

Grid Computing

3rd FCPPL Workshop




Overview

- The LHC & the Grid
- ATLAS & CMS Computing models
- 2009 Activities

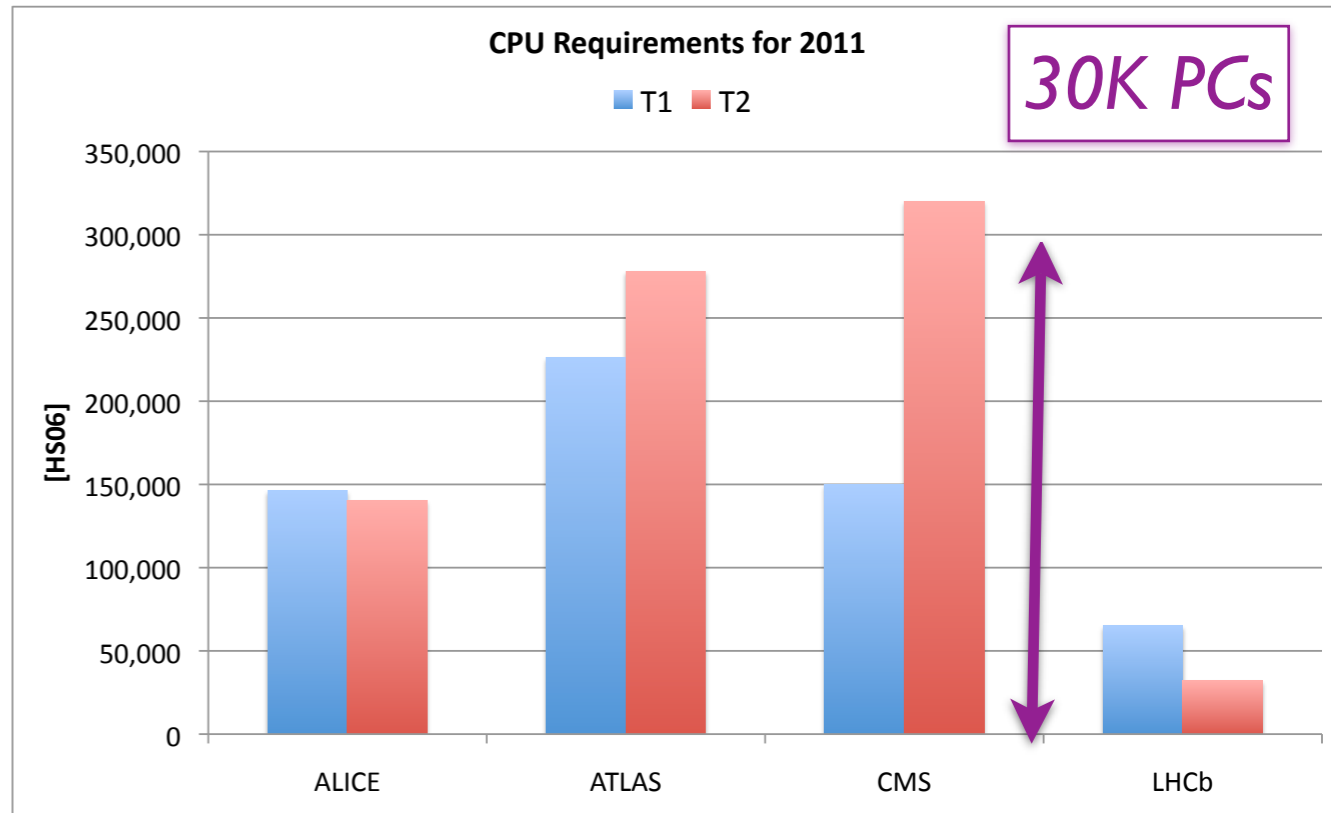


Volume & Power

Disk & CPU requirements for 2011 in *T1s* & *T2s*



1 DVD ~ 5GB

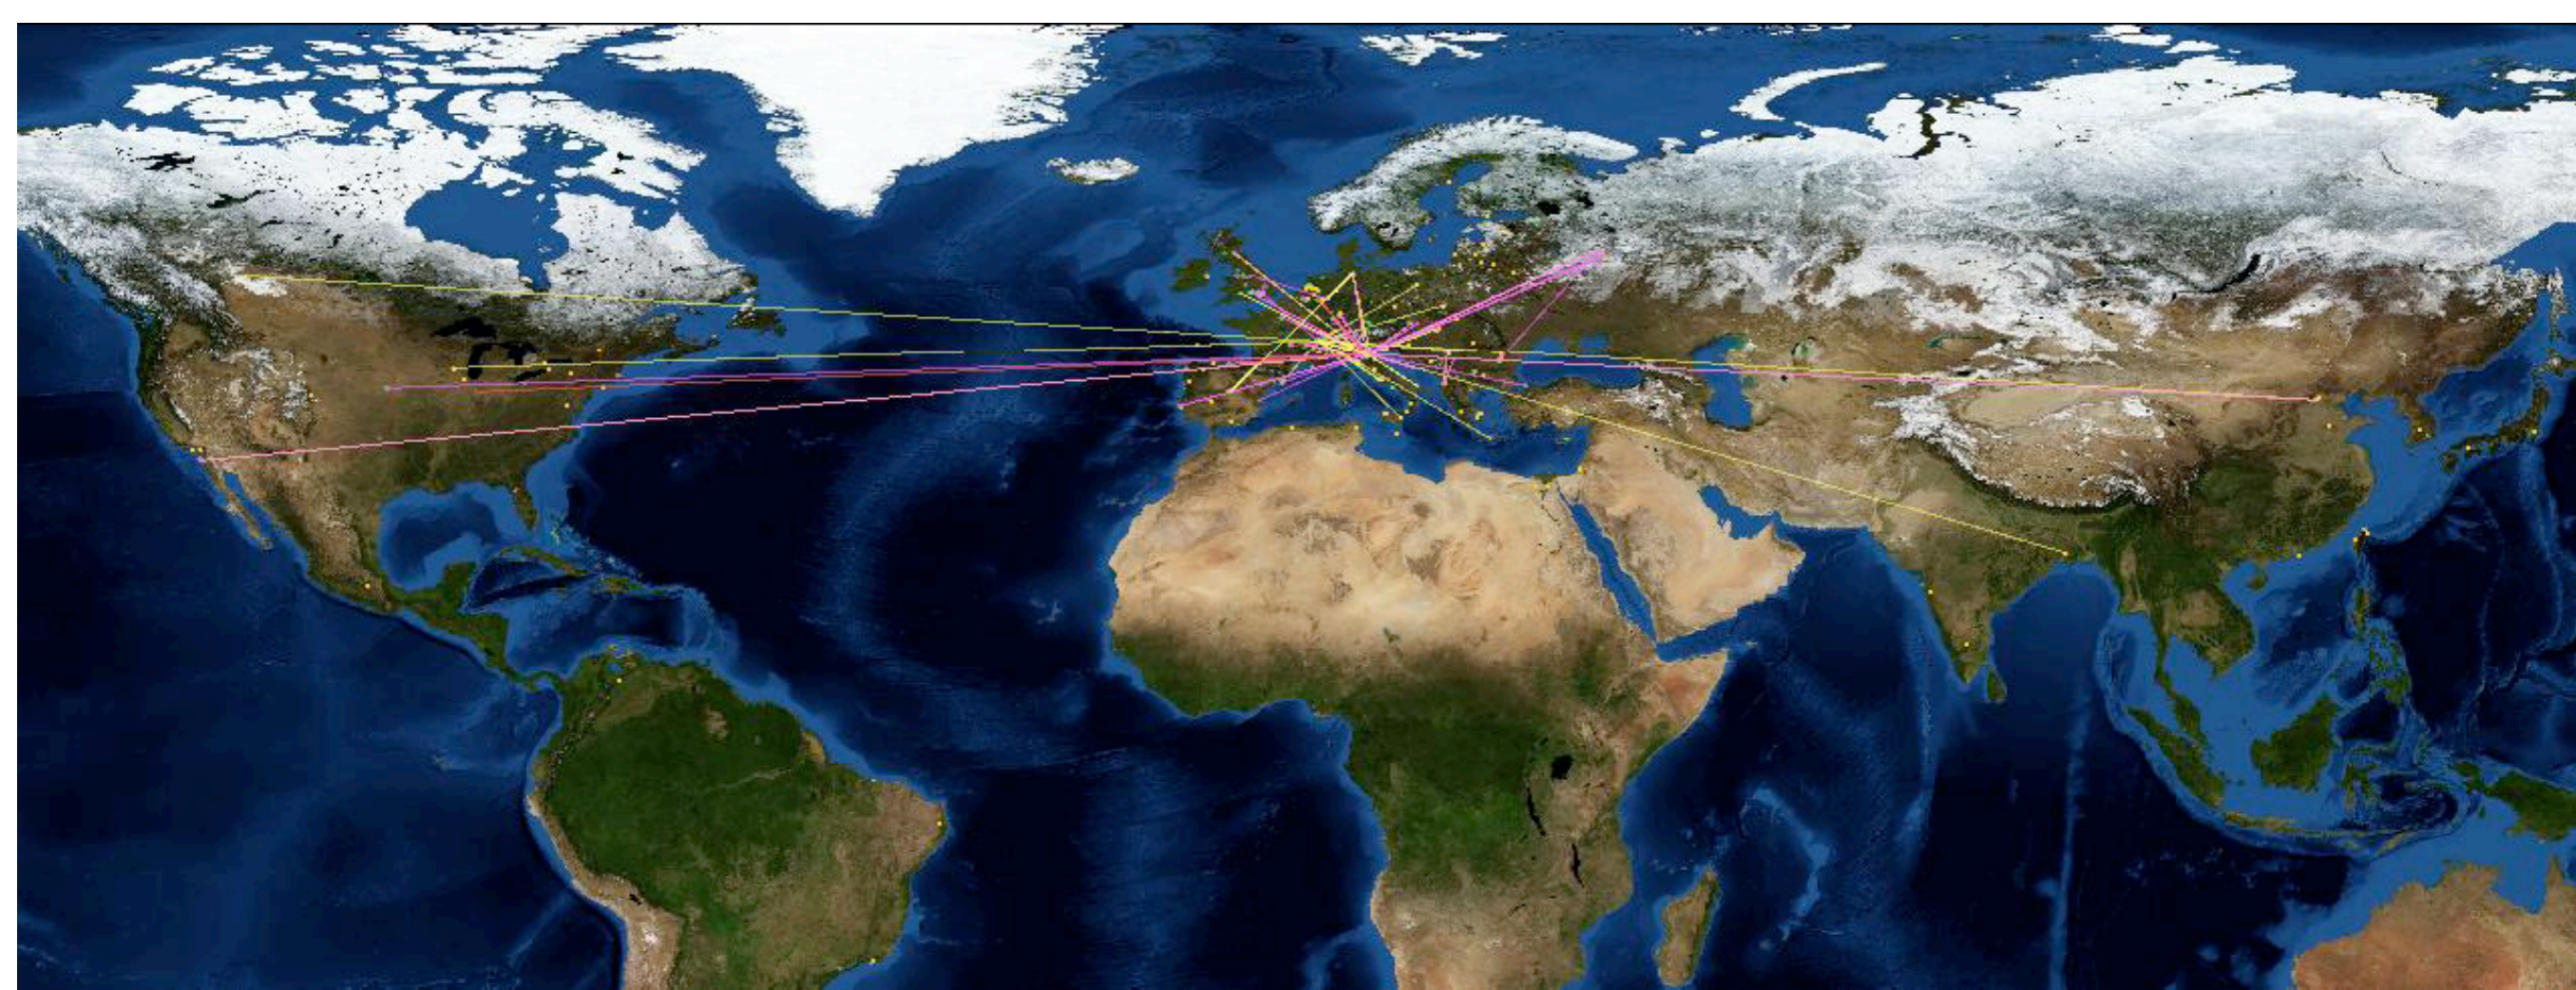


1 regular PC chip ~ 8-10 HS06

2010 - 2011 : ~3G evts expected per experiment
Event sizes : 1.5-2 MB RAW, 1 MB ESD, 0.2 MB AOD

