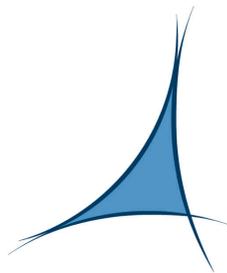


# ATLAS production workflow

Xavier Espinal

Port d'informació científica (PIC)

WLCG 2007 Victoria, BC CANADA 1st-2nd September 2007



**PIC**  
port d'informació  
científica



# The ATLAS Production System

- + Architecture
- + ProdDB
- + Supervisor
- + Executors
- + Job Workflow

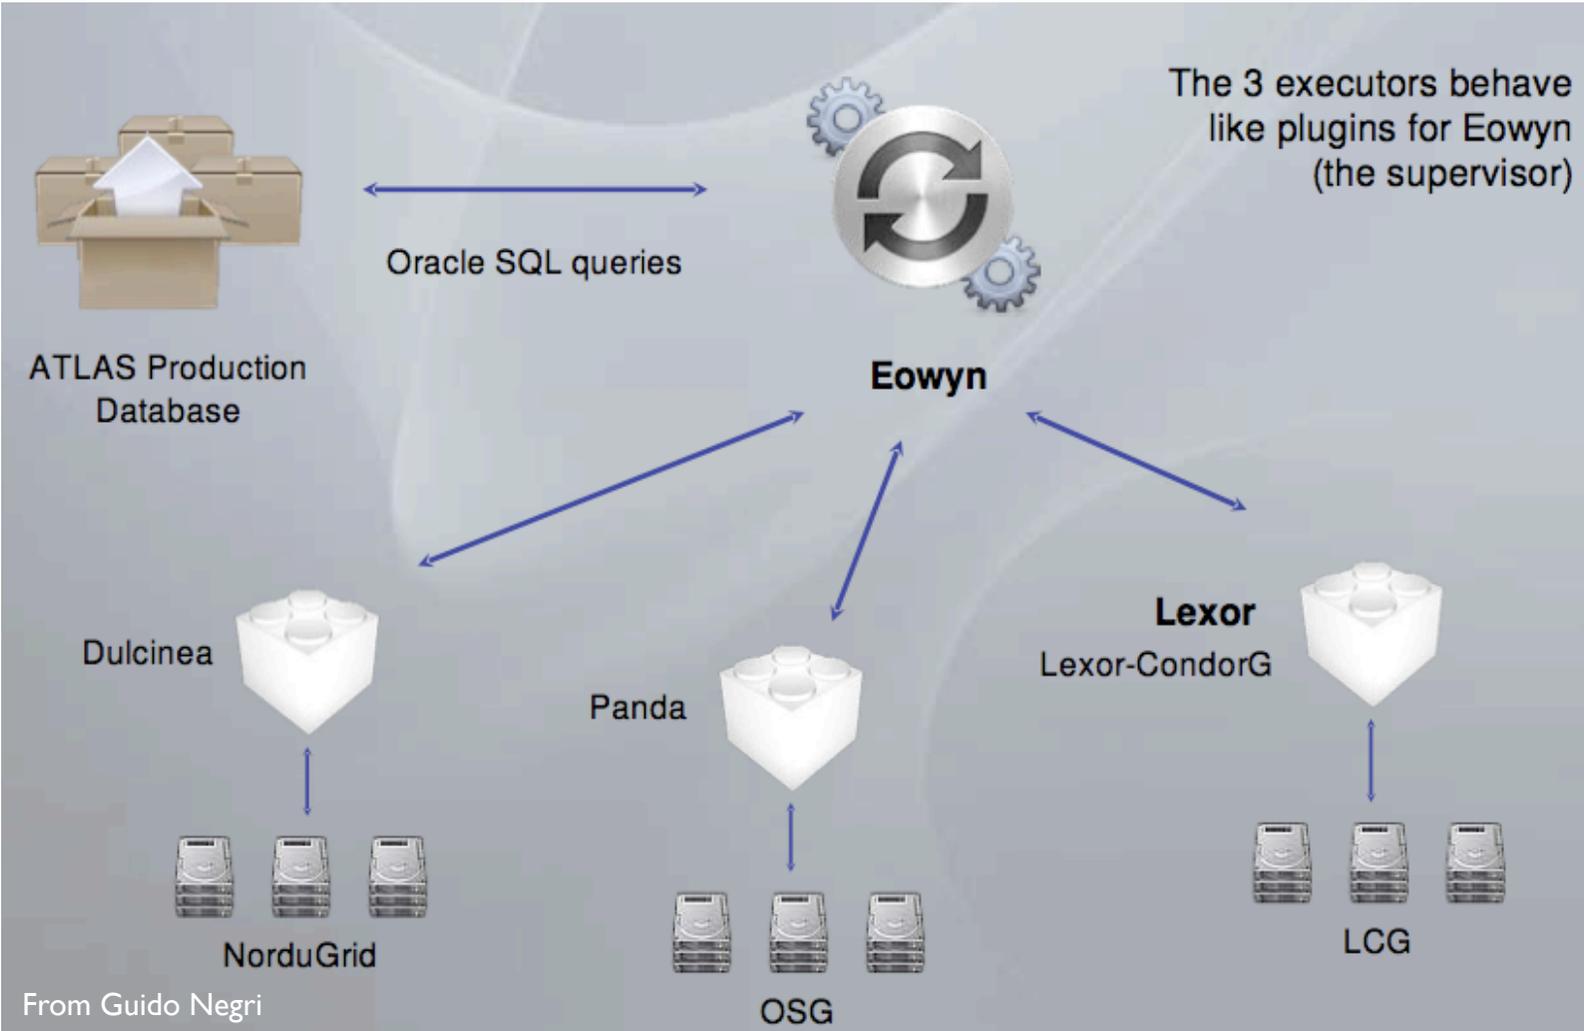


PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow - WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# Architecture



From Guido Negri

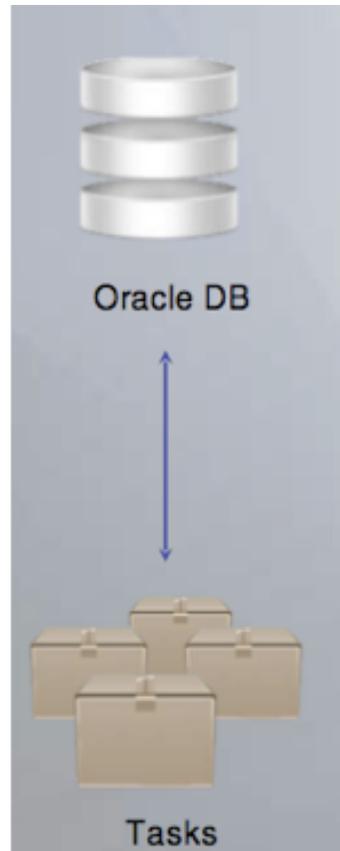


PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# The Production Database (ProdDB)



- ORACLE database at CERN, where jobs for ATLAS are defined.
- Jobs are grouped in datasets and tasks:
  - Datasets: defined the physics content of a collection of jobs.
  - Tasks: Identify all the jobs of the same dataset.
- From the Supervisor point of view only ETASK.TASKID is the relevant info, DATASETNAME is just an attribute.



