Large scale ATLAS simulated production in EGEE

Xavier Espinal (PIC)
on behalf of the ATLAS EGEE production team

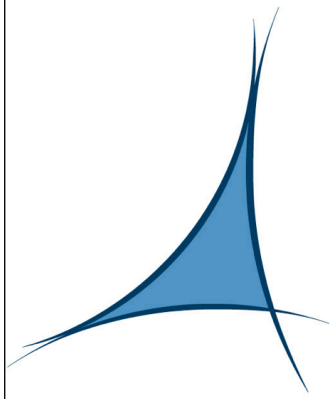

PIC
port d'informació
científica

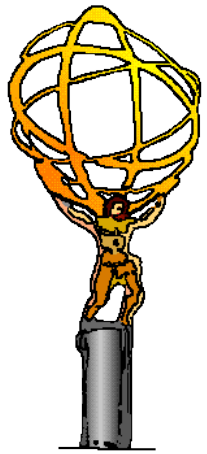

International Conference on Computing in High Energy and Nuclear Physics

2-7 Sept 2007 Victoria BC Canada

OUTLINE

- Introduction
  - The LHC and ATLAS
  - ATLAS computing requirements
- The ATLAS production system
- Experience and results
  - Scope and statistics
  - Simulated events and efficiencies
  - Job errors revision
  - Operations
- The EGEE production operations team
- Next steps
- Summary and conclusions

Introduction: The LHC and ATLAS

- The Large Hadron Collider (LHC) will be worlds largest and most powerful particle accelerator.

- Installed in an underground tunnel of 27 km in circumference astride the border between Switzerland and France.

- Will produce 800 million proton-proton collisions per second, with 14TeV center of mass energy.

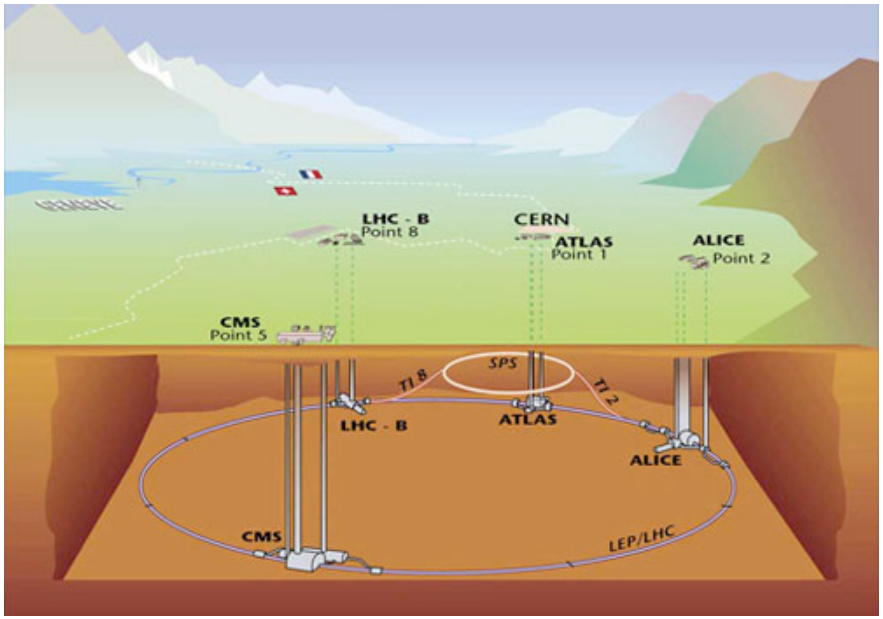

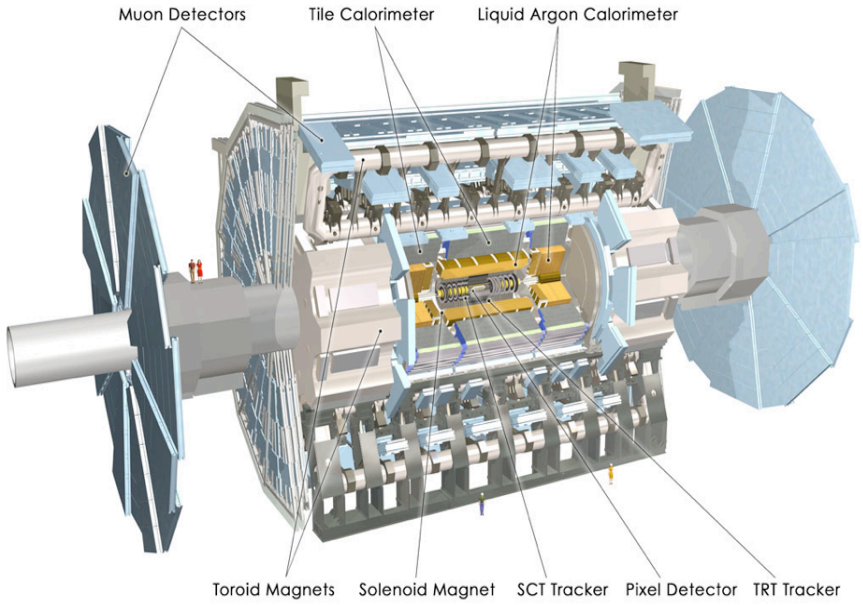

- The ATLAS (A Toroidal LHC Apparatus) detector measures:
  - Diameter: 25 m
  - Barrel toroid length: 26 m
  - Endcap end-wall chamber span: 46 m
  - Overall weight: 7000 Tons

- ATLAS is one of the four LHC detectors, devoted to the study of high-energy proton-proton collisions and heavy ions.

