ATLAS Production System

John Kennedy
LMU Muenchen

Overview

- Who am I
- Prod-Sys Overview – EGEE(LCG)
- DDM
- ProdSys Monitoring
- Recent Production
- Other issues
- Conclusion

Who Am I – and what do I do

- GridKa Cloud Coordinator – Beta Version
  - Responsible for technical coordination (T1/T2's)
  - Organising functional tests
  - Atlas contact – attend many Atlas meetings
  - Organising meetings and ensuring communication
- Primary author of prod-sys monitoring
- Contribute as expert shifter to EGEE production shifts
- DDM-Operations T1 contact for GridKa
- General Grid busy-body

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

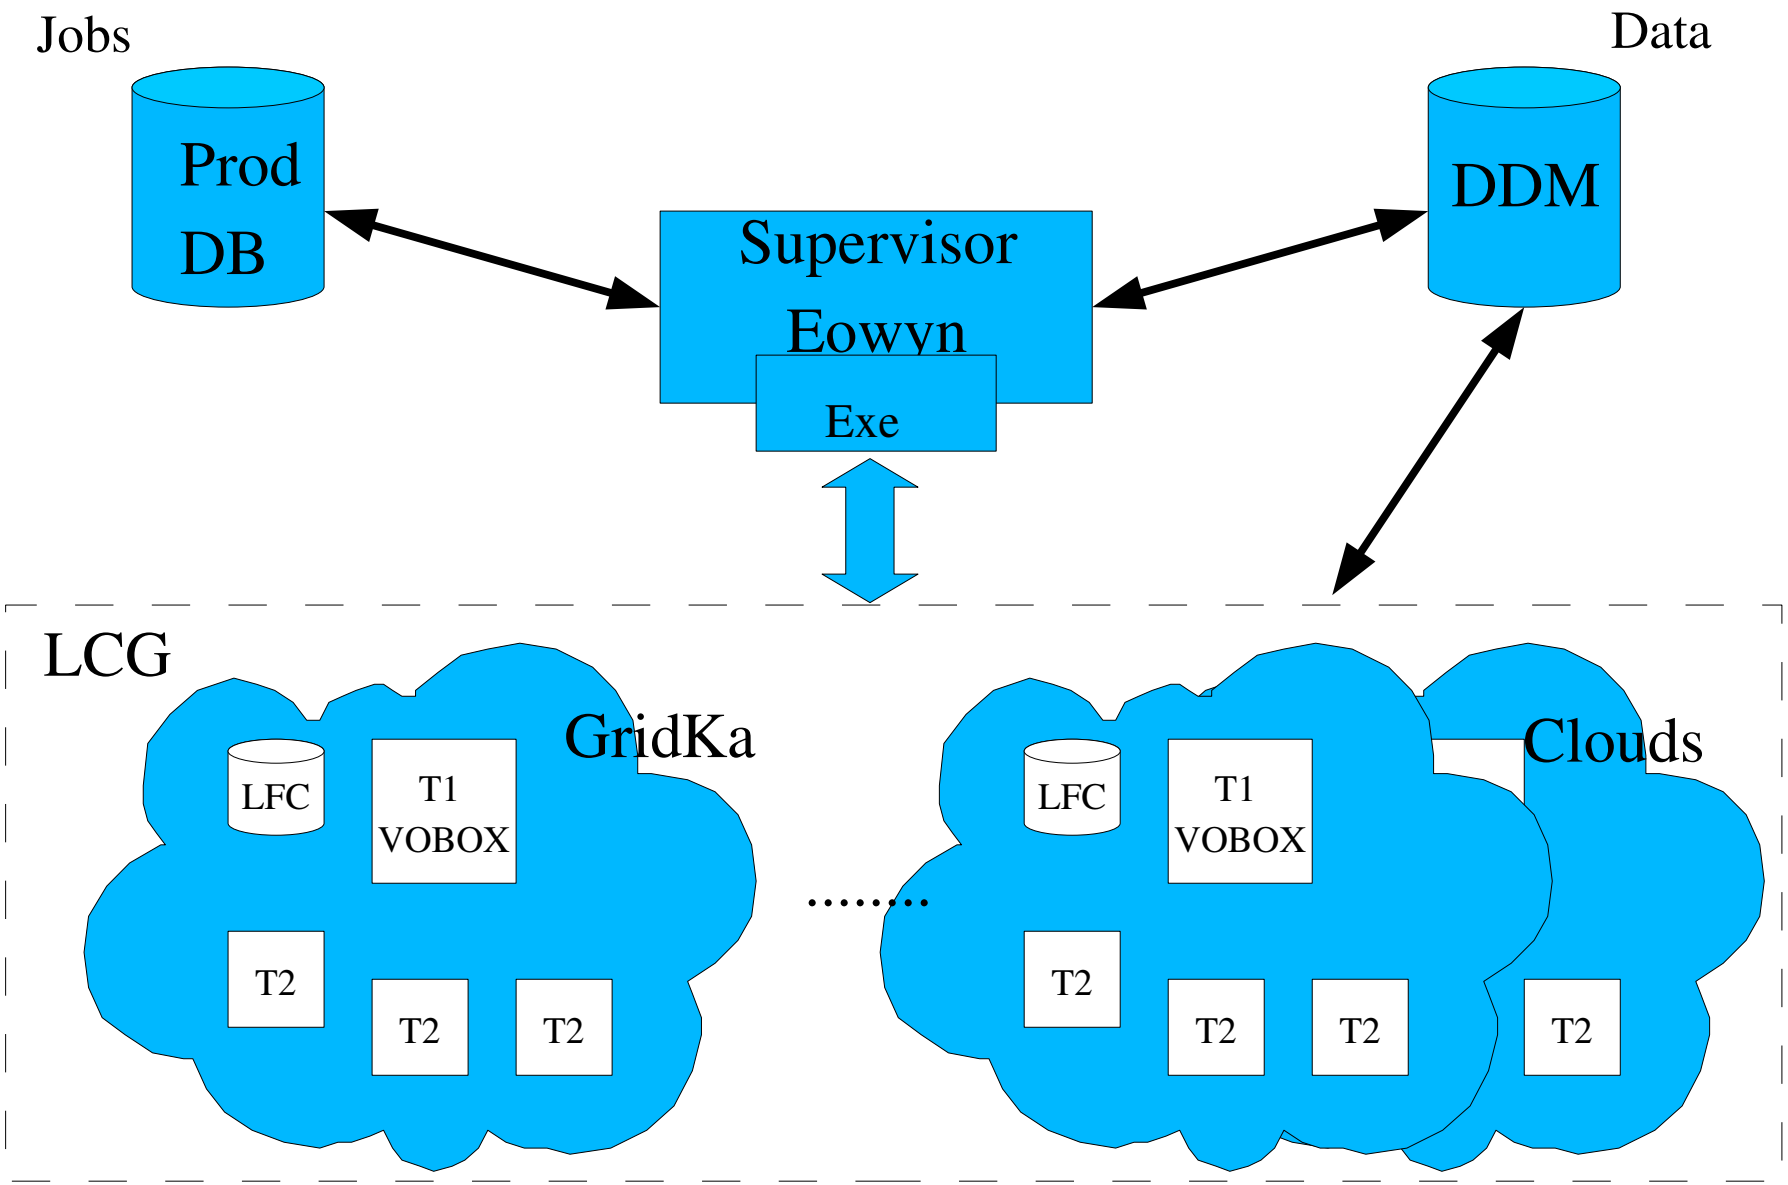
Jobs
Data
Prod DB
Supervisor Eowyn
Exe
DDM
LCG
GridKa
LFC
T1 VOBOX
T2
T2
T2
.....
Clouds
LFC
T1 VOBOX
T2
T2
T2
T2