# ProdDB Schema

- Supervisor only uses 3 tables of the ProdDB:
  - Etask
  - EjobDef
  - EjobExe
- Other tables are used to collect info on tasks and jobs.



# ETASK

TASKID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
TASKNAME	VARCHAR	128
STATUS	VARCHAR	64
ACTUALPARS	VARCHAR	2000
CPUCOUNT	NUMBER	9
CPUUNIT	VARCHAR	32
DISKCOUNT	NUMBER	9
DISKUNIT	VARCHAR	32
RAMCOUNT	NUMBER	9
RAMUNIT	VARCHAR	32
OUTIP	VARCHAR	3
USERNAME	VARCHAR	128
USERGROUP	VARCHAR	32
USERROLE	VARCHAR	32
TASKTYPE	VARCHAR	32
GRID	VARCHAR	32
TRANSFK	NUMBER	9
TRANSUSES	VARCHAR	256
TRANSHOME	VARCHAR	128
TRANSPATH	VARCHAR	128
TRANSFORMALPARS	VARCHAR	2000
TIER	VARCHAR	12
NDONE	NUMBER	9

This table contains all the tasks: each task groups several jobs (from 10 up to ~5000 jobs per each task, depending on the particular nature of the task)

The primary key of the table is **TASKID**: it uniquely identifies a task

Other important fields are:

**ACTUALPARS**: parameters for the ATLAS software which are common to all the jobs of the same task (e.g., needed CPU time, minimum RAM required, ...)

**STATUS**: only jobs of tasks in a RUNNING status will be selected

**TRANSUSES**: defines the ATLAS software version needed by all the jobs of the same task

**GRID**: a string used by the supervisor to select jobs. When configuring a supervisor, you have to specify the value that will match this field, so that only jobs of tasks with GRID=<your\_string> will be picked up

**TIER**: indicates the region (cloud) in which jobs will be run (all the jobs of a same task will run in the same cloud)

**ETASK**

From Guido Negri



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# EJOBDEF

JOBDEFID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
JOBTRANSFK	NUMBER	9
TASKFK	NUMBER	9
CONTROLSTATUS	VARCHAR	64
CURRENTSTATUS	VARCHAR	64
SUPERVISOR	VARCHAR	128
LASTATTEMPT	NUMBER	2
MAXATTEMPT	NUMBER	2
PRIORITY	NUMBER	5
JOBOUTPUTS	VARCHAR	2000
JOBLOGS	VARCHAR	1000
JOBPARS	VARCHAR	4000
LOCKEDBY	VARCHAR	128
TRANSINFO	VARCHAR	512
UPDATETIME	DATE	7

## EJOBDEF

This table contains all the jobs

The primary key of the table is **JOBDEFID**: it uniquely identifies a job

Other important fields are:

**TASKFK**: jobs with the same TASKFK are part of the same task (identified by TASKID=TASKFK in the ETASK table, see the preceding slide)

**CURRENTSTATUS**: only jobs in a TOBEDONE currentstatus will be selected. Jobs can be considered successful only once the supervisor has put them in DONE currentstatus

**MAXATTEMPT**: jobs that fail are resubmitted by the supervisor as many times as its MAXATTEMPT

**LASTATTEMPT**: tells the number of attempt of the last (re)submission. It increases of 1 at each submission of the job. Only jobs with

LASTATTEMPT<MAXATTEMPT will be picked up

**JOBPARS**: lists specific parameters of a job (e.g., input files, random numbers, ...) that, together with ETASK.PARS, will be used by the ATLAS software

From Guido Negri



# EJOBEXE

JOBEXEID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBDEFFK	NUMBER	9
TASKFK	NUMBER	9
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
ATTEMPTNR	NUMBER	2
SUPERVISOR	VARCHAR	128
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
FACILITYID	VARCHAR	256
INFOEXECUTOR	VARCHAR	3000
JOBSTATUS	VARCHAR	64
JOBNATIVESTATUS	VARCHAR	128
STARTTIME	DATE	7
ENDTIME	DATE	7
EXECLUSTER	VARCHAR	128
EXEQUEUE	VARCHAR	128
PROCESSINGHOST	VARCHAR	128
JOBOUTPUTS	VARCHAR	4000
ERRORCODE	NUMBER	9
ERRORACRONYM	VARCHAR	128
ERRORTTEXT	VARCHAR	1000
LOCKEDBY	VARCHAR	128

This table contains all the attempts in running a job

The primary key of the table is **JOBEXEID**: it uniquely identifies an attempt

Other important fields are:

**JOBDEFFK**: entries with the same JOBDEFFK represent resubmissions of the same job, identified by JOBDEFID=JOBDEFFK in the EJOBDEF table

**TASKFK**: entries with the same TASKFK are resubmissions of jobs of the same task, identified by TASKID=TASKFK in the ETASK table

**JOBSTATUS**: shows the status of the job in the executor (submitted, scheduled, running,...). Successful jobs are marked as 'finished'

**ATTEMPTNR**: tells the number of the attempt of resubmission the entry refers to (a job that has been submitted 3 times will have 3 corresponding records in EJOBEXE with 3 different JOBEXEID and 3 different ATTEMPTNR: 1, 2, 3)

From Guido Negri

## EJOBEXE



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# EJOBEXE(2)

JOBEXEID	NUMBER	9
CREATIONTIME	DATE	7
MODIFICATIONTIME	DATE	7
JOBDEFFK	NUMBER	9
TASKFK	NUMBER	9
JOBNAME	VARCHAR	128
PARTNR	NUMBER	9
ATTEMPTNR	NUMBER	2
SUPERVISOR	VARCHAR	128
EXECUTOR	VARCHAR	128
EXECUTORTYPE	VARCHAR	64
FACILITYID	VARCHAR	256
INFOEXECUTOR	VARCHAR	3000
JOBSTATUS	VARCHAR	64
JOBNATIVESTATUS	VARCHAR	128
STARTTIME	DATE	7
ENDTIME	DATE	7
EXECLUSTER	VARCHAR	128
EXEQUEUE	VARCHAR	128
PROCESSINGHOST	VARCHAR	128
JOBOUTPUTS	VARCHAR	4000
ERRORCODE	NUMBER	9
ERRORACRONYM	VARCHAR	128
ERRORTTEXT	VARCHAR	1000
LOCKEDBY	VARCHAR	128

**FACILITYID:** this string represent the ID of the collection of jobs on the gLite Resource Broker

**INFOEXECUTOR:** this string represent the ID of the job on the remote grid (on LCG, it's just the JobID).

The groups different jobs, that are run separately, each with its JobID, but may be queried all in one using the CollectionID

**EXECLUSTER:** this string tells the name of the CE on which the job runs

**ERRORACRONYM:** if an attempt to run a job fails, this acronym is filled with a tag giving the nature of the occurred problem

**ERRORTTEXT:** if an attempt to run a job fails, this field is filled with an error message giving more informations on the error

**JOBOUTPUTS:** this lists the produced output files and their LFC logical file names

## EJOBEXE

From Guido Negri



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# Simple example...