Introduction: ATLAS computing requirements

- The offline computing will have to deal with an output event rate of 200 Hz i.e $2x10^9$ events per year with an average event size of 1.6 Mbyte (320 MB/s).

- ATLAS will produce an amount of data of about 10 Pb per year, to be analyzed by ~2000 physicists from all over the world.

- The design and construction of an experiment like ATLAS requires a large amount of simulated data in order to optimize the detector design, estimate physics performance, and test the software and computing infrastructure.

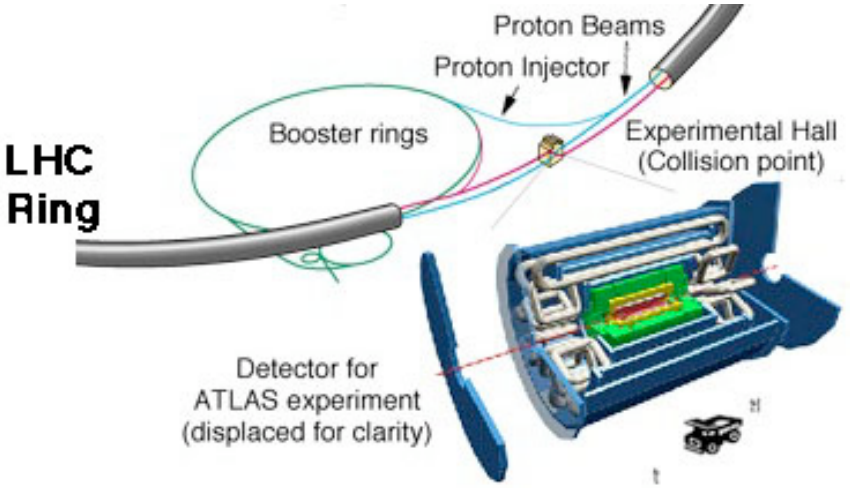

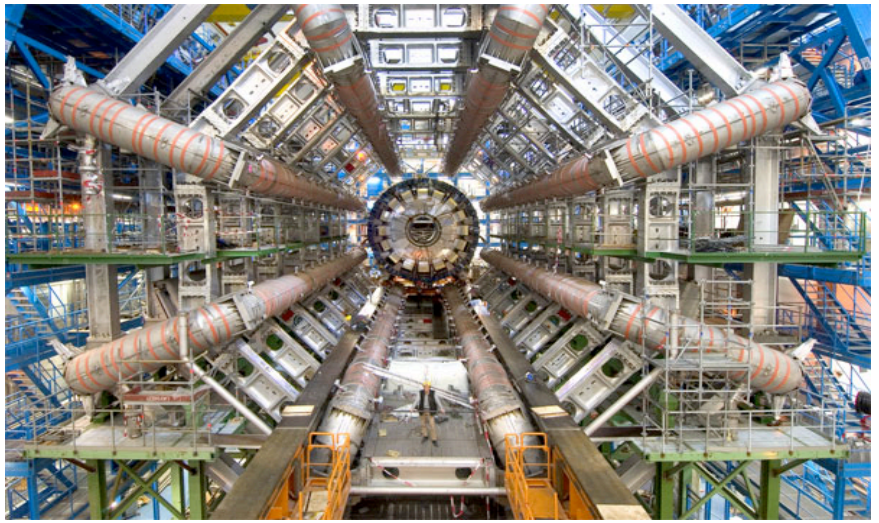

- The funds, electrical power, and human resources necessary for a single, all-purpose computing site would be too great for one laboratory.

- Physicists and computer scientists create a grid-computing system for the experiments, in which more than 100 small and large computing centers share the responsibility for storing, generating, and processing the data.

- Monte Carlo simulated production is performed all over the world both at large computer centers, called Tier-1s, and at smaller sites, called Tier-2s, as well as in institute or university sites, called Tier-3s (ATLAS tiered structure).

The ATLAS Production System

- ATLAS production system provides a common framework in which any grid flavor may be integrated and is a common interface for the three grids.

- Formed from several individual elements which when plugged together provide the required functionality for job submission, tracking, recovery and validation of jobs.

- Composed by: Common database for the production jobs (ProdDB), Common Supervisor (Eowyn), Executors: CondorG, Lexor and Cronus (in EGEE), and Data management system: file cataloging and data transfer (DQ2)