ProdDB

- Hierarchy of tables
  - Task Table - 3,757
    - Contains task definition data
  - JobDef Table - 1,552,615
    - Contains definition of each job within a task
  - JobExe Table - 2,719,133
    - Contains record of each attempt at a specific job

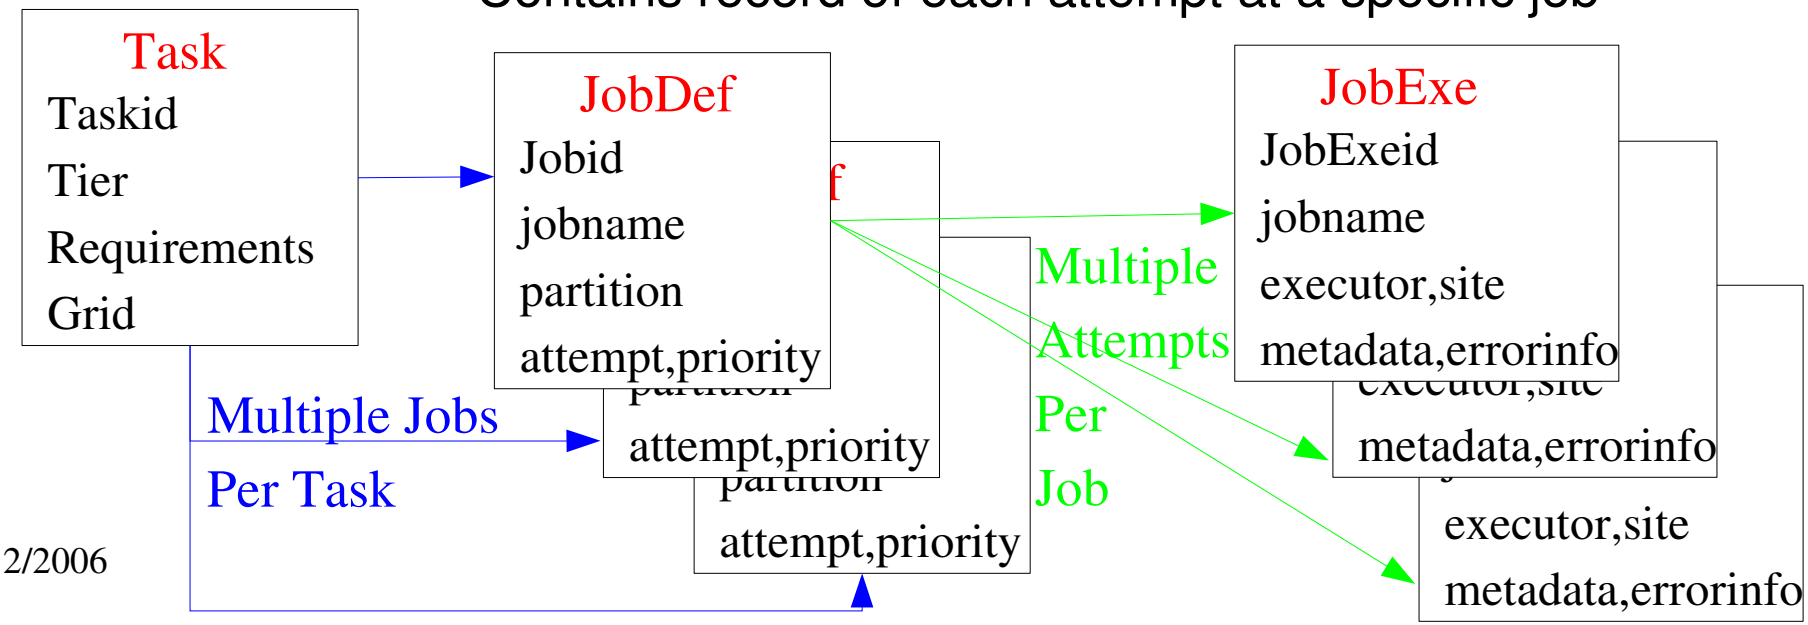

Eowyn - Supervisor

- 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.

Jobs at WN's

- What happens when a job lands lcg
  - Checks environment
    - LFC client,Valid Proxy, ATLAS SW.
  - Install pacman - download from a http cache
  - Install job transforms
  - Stages in inputfiles – loops over all LFC's looking for them
  - Run Athena
  - Stages out outputfiles – to cloud task is assigned to
  - Uploads logfile to webserver
  - Exits cleanly, metadata for executor in pickled file
  - Executor retrieves info

DDM

Distributed Data Management

DDM Intro

- Data Management on a File level cannot scale
- Data Management on a Dataset level rather than a File Level - Sounds Promising
- Central Dataset Catalogs
- Distributed File Catalogs
- By using DDM we expect to make things easier by a few orders of magnitude

DDM Architecture

- Two Main Components
- Dataset Catalogs
  - Contain dataset definitions,locations
  - Single set of central catalogs hosted at cern (currently)
- Site Services
  - Cataloging of data at each site
  - Movement of Data to sites
    - Scheduled transfers (FTS)
  - Removal of old data at sites