An example: task identified by taskid 1000 in the ETASK table consists of 2 jobs, identified by jobdefid 10000 and 10001 in the EJOBDEF table. Job 10001 runs successfully at the very first attempt (one only entry in the EJOBEXE table), while job 10000 fails twice, then it's resubmitted a third time and it succeeds (three entries in the EJOBEXE table)



**TASKID=1000**

**ETASK**

- JOBDEFID=10000  
TASKFK=1000
- JOBDEFID=10001  
TASKFK=1000

**EJOBDEF**

- JOBEXEID=200  
JOBDEFFK=10000  
TASKFK=1000  
status=finished
- JOBEXEID=201  
JOBDEFFK=10001  
TASKFK=1000  
status=failed
- JOBEXEID=202  
JOBDEFFK=10001  
TASKFK=1000  
status=failed
- JOBEXEID=203  
JOBDEFFK=10001  
TASKFK=1000  
status=finished

**EJOBEXE**

From Guido Negri



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# EOWYN: The Supervisor

- Is the brain of the Production system, the only process directly interacting with the ProdDB.
- Eowyn activities are the following ones:
  - Select free jobs form the ProdDB.
  - Pass the job to an executor (who take care of job submission).
  - Ask executor about job status.
  - Validate finished jobs. If job is OK is considered DONE, otherwise will be release for further attempts.
  - Update job entries in the DB.



# Executors

- Modules that behave as plugins to Eowyn, that loads the corresponding executor module in the configuration and is transparent for any grid flavour.
- Executor actions:
  - Tell Eowyn how many jobs can be submitted.
  - Receive jobs from Eowyn and prepare them: *parsing of the job parameters and construction of the JDL.*
  - Submit the jobs to the grid.
  - Ask for the status to the underlying middleware.
  - Retrieve output log files.
- Currently two executor flavors running on EGEE: Lxor, Condor-G.

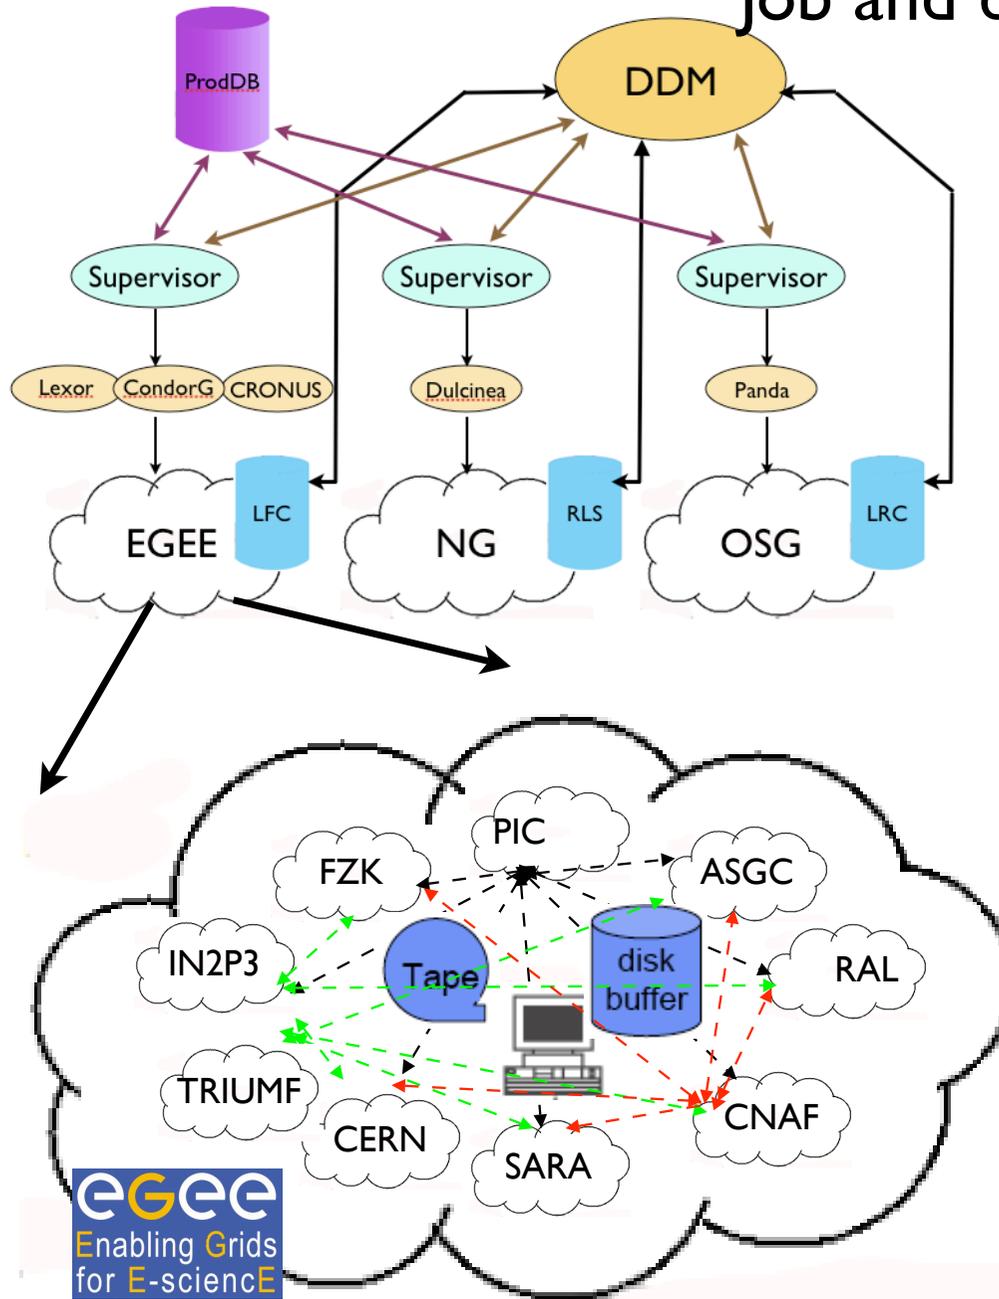


# Job Workflow

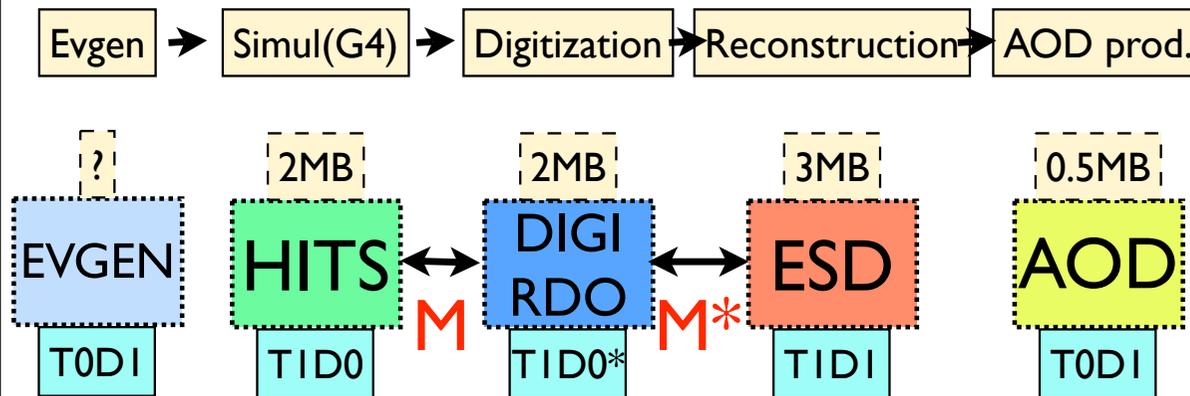
- **Job definition:** jobs entries are created in the ProdDB by the Physics Coordinators (also requests from physics groups).
- **Job selection:** Eowyn pick up jobs from the ProdDB.
- **Job submission:** Eowyn passes the jobs to an executor who submit them to the grid.
- From here on ..... EGEE production team area.