The Challenge : Data Management






**LHC :
Distributed
computing**

**Distributed Data Location
Distributed CPU resources**

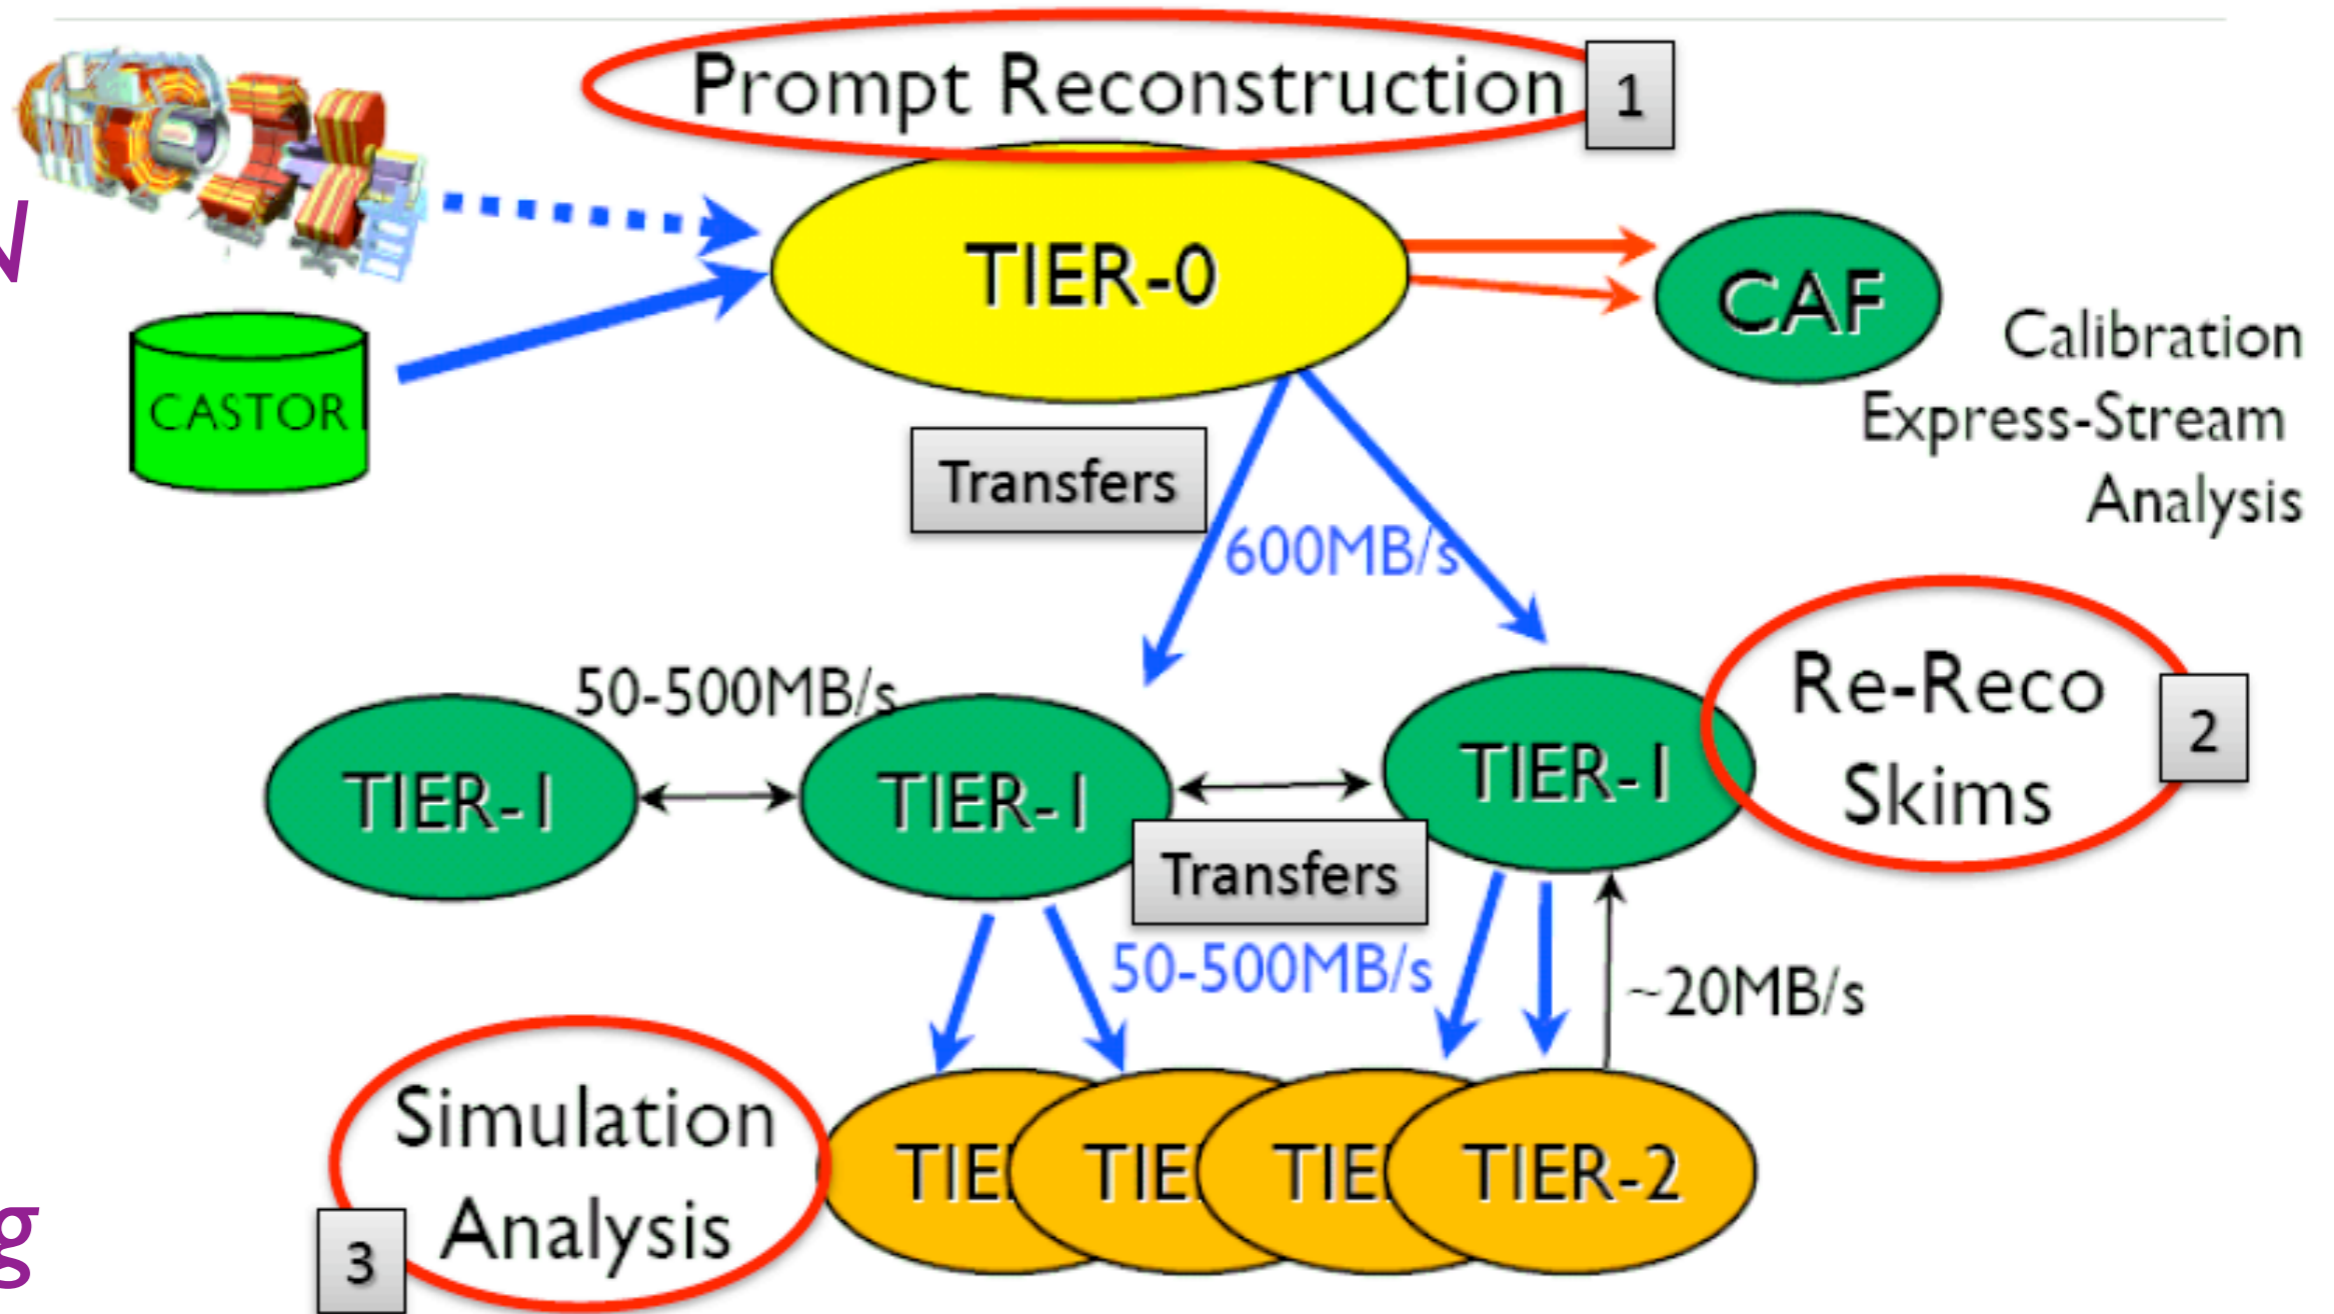
Glued by Network & Resource information











Grid organization

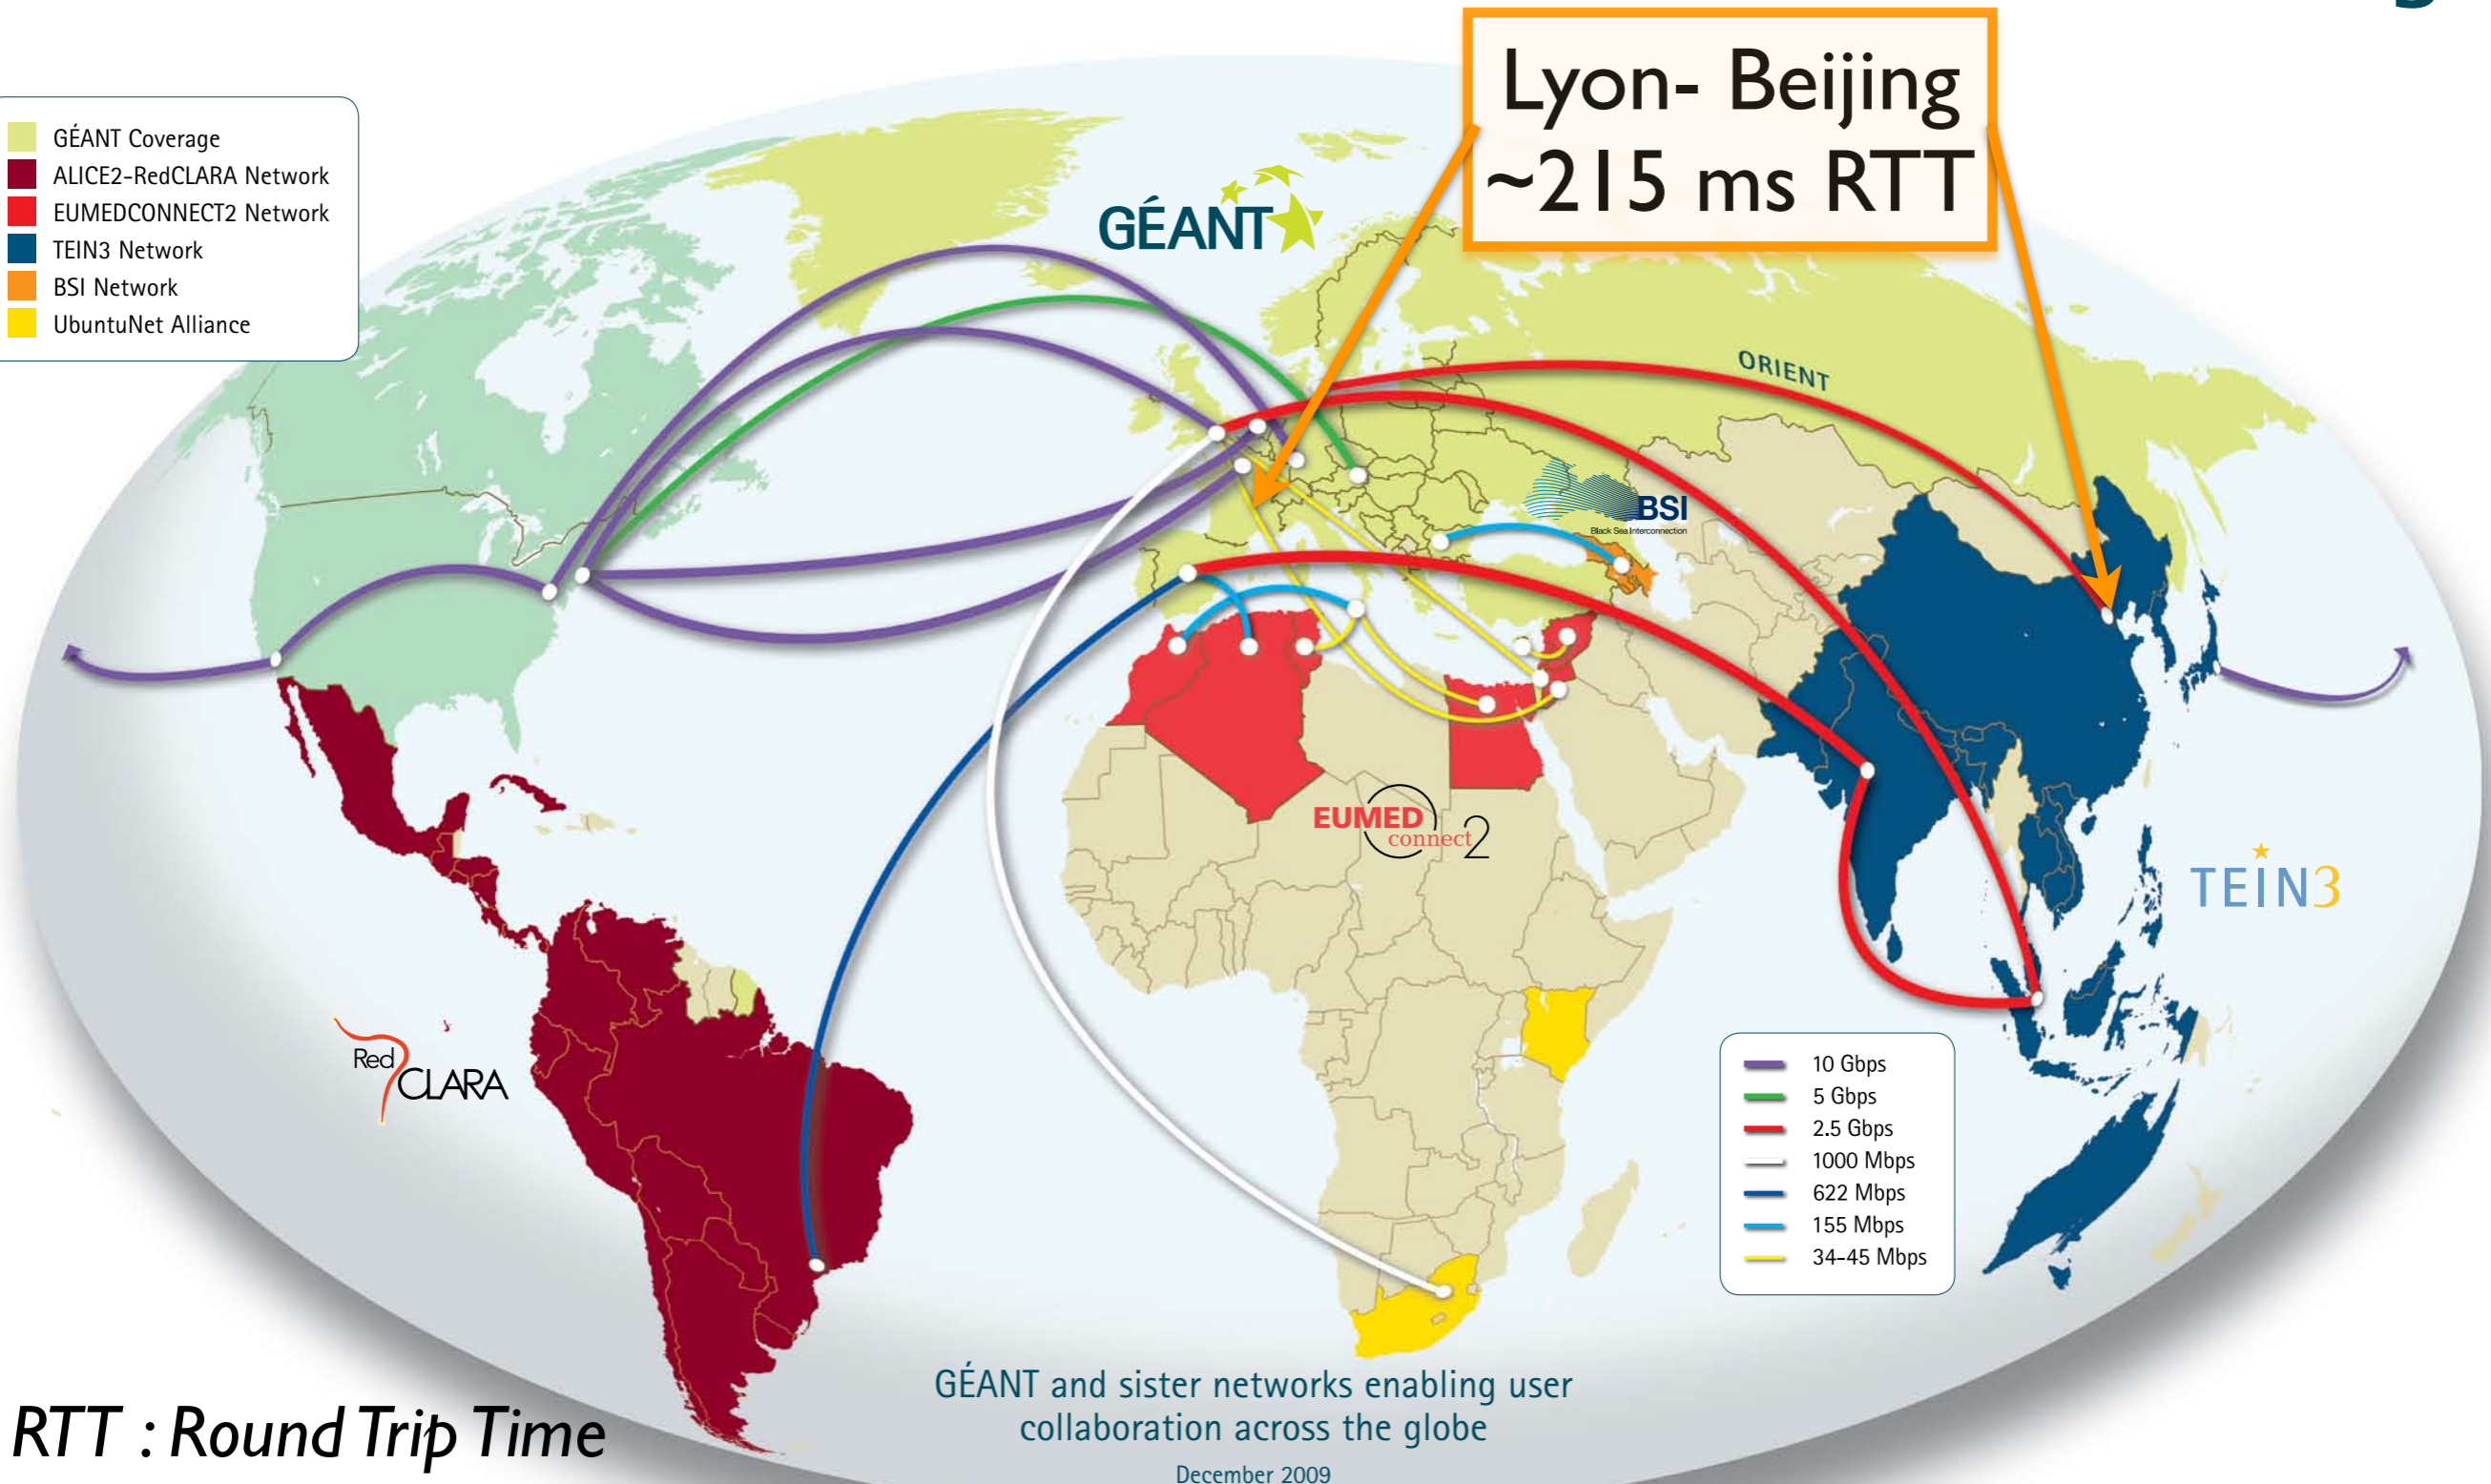
CERN
↓
Lyon
↓
Beijing



Active collaboration between Lyon-T1 and Beijing T2 mandatory

GÉANT At the Heart of Global Research Networking

-  GÉANT Coverage
-  ALICE2-RedCLARA Network
-  EUMEDCONNECT2 Network
-  TEIN3 Network
-  BSI Network
