



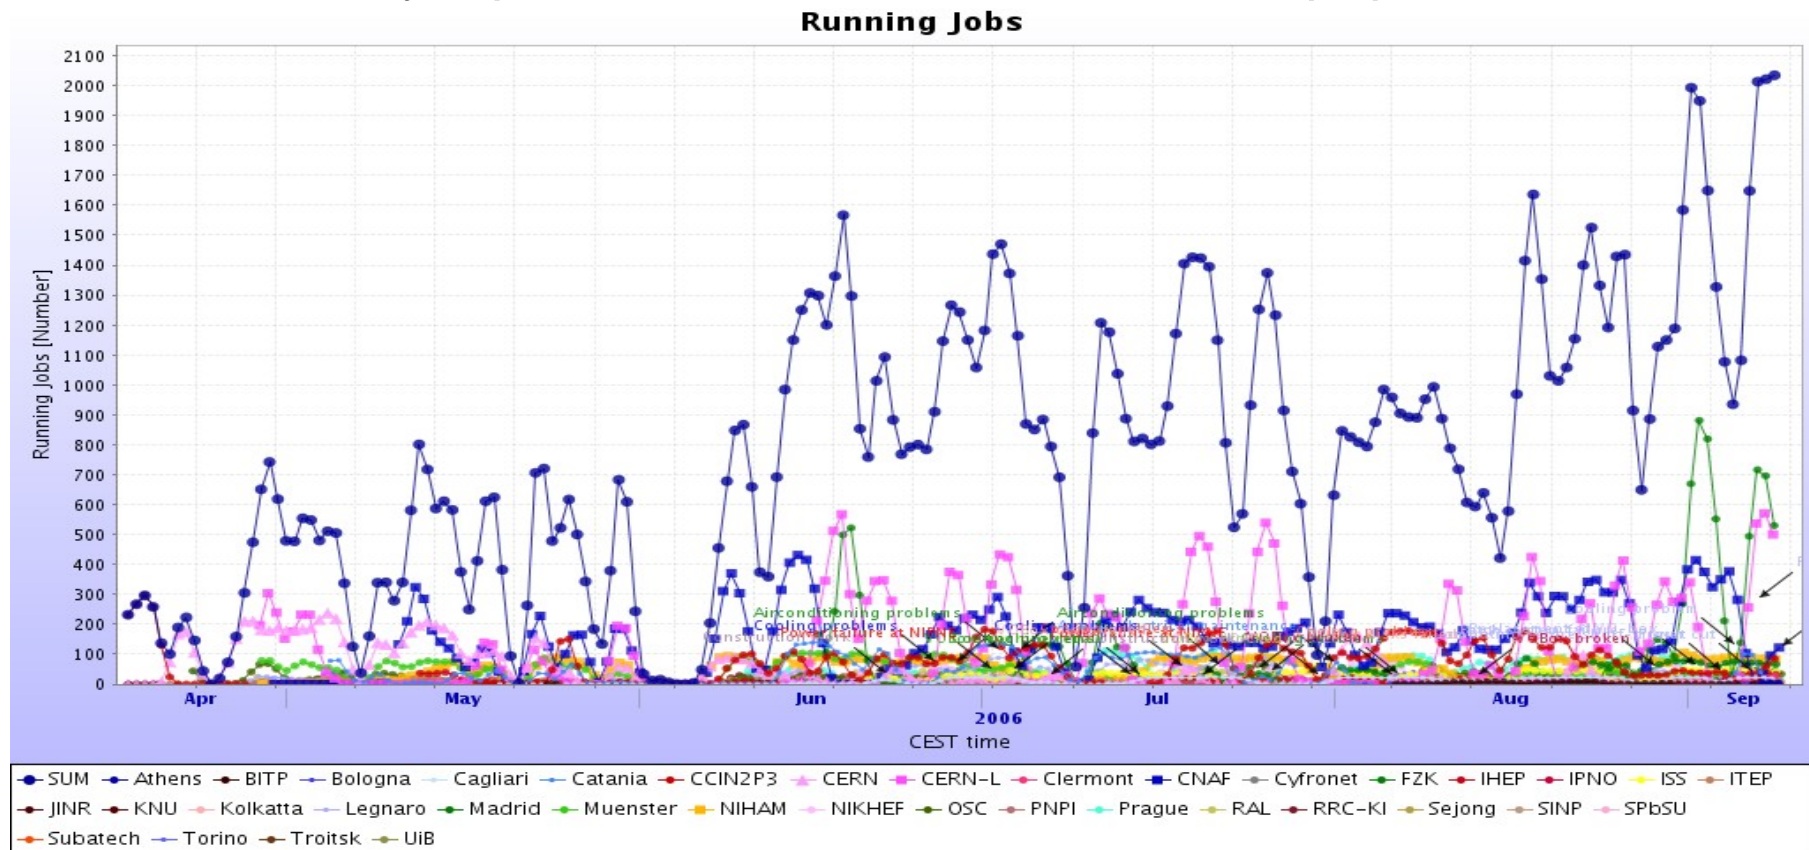
PDC'06

Brief summary of PDC'06 goals

- ❑ Production of MC event for detector and software performance studies
- ❑ Verification of the ALICE distributed computing model
 - Integration and debugging of the GRID components into a stable system
 - ❑ LCG Resource broker, LCG file catalogue, File transfer system, Vo-boxes
 - ❑ AliEn central services – catalogue, job submission and control, task queue, monitoring
 - Distributed calibration and alignment framework
 - Full data chain – RAW data from DAQ, registration in the AliEn FC, first pass reconstruction at T0, replication at T1
 - Computing resources – verification of scalability and stability of the on-site services and building of expert support
 - End-user analysis on the GRID

History of PDC'06 (2)

- Continuous running since April 2006
 - Test jobs, allowing to debug all site services and test the stability
 - From July – production and reconstruction of p+p MC events




History of PDC'06 (2)

- Gradual inclusion of sites in the ALICE Grid - current status:
 - 6 T1s: CCIN2P4, CERN, CNAF, GridKA, NIKHEF, RAL
 - 30 T2s
- Currently available CPU power – 2000 CPUs for ALICE (expected ~4000)
 - Competing for resources with the other LHC experiments
 - Computing centres are waiting for the last moment to buy hardware – will get more for the same price
 - Expect additional resources from Nordic countries and from US (LBL and LLNL)

