



Enabling Grids for E-sciencE

New monitoring applications in the dashboard

Benjamin Gaidioz (CERN)

EGEE User Forum, Clermont-Ferrand 2008

www.eu-egee.org



- **Outline**

- Dashboard
- Updates on job monitoring, framework,
- Production monitoring for ATLAS
- Conclusion

- **The dashboard: a grid monitoring project**
 - Showing VO usage of the grid (jobs, data),
 - Developed at CERN
- **Main applications:**
 - Job monitoring, data transfers monitoring (ATLAS),
 - Efficiency of grid sites,
 - <http://dashboard.cern.ch>
- **Recent developments:**
 - User analysis tasks, HEP production monitoring and shifters help.



Enabling Grids for E-sciencE

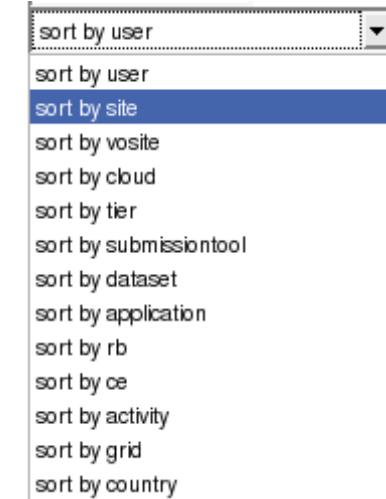
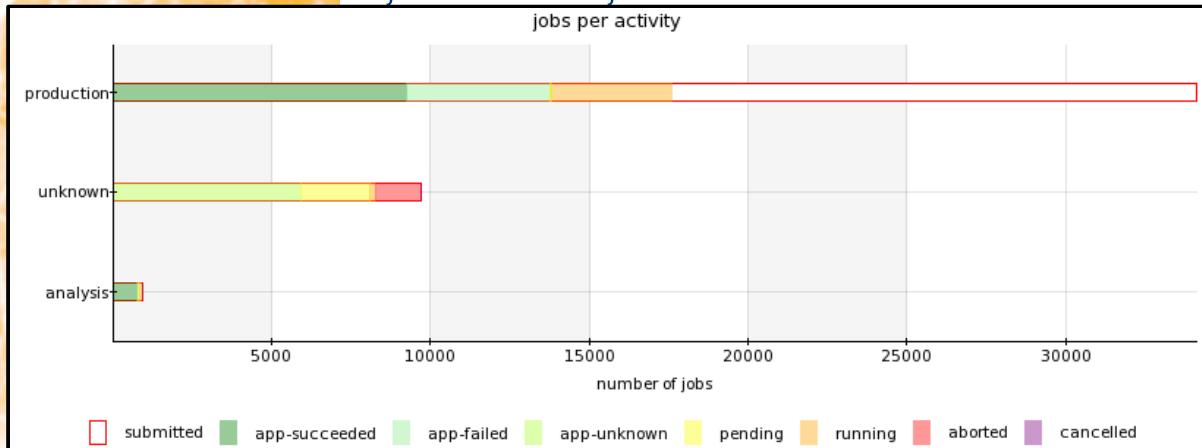
Some updates about existing applications and components