-  UbuntuNet Alliance



Lyon- Beijing
~215 ms RTT

RTT : Round Trip Time

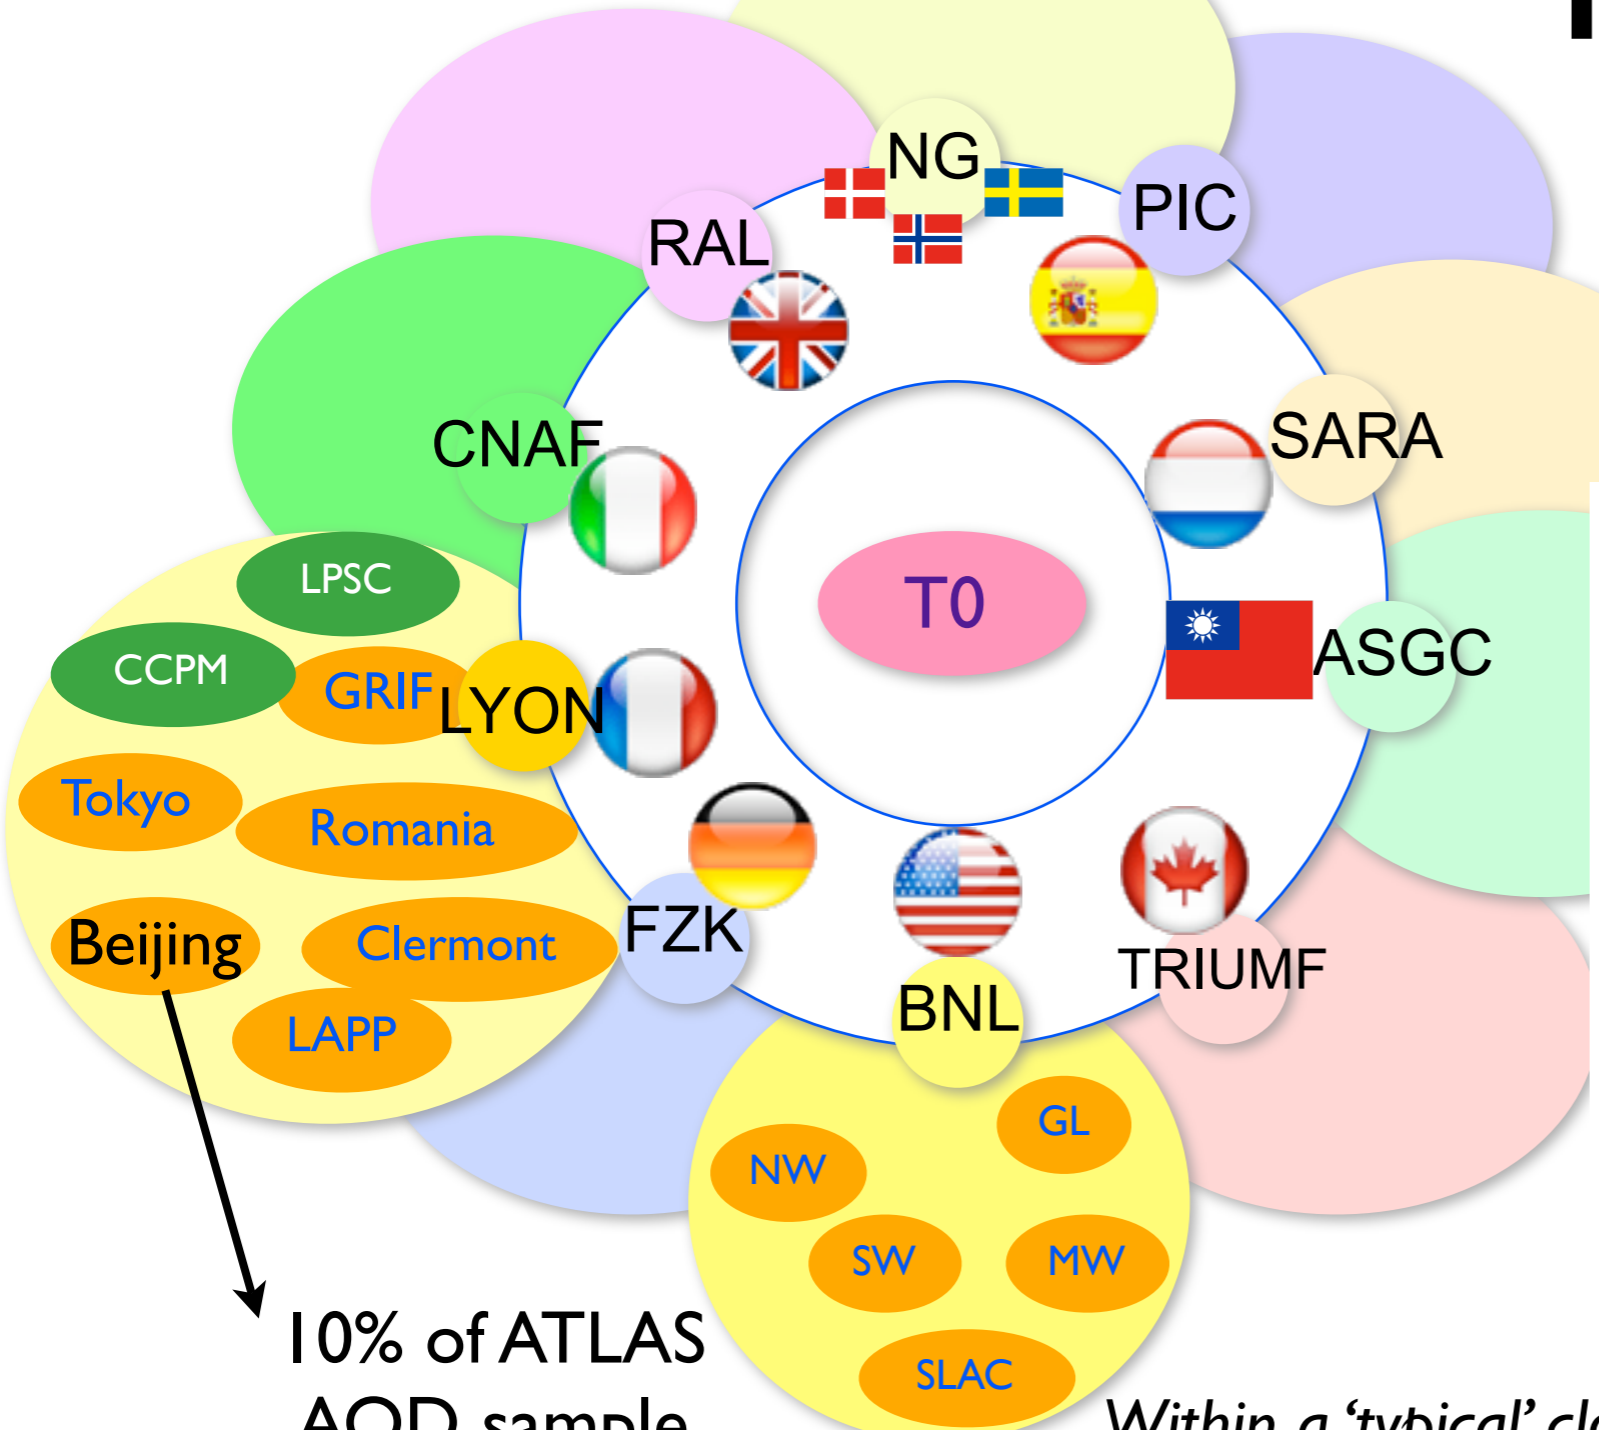
GÉANT and sister networks enabling user collaboration across the globe
December 2009

*Data transfers need careful tuning and monitoring
(time out, parallelism, file size,...)*



ATLAS Topology

“Tier Cloud Model”
Unit: $1 T1 + n T2/T3$




- All T1s have predefined network channel with T0 and with each other T1
- T2s are associated with **one** T1 to form a cloud
- T2s have predefined channel with parent T1 **only**