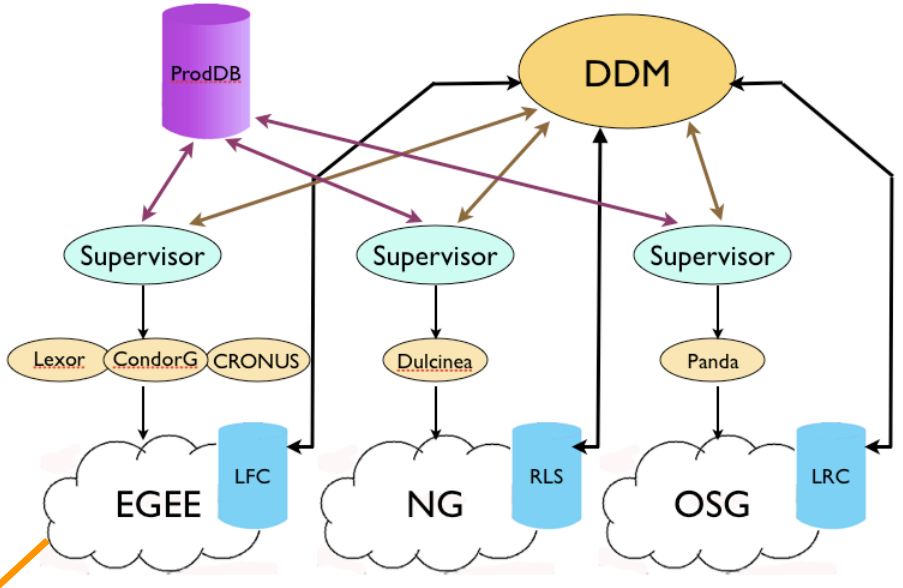

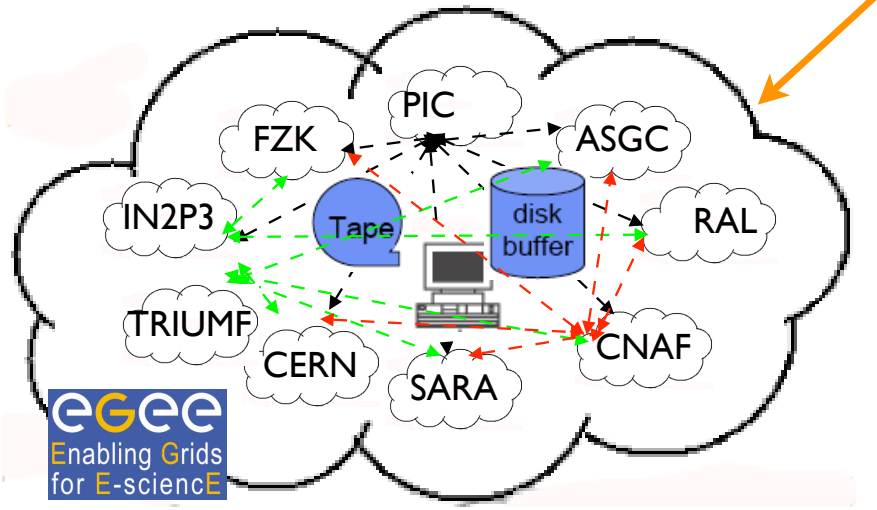

- Executor creates the wrapper files and submit intelligently the jobs to the Grid (free resources, correct software, etc.)

- Within EGEE Grid, the resources of 8 Tier-1s and ~40 Tier-2s yields a power of 26 MSI2k.month and storage disk capacity of ~1.2PB.

- All centers are inter-connected for the data input/output via gLite-FTS (File Transfer Service).

Experience(1/5): Scope and statistics

- One of the targets of the production system is to probe the operability of high level distributed computing, since computing demands of the LHC has no precedents.

- Since November 2006, the ATLAS simulated production in EGEE is supervised by the EGEE production team.

- The current infrastructure successfully covered a ramp-up challenge during late 2006 and early 2007. The target was to finish 20 M events during November and December 2006, and 40M events during the next quarter...

- ...ended earlier than expected as the disk resources of almost all the Tier-1 centers were quickly filled (mid Feb'07)

- Nevertheless ~60M events were produced in the 1st quarter of 2007.

- Hence the production system accomplished the milestone and showed to be able to cope with the ATLAS requirements.

- (Plot) two main zones are clearly seen: ramp-up period (green) and steady state production (red+orange). Finished jobs peaked at 55k jobs finished in a single day (All Grids !).

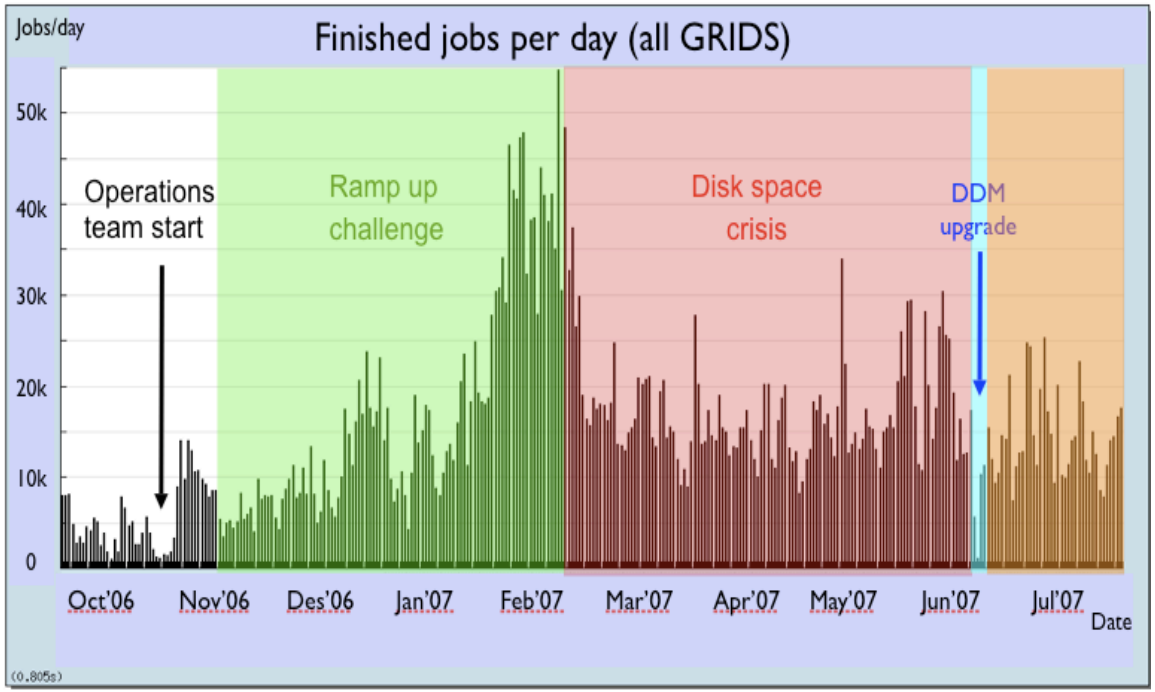

Experience(2/5): Simulated events and efficiencies

- Since the starting of the EGEE production joint operations in November 2006 more than 2.7M jobs finished, yielding an amount of ~135M simulated events.