# Job and data Workflow



# Simulated events chain



M: merged step  
 M\*: To be merged  
 \* or transient

## Data Volume (Status at April'07)

### Full physics, 30M simulated events:

- 30M x (2 MB (HITS)) = 60 TByte HITS
- 30M x (2 MB (RDO)) = 60 TByte RDO
- 30M x (0.5 MB (AOD)) = 15 TByte AOD
- 30M x (3,0 MB (ESD)) = 90 TByte ESD

### Single particle, 20M simulated events:

- 20M x (2 MB (HITS)) = 40 TByte HITS
- 20M x (2 MB (RDO)) = 40 TByte RDO
- 20M x (0.01 MB (AOD)) = 0.2 TByte AOD
- 20M x (0.45 MB (ESD)) = 9 TByte ESD

### ESD:

- ESD file left on disk at the T1 where it is produced
- ESD file archived to tape where it is produced

One full copy of the ESD over the whole grid plus one archived copy on tape.

### AOD:

- AOD file left on disk at the T1 where it is produced
- Each AOD file replicated to all T1 sites
- plus one other full copy at the T2's in the cloud
- Distribution of AOD files in the cloud to be determined by the people in the cloud



PIC  
 port d'informació  
 científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# EGEE Production Operations

- + WhatDoWeDo?
- + Monitoring
- + Tools
- + Workload
- + Data

<https://twiki.cern.ch/twiki/bin/view/Atlas/LcgProduction>



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow - WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# What do we do ?

- Our duty is to look after the jobs submitted to the grid by the executors.
- Investigate its failures and report the problems:
  - Through monitoring and bug/problems reporting.
- Take actions (ProdSys tools):
  - Abort/resubmission of jobs.
  - Data flow: check data at sites, move data (minimally).



# Monitoring

- Our basic tool to spot problems.

- Former ProdSys monitoring web page (JK):

<http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/JobInfo.php>

...To be superseded by the new ProdSys dashboard (next slide).

The screenshot shows a web interface with a navigation bar at the top containing tabs: Home, Datasets, JobInfo, Daily Stats, Site Status, Overviews, and TierViews. Below the navigation bar are five main monitoring panels, each with a 'Help' button on the right:

- Jobs/Errors by Sites:** Contains 'Jobs:' and 'Errors:' sections. Each section has two buttons: 'Last 24 hrs' and 'Last 1 hr'.
- Jobs/Errors by Executor:** Contains 'Jobs:' and 'Errors:' sections. Each section has two buttons: 'Last 24 hrs' and 'Last 1 hr'.
- Jobs/Errors by Task:** Contains 'Jobs:' and 'Errors:' sections. Each section has two buttons: 'Last 24 hrs' and 'Last 1 hr'.
- Stalling(?) Jobs:** Contains three buttons: 'By Site', 'By Executor', and 'By Task'.
- Running Tasks:** Contains one button: 'Running Tasks'.



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# Monitoring

... new ProdSys dashboard (BG):

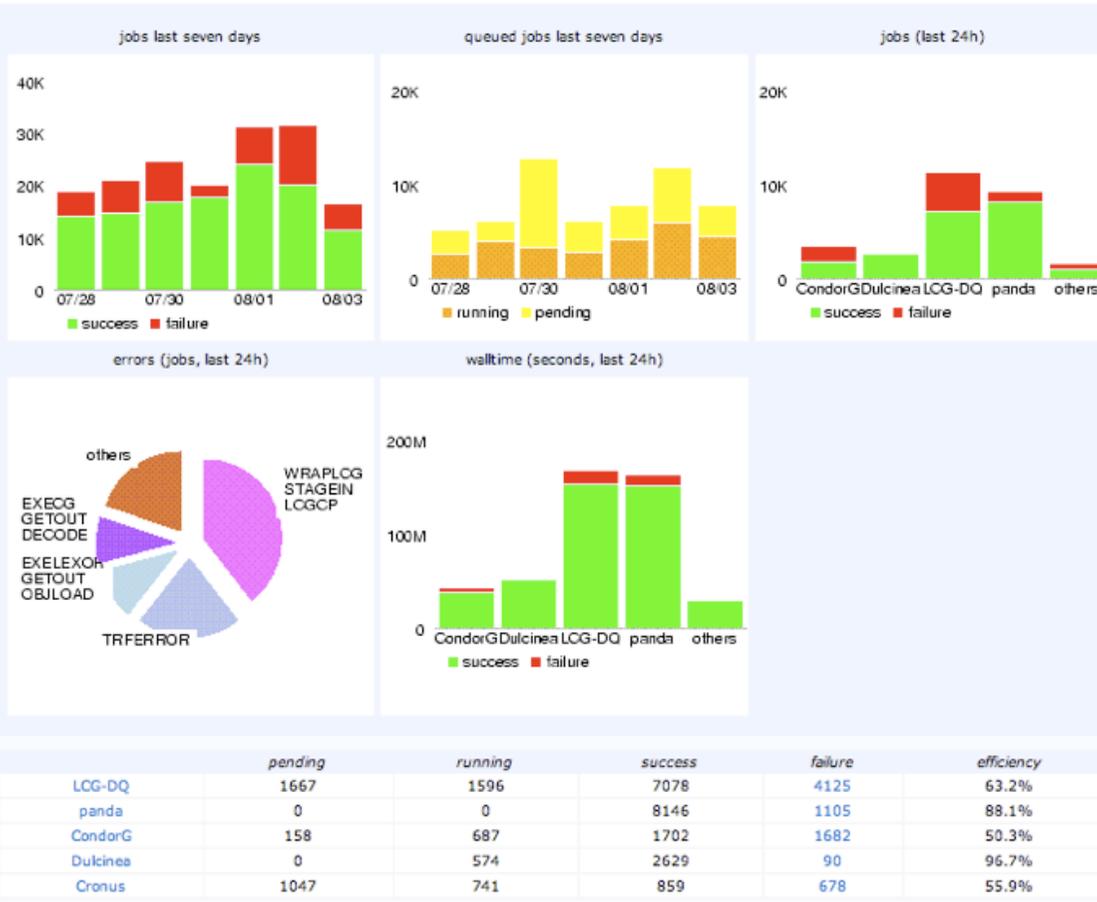
<http://dashb-atlas-prodsys-test.cern.ch/dashboard/request.py/overview>

## ATLAS Production system

Overview User Guide Feedback

### OVERVIEW

- view
- by grid (not avail.)
- by cloud (not avail.)
- by executortype
- by executor
- by cluster
- by task
- select executortype
- LCG-DQ
- panda
- CondorG
- Dulcinea
- Cronus



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



Jobs/Errors by Sites

Jobs: [Last 24 hrs](#)