www.eu-egee.org

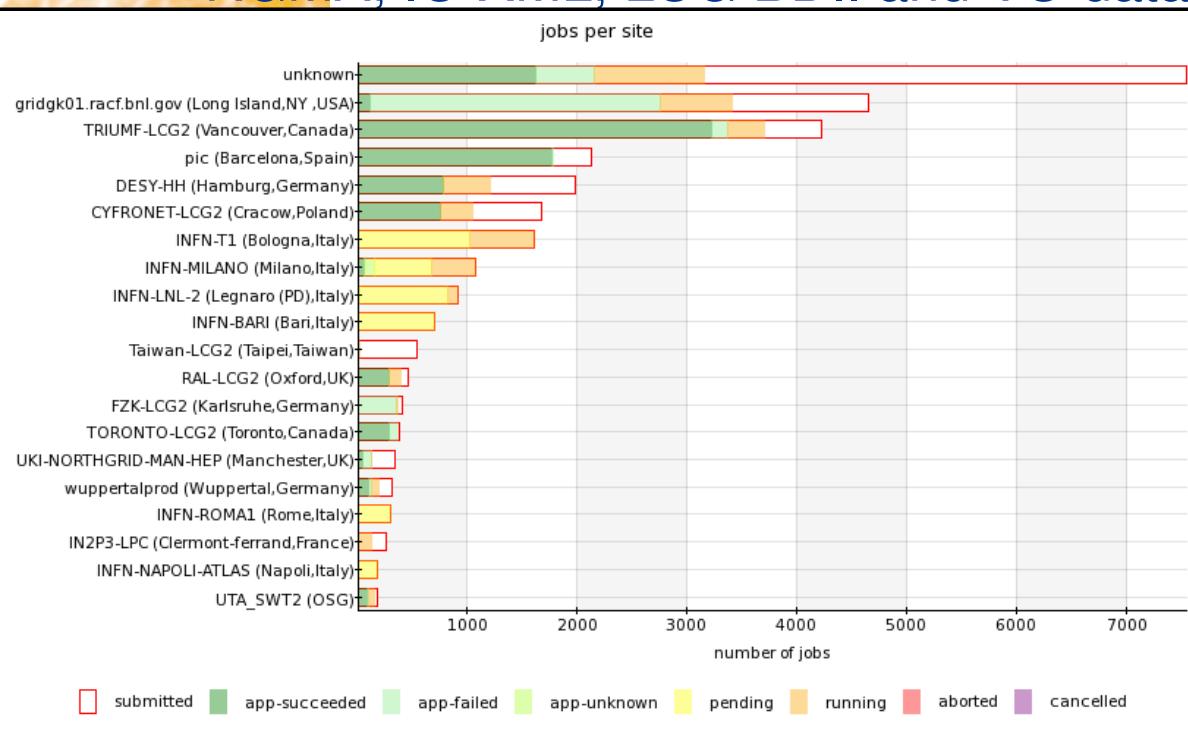


- **The job monitoring application**

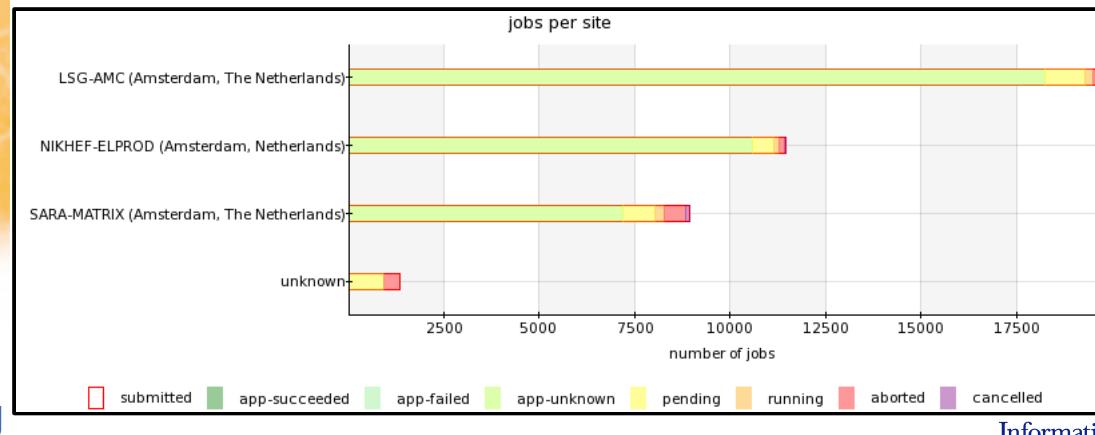
- A tool helping the VO to see its jobs on the grid,
- Multi grid, VO metadata associated to jobs,
- RGMA, IC-XML, LCG BDII and VO data



