



# World-wide daily computing operations in ATLAS

Xavier Espinal and Kaushik De

PIC  
port d'informació  
científica

PIC (Port d'Informació Científica, Barcelona)  
UTA (University of Texas at Arlington)



## ATLAS Distributed Computing shift teams scope:

ATLAS distributed computing spans 10 Tier-1 and about 80 Tier-2 centers, supporting common activities like data distribution, Monte Carlo simulation, data re-processing and data analysis. Such a wide range of activities with cross dependencies between different software systems and site managers require constant monitoring to ensure Quality of Service (QoS) in achieving ATLAS physics goals. Since 2005, we had various shift teams in EGEE, OSG and NG, which turned out to be critically important in supporting MC production. It was clear that we needed to consolidate shifts team for global ADC activities, especially as LHC data approaches.

## ATLAS Distributed shift team status:

Since January 2008 a distributed shift team (ADCoS) has been operating regularly. The group started by combining previous simulation shifts teams and rapidly recruited new people. Its objective is to achieve QoS for ATLAS distributed computing by providing 24h/7d support without anyone taking night shifts (by having people in three time-zones - shifts are taken remotely from home institutes). The main duty of the ADCoS team is to monitor ATLAS production activities: site performance, software problems (feedback to developers) and status of ATLAS services. Since August 2008, ADCoS is covering almost constantly 24h/6d

### - Team Duties:

#### - Simulation and reprocessing activities:

- Spot, report and follow-up of sites issues
- Software bugs
- Data problems: missing files, corrupted files,...

#### - Data export and placement:

- Data export (T0-T1-T2), calibration data streaming
- Simulation data transfers
- Follow-up critical data replication (DB releases)

#### - Critical ATLAS services:

- ATLAS databases
- Central services

### - Manpower:

- 45 people involved: EU(28) and US(10) and ASIA/Pacific(7)

### - Operations mode:

- Regularly covering 24h/6d

### - Communication:

- Joint eLog common for ATLAS distributed computing activities
- Creation of virtual control room (skype based) where shifters on duty, expert and co-ordinators regularly join.
- Regular weekly phone meetings since Jan-2008
- Web based shift summaries interface is in place:
  - Eases time-zone hand-over among shifters and holds repository



## ATLAS Computing Requirements:

ATLAS data rate at 200 Hz, avg. event size: 1.6 MB (320 MB/s)

~2x10<sup>9</sup> events per year and ~ 10 PB of data per year

Data will be analyzed by ~2000 physicist worldwide

Large amount of simulated data required

- physics background and efficiency calculations
- physics performance estimates
- software and computing infrastructure tests