Errors: [Last 24 hrs](#)

Jobs/Errors by Executor

Jobs: [Last 24 hrs](#)

Errors: [Last 24 hrs](#)

Jobs/Errors by Task

Jobs: [Last 24 hrs](#)

Errors: [Last 24 hrs](#)

Clickable:  
 Access to all errors from the task  
 Access to all jobs failed  
 Access to job log files

#	TaskName	TaskId	Grid	Tier	Pending	Running	Finished	Failed	Efficiency (%)
1	valid1.007199.singlepart_mu1.evgen.v12000701.task	11912	CRONUS	NL	0	0	9	1	90
2	valid3_valid1.007199.singlepart_mu1.digit.v13001001.task	11913	CRONUS	NL	0	1	9	0	100
3	valid1.005851.WH120uu_pythia.evgen.v12000605.task	11966	CRONUS	NL	0	0	4	0	100
4	valid3_valid1.005851.WH120uu_pythia.digit.v13001001.task	11967	CRONUS	NL	0	10	0	0	0
5	misal1_mc12.006875.PythiaWH130jgamgam.digit.v12000605.task	12097	LOG-DQ	CA	8	0	2	0	100
6	misal1_csc11.005176.AcerMC_Zbb_3L.digit.v12003103.task	12142	LOG-DQ	CA	8	2	0	18	0
7	pile1s05_misal1_csc11.005300.PythiaH130zz4L.digit.v12000701.task	10808	LOG-DQ	DE	0	2	0	3	0
8	mc12.006509.XX_jimmy_ma_400_tanb_20.evgen.v12000701.task	11016	LOG-DQ	DE	0	7	0	0	0
9	mc12.006501.A0_jimmy_ma_500_tanb_20.evgen.v12000701.task	11036	LOG-DQ	DE	0	1	0	0	0
10	trig1_pile1s05_misal1_mc12.005805.filtered_minbias6.recon.v12000605.task	11212	LOG-DQ	DE	407	664	126	1314	9



# ProdSys Tools

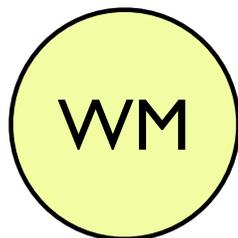
☛ dq2 ans pdbAdmin, installation procedure (thanks to Carl Gwilliam):  
<https://twiki.cern.ch/twiki/bin/view/Atlas/LcgProductionInstallHelp>

- **DQ2**: The ATLAS data management software:
  - List data at the sites.
  - Check/make/delete subscriptions (sites' data request)
  - ...
- **pdbAdmin**: command line interface to interact with the ProdDB through SQL queries:
  - Ask for job status and errors.
  - Abort/release jobs
  - ...



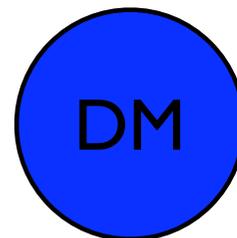
# Hands on

- There are two main working areas, two groups in charge of each one during the shift period:



## Job monitoring

- Site config problems
- Grid related problems
- Executor errors
- Transformation errors\*



## Data flow monitoring

- Data input problems
- Data output problems
- LFC problems
- SRM problems
- BDii problems



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# Workload Management

➤ SHIFTER: to chase failing jobs and notify errors (or abort) when needed, for:

Site config problems  
Grid related problems  
Executor errors  
Transformation errors\*



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# WM ACTIONS (I)

👉 Site related problems: **GGUS ticket** (<https://gus.fzk.de/pages/home.php>)

```
ERROR:Setup file not found: /opt/exp_software/atlas/software/12.0.31/cmtsite/  
setup.sh'
```

👉 Executor related problems: mail to executor handlers and production mailing list ([atlas-lcg-production\(at\)cern.ch](mailto:atlas-lcg-production(at)cern.ch))

```
WRAPLCG_WNCHECK_PROXY  
['2007-01-15 07:27:03,637 isProxyValid: ERROR Valid time 16 hours less than  
required 22 hours\n']
```

👉 Transformation error\*, Savannah bug report to validation team (<https://savannah.cern.ch/bugs/?group=validation>)

```
TRFERROR - The VxTrackAtVertex returned not refitted  
LArG4HitManagement ERROR BuildHitCollections: CaloDM_ID error code 4  
CaloID_Exception - Error code: 4LArG4HitManagement ERROR Run=0 Event=1,  
JobOptions file SimuJobTransforms/Calhits.py not found
```

👉 Periodically check for paused tasks and status (unpause if former problem is over):

[http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/paused\\_tasks.html](http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/paused_tasks.html)

👉 Each action is written in a log web page (eLOG: <http://carnage.triumf.ca:8080/>)



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# WM ACTIONS (II)

pdbAdmin

Used as an interface to the ProdDB:

- Abort jobs
- Release jobs
- List job errors for a task
- ...

## ACTIONS:

Only one of these is allowed

- assign used to assign a single task to a T1 cloud. Determines input dataset locations, subscribes input and output datasets to the chosen T1, and releases jobs.
- release Set TOBEDONE jobs in ('WAITINGINPUT','WAITING','WAITINGCOPY')
- autoabort Set TOBEDONE jobs in AUTOABORT if last error not TRF error
- forceabort Force abort jobs with TRFERROR
- max Increase maxattempt if TOBEDONE and last=max
- abort Abort jobs matching criteria AND in ('AUTOABORT','TOBEDONE')  
i.e. will not abort pending/running/waiting jobs.
- setprio= Set priority
- subs Subscribe the output datasets of matching tasks. Requires --tier to be set.
- reasons Prints the errors for stuck jobs,AUTOABORT or reached maxattempt
- stats Summarise status. No. running, waiting, done,...
- errors List trf errors for validation
- getinputs Print information from the jobpars for first partn
- getoutputs in each matching task
- getparams
- pause Pause matching tasks (often tiers due to lfc/srm problem)
- listtrf To list jobs with TRFERROR of a task.
- taskrelease To list tasks that to be released.
- missingfiles To list jobdefids that missing files.
- jobstatus To print quick info about jobs status and status of the corresponding tasks that the jobs belong.
- abortandrelease To abort or release jobs according to ERRORACONYM etc.A helpful bug report file will be created at the same time.
- evgen To list all evgen tasks.
- jobsremains To list remaining jobs in RUNNING tasks. --percentdone can be specified to filter tasks with percentdone greater than certain value.

