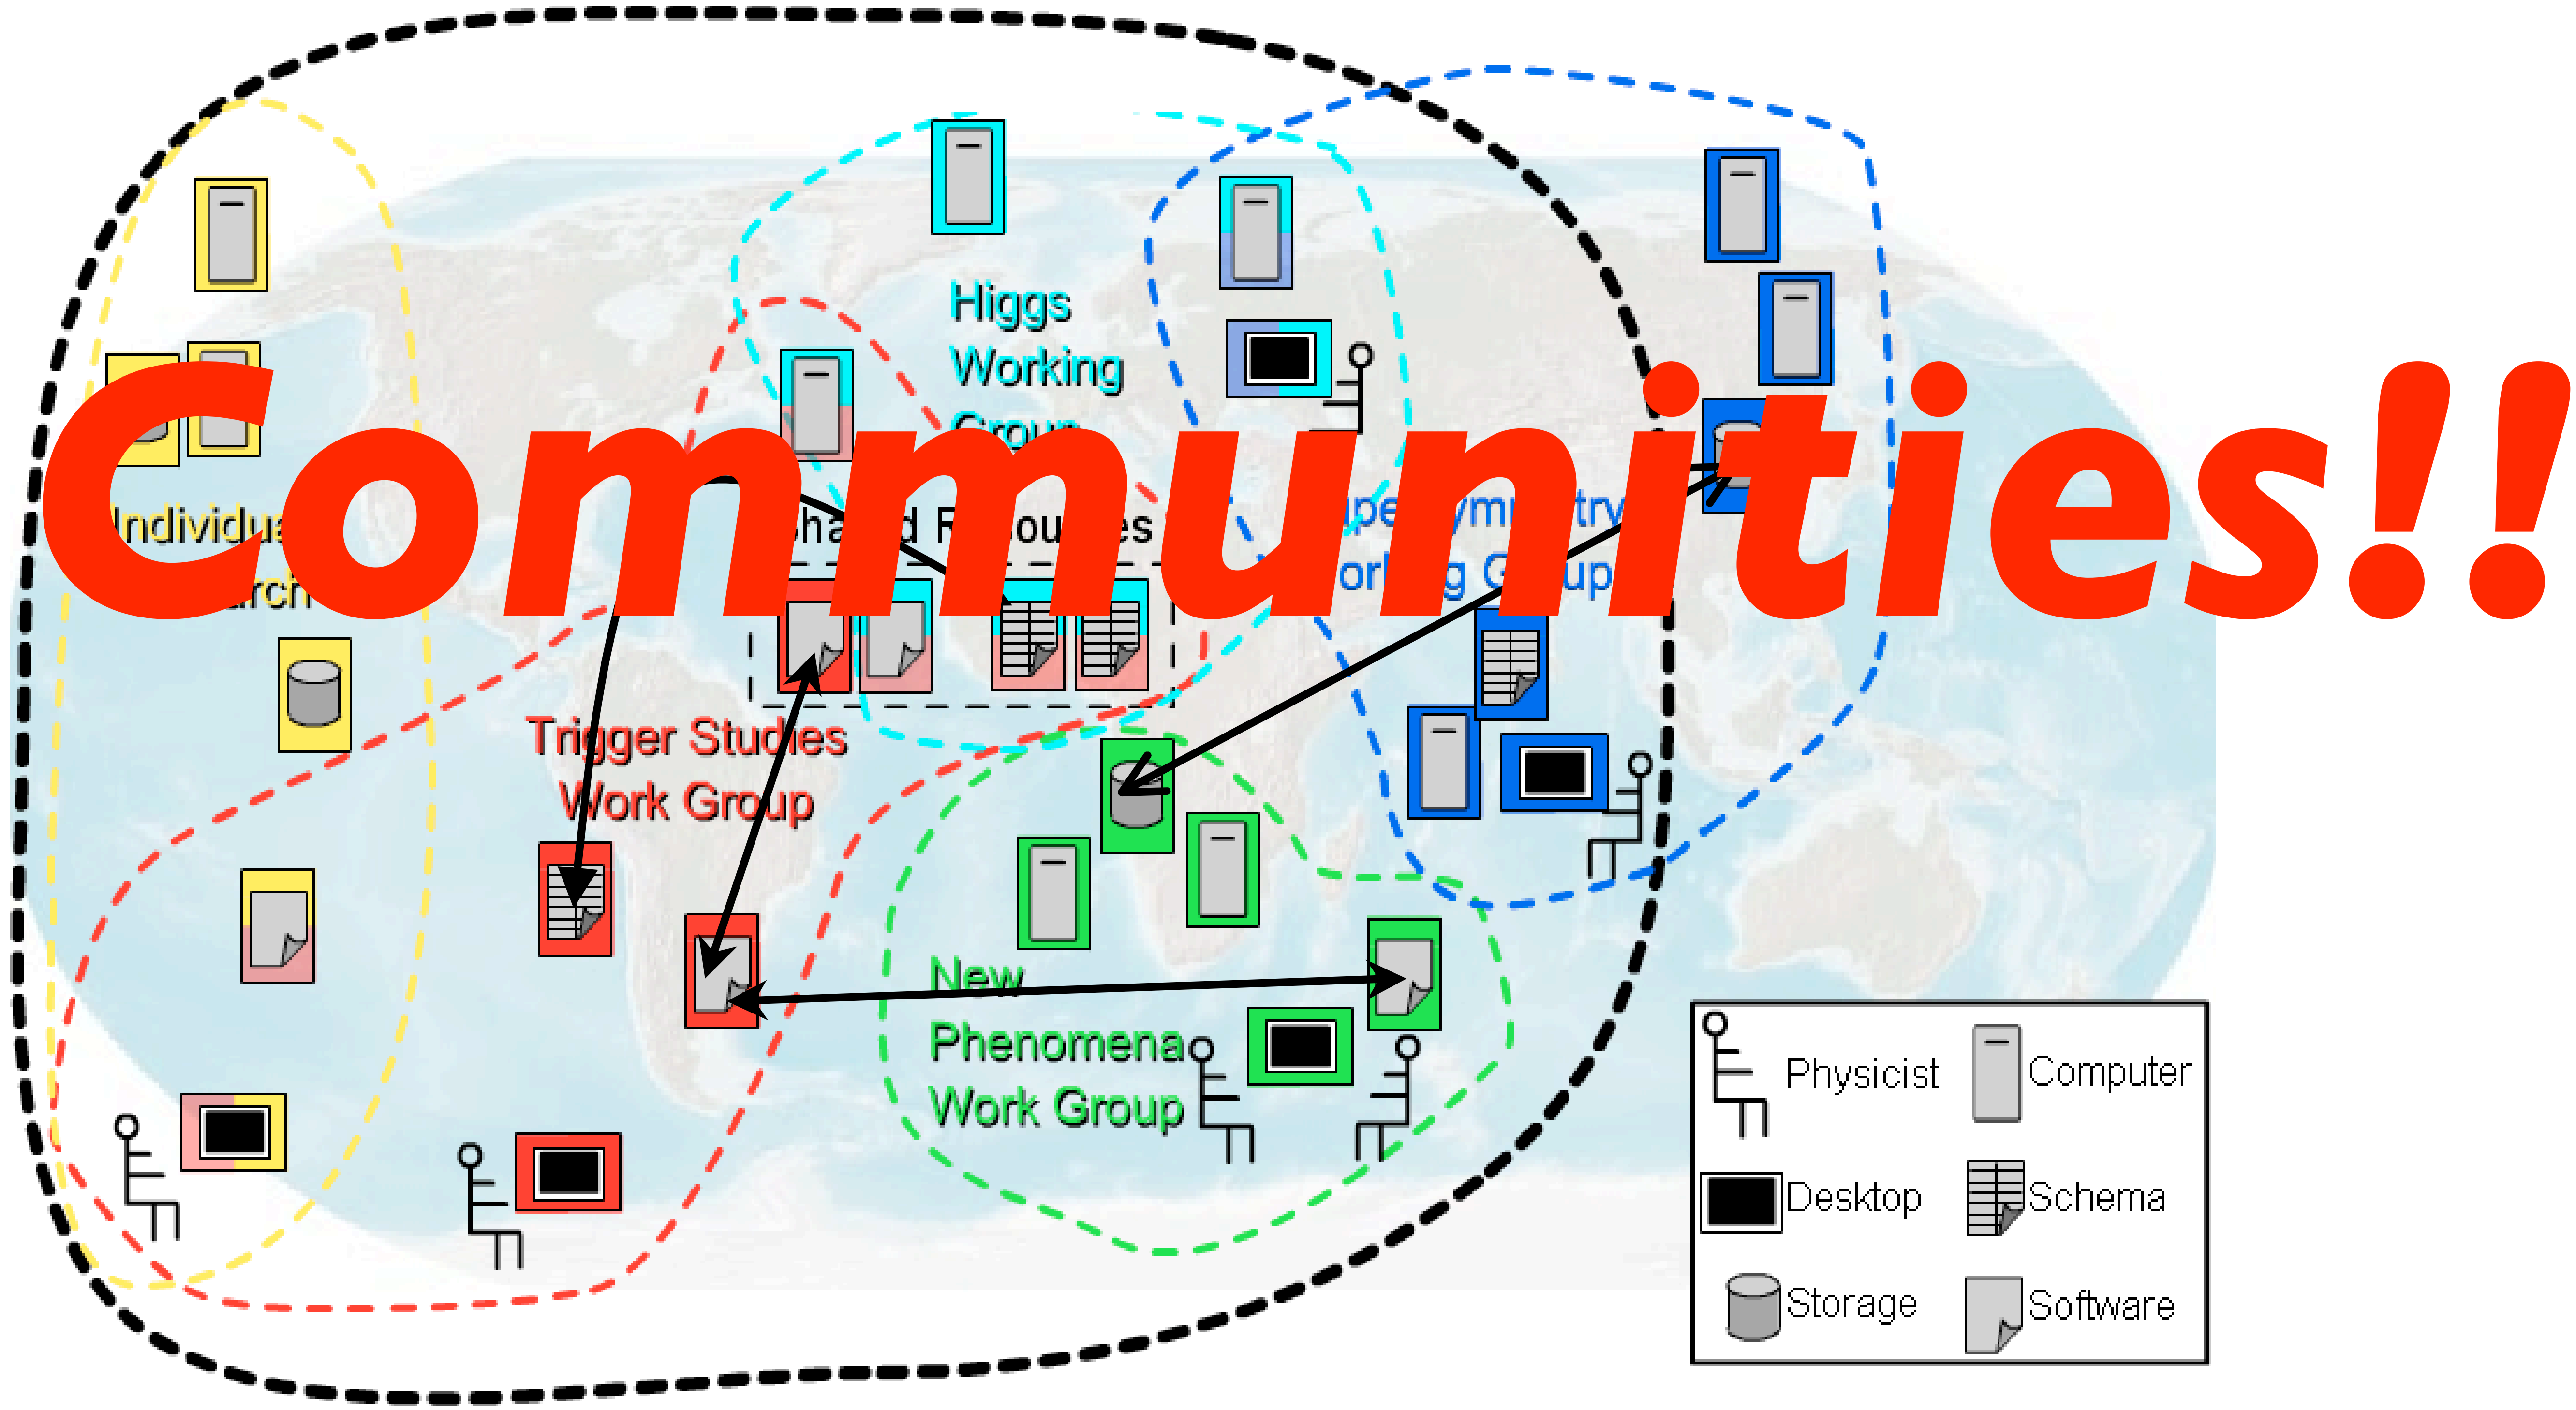

“Work Packages” for LHC Computing

Facilities and Fabric Infrastructure

- Tier-1 and Tier-2 centers, University infrastructure

Distributed Computing Infrastructure

- Networks, throughput, servers, catalogs

Grid Services

- Middleware, “Virtual Organizations” support, end-to-end and higher level services, trouble shooting and fault tolerance, distributed science environment

Experiment Specific Software

- Core software, frameworks, architectures, applications physics and detector support

Collaboratory Tools and Support

- Communication, conferencing, sharing, Virtual Control Room

Support Services

- Training, info services, help desk

Building an LCG Grid

LCG

Summary

I.Bird, LCG

- LCG-1 (EDG 2.0) is likely to be late …
- … and untested
- MDS is known to have serious problems (LDAP lookup)
- R-GMA is unknown
- EDG 2.0 has assumed that R-GMA will be used
- MDS can be used but RLS will need static config files (OK for ~10 sites)
- SE is not delivered
- Work is needed to deliver file access solution
- Scalability of RB in this version is still unknown
  - MDS fundamental problems have not changed
  - RB architecture has changed
  - Condor-G limitations have been addressed

CERN

last update: 25/04/2003 03:12 *Ian Bird 8 april 03*

If we can make Grid work, large resources become available to experiments

- ➡ But a long way to go even for basic Grid software
- ➡ We will manage for next round of production and Data Challenges

so, how about LHC Data Analysis?