DDM Architecture

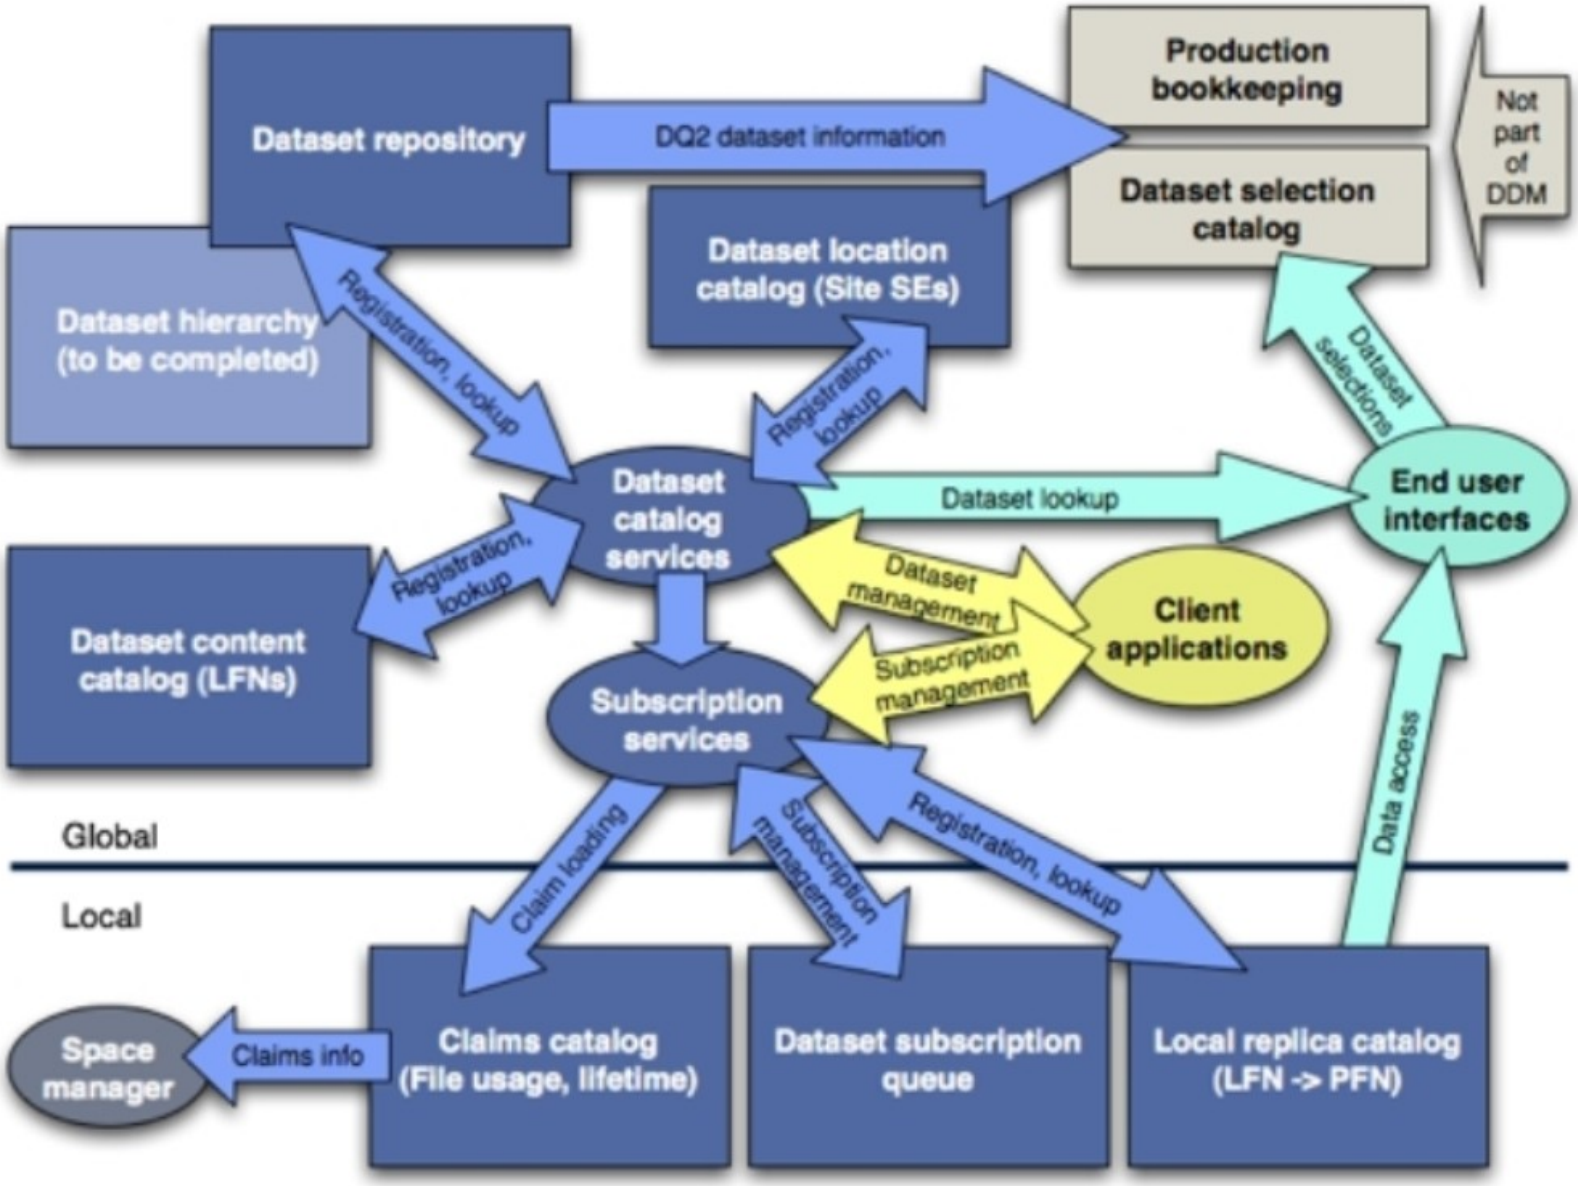

Warning: No matter how confusing you make your figures people will re-use them!

Datasets

- Datasets – what are they:
  - Contain files
    - Logical File Name, GUID, metadata
  - Have Versions
    - Files can be added/removed between versions
  - Has Unique Identifiers
    - Per dataset, per Version
  - Hierarchies
    - Datasets which contain other Datasets

Datasets Catalogs

- Dataset Repository Catalog
  - What Data exists in the system
- Dataset Content Catalog
  - What files are in a given dataset (version)
- Dataset Location Catalog
  - Where is the Data located
- Dataset Subscription Catalog
  - Keeps track of all requests for data (sites)
- Dataset Selection Catalog Not Implemented/Not Part of DQ2
  - What Datasets match a certain query
- Dataset Hierarchy Catalog Not Implemented
  - Records hierarchical organisation between datasets

Dataset states/versions

- States
  - Open
    - Files can be added
  - Closed
    - Files cannot be added (new version can be made)
  - Frozen
    - No more files/Versions (REAL DATA)
- Versions
  - Can track changes in data
  - Files can be added removed between versions
  - Subscriptions will get latest version unless otherwise instructed.

Site Services

- Subscription handling
  - Ensure that contents of a subscribed dataset are moved to the site
  - Automatic handling of dataset updates
- Cache turnover
  - Allow for deletion of unwanted data
- Cron Jobs run at T1's providing services for associated centers, you need to be supported to use DDM

A subscription Explained

- Subscribe a dataset to a site
- Services at T1 will pickup subscription and submit transfer request (FTS)
- T1 services monitor transfer status doing re-trys if needed
- If dataset is closed once completely at site marked as COMPLETE, otherwise remains INCOMPLETE (even though all files there)

A subscription Explained

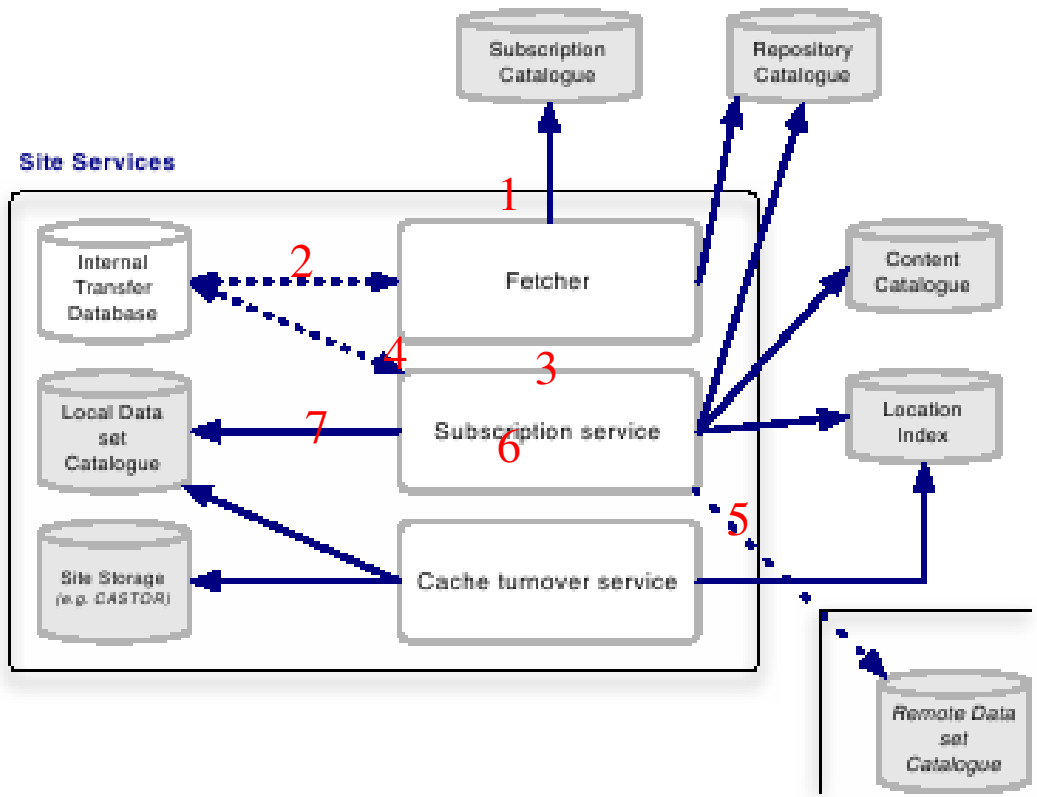

1) Fetcher- picks up new subscriptions from subscription catalog
2) Queues requests to local database of ongoing transfers
3) Subscription service finds missing files (repo,content cat,local cat)
4) Adds missing files to local database
5) Finds replicas of missing files
6) Interacts with data transfer layer (FTS), moves files in blocks
7) Bookkeeping of transfered files, registering in local dataset catalog