- Operations team exchanged more than ~3.5k mails, reported ~600 bugs and wrote more than 100 reports during this 9 months of operations.

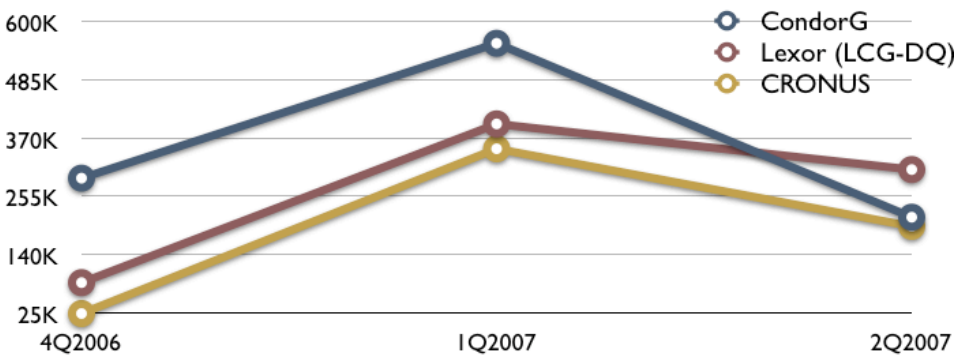

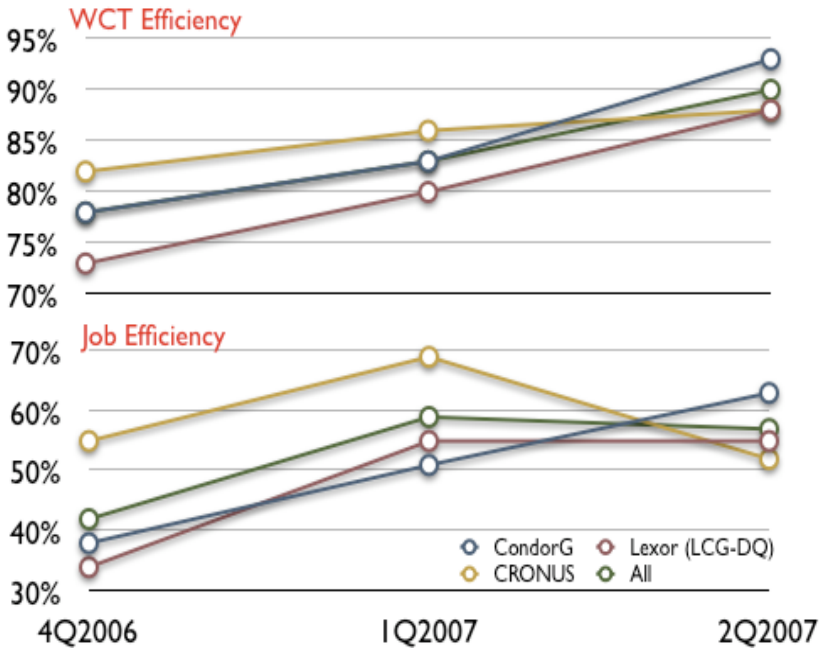

- Job and Wall Clock Time (WCT) efficiency has been continuously improving since the start of the operations.

- Although simulated production is a vivid body, which suffer permanent upgrades with new releases/patches and software validation periods, both producing short timed failures.

- WCT is well controlled with an average of 90% efficiency for the three EGEE executors.

Experience(3/5): Job errors revision

- Workload management englobe the intrinsic job related problems: software at the sites and VO specific, Grid related, task and job definitions, Executors, etc.

- Data management comprises stage-in and stage-out failures: site problems (SE outages, LFC time-outs, BDii errors, etc.) or global Data Management problems: missing files, corrupted files, etc.