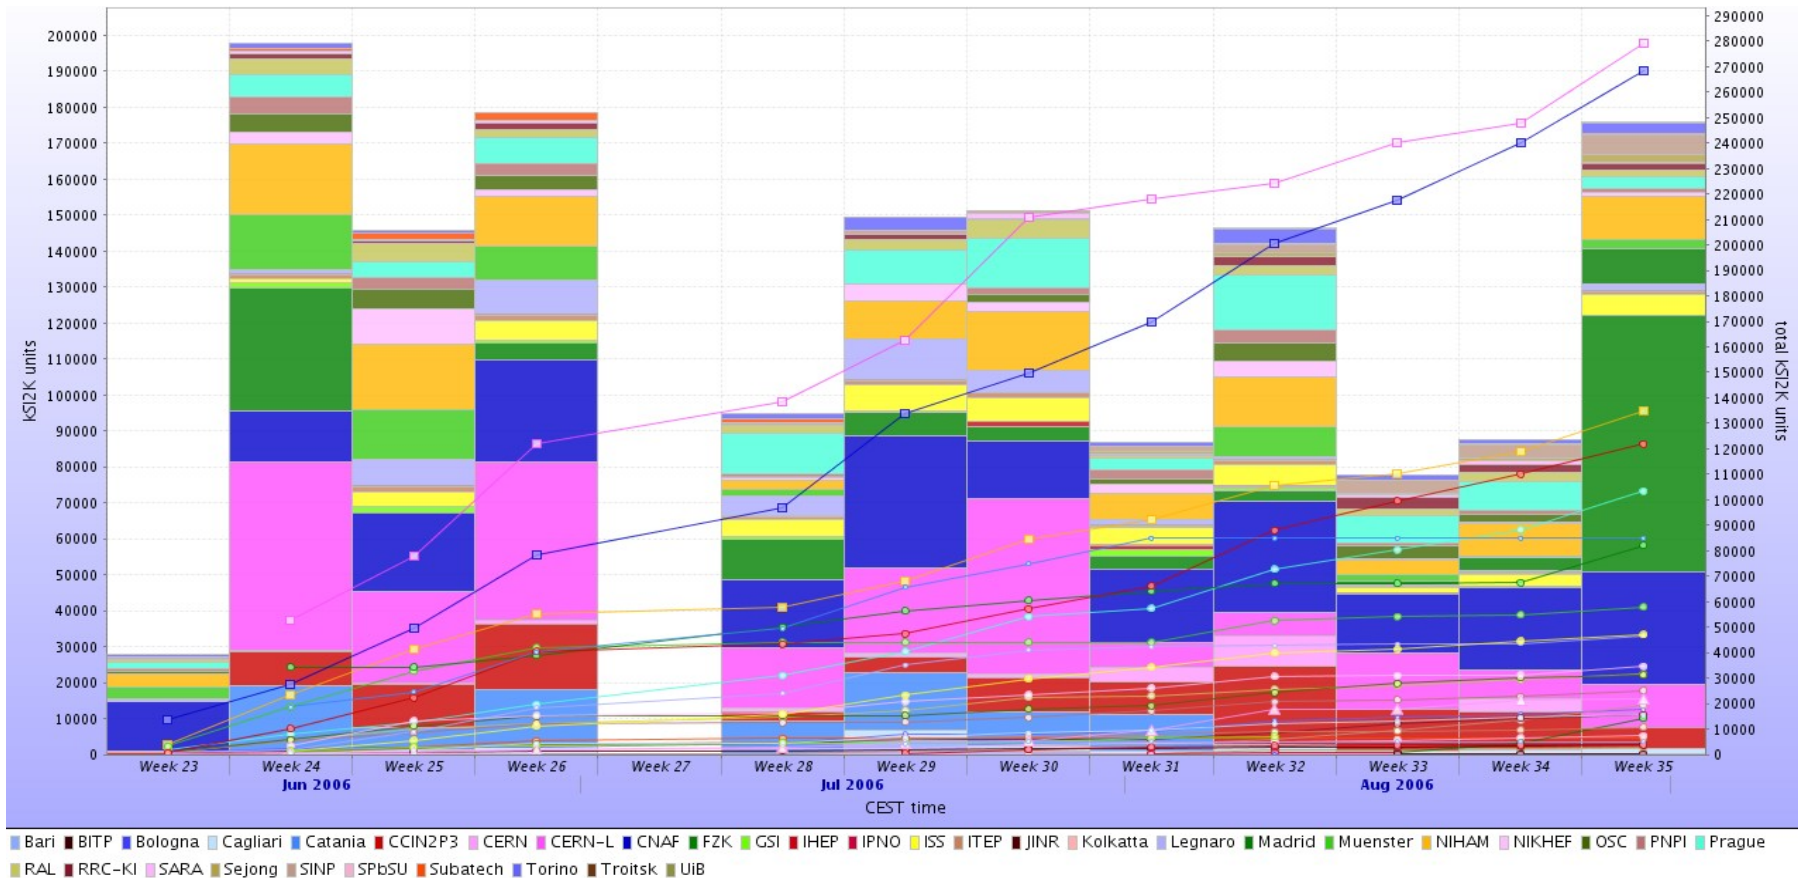
Resources statistics

- Resources contribution (normalized Si2K units): 50% from T1s, 50% from T2s
 - The role of the T2 remains very high!