Client and End User are Tools

- Client tool exists for interface to central catalogs
  - Registering subscriptions, select source/destination...
  - Getting info about datasets
  - Listing subscriptions at sites
- A Rich Sample of end user tools.
  - dq2_ls
  - dq2_get
  - dq2_register
  - dq2_put
  - dq2_cleanup
  - dq2_sample
- End user tools may be extended

Prod-Sys Monitoring

The Monitoring Framework

Prod-Sys Monitoring

- Single point of info for users
- Monitoring
  - Enable people to spot problems with Sites,Tasks,Executors...
  - Allow for tracking of Task status
  - Present Info in a simple way
  - Allow for configuration
- Accounting
  - Use info from proddb to provide views of jobs/walltime/cputime

Prod-Sys Monitoring

- Attempt to abstract Monitoring away from production db.
- Moving snapshot windows and summary tables

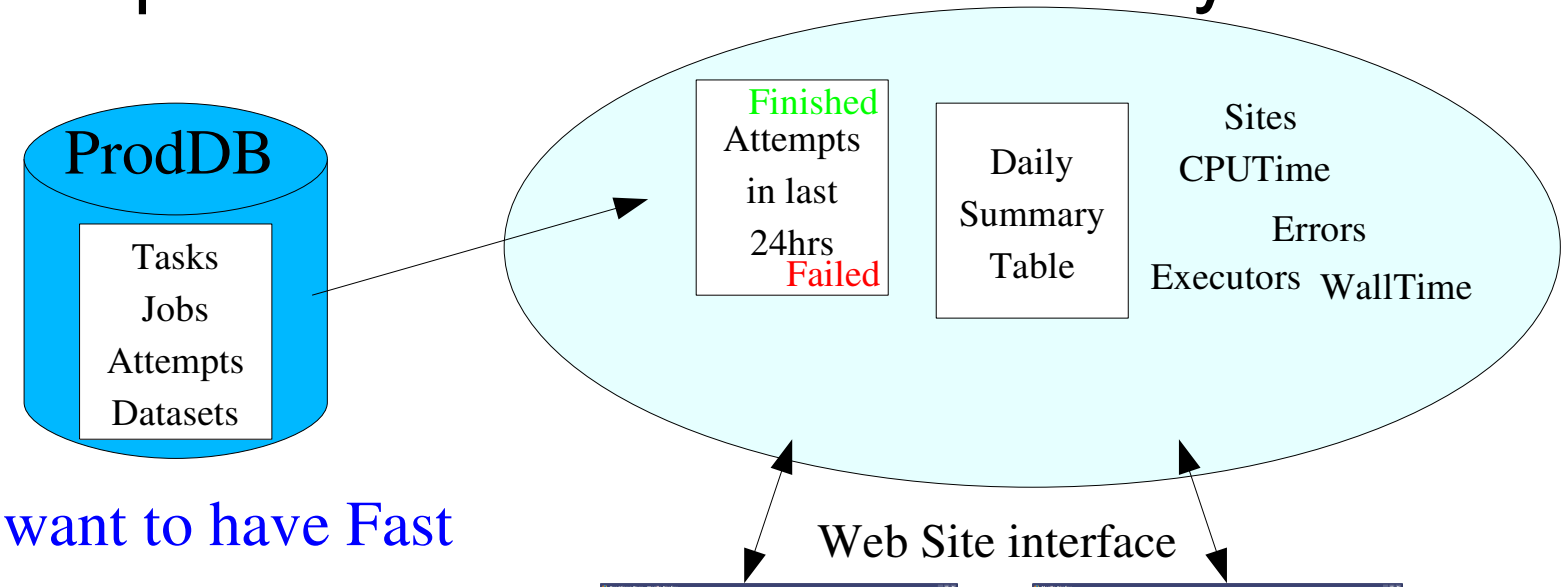

People want to have Fast Simple Monitoring.
Problems should jump out at people

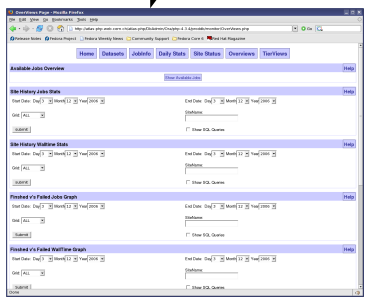

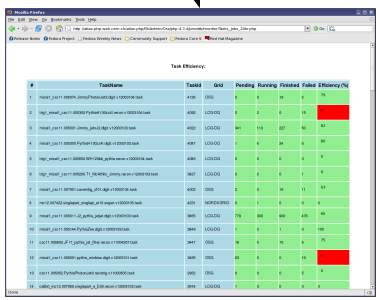

Monitoring-example

Sites