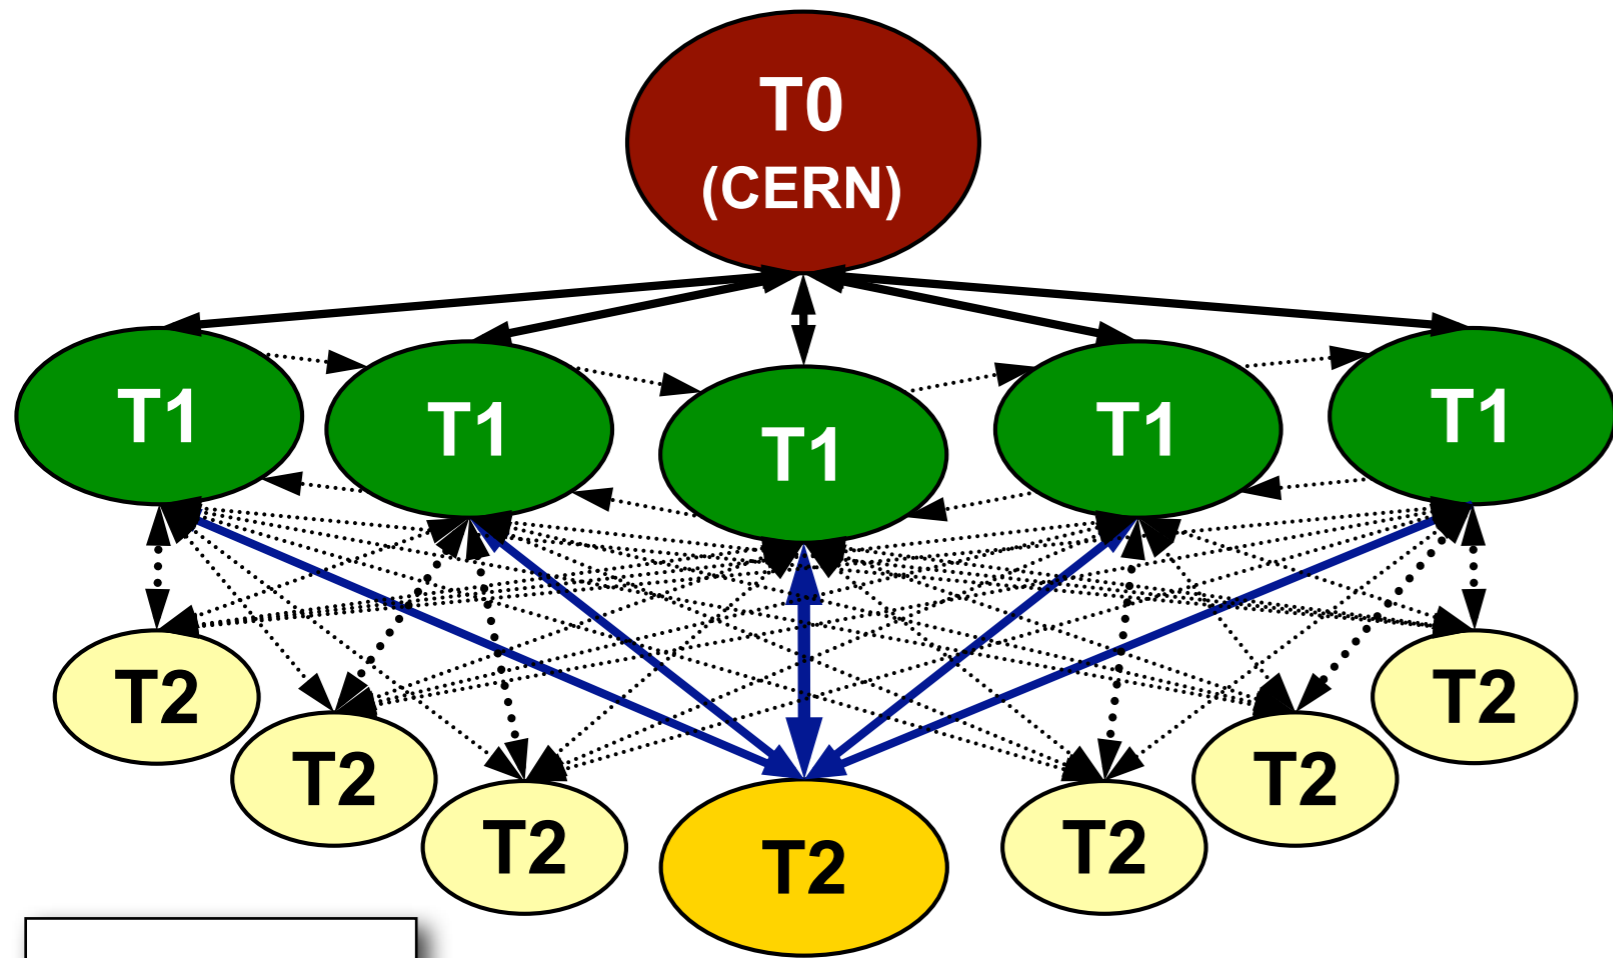
10% of ATLAS AOD sample

Within a ‘typical’ cloud
 T1: 10 % RAW, 20 % ESD, 100% AOD
 $\Sigma T2$: 100 % AOD, small fraction ESD, RAW





CMS Topology




7 T1s
~50 T2s

Very flexible transfer topology

- Any Tier-2 downloads and uploads data from any Tier-1
- Tier-2 to Tier-2 transfers are also allowed (though not encouraged)

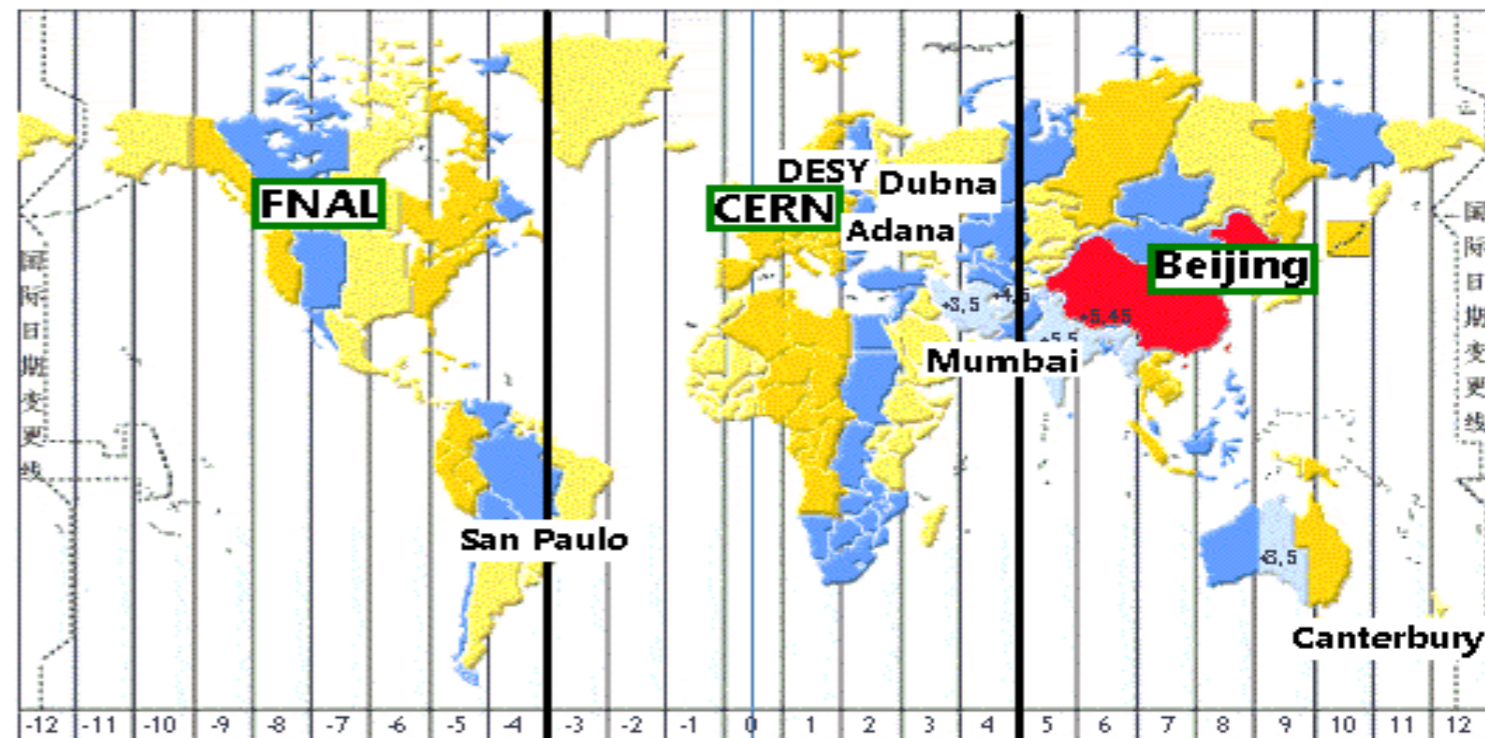

Different Topologies → Different organization

CMS Remote Operation Center




CMSROC@Beijing

- Established in July 2009
- Operational in August
- CMS Computing shift in Beijing
- 24 hours coverage of CMS computing shift



Eric Lançon / 2010 FCPPL Workshop, Lyon





ATLAS cloud-operation

- ‘Squad-FR’ Team
 - Interface between T0 - sites
 - Monitor all Grid activities
 - Data transfers
 - MC production
 - Analysis

IHEP involvement
Erming Pei
Xiaofei Yan

MC Job efficiency	2008	2009
FR-Cloud	71%	88%
Beijing	62%	90%

3.5 more jobs in 2009

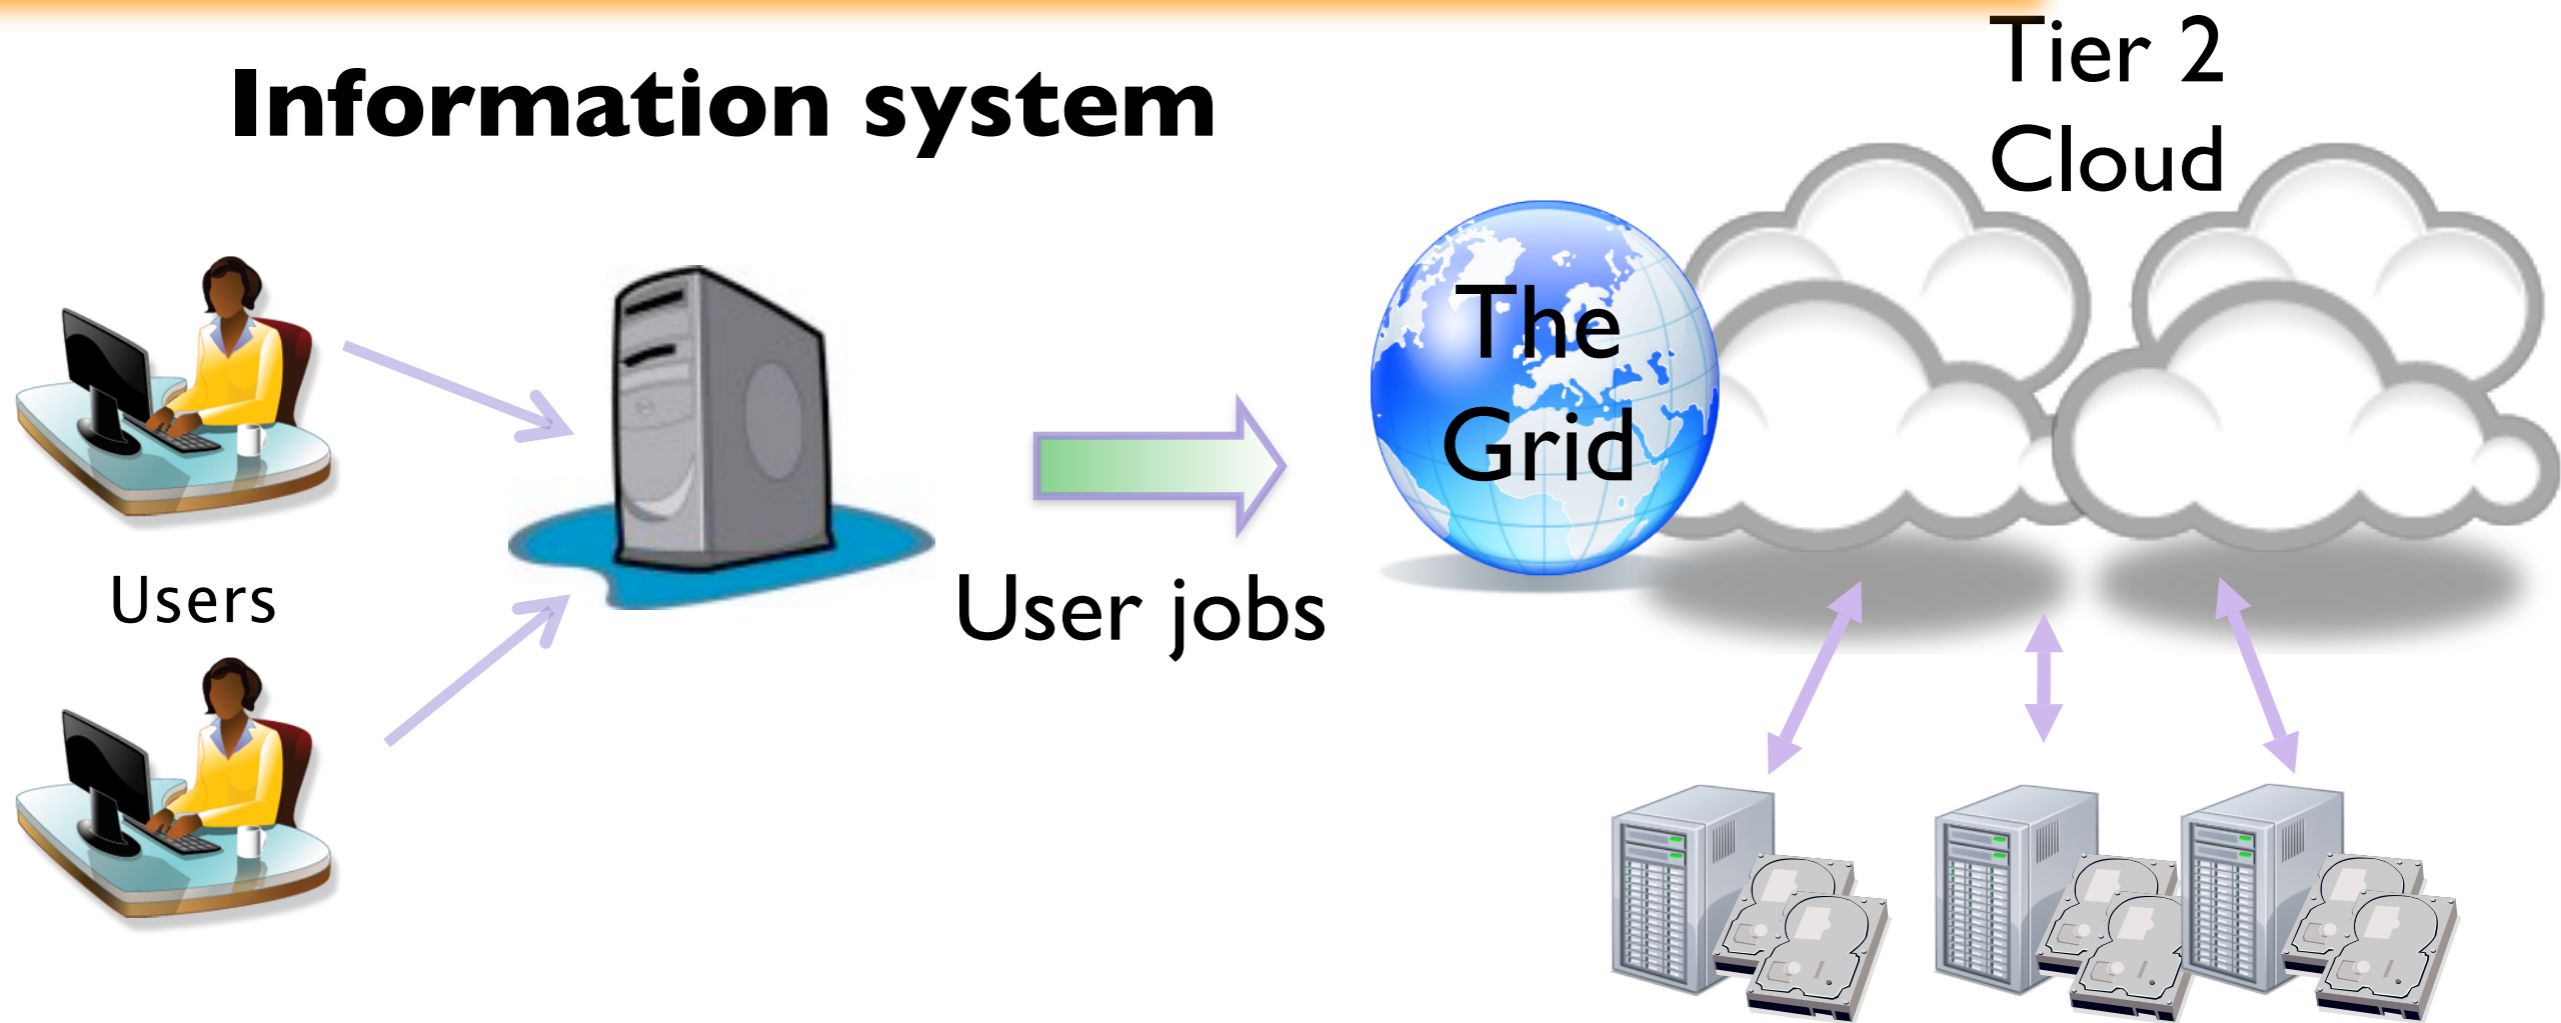
- Dedicated Technical mailing list and meetings