- **The job monitoring application**
 - A tool helping the VO to see its jobs on the grid,
 - Multi grid, VO metadata associated to jobs,
 - RGMA, IC-XML, LCG BDII and VO data



- **Update about the job monitoring application**
 - Standalone installation by NIKHEF for VL-eMed (VO),
 - Running since a year,
 - Now in collaboration with them to upgrade and include vlemed metadata in the monitoring information
 - More info: dashboard poster, fMRI analysis demo



any user

any user

Andreas C. van Kan
Coen Schrijvers
KAMEL BOULEBIAR
Kamel Boulebiar
Maarten Hendrik Van Ingen
Maarten van Iterson
Marques Ricardo
Piter T. De Boer
Sanna Gevers
Silvia Delgado Olabarriaga
Thomas Flanitzer
Tristan Glatard
unknown user
unknown vlemed

- **User task monitoring**
 - A user view for the job monitoring information,

| TaskMonitorId | Num of Jobs | Pending | Running | Successful | Failed |
|--|-------------|---------|---------|------------|--------|
| ganga_314_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 300 | 0 | 0 | 232 | 68 |
| ganga_315_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 300 | 0 | 0 | 231 | 68 |
| ganga_316_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 128 | 22 |
| ganga_317_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 132 | 18 |
| ganga_318_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 100 | 0 | 0 | 93 | 7 |
| ganga_319_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 149 | 1 |
| ganga_320_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 100 | 0 | 0 | 99 | 1 |
| ganga_321_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 149 | 1 |

- Hopefully more practical for distributed analysis users.

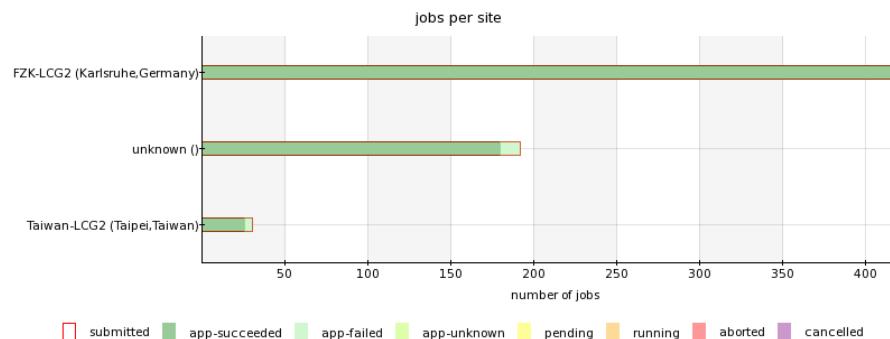
- **User task monitoring**
 - A user view for the same data (show progress of tasks)

| TaskMonitorId | Num of Jobs | Pending | Running | Successful | Failed |
|--|-------------|---------|---------|------------|--------|
| ganga_314_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 300 | 0 | 0 | 232 | 68 |
| ganga_315_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 300 | 0 | 0 | 231 | 68 |
| ganga_316_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 128 | 22 |
| ganga_317_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 132 | 18 |
| ganga_318_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 100 | 0 | 0 | 93 | 7 |
| ganga_319_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 149 | 1 |
| ganga_320_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 100 | 0 | 0 | 99 | 1 |
| ganga_321_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 149 | 1 |