Efficiency

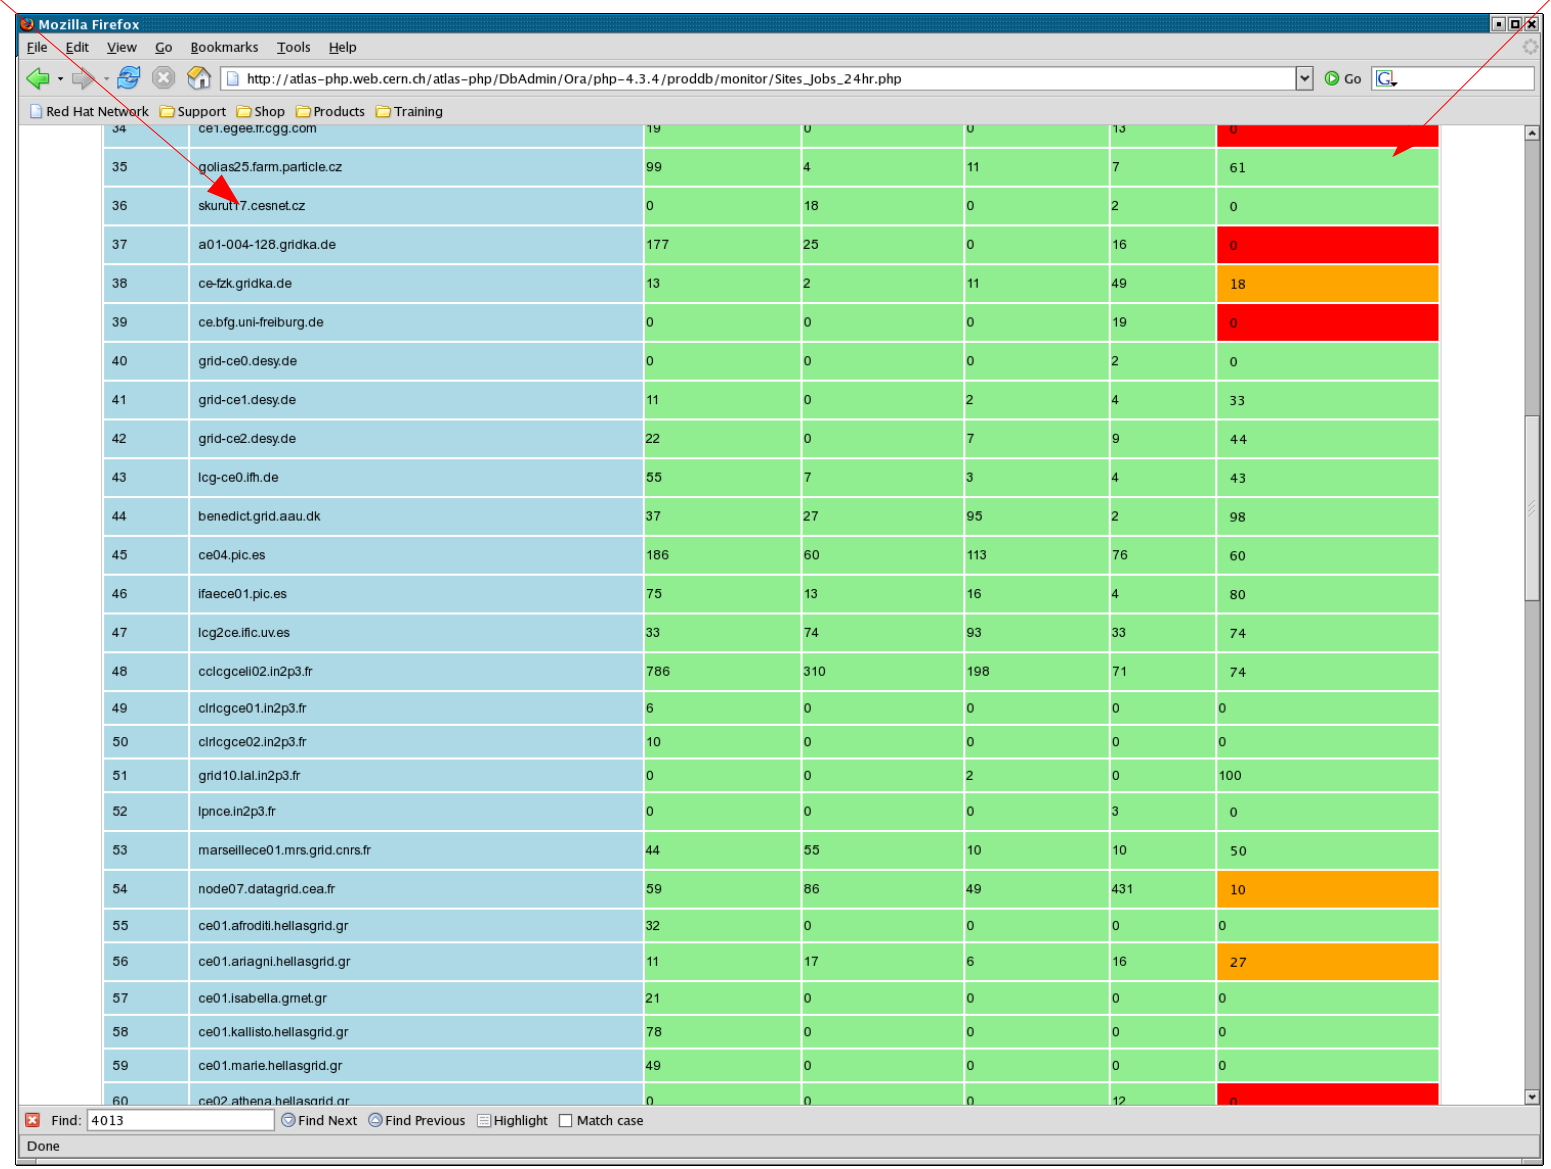

34	ce1.egee.ir.cgg.com	19	0	0	13	0
35	golias25.farm.particle.cz	99	4	11	7	61
36	skuruti7.cesnet.cz	0	18	0	2	0
37	a01-004-128.gridka.de	177	25	0	16	0
38	ce-fzk.gridka.de	13	2	11	49	18
39	ce.bfg.uni-freiburg.de	0	0	0	19	0
40	grid-ce0.desy.de	0	0	0	2	0
41	grid-ce1.desy.de	11	0	2	4	33
42	grid-ce2.desy.de	22	0	7	9	44
43	lcg-ce0.ifh.de	55	7	3	4	43
44	benedict.grid.aau.dk	37	27	95	2	98
45	ce04.pic.es	186	60	113	76	60
46	ifaece01.pic.es	75	13	16	4	80
47	lcg2ce.ific.uv.es	33	74	93	33	74
48	cclcgceli02.in2p3.fr	786	310	198	71	74
49	clrlcgce01.in2p3.fr	6	0	0	0	0
50	clrlcgce02.in2p3.fr	10	0	0	0	0
51	grid10.lal.in2p3.fr	0	0	2	0	100
52	lpnce.in2p3.fr	0	0	0	3	0
53	marseillece01.mrs.grid.cnrs.fr	44	55	10	10	50
54	node07.datagrid.cea.fr	59	86	49	431	10
55	ce01.afroditi.hellasgrid.gr	32	0	0	0	0
56	ce01.arlagni.hellasgrid.gr	11	17	6	16	27
57	ce01.isabella.grnet.gr	21	0	0	0	0
58	ce01.kallisto.hellasgrid.gr	78	0	0	0	0
59	ce01.marie.hellasgrid.gr	49	0	0	0	0
60	ce02.athena.hellasgrid.gr	0	0	0	12	0

Color coding helps spot sinners!