Analysis on the Grid

Each experiment has different implementation

Information system



Key element : information system

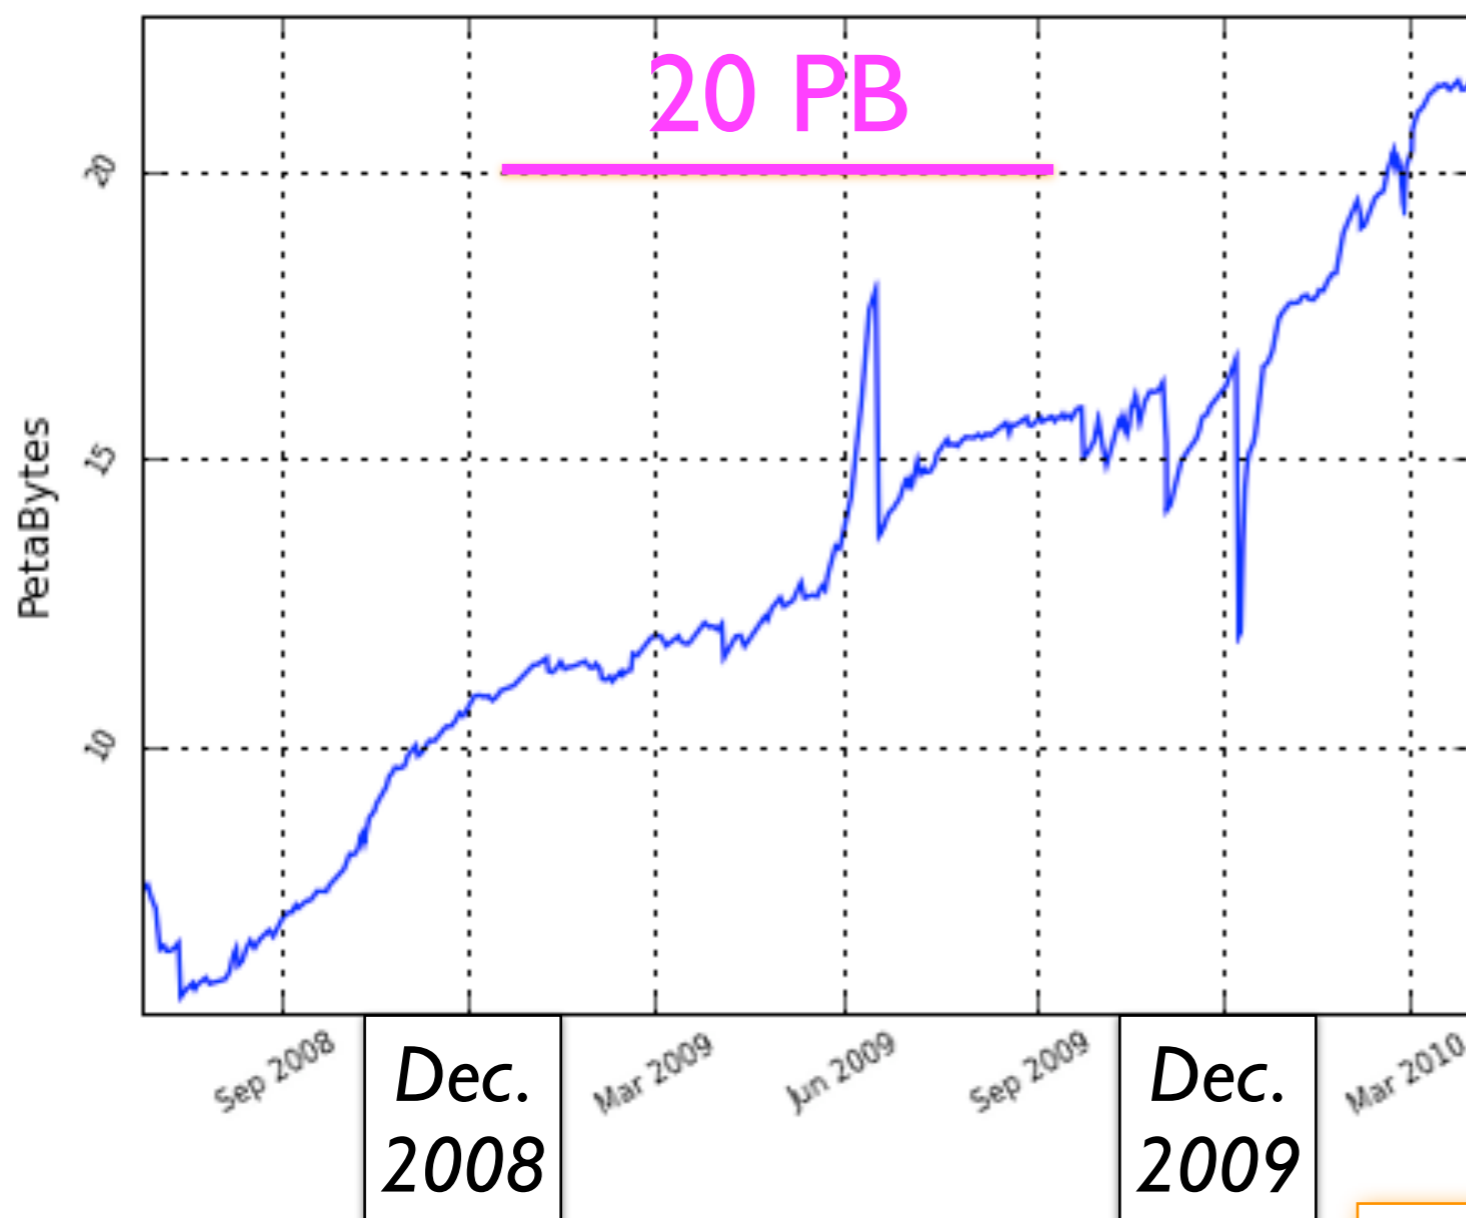
- Site availability & performance
- Data location





ATLAS data Volume


Total GRID disk usage according to dq2



- > 1,000 users
- > 500 endpoints (storage elements)
- 3 Grids
- ~2,000,000 dataset replicas
- ~80,000,000 file replicas

Catalog of >80M identities

Eric Lançon / 2010 FCPPL Workshop, Lyon



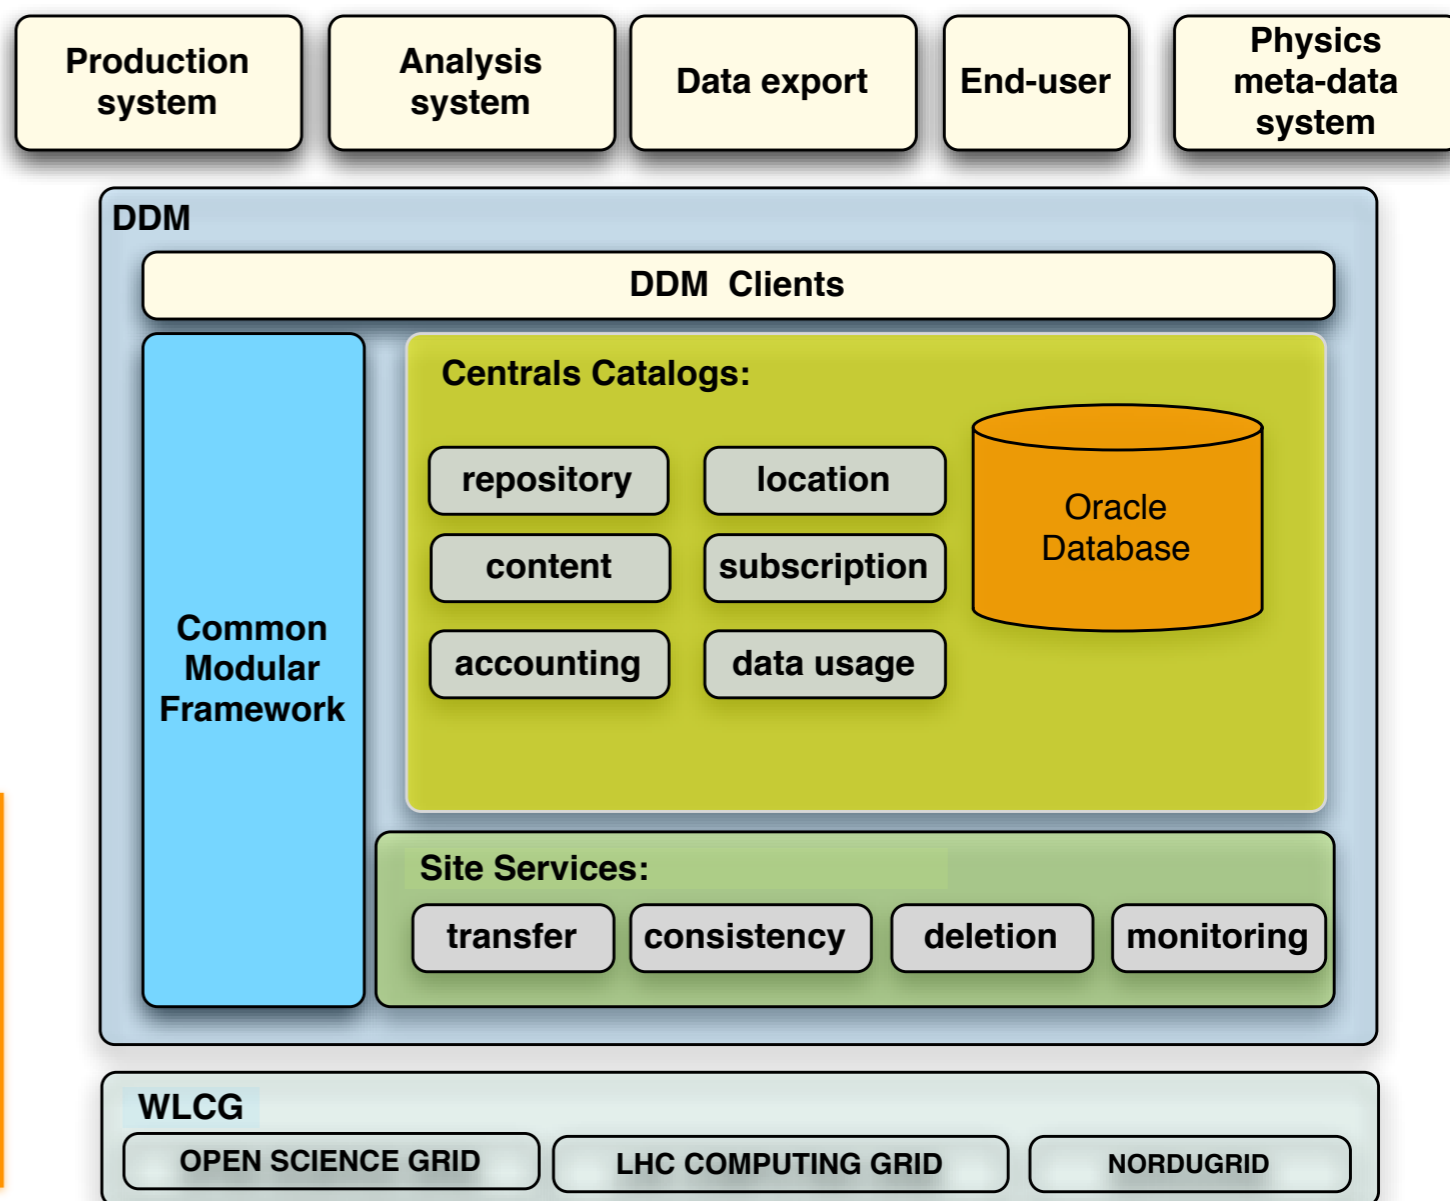


ATLAS Data Distribution Management

- Central link between all ATLAS Grid components
- Manage the experiment's data
 - Taking care of the Data movement between sites
 - Providing data access to
 - Analysis systems
 - Production systems
 - Physics meta-data systems
 - End users
 - Bookkeeping & accounting

*Proposal for 2010-2011:
PhD student from IHEP Comp. Center
to join DDM Development team
(leader V. Garonne, LAL)*

Baseline service for ATLAS data organization, access, placement and deletion



From YES!




**Excitement in the ATLAS Detector Control Room:
The first LHC event on 10th September 2008**

LHCC, 24-Sep-2008, PJ

Back to fall 2008


To NO!



Incident on 19th September



- During commissioning of the last main bend circuit to 5 TeV an incident occurred resulting in the triggering of quench heaters of about 100 magnets and a large He discharge into the tunnel.
- The most probable cause is a faulty electrical connection between two magnets. The sector is being brought to room temperature for repair.
- The time needed for warmup, repair and cooldown precludes a restart before CERN's obligatory winter shutdown.
- The shutdown schedule is being modified to gain ~ 1 month of LHC operation in 2009.



'Details' were not yet known at that time



We did not sleep for a year!

- Improvements of
 - Monitoring tools
 - Procedures
- Test, Test, Test
 - Example : STEP09 see Next Slides
(Apologizes for the ATLAS bias...)