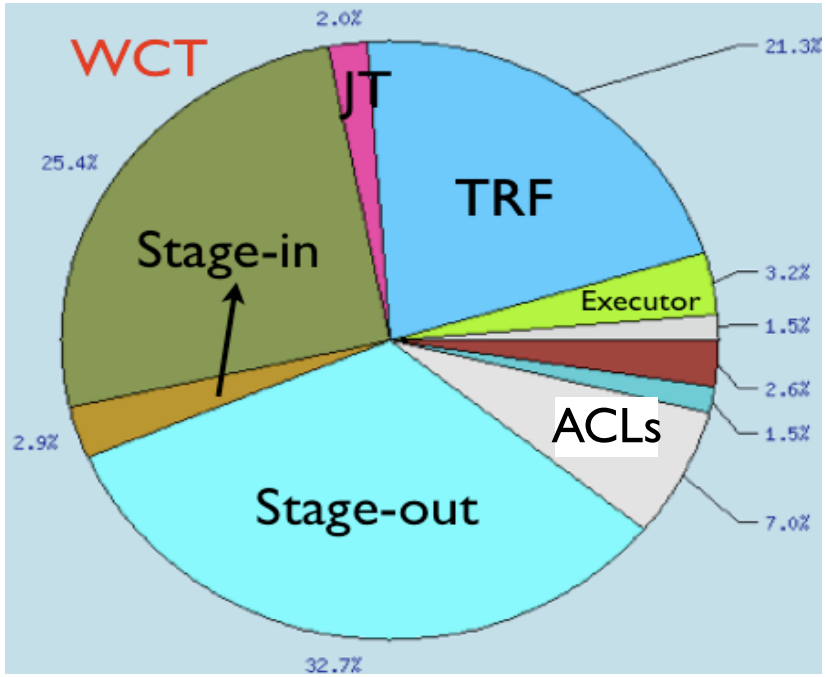

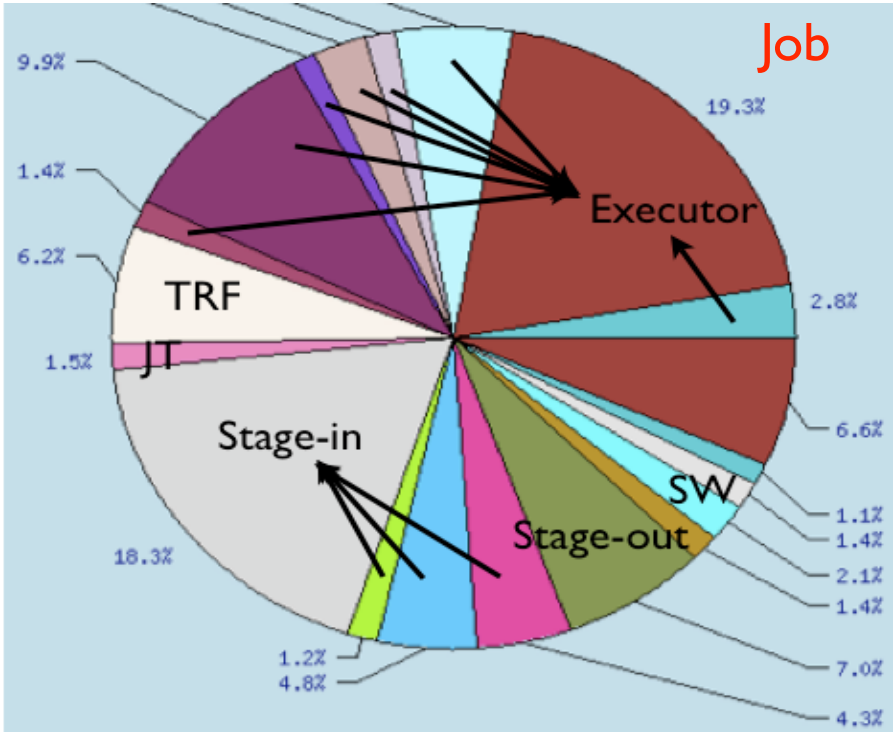

- Main impact in WCT is due to failures in data I/O: **65%** of the total.

- Losses in WCT due to ATLAS software is **32%** (highly variable between validation and full-production phases)...

- ...but as the global WCT efficiency is **~90%**, the net impact for the total CPU losses is low.

Experience(4/5): Production operations deployment

- Operations team take care of the ATLAS simulated event production in EGEE.
- Operation is based on a shift system:
  - Group of shifters work together during a whole week: production coordinators, 2 "senior" shifters and 2 "trainees" are on duty:
    - Production coordinators: control the task assignation, cloud production and monitor the overall production activity.
    - Shifters are separated in two working groups:
      - 1) Workload management: perform job monitoring.
      - 2) Data management: control the correct data flow (job inputs and outputs).
- Meetings:
  - Operations meeting is performed weekly by phone conference.
  - F2F every three month (ATLAS software week framework).
- Active dedicated mailing list among the members and ATLAS computing community.
- eLog web system is used to track all the incidences.
- **Extremely useful and fruitful experience !**

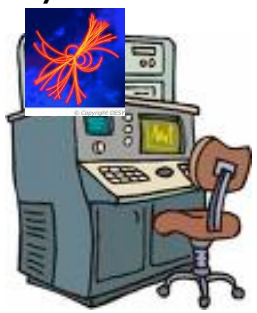

EGEE production team

- **Many people involved, but not entirely dedicated to EGEE production.**
- **Shifters are contributing about 1week/month, yielding a manpower of ~1.5 FTE's per week.**
- **EGEE production team:**

  - Production coordinators: Simone Campana, Rod Walker and Xavier Espinal.
  - Monitoring: John Kennedy and Benjamin Gaidioz (New dashboard)
  - Database tools: Suijian Zhou
  - Shift coordination: Xavier Espinal
  - Senior Shifters: Silvia Resconi, Mei Wen, Alessandra Doria, John Kennedy, Luis March, Xavier Espinal, Suijian Zhou, Carl Gwilliam, Guido Negri
  - Trainee shifters: Elisabetta Vilucchi, Agnese Martini, Marcel Schroers, Jaroslav Guenther, Miroslav Jahoda, Jordi Nadal and Kendall Reeves.
  - French cloud shifters team: Sandrine Laplace, Frederic Derue, Jerome Schwindling, Karim Bernardet,Terront Trujillo
  - Executor handlers: David Rebatto, Guido Negri, Silvia Resconi, French team, Rod Walker, Sanjay Padhi.