Clicking gives error summary

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

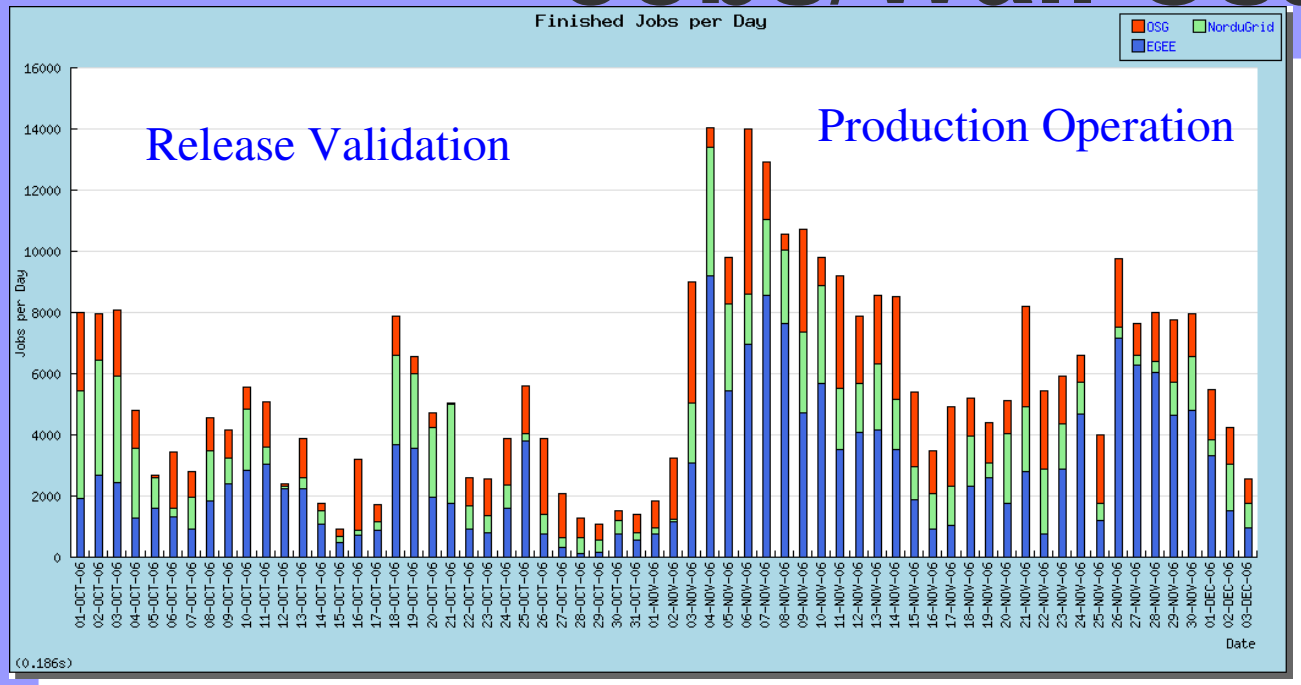

Successful jobs Oct-Dec
All Grids perform better
after Validation Phase

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

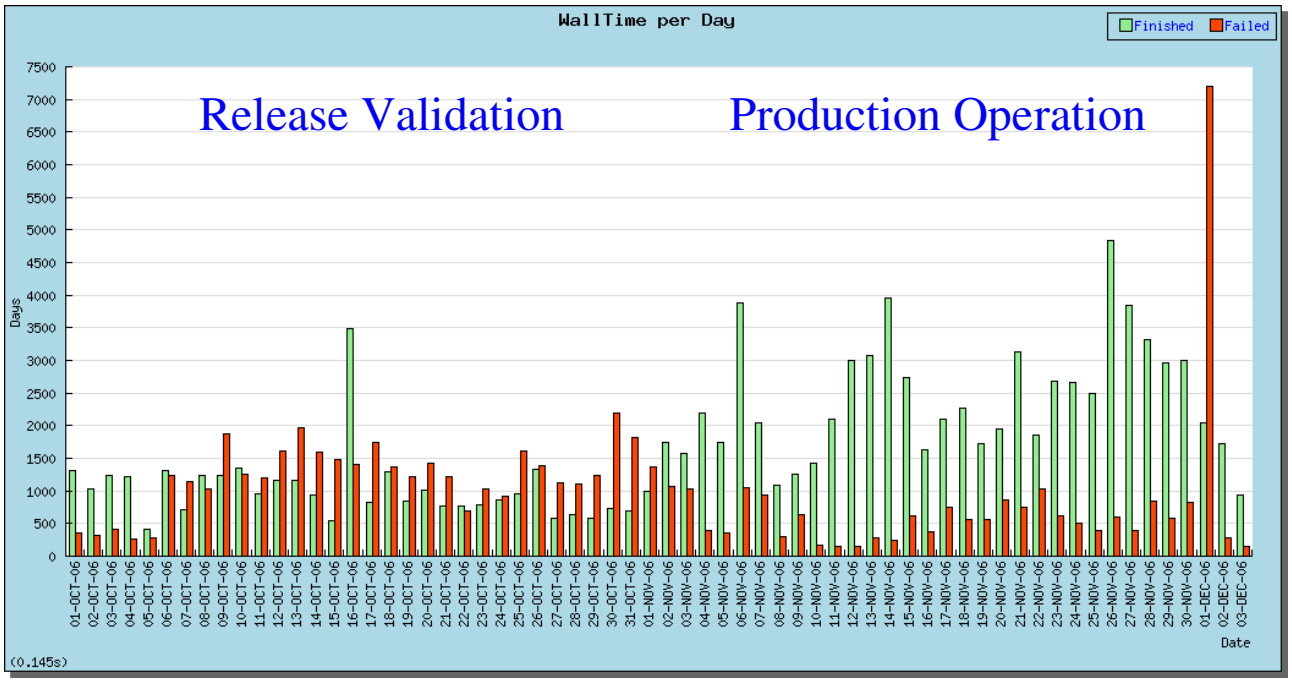

05/12/2006

Errors in a Day for EGEE

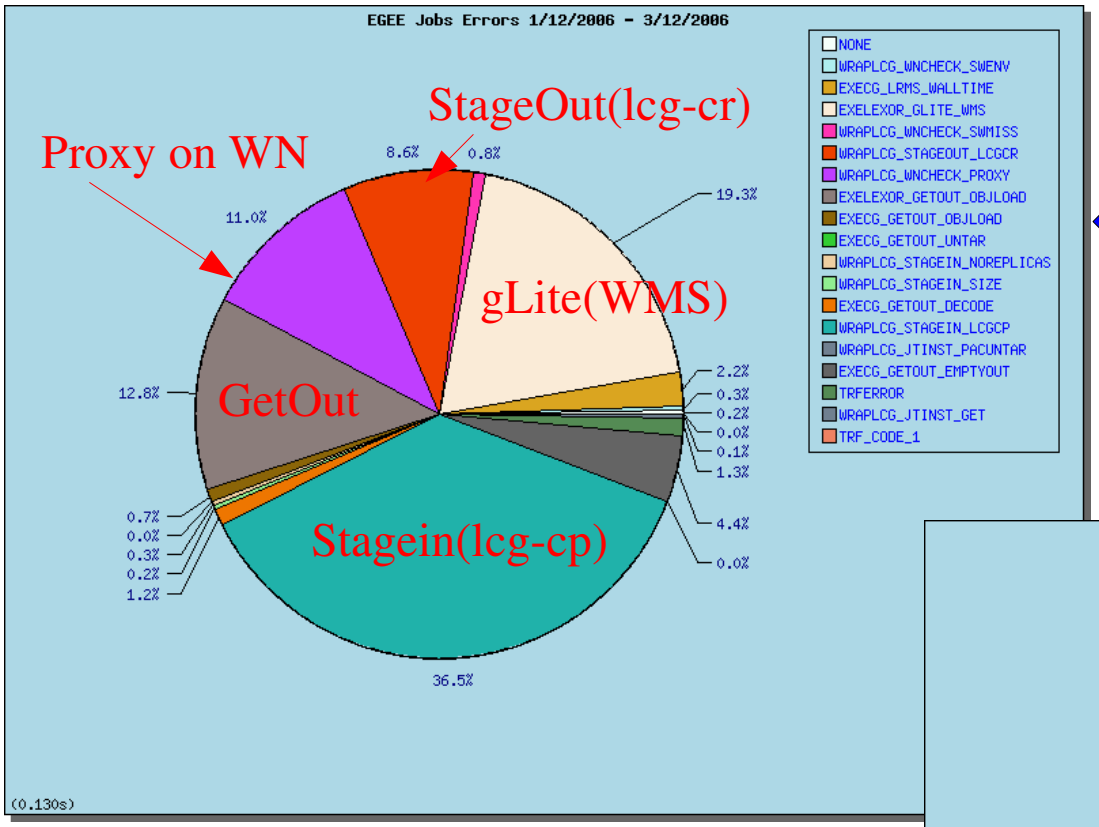

Looking at Jobs to highlight problems and reduce load

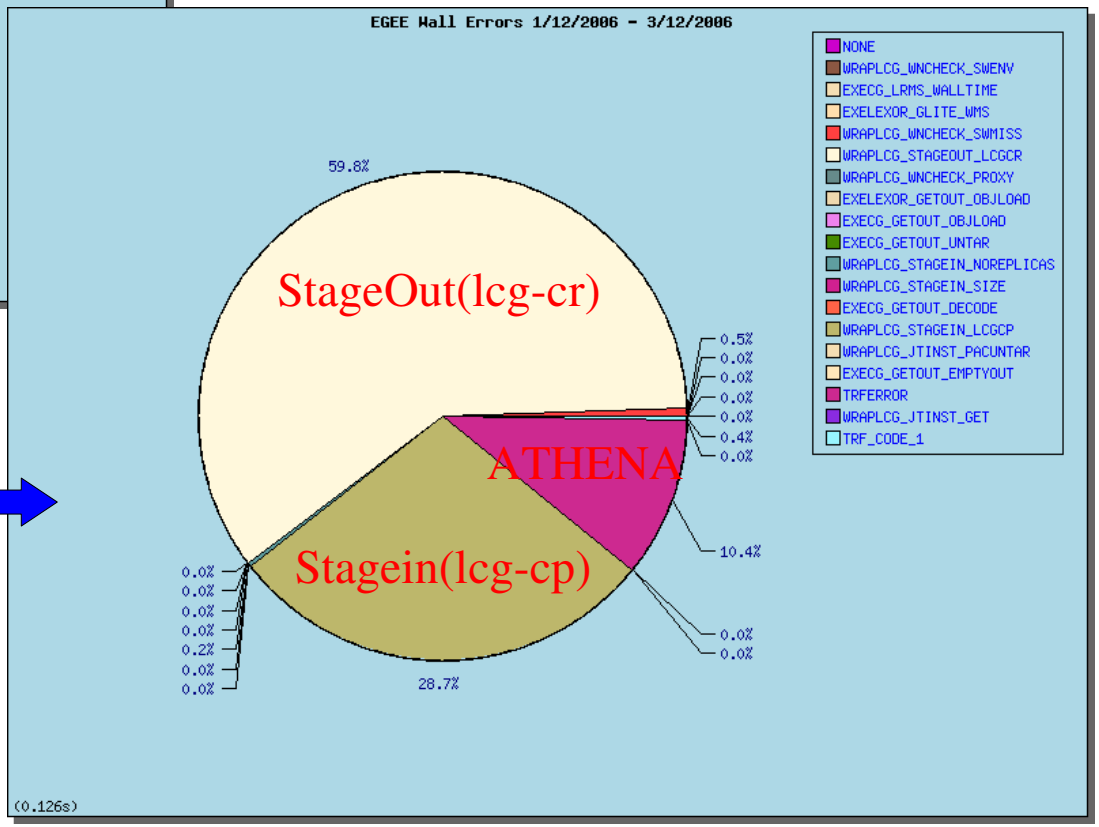

Looking at Walltime to find where/why we waste resources

Attempted Jobs

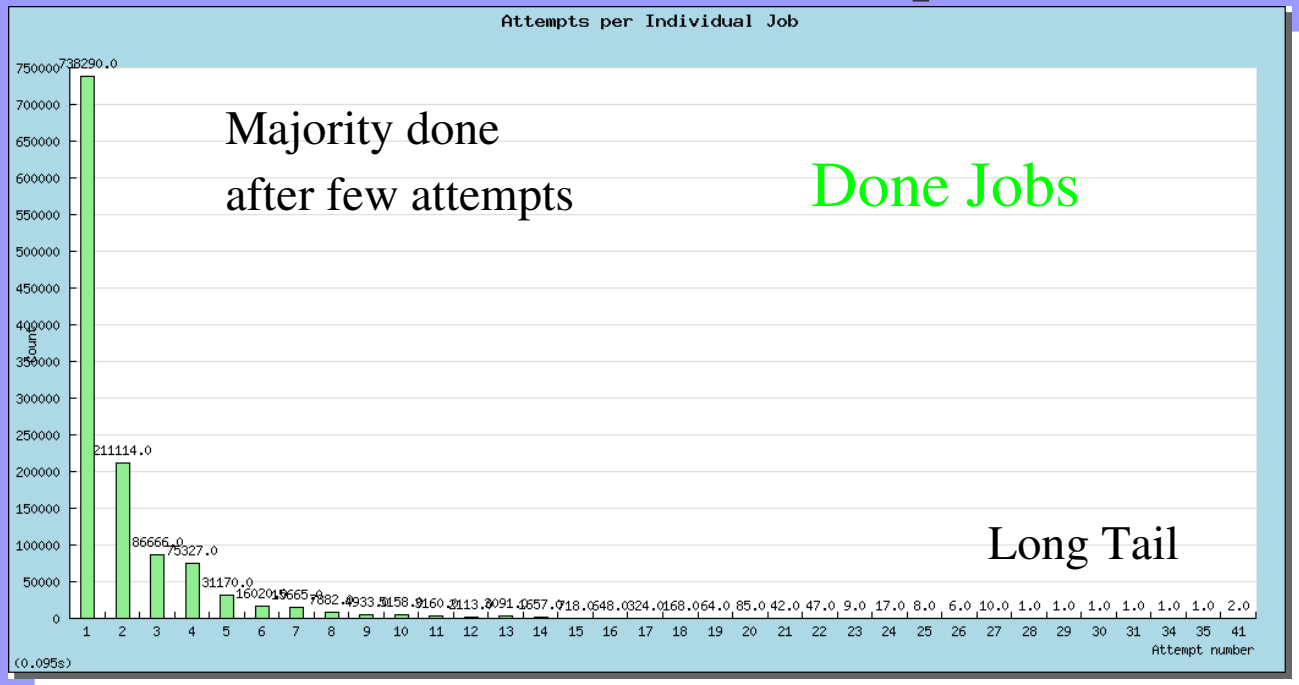

Num of attempts for successful jobs

Num of attempts for all jobs

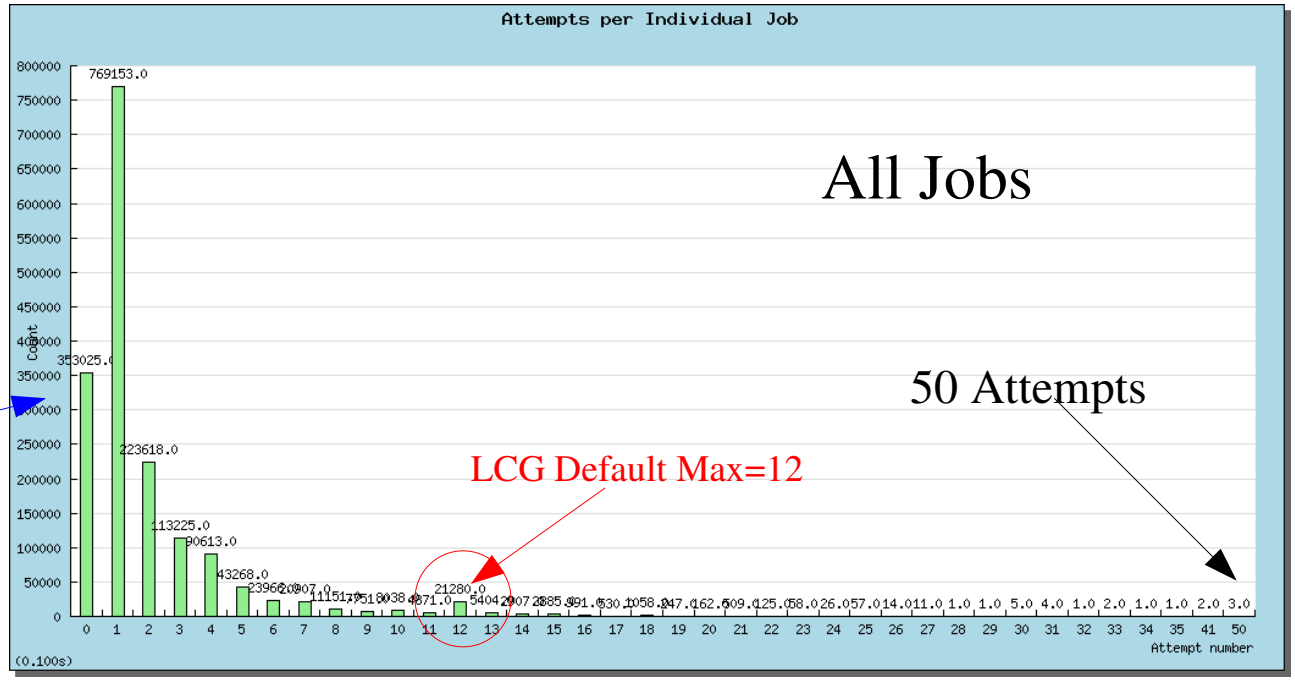

0 attempts for
jobs from aborted tasks
Thus stops false submission

05/12/2006

Zeuthen

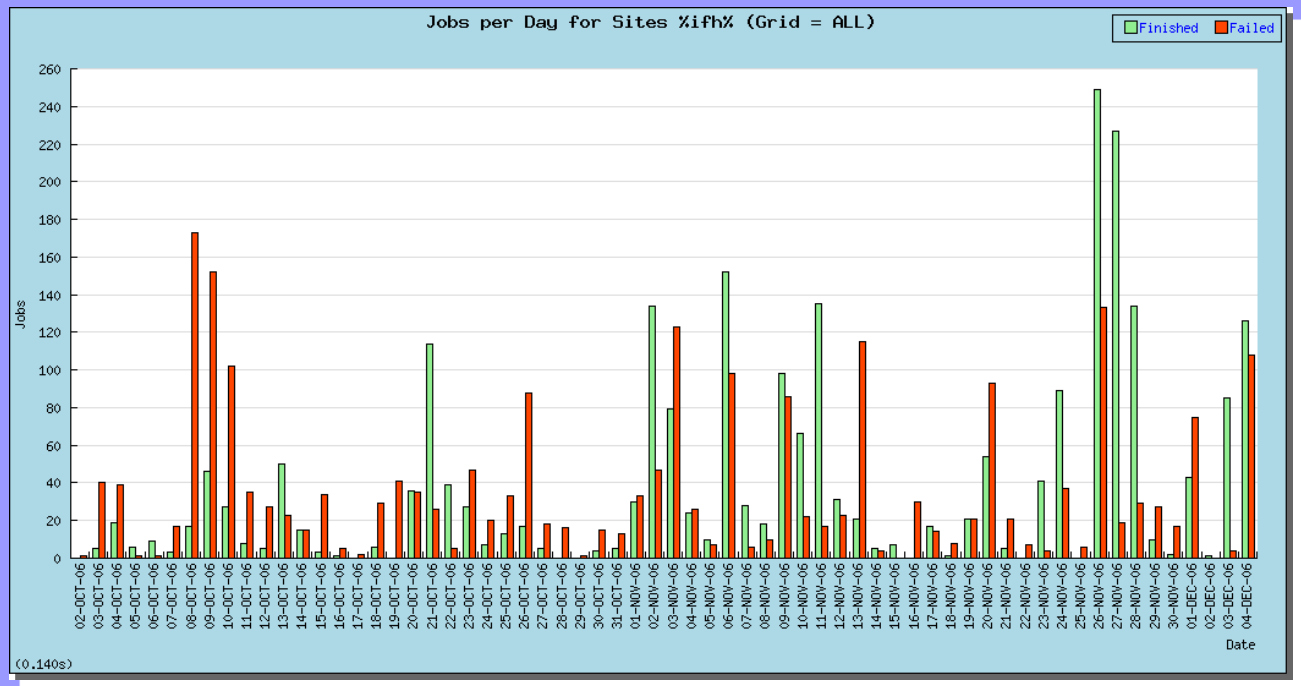

Jobs

Oct-Dec:

2430 Successful Jobs (51% Eff)

Whole year:

5966 Successful Jobs (51% Eff)

Walltime

Oct-Dec:

1079 Days (90% Eff)

Whole year:

2374 Days (88% Eff)

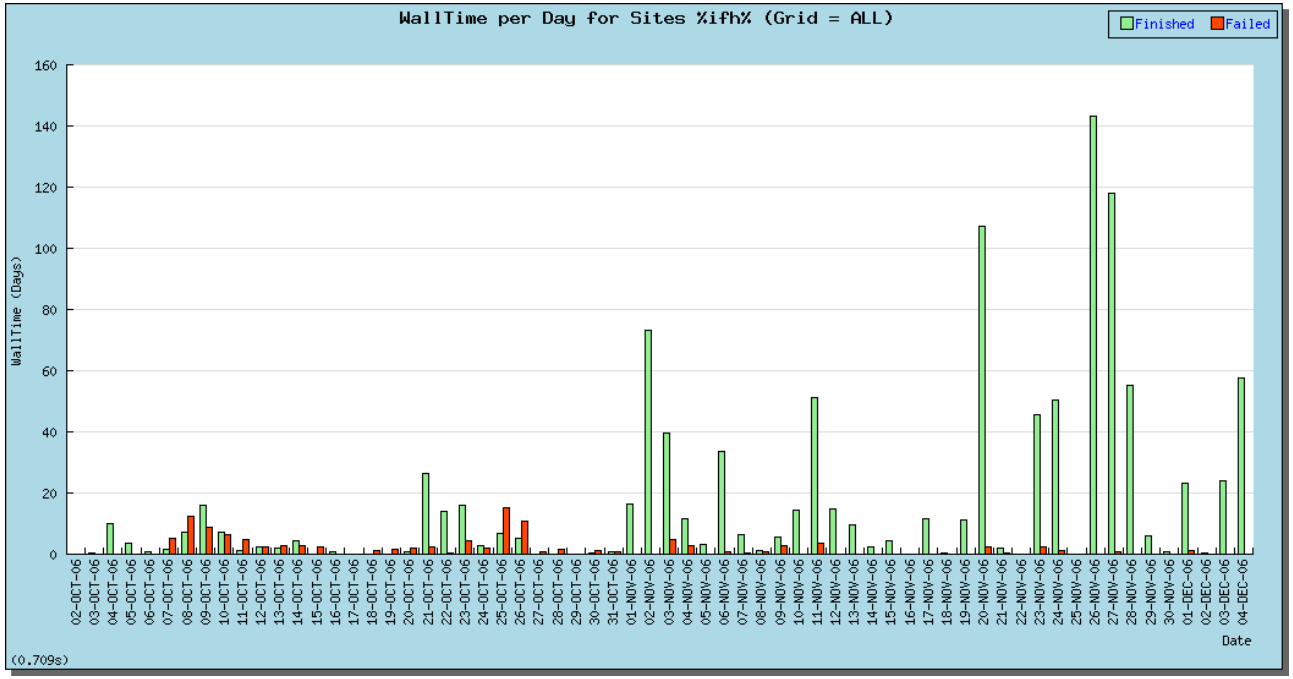

05/12/2006

Some Leftovers

Important but not mentioned yet

Production Shifts

Trying to split responsibilities to make life easier

- Assigning Jobs
  - Assigning tasks to clouds
- Submitting Jobs
  - Running an executor
- Shifts (experts and trainees)
  - Job Management:
    - Following tasks, reporting errors, stopping failing tasks
  - Data Management:
    - Following data flow, reporting+fixing problems

Cloud Services

- Much of the dependency is centered around the T1 (DDM)
  - VObox/LFC/FTS at T1
- Need ATLAS sw installed
- Data management within the cloud is V important
  - Ensuring Input Datasets available
  - DB releases replicated
  - Output Dataset Aggregation on T1
- We need to make sure we have people taking care of our services

Conclusion/Outlook

- Main Constituent parts of Prod-Sys exist
  - Supervisor
  - Executor
  - DDM
  - Monitoring
- And they are being tested and improved
- 10 fold increase in rate required next year
  - Expected increase in steps x2 every 2-3 Months
- We are confident that this can happen but it is a big challenge
  - Need manpower in prod-sys itself and also in Clouds to ensure services run smoothly

More Info

- The Wiki pages are a good source of info:
  - https://twiki.cern.ch/twiki/bin/view/Atlas/ProdSys
  - https://twiki.cern.ch/twiki/bin/view/Atlas/DDM
- Production Shifts
  - https://twiki.cern.ch/twiki/bin/view/Atlas/LcgProduction
- Production System Monitoring
  - http://atlas-php.web.cern.ch/atlas-php/DbAdmin/Ora/php-4.3.4/proddb/monitor/Home.php