*Would we have been ready in 2008?
Probably not... but we will never know*



June '09


What STEP09 was for us:

- STEP09 is *Scale Testing for the Experiment Program 09*, i.e., Offline computing systems commissioning test
- It involved all major offline activities:
 - Monte Carlo Production
 - Full Chain Data Distribution
 - Reprocessing at Tier-1s
 - User Analysis Challenge: Hammercloud
 - ATLAS Central Services Infrastructure
- Done in conjunction with other LHC experiments
 - Combined tape access with CMS at T1s critical



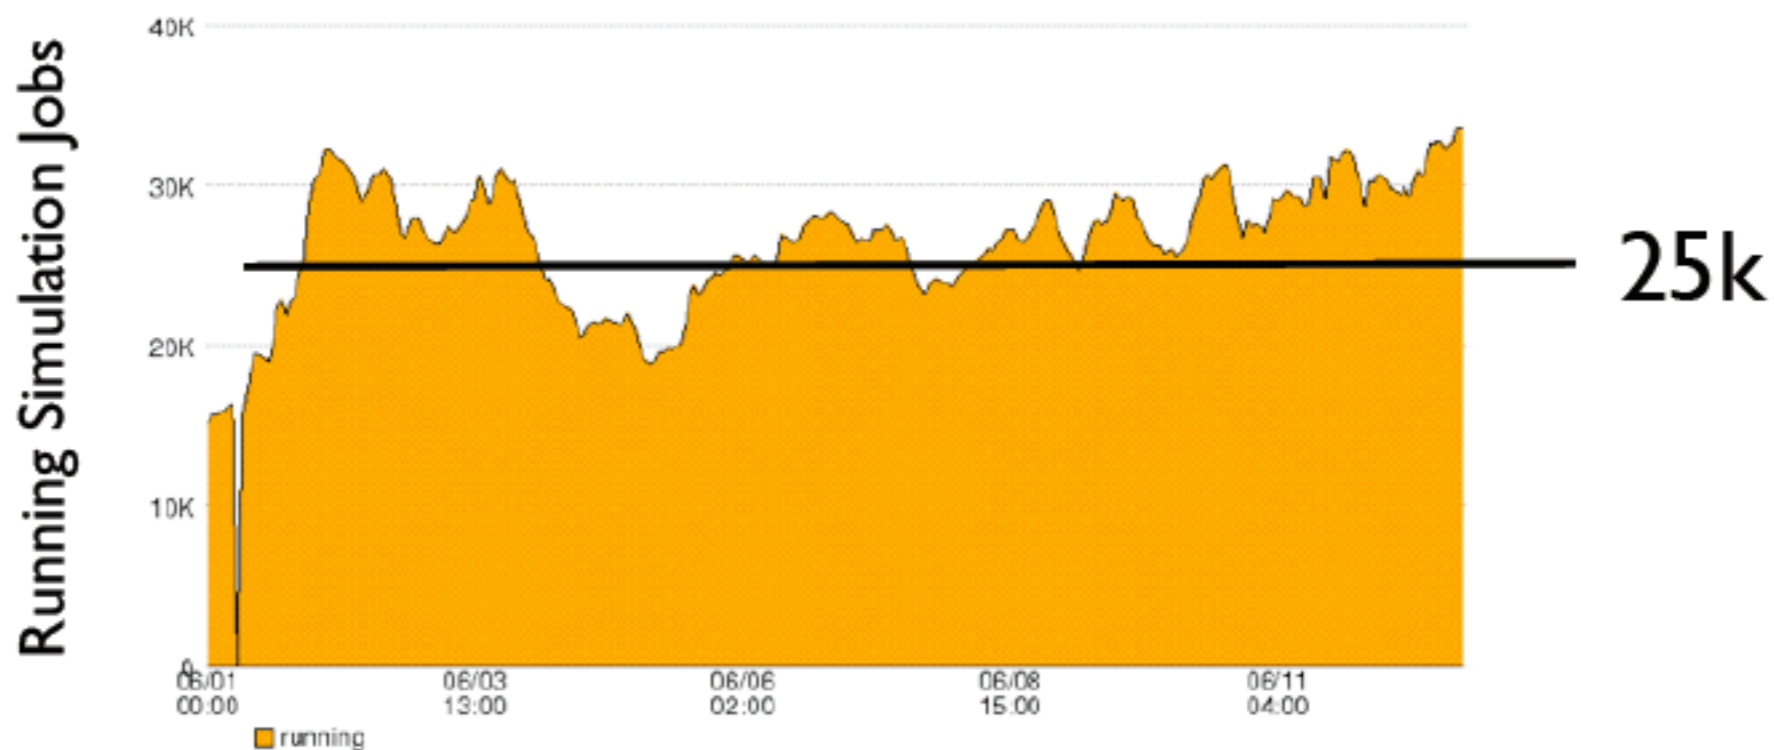
A Very Intensive Period
for all of us...

Eric Lançon / 2010 FCPPL Workshop, Lyon




MonteCarlo Production

- Millions of hours of simulation production done
 - Production already well validated by increasingly large production runs
 - Operationally this is a solved problem
- N.B. Simulation filled all free resources to produce 12M events during STEP which matches ATLAS' mc09 requirements



- *All Sites full with MC production*
- *But we knew already it was feasible*

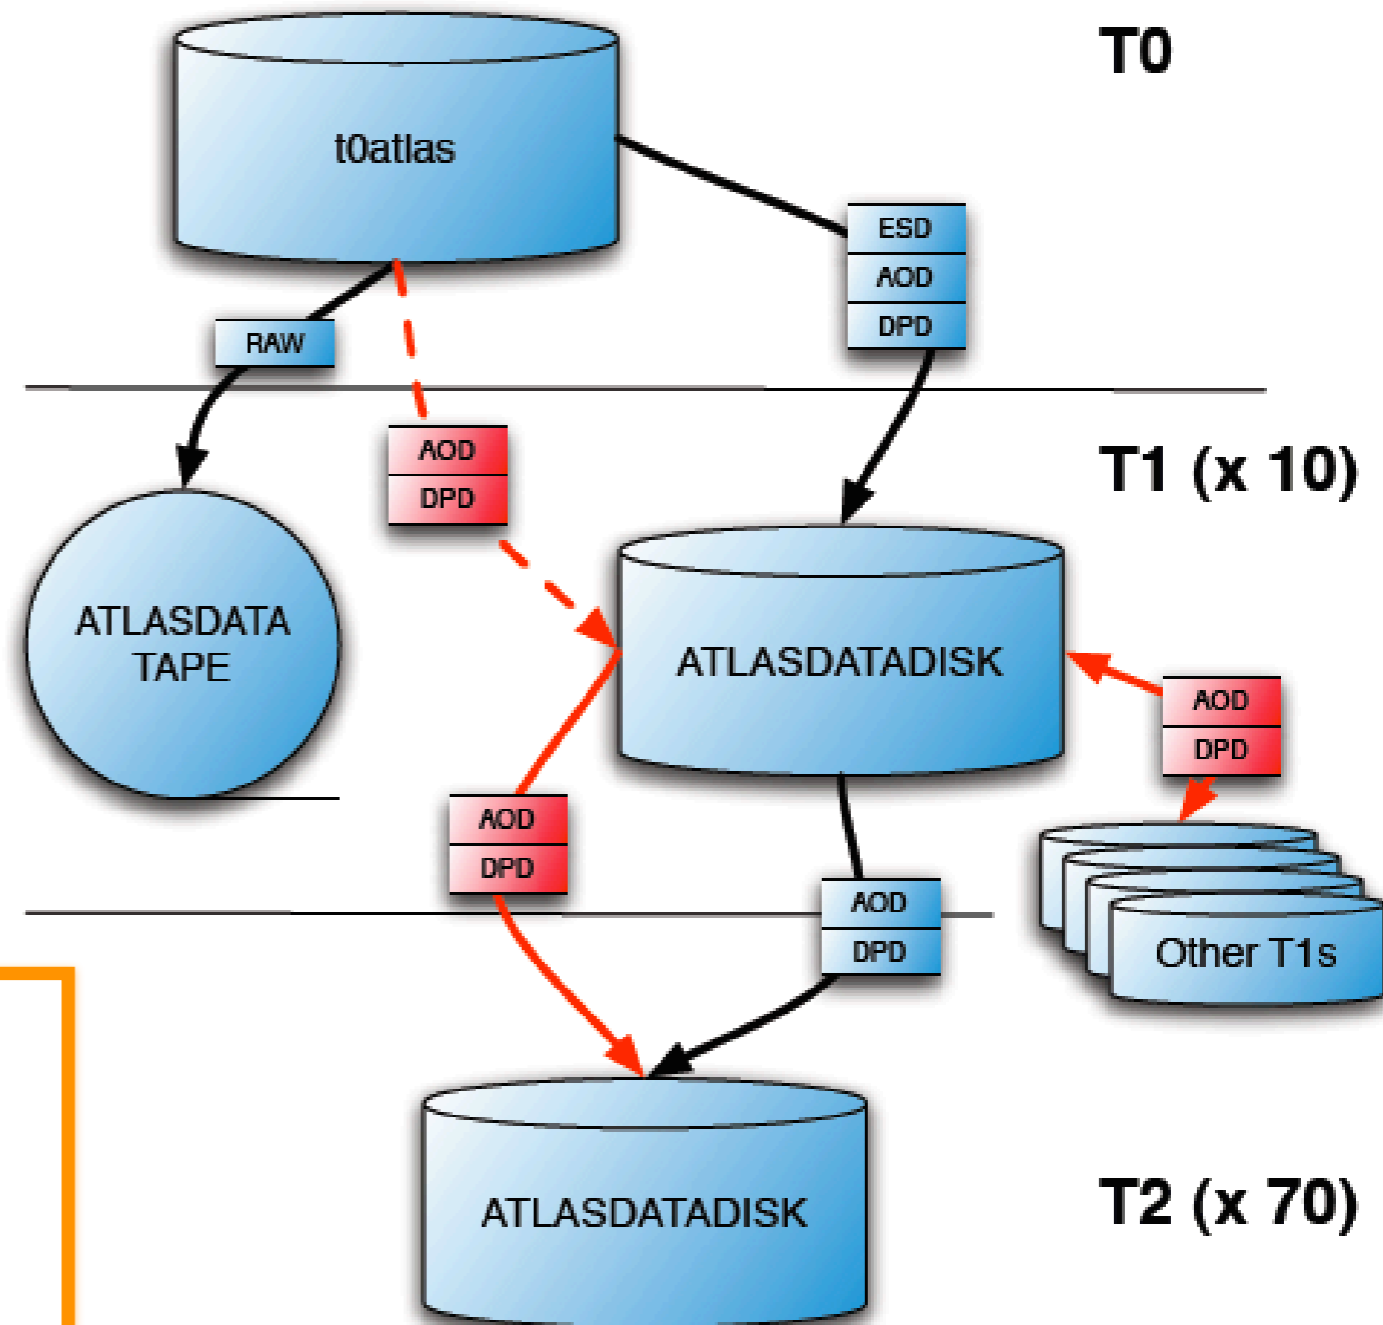




Data Distribution

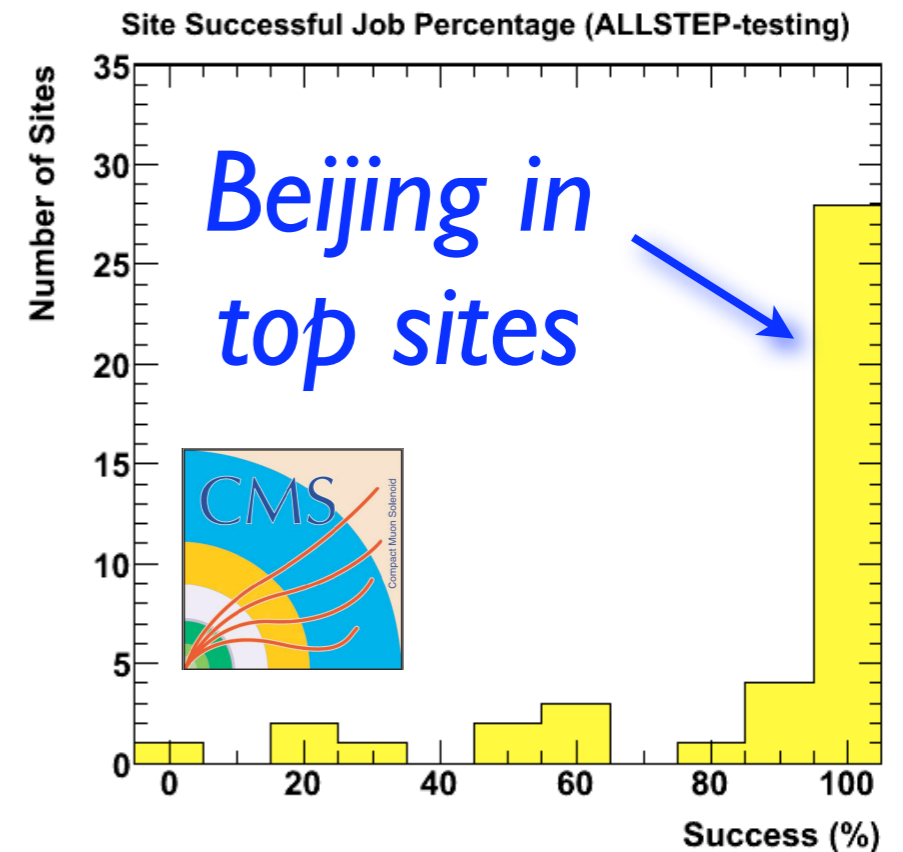
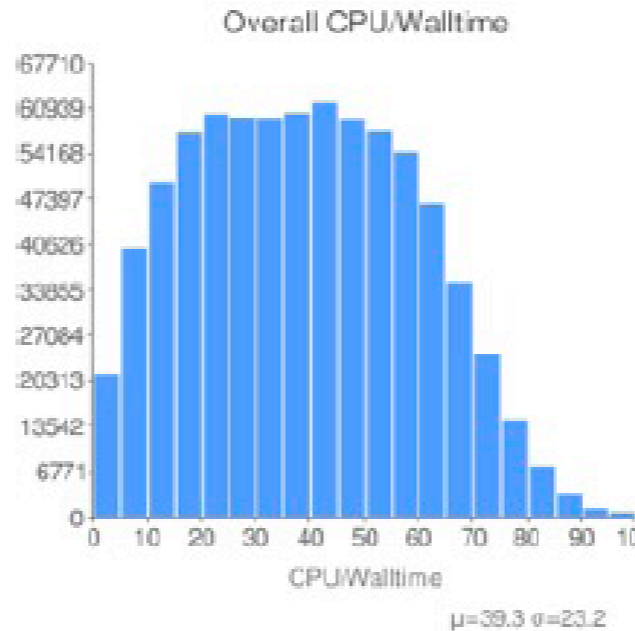

- Data taking and first reconstruction passes
 - RAW and ESD from CERN → distributed to T1 sites (1, 2 copies respectively, RAW to tape)
 - AOD and DPD from CERN distributed to all T1 sites (10 copies)
 - AOD and DPD from CERN distributed to T2 from their parent T1 (1 to 2.7 copies per cloud)
- Reprocessing at Tier-1s
 - AOD and DPD distributed from all T1s to all other T1s
 - AOD and DPD from all T1s

*The Complete Chain:
what real life would
look like...*



'User' Analysis

- ~1M jobs submitted, 83.4% success rate *Beijing 88%*
- 26.3B events processed, 28.6kHz across grid
- N.B. This sounds impressive, but is actually only a few 10s of power users
- Average job's event rate: 7.7Hz
- Average job's cpu efficiency: 0.39





Analysis jobs systematic tests

Hammercloud

Home Clouds Tests Last Tests Time HC Stats Administration

Test 1252

<< Back

Overall Sites Metrics Other




Regular Tests performed by Erming Pei on FR-cloud

Summary

state	id	host	clouds	start time (CET)	end time (CET)	submitted jobs
completed	1252	voatlas49.cern.ch	FR_PANDA	2010-03-31 18:35:00	2010-04-01 18:35:02	14945

