

ATLAS Production System



John Kennedy
LMU Muenchen
DPG Frühjahrstagung
09/03/2007



Overview

- Production System
 - High level View
 - Components
 - Some stats from production
- Some new activities
- Conclusion/Outlook

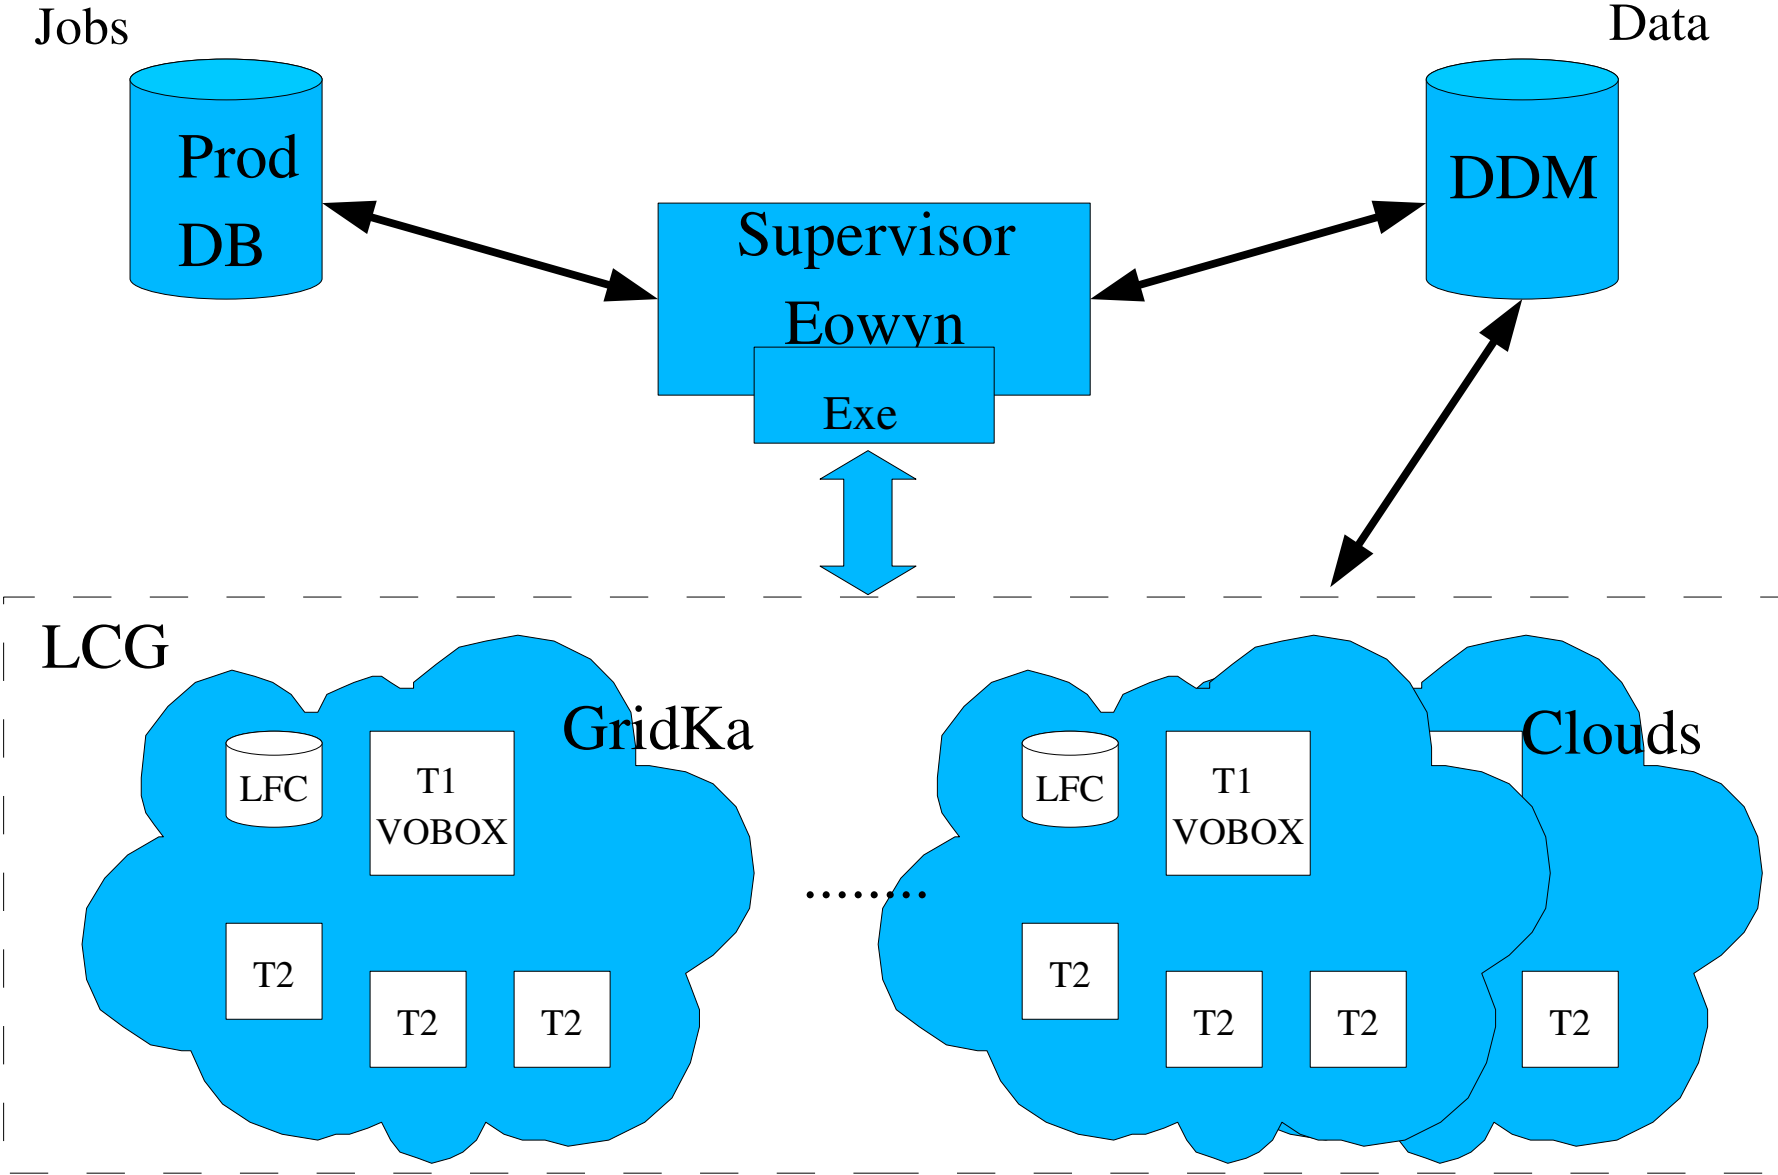
Prod-Sys

The ATLAS Production System

Prod-Sys

- Looking from High above
 - Jobs are Retrieved from central database and sent to Grid
 - Datasets subscribed(DDM) to a T1 at assignment
 - Jobs can go to any cloud but data is confined to a specified cloud (in the main)
 - A job starts at a site gets data then when finished puts data on assigned Cloud (T1 then T2's as fallback)
 - Jobs are checked for errors and re-submitted if required