Experience(5/5): Operations

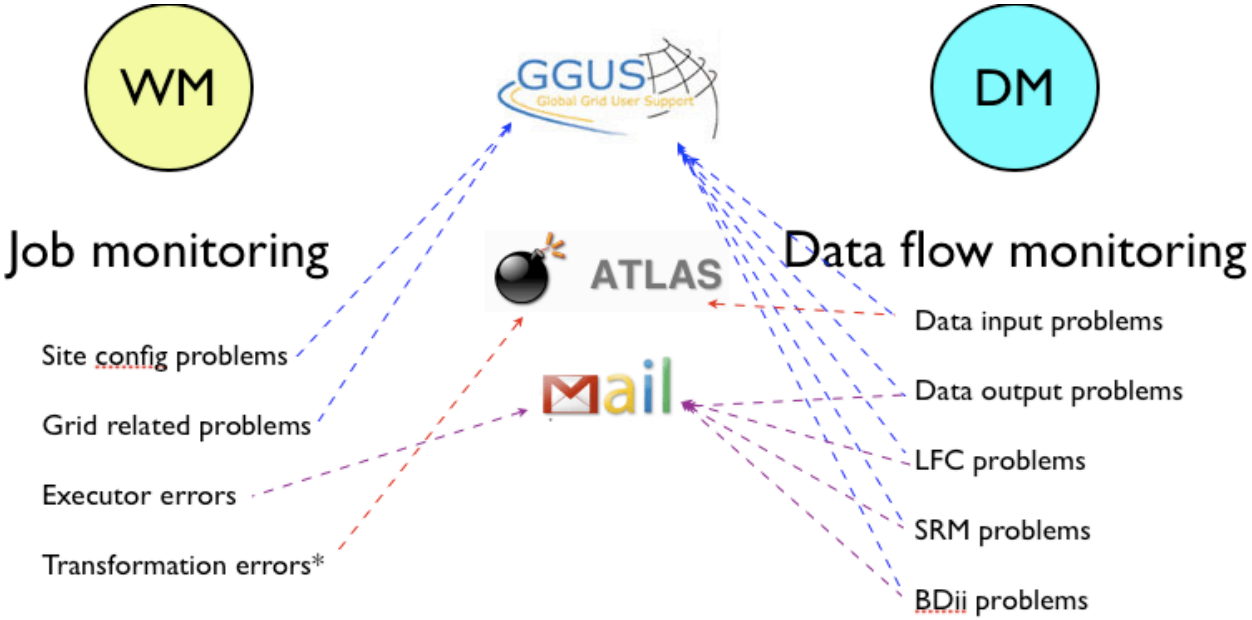

Every action recorded in:

ELOG

...and discussed in weekly meetings

Future steps

- **New ramp-up of production is expected: need for automation.**
  - Not human scalable, begin to automate error spotting and reporting.
  - First action done: transformation errors pseudo-automatically reported.
  - Next (September): Limit the job attempts once stage-out fails and control the re-release.

- **Keep working in monitoring:**
  - For Workload Management has improved a lot and is really useful.
  - For ProdSys Data Management is still a bit “dark”.
  - New dashboard for job monitoring is coming !

- **Automatic task assignment.**
  - Production coordinators overflooded by the requests.
  - After task is defined by the physics coordinators this has to be assigned to a certain cloud.
  - Task assignment has to be done in an intelligent way to minimize data movement.

- **Production and executor tutorial at next ATLAS Software Week (Oct’07, TBC)**

Summary and conclusions

- **Production system showed that could cope with the requirements during the production ramp up challenge (Nov06-Mar07)**
  - 60M events produced in first quarter of 2007 (coexistent with the disk shortage)

- **Job and WCT efficiencies has been improving almost continuously:**
  - **~90%** efficiency in WCT and **~60%** for job efficiency
  - More experienced team.
  - Infrastructure and system better known and more debugged.
  - Improvements in LFC, lcg-utils, SRM.

- **Monitoring pages are the key for successful job monitoring:**
  - Former ProdSys has been continuously revised and updated.
  - New dashboard is under testing and will be deployed soon.

- **Next phase is to start automation at higher level.**

- **Operations group and shift system has demonstrated to be extremely necessary.**
  - Enriching experience for everyone.
  - Some trainee shifters already promoted to senior.


Thanks for your attention !