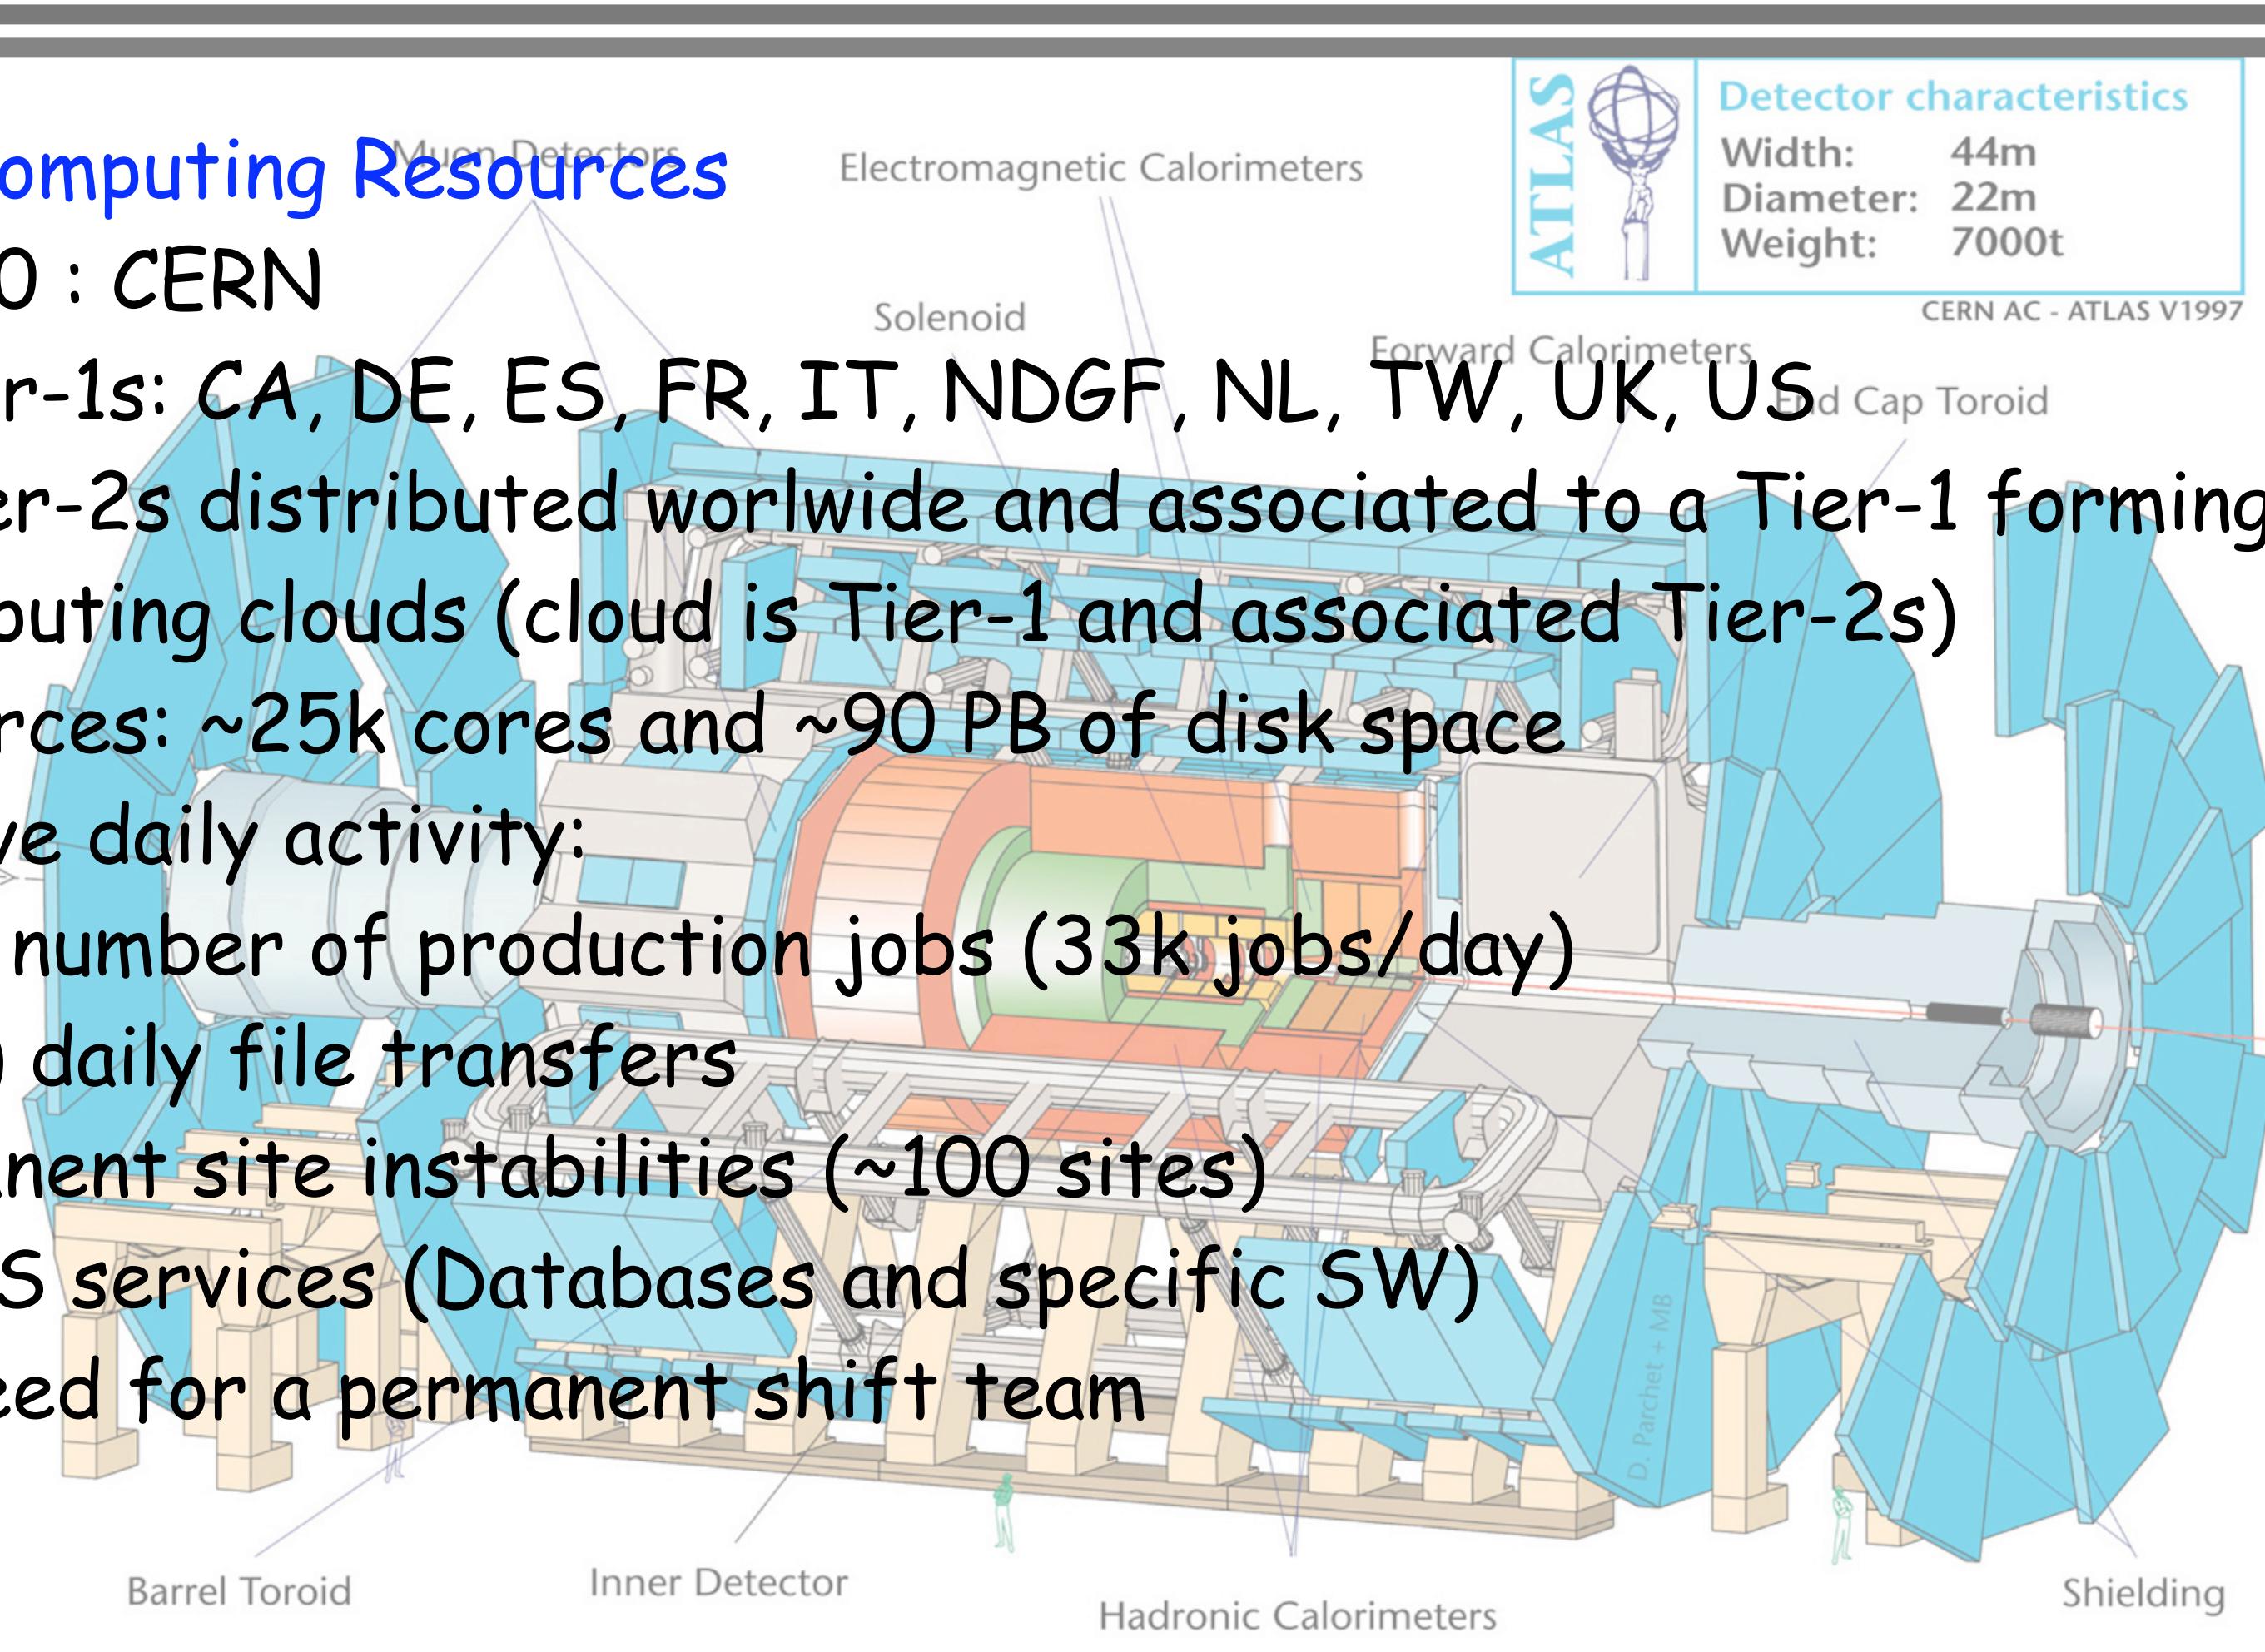
Funds, electrical power, HR and scaling issues for a single all-purpose lab too high:

- Need for grid-computing technology

Responsibility for storage, generation, re-processing of data is shared

## ATLAS Computing Resources

- 1 Tier0 : CERN
- 10 Tier-1s: CA, DE, ES, FR, IT, NDGF, NL, TW, UK, US
- 80 Tier-2s distributed worldwide and associated to a Tier-1 forming 10 computing clouds (cloud is Tier-1 and associated Tier-2s)
- Resources: ~25k cores and ~90 PB of disk space
- Intensive daily activity:
  - Large number of production jobs (33k jobs/day)
  - O(1M) daily file transfers
  - Permanent site instabilities (~100 sites)
  - ATLAS services (Databases and specific SW)
- Clear need for a permanent shift team



### - Monitoring tools:

- Panda monitor, ATLAS production and data management dashboards

- ATLAS Central services (CERN SLS)

- ATLAS databases monitoring

### - Consolidated layered structure:

- Trainee shifters (2 day shift), Senior shifters (2 day shift), Expert shifter (one week shift)
- Triage and escalate

### - ADCoS in numbers:

- GGUS and US/RT used for site related issues:

- 738 GGUS tickets in the past year (>2 tickets/day)

- Savannah portals :

- data replication (>1k tickets past year)
- re-processing and SW issues (>2.2k tickets past year)