| SchedulerJobId | EventRange | Submission Attempts | Appl Exit Code | Grid End Status | Site | Submitted | Started | Finished |
|---|------------|---------------------|----------------|-----------------|-------------|---------------------|---------------------|---------------------|
| https://rb106.cern.ch:9000/OomR5zGUBEMvKV2c3dw-Mw | 111 | 1 | 134 | DONE | Taiwan-LCG2 | 2008-02-09 23:00:52 | 2008-02-10 05:08:42 | 2008-02-10 05:27:17 |
| https://rb106.cern.ch:9000/W3YsFNM4YSUuwGDhYoWSnw | 113 | 1 | 188 | DONE | Taiwan-LCG2 | 2008-02-09 23:01:14 | 2008-02-10 05:06:43 | 2008-02-10 05:13:20 |
| https://rb121.cern.ch:9000/GbqyvuDn1SBxqEv5UyQ_dw | 119 | 1 | 188 | DONE | Taiwan-LCG2 | 2008-02-09 23:02:31 | 2008-02-10 05:12:34 | 2008-02-10 05:18:54 |
| https://rb106.cern.ch:9000/_PAKgX0t8Ql3mkrwWjimg | 12 | 1 | 188 | DONE | Taiwan-LCG2 | 2008-02-09 22:43:03 | 2008-02-10 04:00:07 | 2008-02-10 04:06:35 |
| https://rb106.cern.ch:9000/30GX9_jmRMmcBm_Xvi5-n | 120 | 1 | 134 | DONE | Taiwan-LCG2 | 2008-02-09 | 2008-02-10 | 2008-02-10 |

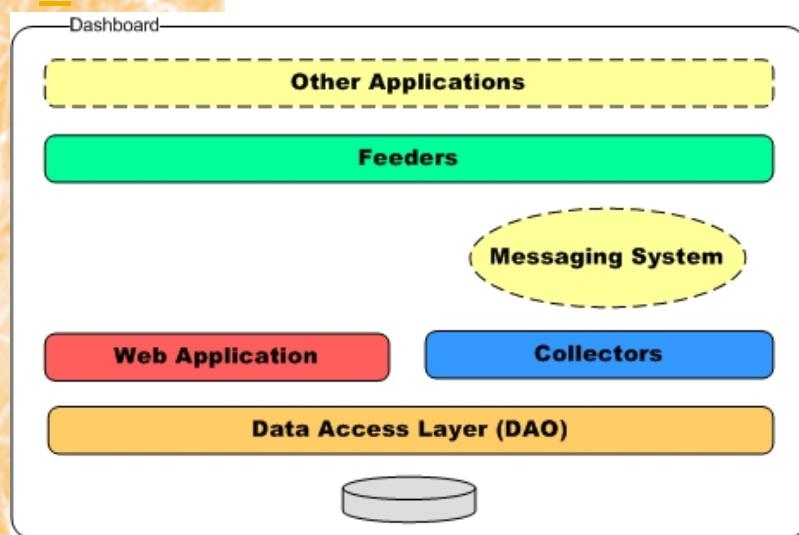
- **User task monitoring**
 - A user view for the same data (show progress of tasks)

| TaskMonitorId | Num of Jobs | Pending | Running | Successful | Failed |
|--|-------------|---------|---------|------------|--------|
| ganga_314_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 300 | 0 | 0 | 232 | 68 |
| ganga_315_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 300 | 0 | 0 | 231 | 68 |
| ganga_316_dprice@lxplus218.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 128 | 22 |
| ganga_317_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 132 | 18 |
| ganga_318_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 100 | 0 | 0 | 93 | 7 |
| ganga_319_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 149 | 1 |
| ganga_320_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 100 | 0 | 0 | 99 | 1 |
| ganga_321_dprice@lxplus221.cern.ch:/afs/cern.ch/user/d/dprice/gangadir/repository/ | 150 | 0 | 0 | 149 | 1 |



Integrated with the job monitoring application

- **dashboard framework: common base for applications:**



- Apache, mod_python, XSL/XML,
- Build system, release system, documentation (dev guide, demo application)

- ✓ arda.dashboard.dao-oracle-job-productII
- ✗ >arda.dashboard.demo
 - ▷ build
 - ▷ config
 - ▷ doc
 - ✗ >lib
 - ▷ >dashboard
 - ✗ >templates
 - ▷ css
 - ▷ js
 - ▷ >xsl
 - 📄 .cvsignore 1.1.1.1 (ASCII -kkv)
 - 📄 LICENSE 1.1.1.1 (ASCII -kkv)
 - 📄 MANIFEST.in 1.2 (ASCII -kkv)
 - 📄 module.cfg 1.3 (ASCII -kkv)
 - 📄 README 1.1.1.1 (ASCII -kkv)