## MODIFIERS

- force All update actions on the Db are preceded by a y/n prompt unless this is set
- force\_autoabort Special option to release jobs with "TRF\_GBB\_CPU".
- hardmax= Jobs exceeding this number of attempts, default 12, are not released by --max or --autoabort

## SELECTION:

- These select the tasks/jobs operated on, i.e. builds the where clause
- grid= Mandatory: LCG-DQ,LCG-CG-DQ,OSG,NORDUGRID
- task= Task id in form 1-10 for inclusive range or 1,3,5 for list
- tier= ETASK tier, e.g. CA,UK,ES
- partnr= Partition number, 1-10 and 1,3,5 formats
- jobdef= Jobdefinition number, 1-10 and 1,3,5 formats
- taskname= Match to taskname, usually containing SQL wildcards, eg.  
--taskname="%v120002%"

## SET VALUE

- setier= Set the ETASK tier field



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# WM ACTIONS (II)

## pdbAdmin commands

Abort jobs.:

```
./pdbAdmin --grid='LCG-DQ' --jobdef=6465 --abort
```

To list remaining jobs in RUNNING tasks:

```
./pdbAdmin --grid='LCG-DQ' --task=9709 --jobsremains
```

```
percentdone=95%
```

```
(9709, 'misall_mc12.008095.PythiaPhotonJet1_FIXED.digit.v12003103.task', 'ES', 'RUNNING') 99% Done  
(5446955, 'RUNNING', '2007-08-03', None, 'CERN_EGEE_cronus', 11)  
(5446955, 'RUNNING', '2007-08-03', 'EXELEG_GETOUT_OBJLOAD', 'GLOW_EGEE_cronus', 10)  
(5446955, 'RUNNING', '2007-08-03', 'EXELEXOR_GLITE_MAXRETRYCOUNT', 'david_test', 9)  
(5448830, 'RUNNING', '2007-08-03', None, 'CERN_EGEE_cronus', 9)  
(5448830, 'RUNNING', '2007-08-03', 'EXELEG_GETOUT_DECODE', 'CERN_EGEE_cronus', 8)  
(5448830, 'RUNNING', '2007-08-03', 'EXELEXOR_GLITE_MAXRETRYCOUNT', 'david_test', 7)
```

Prints the errors for stuck jobs, AUTOABORT or reached maxattempt:

```
./pdbAdmin --grid='LCG-DQ' --task=10029 --reasons
```

To list jobs with TRFERROR of a task:

```
./pdbAdmin --grid='LCG-DQ' --task=10029 --listtrf
```

To pause (Tasks, Clouds):

```
./pdbAdmin --grid='LCG-DQ' --task=10029 --pause
```

```
./pdbAdmin --grid='LCG-DQ' --tier=ES --pause
```



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# Data Management

➔ SHIFTER: to chase failures concerning data input/output and notify errors, pause tasks (or clouds) when needed, for:

Data input problems  
Data output problems  
LFC problems  
SRM problems  
BDii problems



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# DM ACTIONS (I)

## dq2 tools

Locate and query the data stored at the sites.  
Check Subscriptions info, ...

dq2-close-dataset	dq2-list-dataset	dq2-ping
dq2-delete-files	dq2-list-dataset-by-creationdate	dq2-register-dataset
dq2-delete-replicas	dq2-list-dataset-replicas	dq2-register-files
dq2-delete-subscription	dq2-list-dataset-site	dq2-register-location
dq2-destinations	dq2-list-files	dq2-register-subscription
dq2-erase	dq2-list-subscription	dq2-register-version
dq2-freeze-dataset	dq2-list-subscription-info	dq2-reset-subscription
dq2-get-metadata	dq2-list-subscription-site	dq2-reset-subscription-site
dq2-get-number-files	dq2-metadata	dq2-sources

## Staging problems (lcg-cp or lcg-cr), quick overview info in ProdSys monitoring:

Stage In&Out task problems:

[http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/Tasks\\_Stage\\_Errors\\_24hr.php](http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/Tasks_Stage_Errors_24hr.php)

Tier Stageout problems:

[http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/Tier\\_Stageout\\_View.php](http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/Tier_Stageout_View.php)

Staging errors per site:

<http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/php/DbAdmin/Ora/php-4.3.4/proddb/monitor/staging.html>

 Submit a GGUS to the site in case of staging problems.

 In case of permanent problems, pause the cloud.

 Each action is written in a log web page (eLOG: <http://carnage.triumf.ca:8080/>)