Input type: PANDA
 Output DS: user09.JohannesElmsheuser.ganga.sitetest.FRPANDA.20100330.1.[sitename]
 Input DS Patterns: mc09*merge.AOD*.e*r1*
 Ganga Job Template: /data/gangarobot/hammercloud/inputfiles/muon1566/muon1566_panda.tpl
 Athena User Area: /data/gangarobot/hammercloud/inputfiles/muon1566/MuonTriggerAnalysis_1566.tar.gz
 Athena Option file: /data/gangarobot/hammercloud/inputfiles/muon1566/MuonTriggerAnalysis_1566.py

[View Test Directory \(for debugging\)](#)



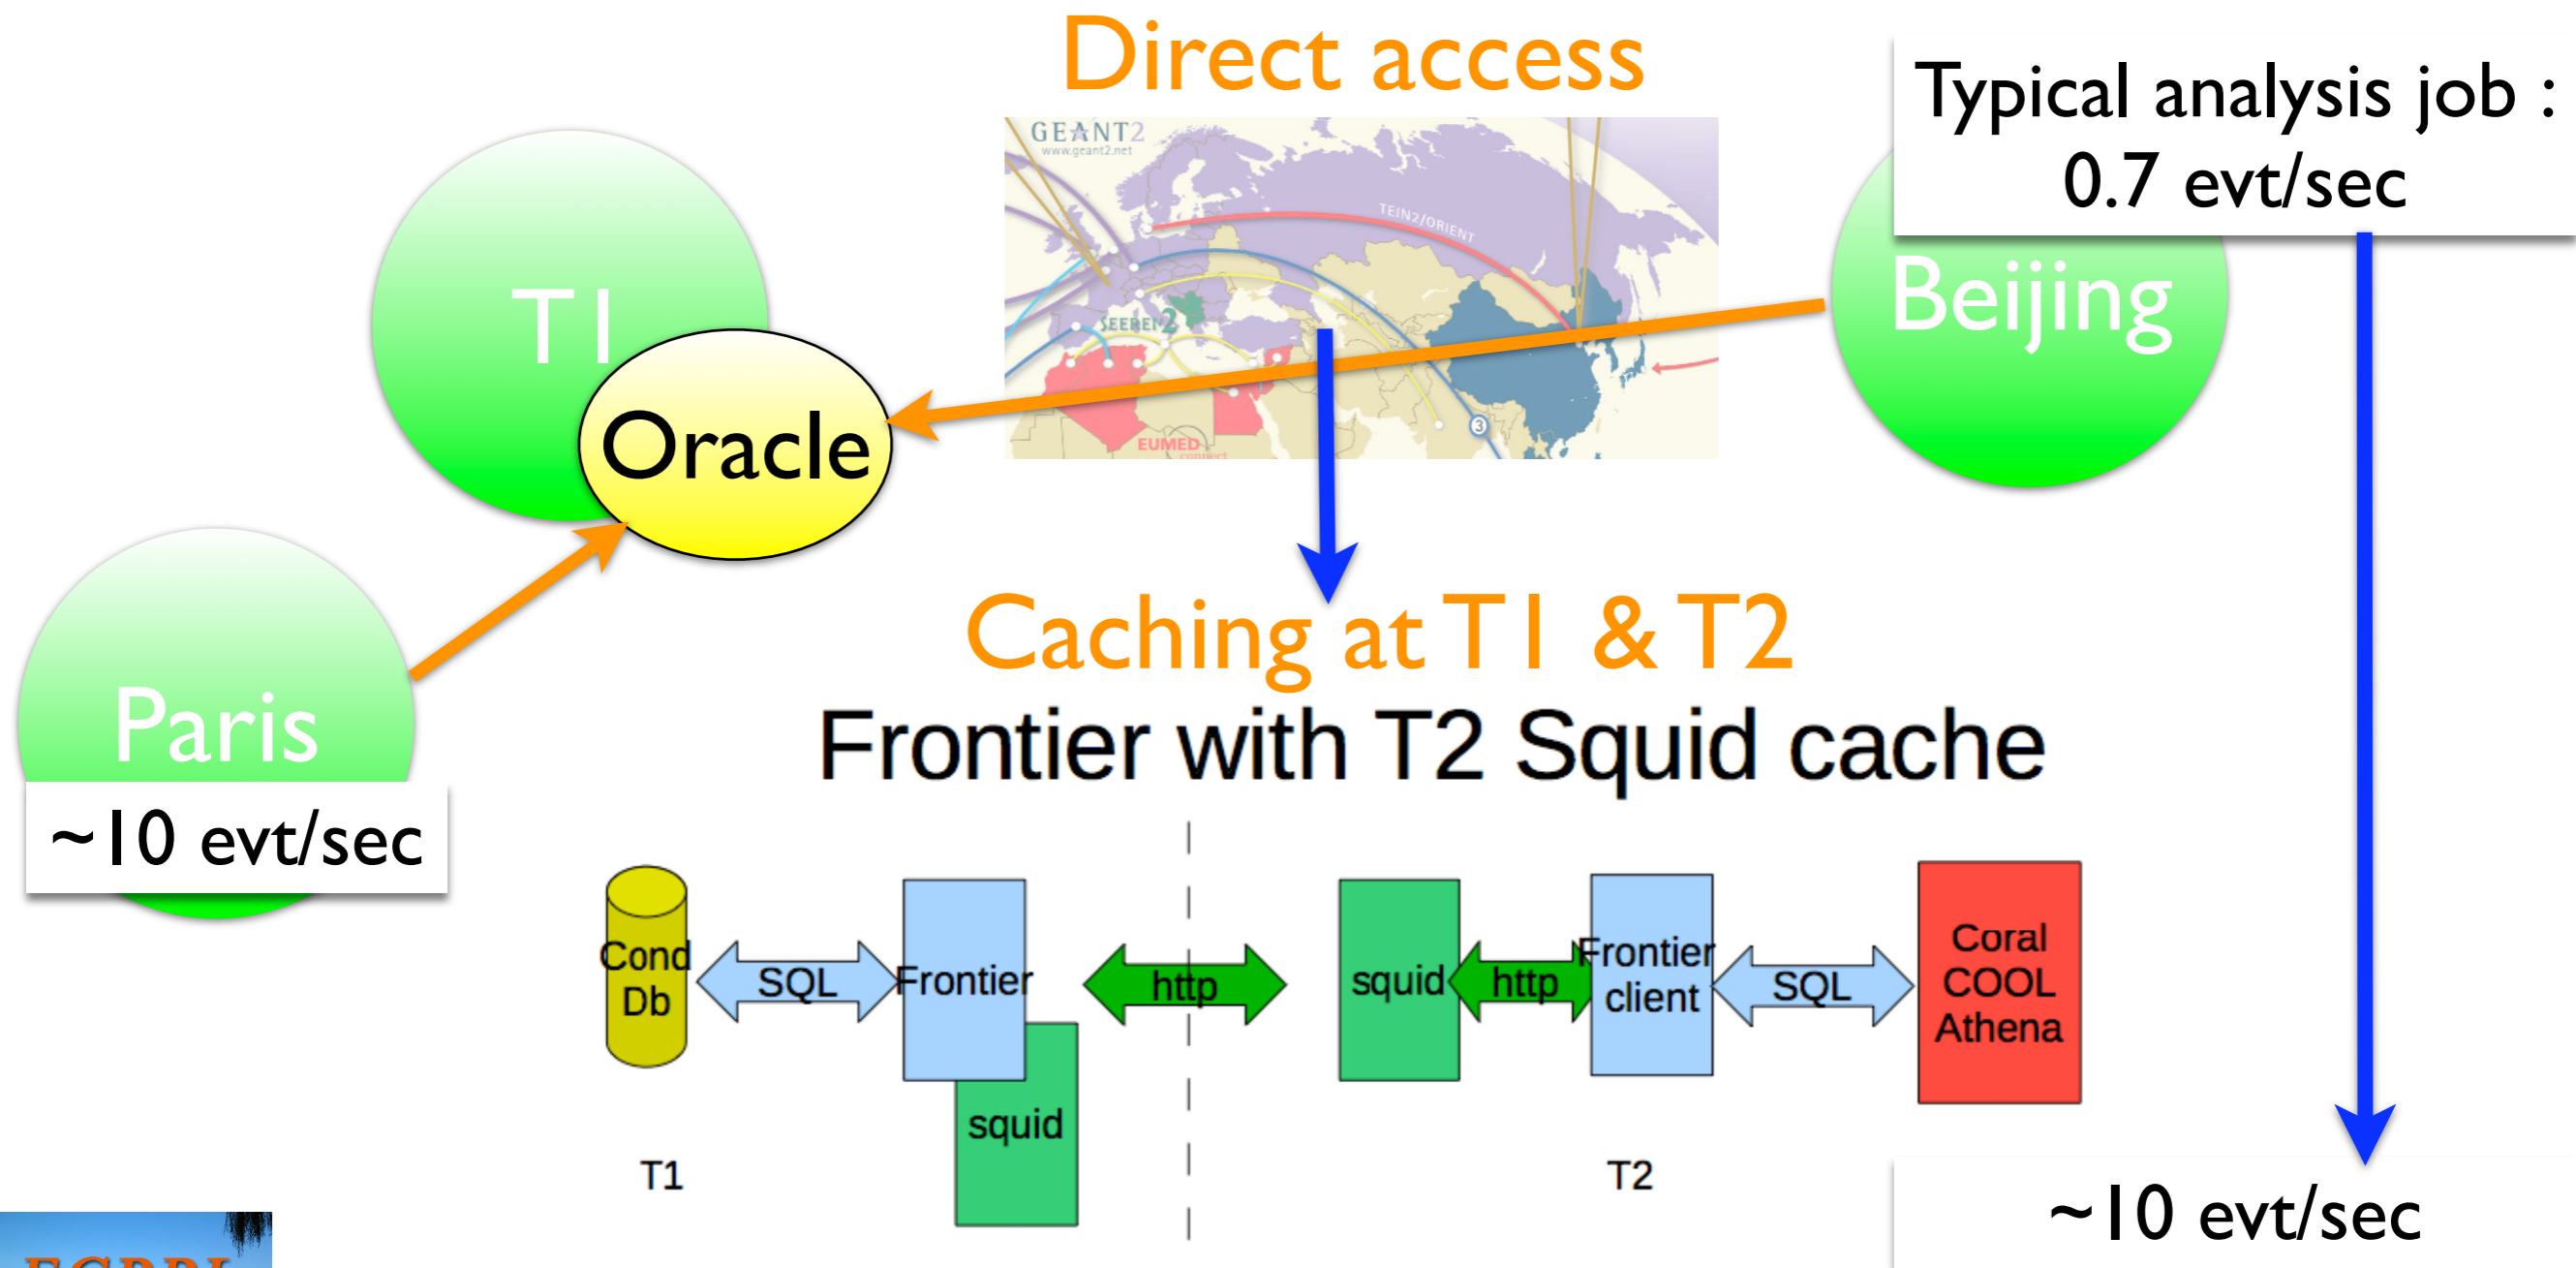
Contribution to monitoring tool from Erming Pei (IHEP)

Eric Lançon / 2010 FCPPL Workshop, Lyon





Calibration constants access





Some details of ATLAS activities in 2009


- One week stay at Beijing of a French engineer : deployment of French-Cloud T2s common storage setup
- 2 month stay of Erming Pei at CERN (fall 2009) work with FR-cloud team on
 - Analysis Stress Tests
 - Squid server deployment and test on FR-cloud
 - Data distribution monitoring
- 3 people from FCPPL participated to the ATLAS Asia-Pacific Computing workshop in Tokyo
- Monthly phone meetings