Resources statistics (2)

- Total amount of CPU work: 2MSi2K units, equivalent to 500 CPUs working continuously since the beginning of the exercise (142 days)




Data movement

- Step 1: produced data is sent to CERN
 - Up to 150 MB/sec data rate (limited by the amount of available CPUs) – ½ of the rate during Pb+Pb data export



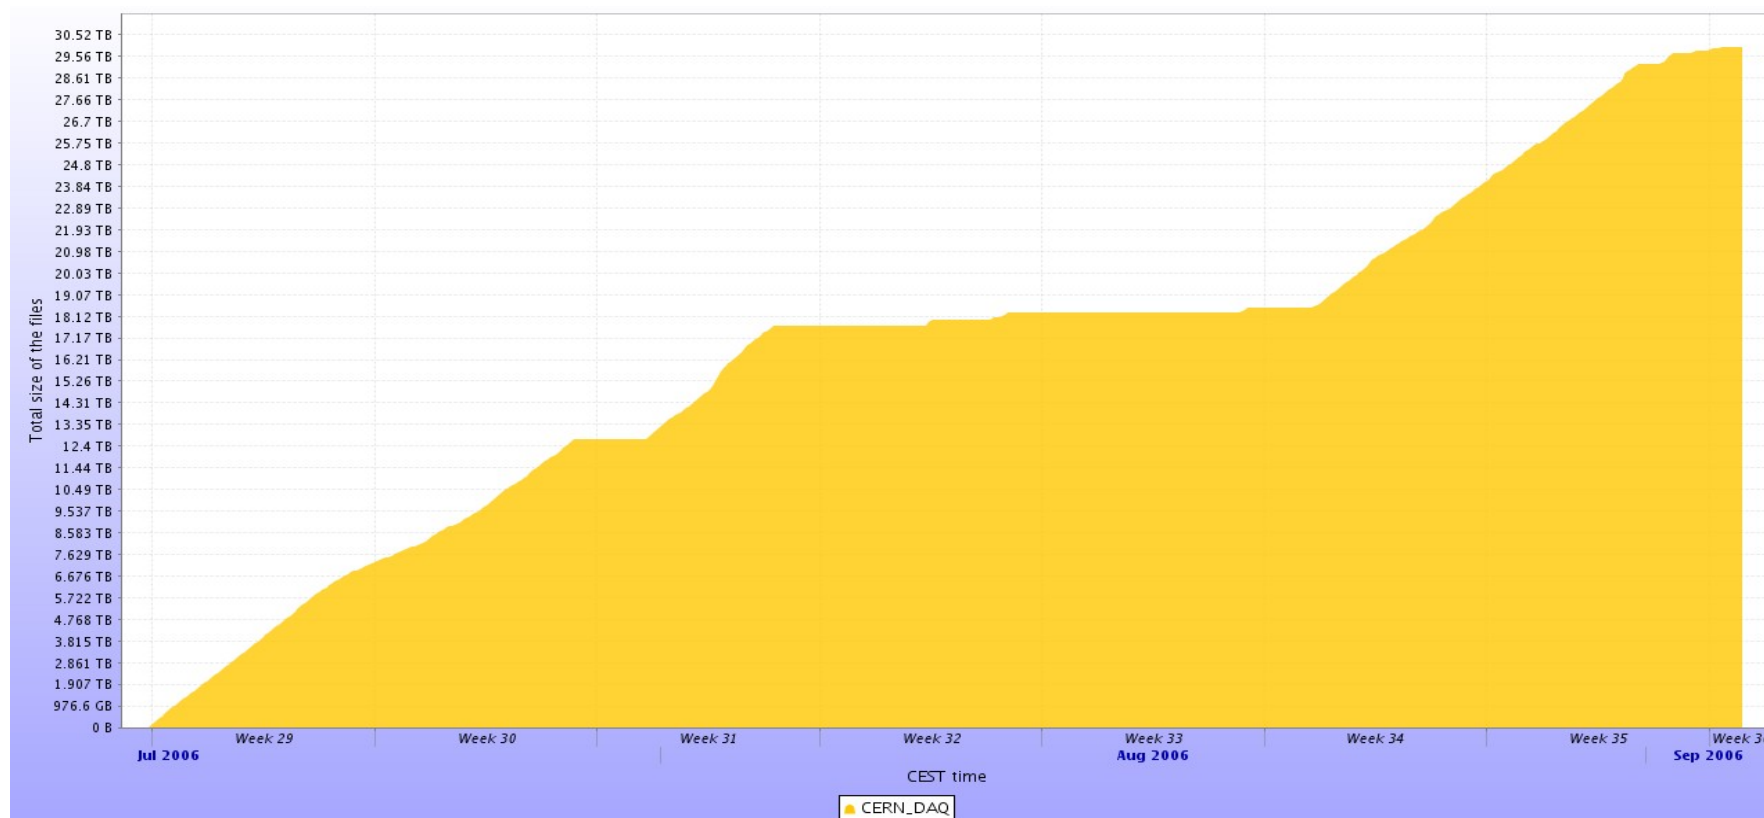
Data movement (2)