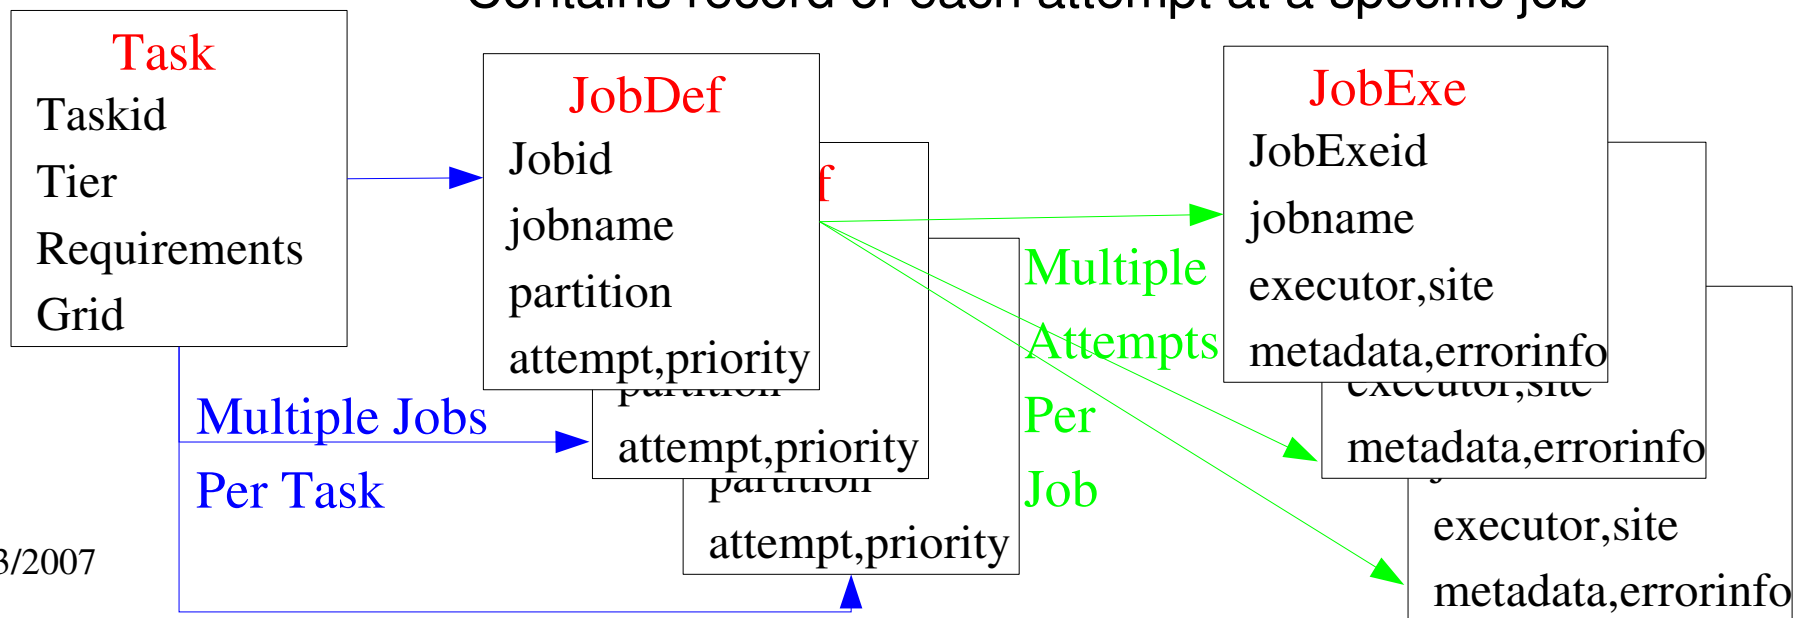
Prod-Sys



ProdDB

- **Hierarchy of tables**

- **Task Table** - 5,865
 - Contains task definition data
- **JobDef Table** - 4,032,571
 - Contains definition of each job within a task
- **JobExe Table** - 5,790,636
 - Contains record of each attempt at a specific job



Eowyn - Supervisor

- The Supervisor component
 - Retrieve Jobs from ProdDB
 - Submit to executor – communication via python objects
 - Maintain in memory database of associated jobs, modify ProdDB on state change
 - Manages failed jobs, re-releasing to be tried again or **autoaborting** job if failing persistently for same reason
 - Logging of transactions in proddb
 - Post Processing, filling output files into DDM datasets

Executor(s)

- **Interface to the Grids**
 - One For Each Grid Flavour, written by experts associated to corresponding grid
 - Determining number of free slots
 - Creating wrapper files
 - Submit job to grid
 - Mechanism for monitoring job status
 - Retrieving job after execution
 - Able to interpret errors, grid specific problems.


Recent Production

A look at some Production Stats

Recent Production


- Production tends to run in phases
- **Release Validation**
 - Standard release validation routine
 - Many errors in tasks, need reporting and fixing
 - Often seen that Grid related errors are not main problem during this period
- **Production**
 - Once a Release have been validated and probs fixed/understood
 - Errors switch to being Grid Dominated

Jobs/Wall Usage




Successful jobs Oct-Dec
All Grids perform better
after Validation Phase

High Failure rates during
Validation Phase.
Savannah bug reports etc
from operation teams