Some details of CMS activities in 2009

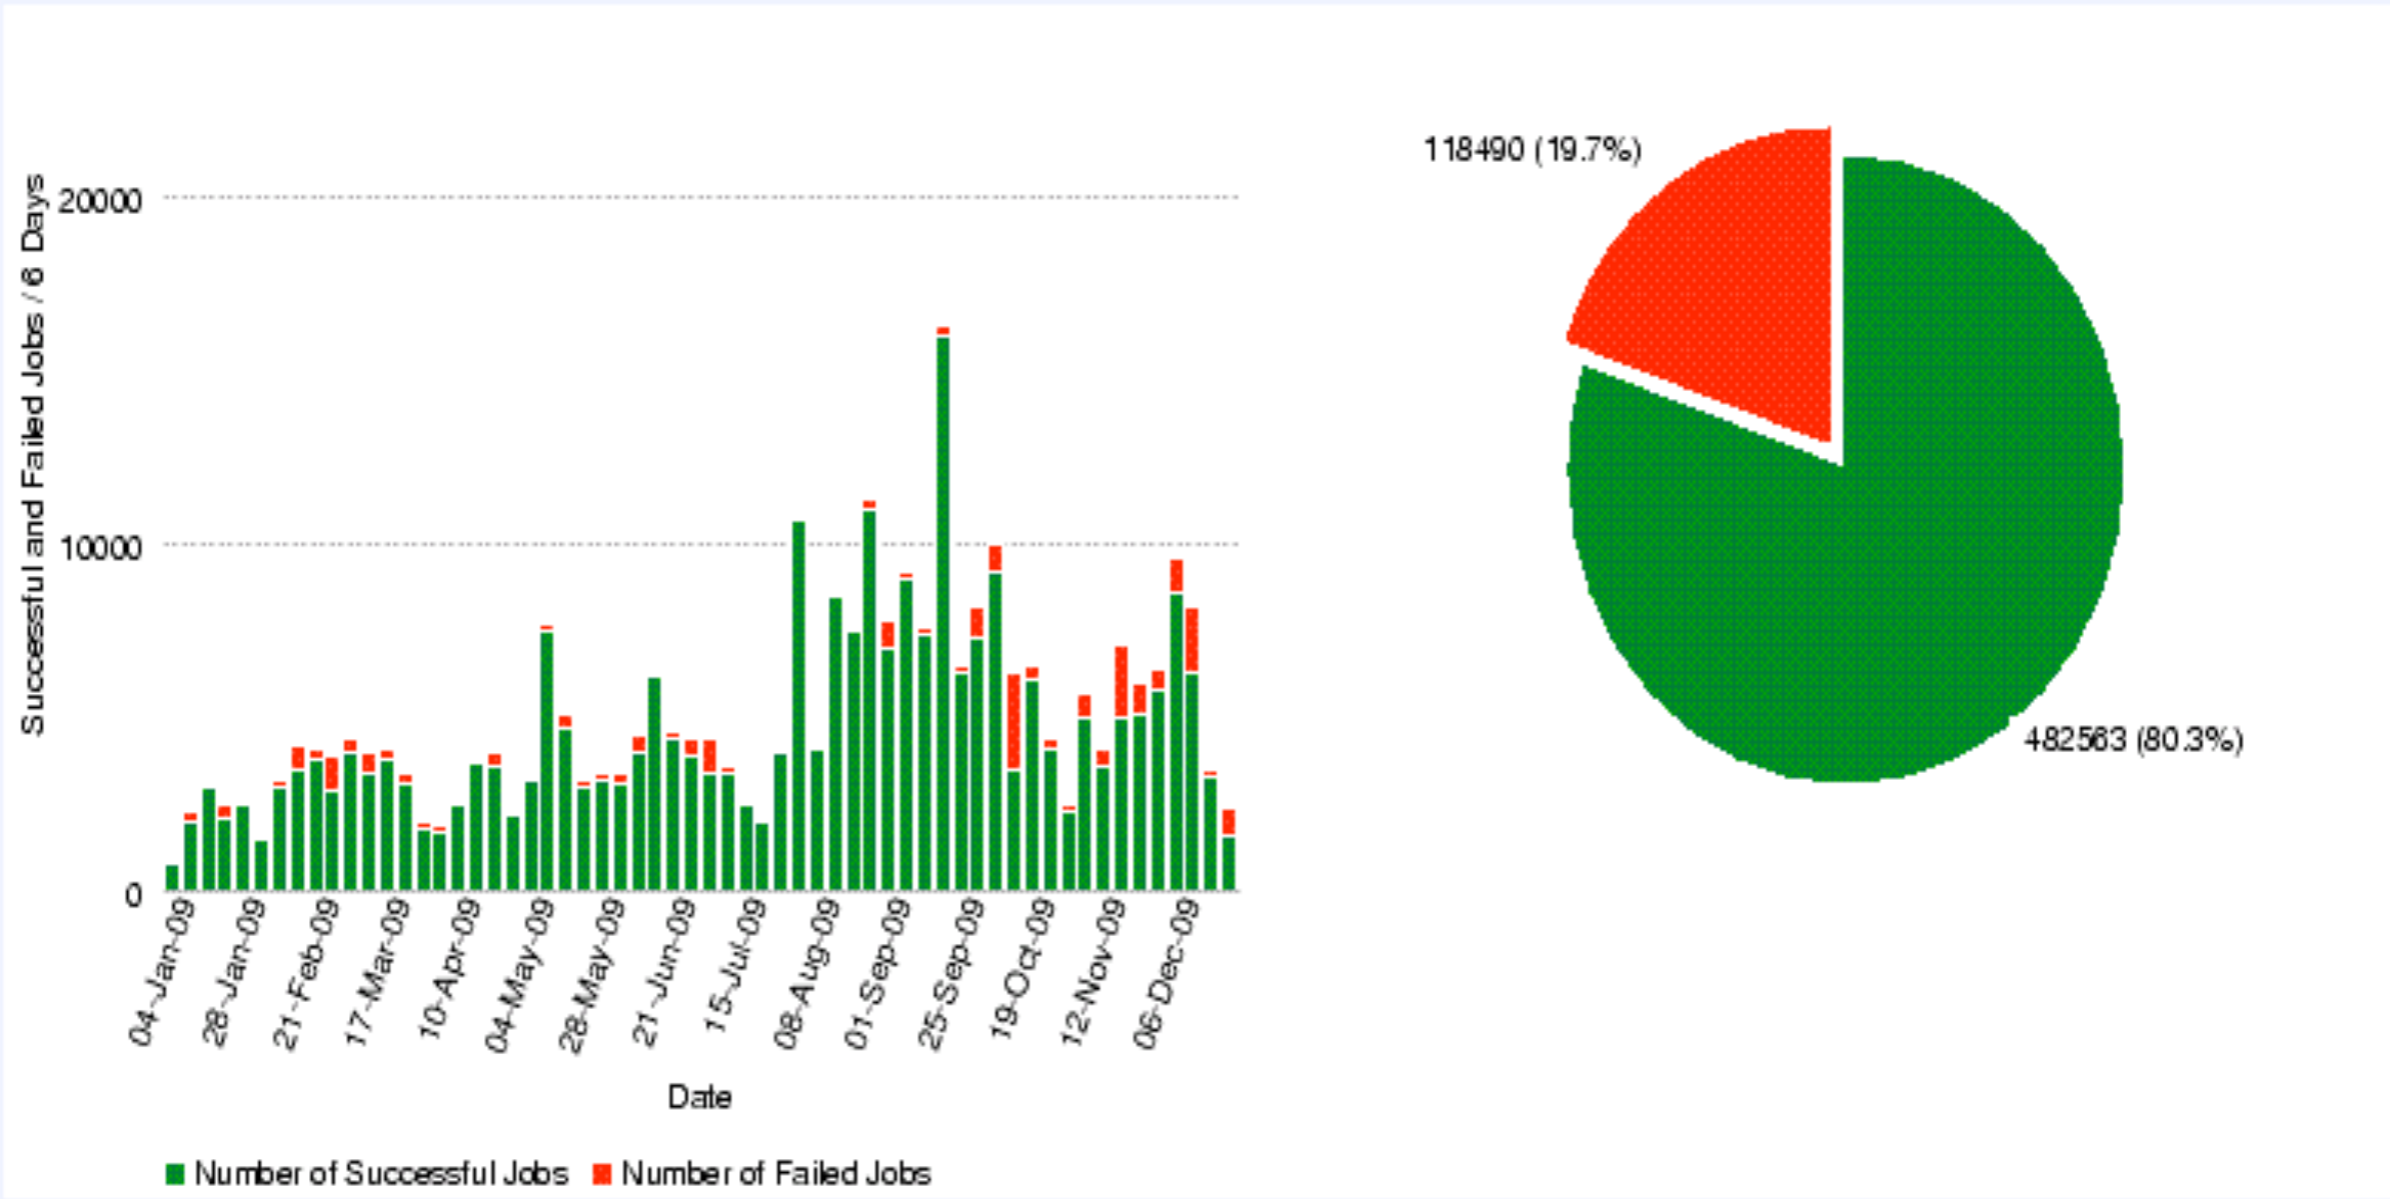
- Establish transfer links between Beijing, CSCS, HIP, MIT sites
- Establish real time video links between CERN CMS computing center, Point 5, FNAL, Beijing Centers
- Deploy ispy real-time event display at cmsRoc@beijing
 - Take on part of CMS computing shift on cmsRoc@beijing
- Deploy CERNVM server in Asian region issue
- Visit to CC-IN2P3
 - Sharing experience on the dCache system
 - Communicating and Sharing experience on the xrootd/proof system
- Participation to :
 - Weekly CMS facility meeting for T1 and T2 sites
 - CMS offline computing workshop






CMS Jobs at IHEP in 2009

Number of Successful and Failed Jobs (Show Percentage of Successful Jobs)




*More than 600k jobs
Success rate 80% (including user application errors)*





Beijing - CMS connectivity

CMS PhEEx - Cumulative Transfer Volume
52 Weeks from Week 51 of 2008 to Week 51 of 2009




- T2_CN_Beijing to T1_US_FNAL_Buffer
- T2_CN_Beijing to T1_DE_FZK_Buffer
- T2_CN_Beijing to T1_FR_CCIN2P3_Buffer
- T2_CN_Beijing to T1_TW_ASGC_Buffer
- T2_CN_Beijing to T1_UK_RAL_Buffer
- T2_CN_Beijing to T1_ES_PIC_Buffer
- T2_CN_Beijing to T2_US_MIT
- T2_CN_Beijing to T2_CH_CSCS
- T2_CN_Beijing to T2_FI_HIP
- T2_CN_Beijing to T1_DE_KIT_Buffer
- T2_CN_Beijing to T1_IT_CNAF_Buffer

Total: 119.95 TB, Average Rate: 0.00 TB/s

Link with 8 T1s & 3 T2s established

52 Weeks from Week 51 of 2008 to Week 51 of 2009



- T1_US_FNAL_Buffer to T2_CN_Beijing
- T1_TW_ASGC_Buffer to T2_CN_Beijing
- T1_CH_CERN_Buffer to T2_CN_Beijing
- T1_IT_CNAF_Buffer to T2_CN_Beijing
- T1_UK_RAL_Buffer to T2_CN_Beijing
- T1_ES_PIC_Buffer to T2_CN_Beijing
- T1_FR_CCIN2P3_Buffer to T2_CN_Beijing
- T1_DE_FZK_Buffer to T2_CN_Beijing
- T2_US_MIT to T2_CN_Beijing
- T2_FI_HIP to T2_CN_Beijing
- T2_CH_CSCS to T2_CN_Beijing
- T1_DE_KIT_Buffer to T2_CN_Beijing

Total: 153.60 TB, Average Rate: 0.00 TB/s




Last Week...

LHC Page1 Fill: 1005 E: 3500 GeV 30-03-2010 14:37:02

PROTON PHYSICS: STABLE BEAMS

Energy: 3500 GeV I(B1): 1.54e+10 I(B2): 1.33e+10




Comments 30-03-2010 14:36:55 :


More than 1h of stable beams!

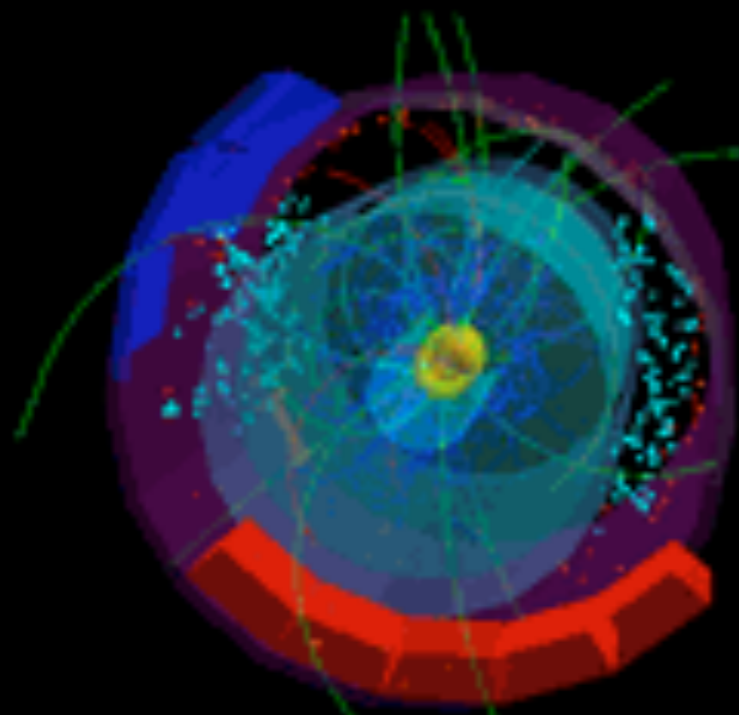
No black holes

...yet




Eric Lançon / 2010 FCPPL workshop, Lyon



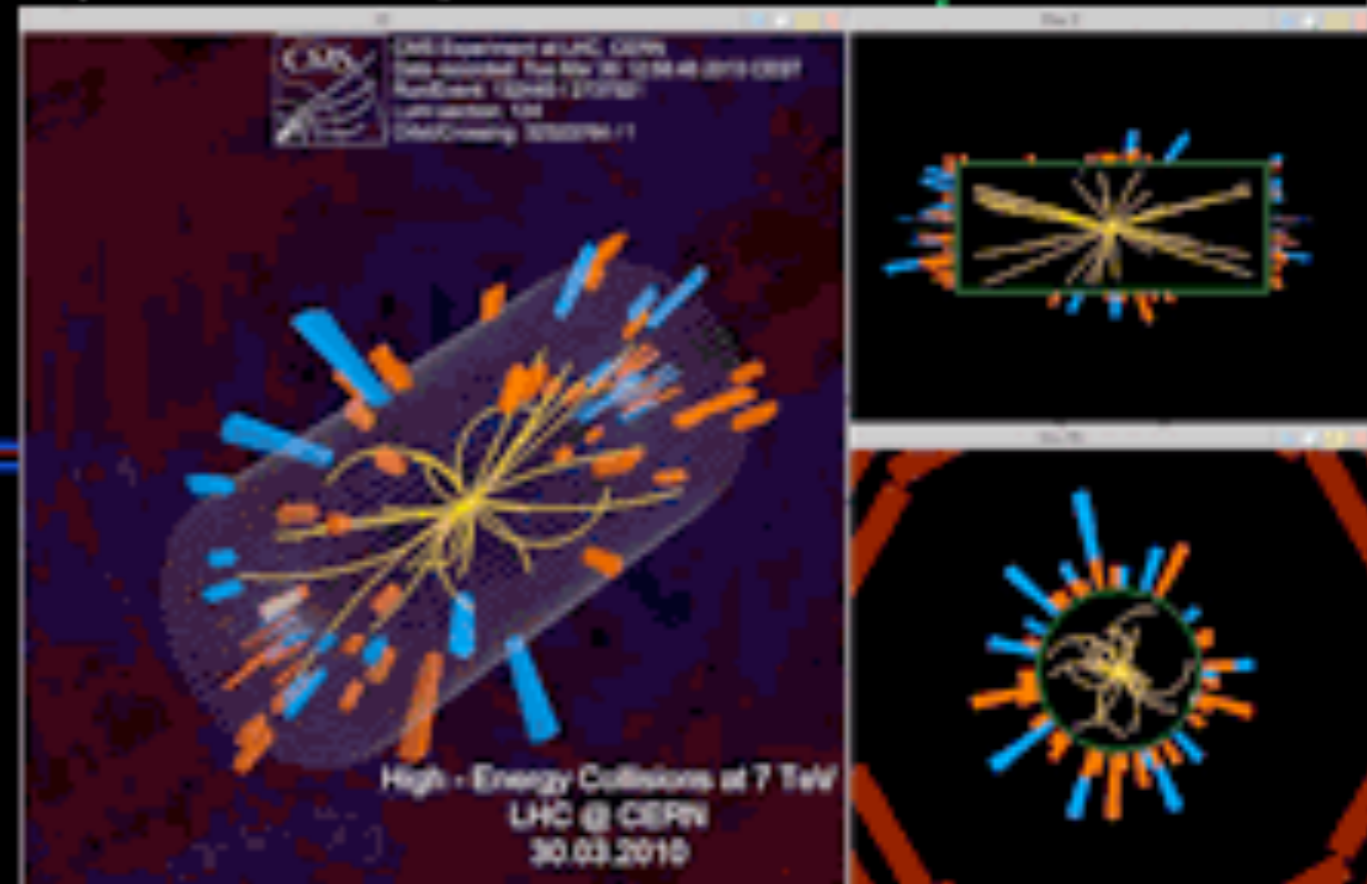



Collision Event at 7 TeV



ATLAS EXPERIMENT
2010-03-30, 12:58 CEST
Run 152366, Event 368299
<http://atlas.web.cern.ch/Atlas/public/EVTDISPLAY/events.html>

Physics can start!



High - Energy Collisions at 7 TeV
LHC @ CERN
30.03.2010



Prospect for 2010

- Distribute & Analyze Data
 - ATLAS & CMS : cooperation on analysis hardware setup
- ATLAS :
 - Tighten collaboration for more integrated & efficient cloud co-operation
 - PhD student on Data Distribution system