- Step 2: data is replicated from CERN to the T1s
 - Test of LCG File Transfer Service
 - Goal is 300 MB/sec – exercise is still ongoing



Registration of RAW data from DAQ

- Dummy registration in AliEn and CASTOR2 – continuous since July 15
- Registration and automatic reconstruction of TPC test data



MC production statistics

- Total of 0.5 PB of data registered in CASTOR2
 - 300K files – 1.6 GB/file
 - Files are combined in archives for optimal load on MSS
- 7M p+p events combined in 70 runs (production is continuing)
 - ESDs, simulated RAW and ESD tag files
- 50K Pb+Pb event in 100 runs
 - ESDs and ESD tag files

Elements of the GRID machinery

- AliEn - Entry point of ALICE to the GRID – both for data production and user analysis
 - Steady improvements in stability and capacity
 - Central services are now at 95% availability and are practically at production level
 - Quick advances on the end-user interface (gshell)
 - Software improvements are coupled with user training – tutorials on analysis are given monthly
 - Analysis job submission and tracking is very robust
- Remaining major issues
 - Access to data – standard storage solutions (CASTOR2, DPM, dCache) with the require by ALICE xrootd interface are still in active development
 - Blockages and high latency of storage is quite frustrating for the end user
 - This has to be solved by the October phase (user analysis) of PDC'06
 - Stability of participation and level of provided resources at the computing centres
 - The resources commitment was overestimated almost everywhere

Elements of the GRID machinery (2)

- LCG components
 - Steady improvement in quality of service
 - ALICE is using fully all stable services of LCG
 - The VO-box infrastructure (required at all sites) has been tested and found to be a scalable and secure solution
 - Plays the role of an interface between the AliEn services and the local LCG/site services
 - Was a highly controversial issue initially, now accepted standard solution at all ALICE sites
- Interfaces to other GRID flavours
 - ARC (NFGF) – prototype exists and is being tested at Bergen
 - OSC (USA) – discussion started
- Remaining major issues:
 - Standard storage solutions, allowing for the sites to install only LCG services (with the ALICE required xrootd)
 - Still in development

Elements of the GRID machinery (2)

- Operational support
 - ALICE's goal is to automatize as much as possible the GRID operations - small team of experts take care of everything
 - Core team (ARDA, SFT, LCG GD, INFN, Offline) are responsible for the software and central services operations, training of regional experts and general dissemination of GRID-related information
 - Regional experts (1 per country) are responsible for the site operations and interactions with the local system administrators
 - Total of 15 people are responsible for the software development, daily operations and support of the ALICE GRID
 - Prolonged periods of running (already 5 months) only possible due to the pool of experts with overlapping knowledge and responsibilities
 - Still, this is quite a strain on very few people – expecting that with the more mature software, the load will go down
- GRID monitoring - MonALISA
 - Almost complete monitoring and history of running on the GRID available at <http://alimonitor.cern.ch:8889>

ALICE sites on the world map



Conclusions and outlook

- Production is running on the GRID continuously since April
 - Testing of the ALICE computing model with ever increasing complexity of tasks
 - Gradual build up of the distributed infrastructure in preparation for data taking in 2007
 - Training of experts and collection of operational experience
 - Improvements of the AliEn software - hidden thresholds are only uncovered under high load
- We are reaching a 'near production' level of service
 - Storage still requires a lot of work and attention
- Next big event – user analysis of the produced data
 - Pilot users help a lot in finding the weak spots and push the development forward