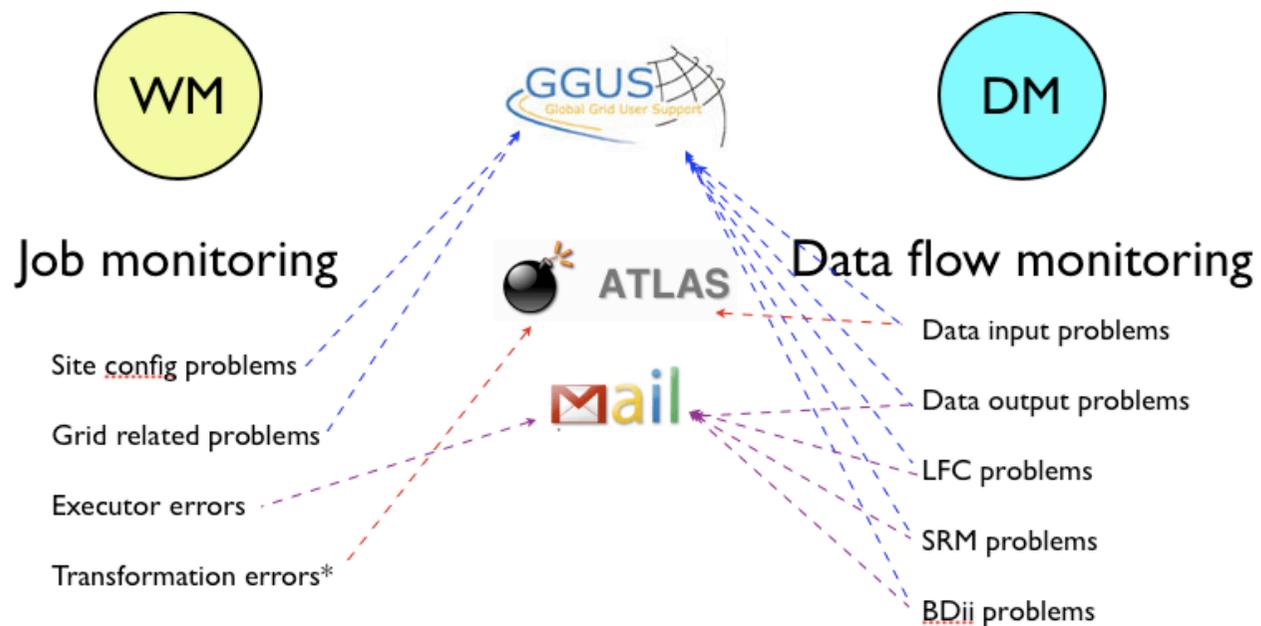


PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007



# ProdSys Flow diagram



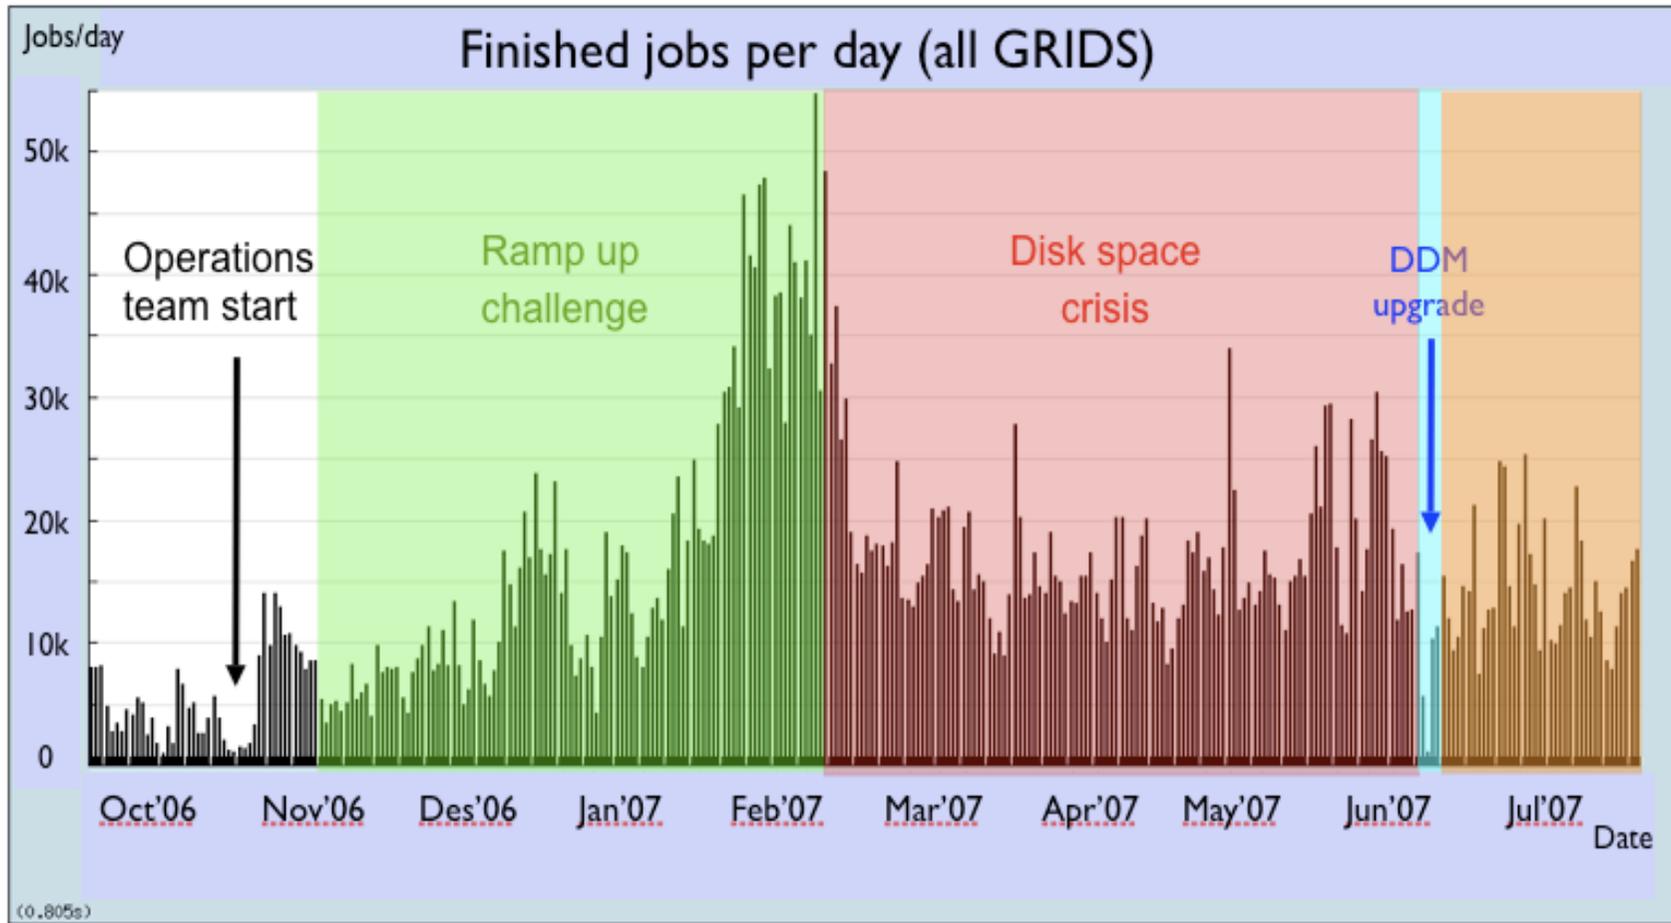
Every action recorded in:



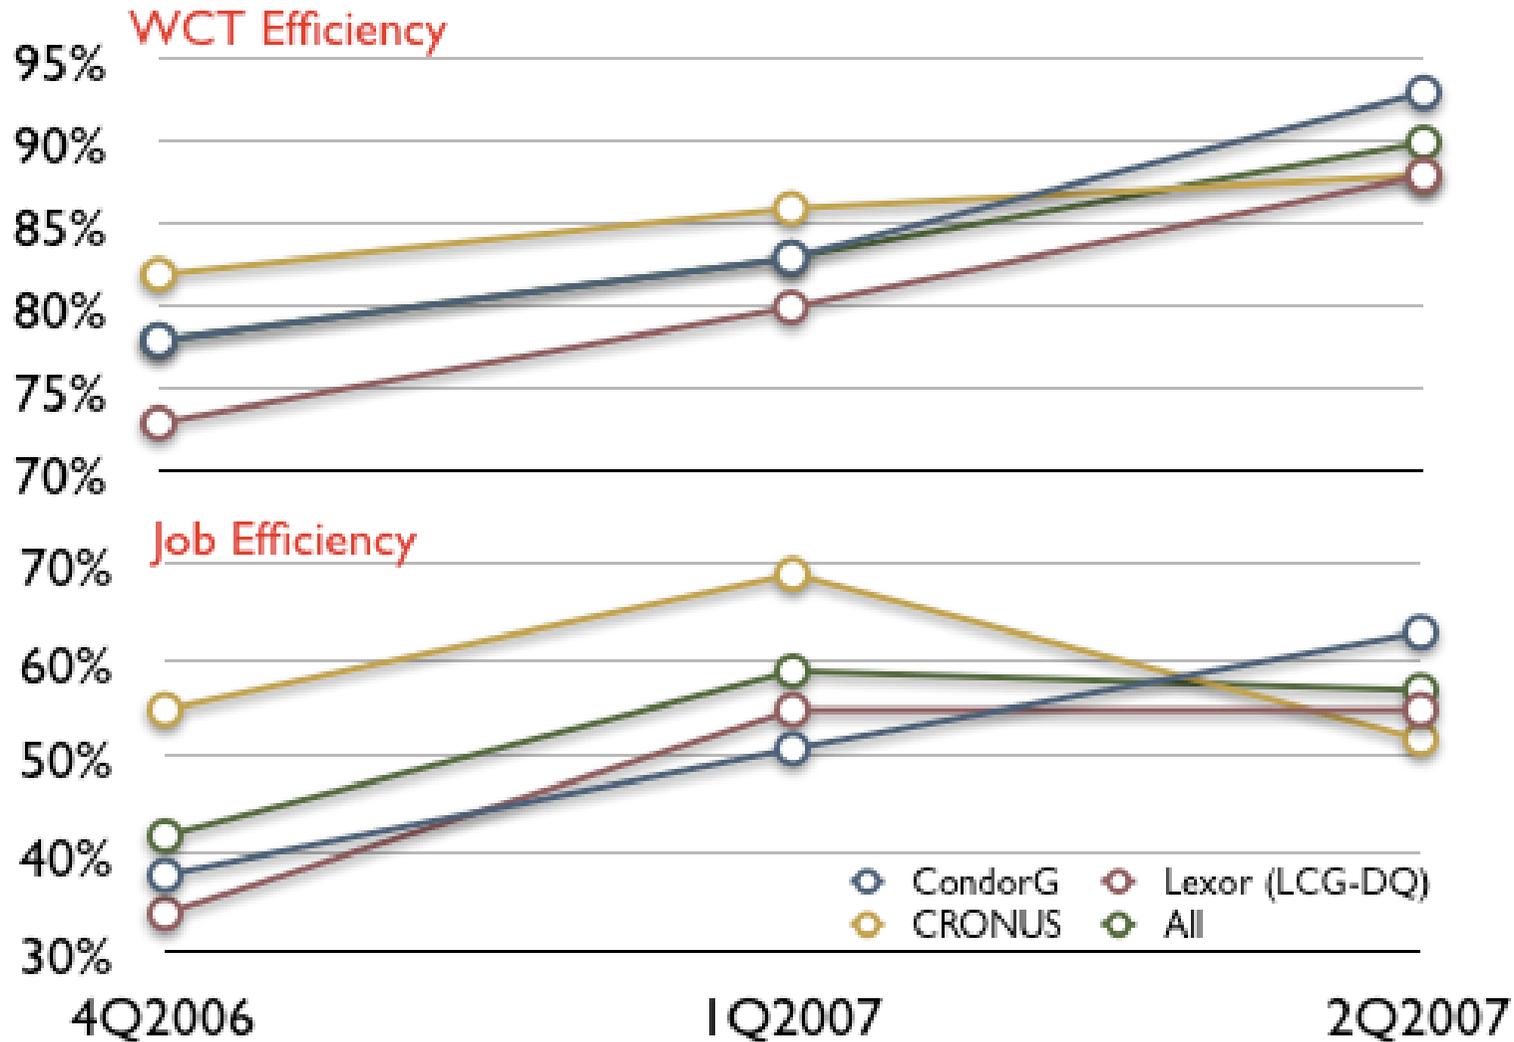
...and discussed in weekly meetings



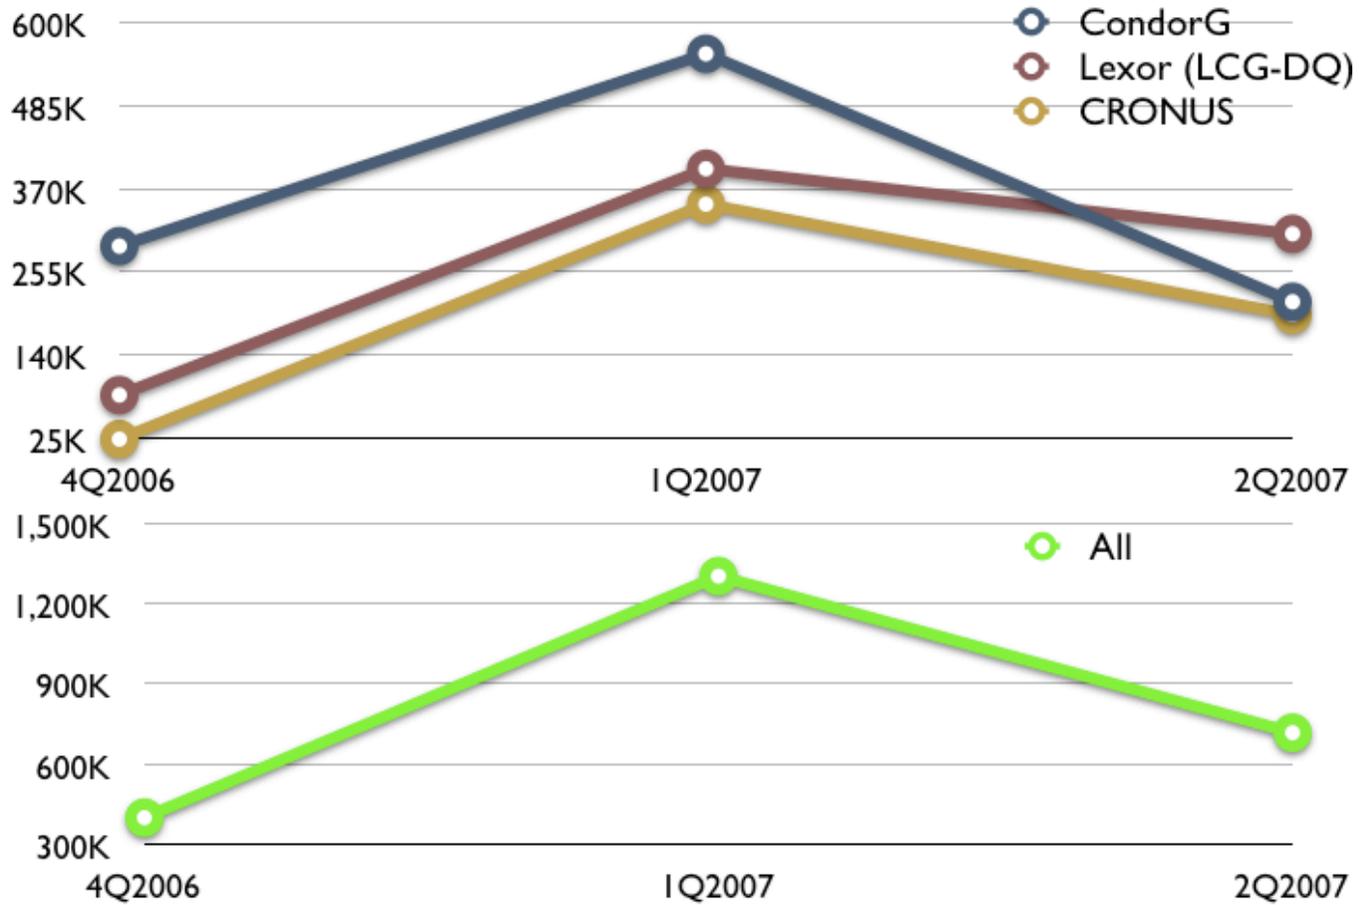
# EGEE production in numbers



# EGEE production in numbers



# EGEE production in numbers



# The EGEE production team

## Coordinators

Rod Walker Simone Campana Xavier Espinal

## Senior shifters

John Kennedy Sujian Zhou Silvia Resconi Xavier Espinal Luis March Guido Negri Carl Gwilliam Alessandra Doria Mei Wen

## French cloud

Jerome Schwindling Sandrine Laplace Frederic Derue

## Trainee shifters

Elisbeta Vilucchi Agnese Martini Marcel Schroers Jaroslav Günther Miroslav Jahoda Kendall Reeves Jordi Nadal

## Database Tools

Sujian Zhou

## Monitoring

John Kennedy Benjamin Gaidioz

## Shift Coordinator

Xavier Espinal

## Executor Instances

CondorG	Lexor	CRONUS
FR team	David Rebatto	Sanjay Padhi
Rod Walker	Guido Negri	
	Silvia Resconi	

- Weekly phone meeting.
- Shift rate: one\_week/month.
- Begin as trainee shifter associated to a senior.
- Good chance to discover ATLAS computing world.



Thanks for the attention !



PIC  
port d'informació  
científica

Xavier Espinal - ATLAS production workflow- WLCG 2007 Victoria, BC CANADA - 1st-2nd September 2007