Backup Slides

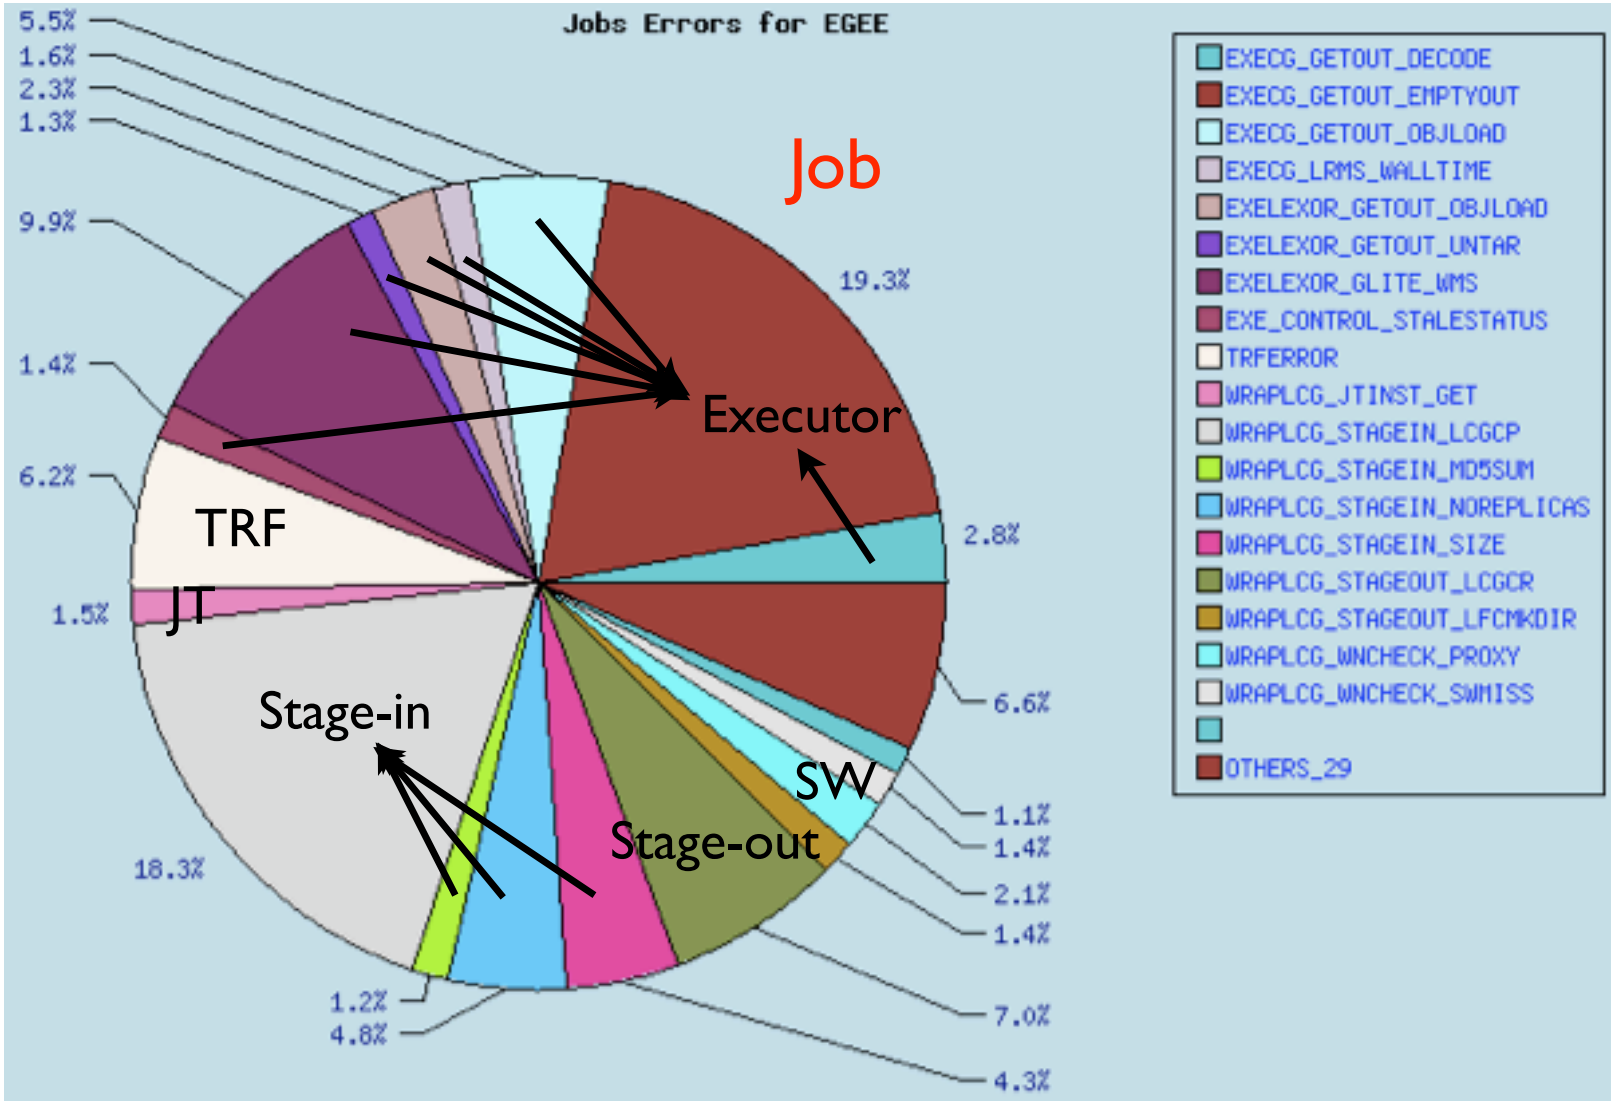
Jobs Errors for EGEE
5.5%
1.6%
2.3%
1.3%
9.9%
1.4%
6.2%
1.5%
18.3%
1.2%
4.8%
19.3%
2.8%
6.6%
1.1%
1.4%
2.1%
1.4%
7.0%
4.3%
Job
Executor
TRF
JT
Stage-in
Stage-out
SW
EXECG_GETOUT_DECODE
EXECG_GETOUT_EMPTYOUT
EXECG_GETOUT_OBJLOAD
EXECG_LRMS_WALLTIME
EXELEXOR_GETOUT_OBJLOAD
EXELEXOR_GETOUT_UNTAR
EXELEXOR_GLITE_WMS
EXE_CONTROL_STALESTATUS
TRFERROR
WRAPLCG_JTINST_GET
WRAPLCG_STAGEIN_LCGCP
WRAPLCG_STAGEIN_MD5SUM
WRAPLCG_STAGEIN_NOREPLICAS
WRAPLCG_STAGEIN_SIZE
WRAPLCG_STAGEOUT_LCGCR
WRAPLCG_STAGEOUT_LFCMKDIR
WRAPLCG_WNCHECK_PROXY
WRAPLCG_WNCHECK_SWMISS
OTHERS_29