09/03/2007


Breaking Records



45000 jobs per day completed




Still O(50%) Job efficiency
But seems it could be getting better




09/03/2007

Breaking Records




45000 jobs per day completed

Walltime efficiency looks much better $O(80\%)$



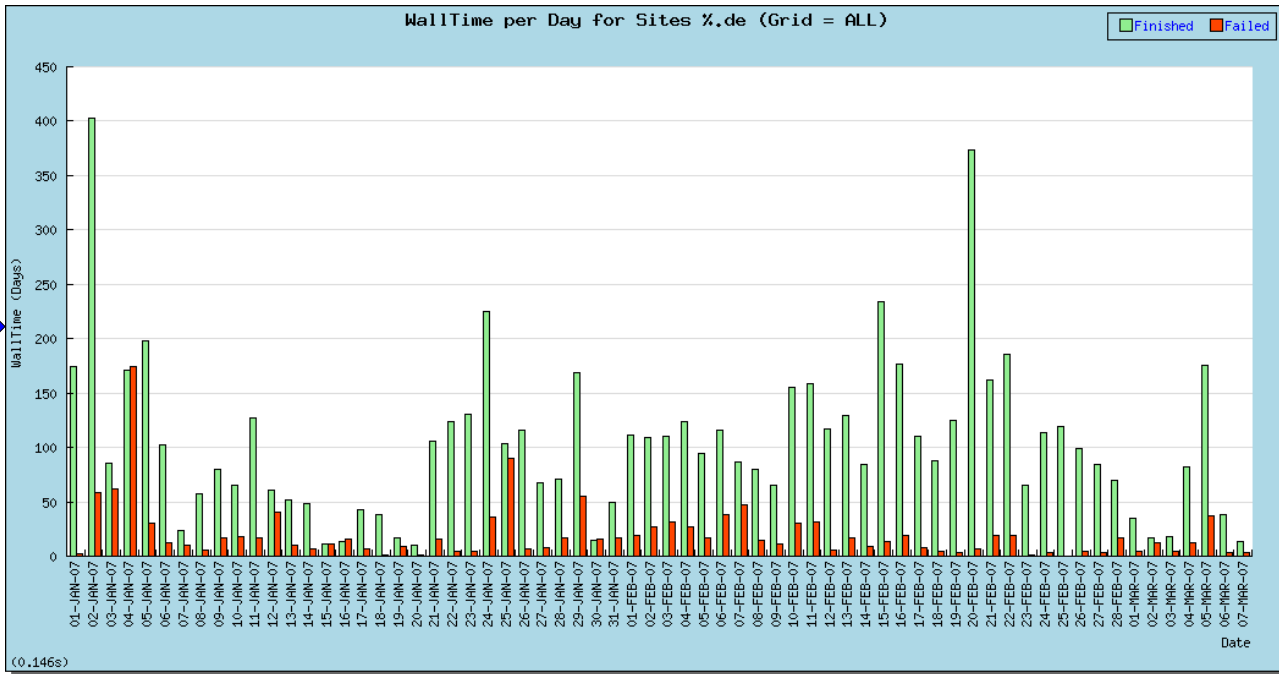
09/03/2007

Production in DE (This year)



Jobs with varying rates
eff O(60%)

Walltime efficiency looks
much better O(85%)



Production in DE (This year)

Site Efficiency:

#	Site	Finished	Failed	Efficiency (%)
1	a01-004-128.gridka.de	485	758	39
2	ce-fzk.gridka.de	28034	23364	55
3	ce.bfg.uni-freiburg.de	345	704	33
4	grid-ce.physik.uni-wuppertal.de	6388	4813	57
5	grid-ce0.desy.de	628	412	60
6	grid-ce1.desy.de	39984	10071	80
7	grid-ce2.desy.de	7002	3388	67
8	lcg-ce0.ifh.de	10261	9010	53
9	lcg-lrz-ce.lrz-muenchen.de		7	0
-	Total	93127	52527	64

← Jobs spread over sites
eff O(60%)

Site Efficiency:

#	Site	Finished (days)	Failed (days)	Efficiency (%)
1	a01-004-128.gridka.de	101.48	11.26	90
2	ce-fzk.gridka.de	2008.73	530.41	79
3	ce.bfg.uni-freiburg.de	35.05	8.63	80
4	grid-ce.physik.uni-wuppertal.de	345.66	102.11	77
5	grid-ce0.desy.de	139.74	20.88	87
6	grid-ce1.desy.de	2574.18	401.50	87
7	grid-ce2.desy.de	738.52	117.97	86
8	lcg-ce0.ifh.de	935.38	112.34	89
9	lcg-lrz-ce.lrz-muenchen.de	0.00	0.00	0
-	Total	6878.74	1305.1	84

Walltime efficiency looks
much better O(85%) →

What's New

what are **WE** doing

Executor and Shifts

- Running an Executor
 - CondorG instance running from LMU – Tariq Mahmoud (going well so far!)
 - Increasing # jobs as we learn (Now 250)
- Shifts
 - Increased contributions to production shifts
 - Marcel Schrörs from Wuppertal – here next week!
 - Manpower from CZ to help
- Help DE Cloud
 - Increased involvement and knowledge will help our sites to be more efficient

Conclusion/Outlook

- Production System
 - Jobs per day reached 45,000
 - Continued increase expected x2 each 2-3 Months – Big Challenge!
- DE in prod-sys
 - Our contribution increasing, more jobs = more fun!
 - Better monitoring of our sites and tasks – reporting and fixing problems!
- **Lots to be done this year!!!**