- dashboard framework: common base for applications:

```
[lxplus237] /afs/cern.ch/user/b/bgaidioz > dashb-job-summary -f text/plain
name | sub | unk | pend | run | term | done | canc | abort | g-unk | grid% | succ | fail | a-unk | app% | d/
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
production | 45621 | 15577 | 0 | 6130 | 23914 | 0 | 0 | 0 | 23914 | 0.00 | 19588 | 4326 | 0 | 81.91 | 0
analysis | 1443 | 0 | 565 | 70 | 808 | 327 | 9 | 9 | 463 | 97.32 | 529 | 262 | 17 | 66.88 | 28
unknown | 1111 | 0 | 155 | 162 | 794 | 720 | 44 | 27 | 3 | 96.39 | 3 | 0 | 791 | 100.00 | 0
[lxplus237] /afs/cern.ch/user/b/bgaidioz > dashb-job-summary -f text/plain --sortby=rb --activity=unknown
name | sub | unk | pend | run | term | done | canc | abort | g-unk | grid% | succ | fail | a-unk
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
rb01.pic.es | 452 | 0 | 99 | 100 | 253 | 252 | 0 | 1 | 0 | 99.60 | 0 | 0 | 0 | 253
rb121.cern.ch | 222 | 0 | 14 | 38 | 170 | 169 | 0 | 1 | 0 | 99.41 | 0 | 0 | 0 | 170
rb106.cern.ch | 128 | 0 | 4 | 5 | 119 | 114 | 0 | 5 | 0 | 95.80 | 0 | 0 | 0 | 119
lcgrb01.gridpp.rl.ac.uk | 124 | 0 | 15 | 1 | 108 | 52 | 43 | 13 | 0 | 80.00 | 0 | 0 | 0 | 108
lcgrb02.gridpp.rl.ac.uk | 89 | 0 | 19 | 7 | 63 | 62 | 1 | 0 | 0 | 100.00 | 0 | 0 | 0 | 63
lcgrb03.gridpp.rl.ac.uk | 70 | 0 | 1 | 0 | 69 | 69 | 0 | 0 | 0 | 100.00 | 0 | 0 | 0 | 69
lcg005.ihep.ac.cn | 10 | 0 | 0 | 2 | 8 | 1 | 0 | 7 | 0 | 12.50 | 0 | 0 | 0 | 8
lapp-rb01.in2p3.fr | 10 | 0 | 1 | 9 | 0 | 0 | 0 | 0 | 0 | 0.00 | 0 | 0 | 0 | 0
lb103.cern.ch | 5 | 0 | 2 | 0 | 3 | 0 | 0 | 0 | 3 | 0.00 | 3 | 0 | 0 | 0
rb-fzk.gridka.de | 1 | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 100.00 | 0 | 0 | 0 | 1
[lxplus237] /afs/cern.ch/user/b/bgaidioz >
```

- <item>

<terminated>581</terminated>

<name>SARA-MATRIX (Amsterdam, The Netherlands)

<unknown>0</unknown>

<app-failed>0</app-failed>

<submitted>584</submitted>

<app-succeeded>0</app-succeeded>

<running>1</running>

<done>581</done>

<aborted>0</aborted>

- Very practical flexibility in output (text/xml, text/csv, HTML, png)
- Common base for API and CLI.

- **dashboard framework: current developments:**
 - Monitoring of components integrated with Lemon, Nagios,
 - Integration with SAM (service monitoring in EGEE):
 - e.g. alert if ATLAS production is bad on site
 - Improve the flexibility of web applications for faster developments (more component based, AJAX, etc.).



Enabling Grids for E-sciencE

Monitoring of the production for ATLAS

www.eu-egee.org



Information Society
and Media



- **ATLAS production activities:**

- 75,000 jobs per day, on three middlewares (EGEE, OSG, Nordugrid),
- A central “queue” of jobs to run, dequeued by several submission systems (now converging to one),
- A couple of experts and many volunteers shifters,
- Knowledge transfer from experts to trainee shifters:
 - Procedures, workbook,
 - Tutorials during software weeks,
 - Permanent mailing-list, logbook, chat, etc.