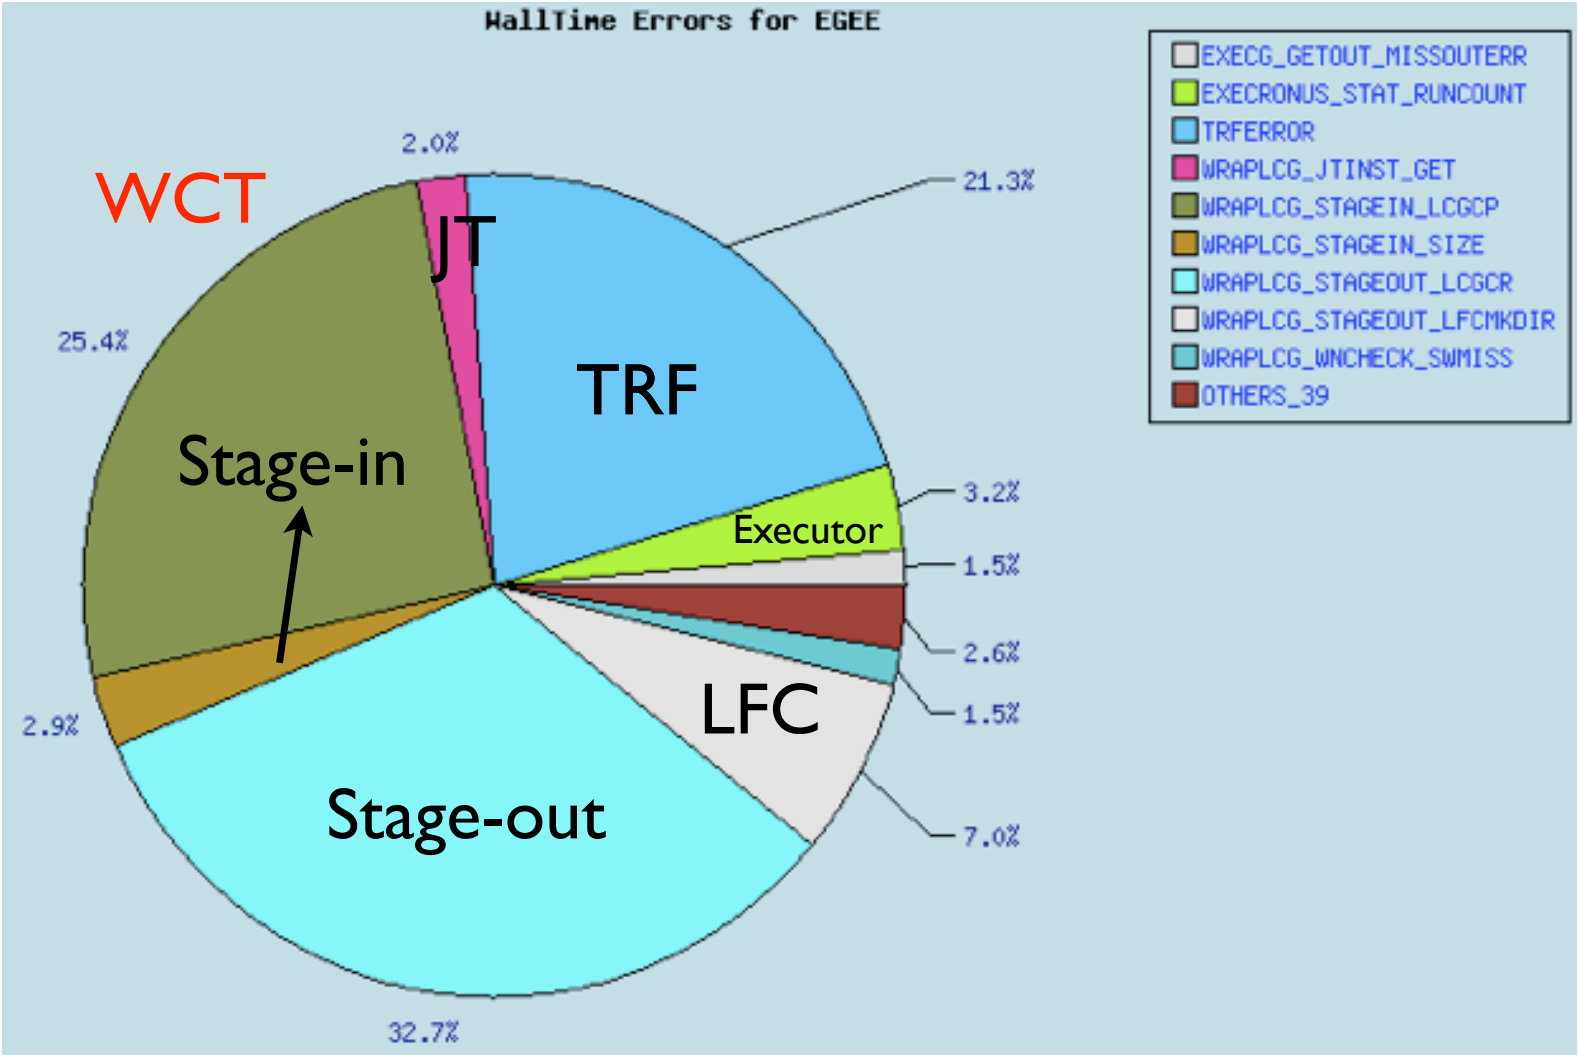
WallTime Errors for EGEE
EXECG_GETOUT_MISSOUTERR
EXECRONUS_STAT_RUNCOUNT
TRFERROR
WRAPLCG_JTINST_GET
WRAPLCG_STAGEIN_LCGCP
WRAPLCG_STAGEIN_SIZE
WRAPLCG_STAGEOUT_LCGCR
WRAPLCG_STAGEOUT_LFCMKDIR
WRAPLCG_WNCHECK_SWMISS
OTHERS_39
WCT
JT
TRF
Stage-in
Executor
LFC
Stage-out
2.0%
21.3%
25.4%
3.2%
1.5%
2.6%
1.5%
2.9%
7.0%
32.7%