- eLog: spread information among the ADCoS Team

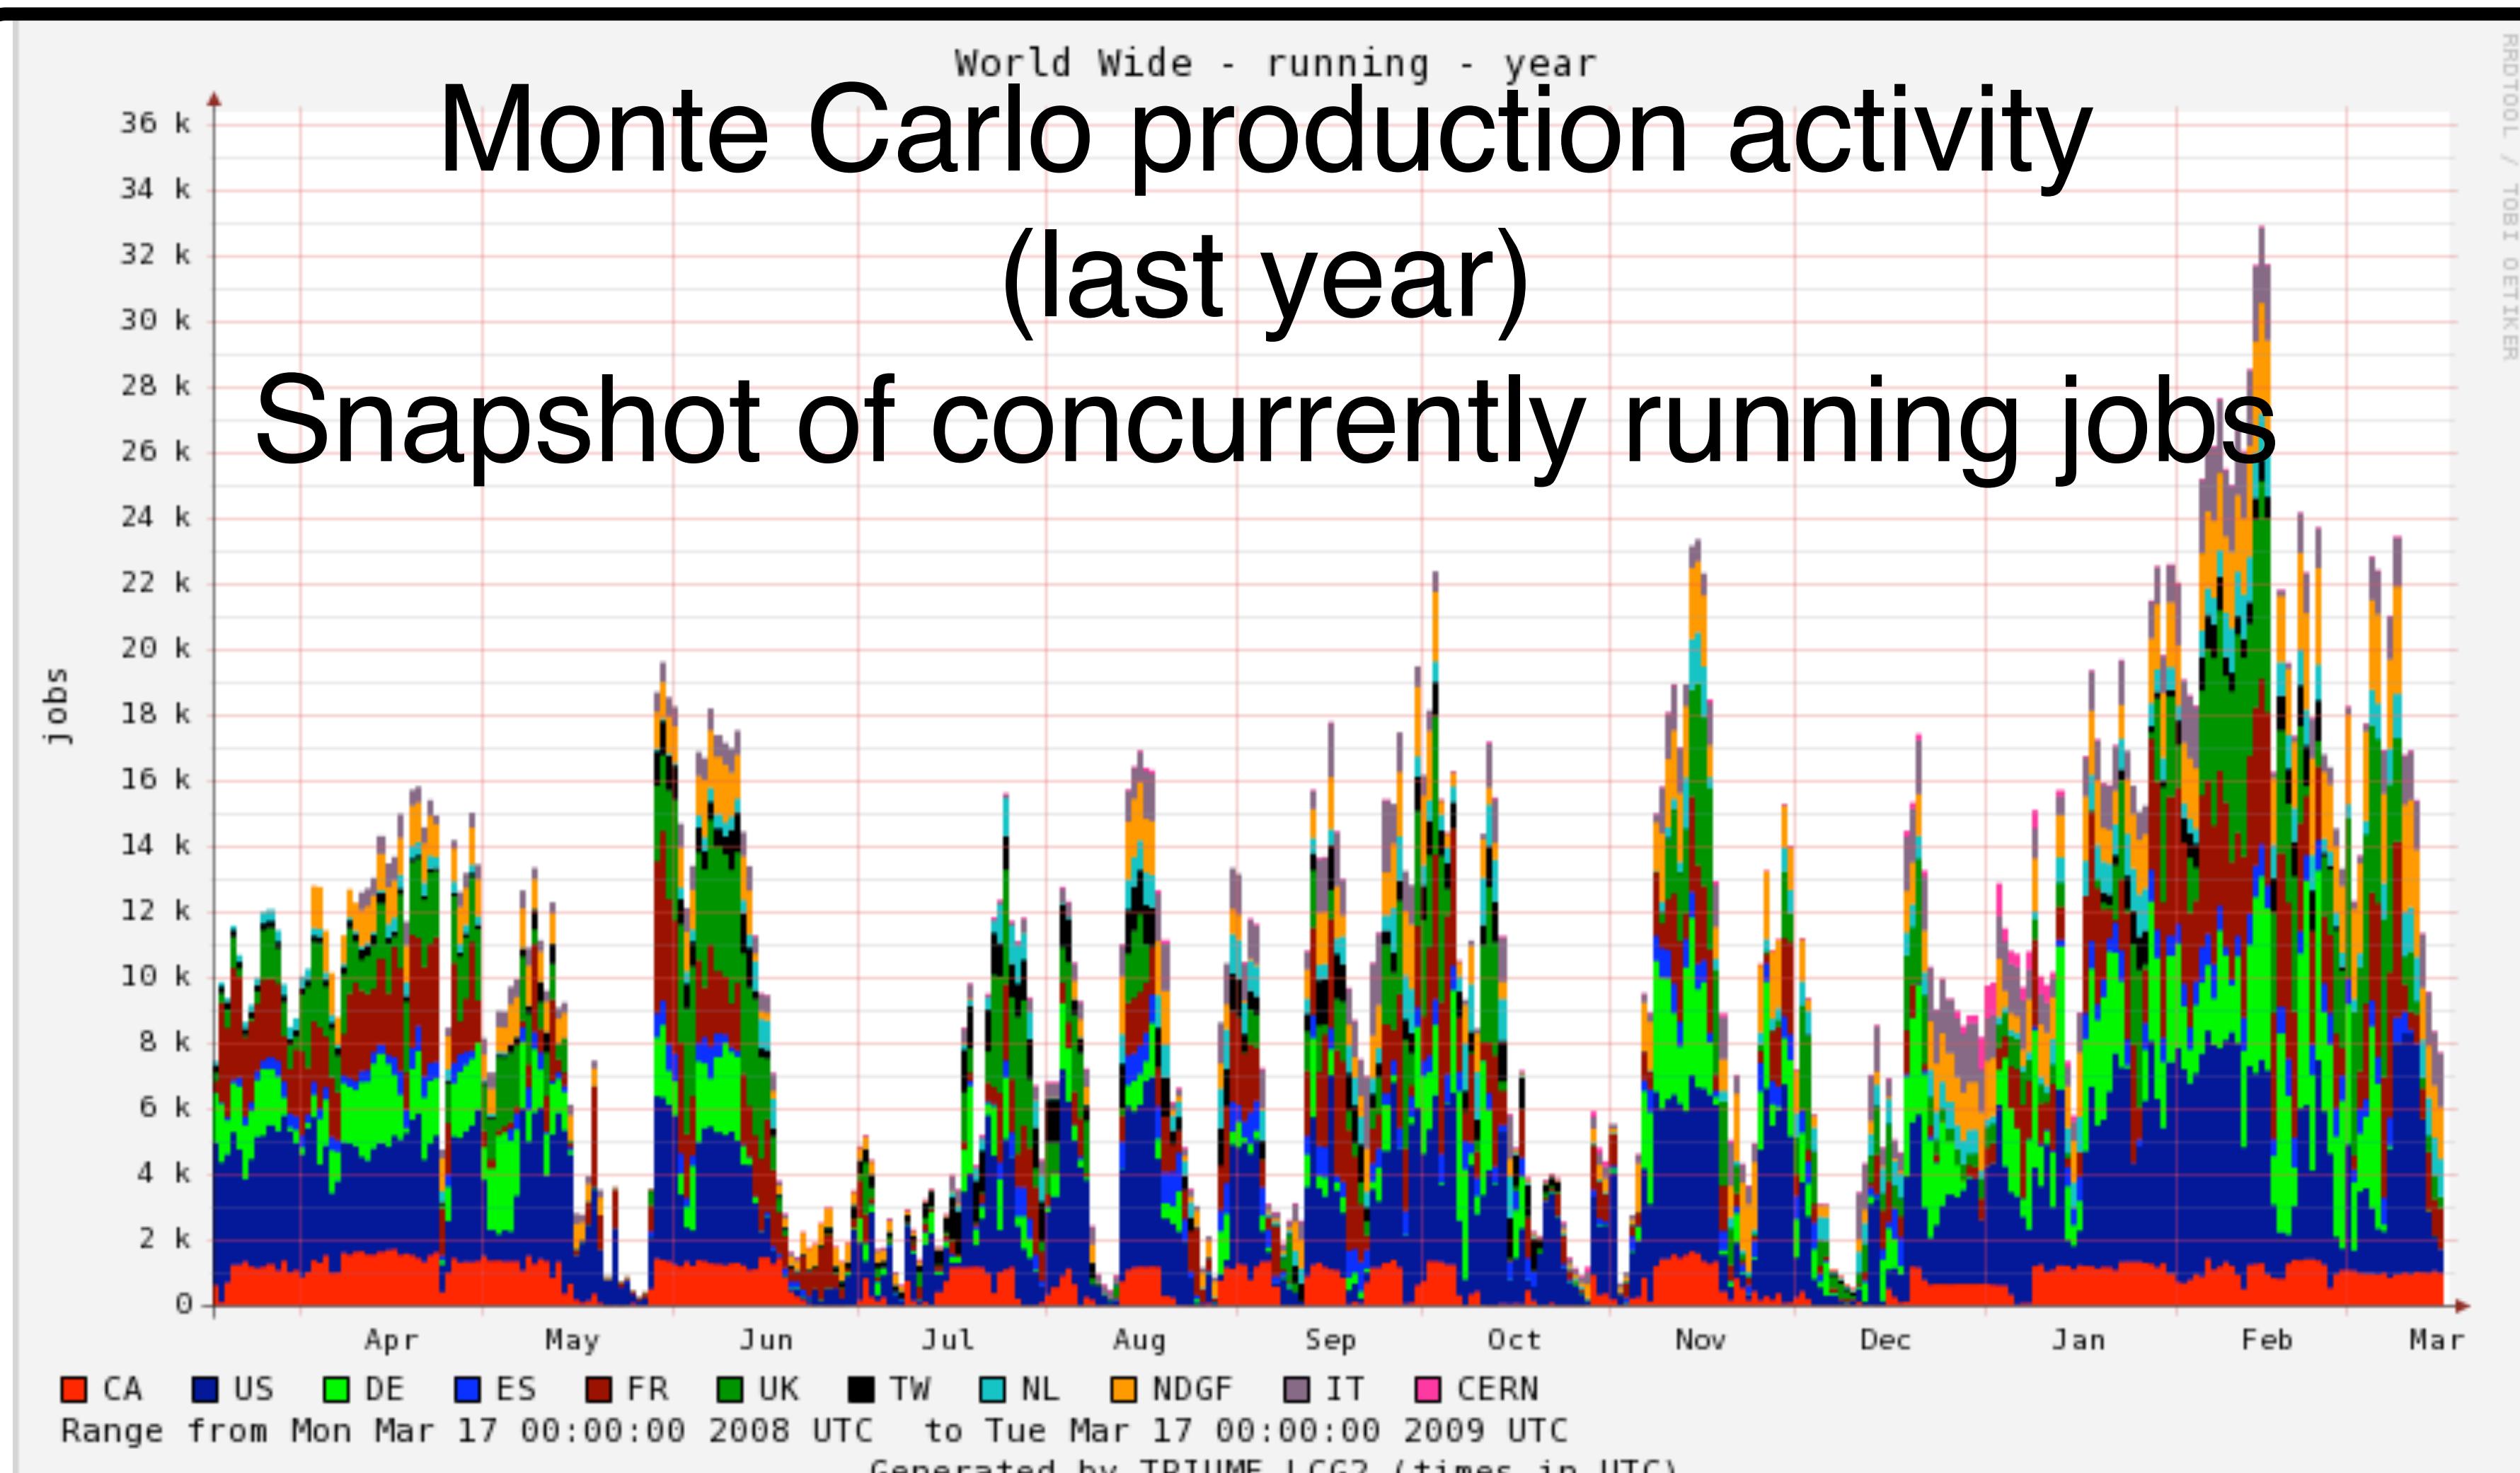
- ~3k eLog entries

- Alarms to Experts

## Monte Carlo production activity

(last year)

Snapshot of concurrently running jobs



More information: <https://twiki.cern.ch/twiki/bin/view/Atlas/ADCoS>