- **Monitoring:**

- ATLAS is using a dashboard application for data transfers (DDM monitoring, for DMM see their presentation and poster),
- Had a working monitoring application for production (John Kennedy, ATLAS),
- Asked to have a dashboard application instead:
 - Maintenance was becoming an issue,
 - Potential integration with data transfer monitoring,
- Prototype developed from feb to sep 2007

- **Developments:**

- Several meetings with production system developers and expert shifters,
 - Private discussions with expert shifters,
 - Identification of missing features in the current system, etc.

- **Observations:**

- People were really willing to explain and give advices, but it's very time consuming for them (explanations, e-mails, etc.).

- **Prototype:**

- Provided practical browsing functionalities (missing in the original application),
- Overview displays (number of jobs, errors today, etc.),
- User's guide, FAQ

- **Observations:**

- In principle all possible views were accessible in three clicks,
- Everyone seemed to like.

ATLAS Production system

[Overview](#) [User Guide](#) [Feedback](#)

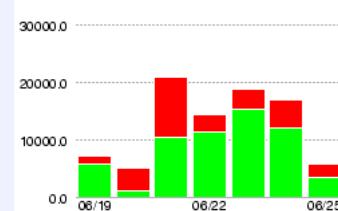
OVERVIEW

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

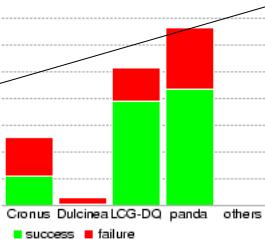
select a view

[others](#)
[GRIDARC](#)
[FILE DOWNLOAD](#)
[TRFERROR](#)
[WRAPSG](#)
[DD2](#)
[SERVERERROR](#)

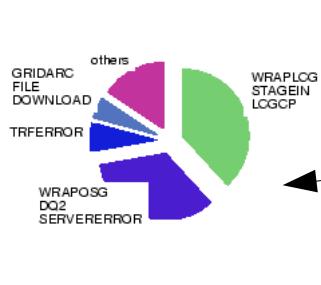
jobs last seven days



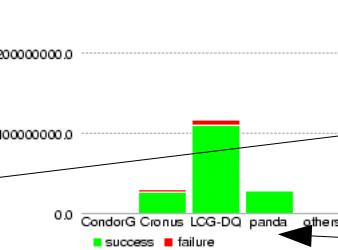
jobs (last 24h)



errors (last 24h)



walltime (last 24h)



terminated jobs

today (main four)

main four errors today

today, walltime

and Media

| | pending | running | success | failure | efficiency |
|----------|---------|---------|---------|---------|------------|
| panda | 0 | 0 | 4350 | 2253 | 65.9% |
| LCG-DQ | 479 | 207 | 3871 | 1257 | 75.5% |
| Cronus | 1 | 330 | 1055 | 1481 | 41.6% |
| Dulcinea | 0 | 0 | 2 | 273 | 0.7% |
| CondorG | 0 | 2 | 25 | 8 | 75.8% |

go to the error page

ATLAS Production system

[Overview](#) [User Guide](#) [Feedback](#)

OVERVIEW

- view
 - [by grid \(not avail.\)](#)
 - [by cloud \(not avail.\)](#)
 - [by executortype](#)
 - [by executor](#)
 - [by cluster](#)
- [select executortype](#)

- [panda](#)
- [Cronus](#)
- [LCG-DQ](#)
- [Dulcinea](#)
- [CondorG](#)



go to the job details display

first three errors of each



ATLAS Production system

[Overview](#)
[User Guide](#)
[Feedback](#)

OVERVIEW

view

| | <i>jobexecid</i> | <i>jobdefk</i> | <i>taskfk</i> | <i>jobname</i> | <i>error</i> | <i>message</i> |
|-----------------------|------------------|----------------|---------------|--|------------------------|-------------------------------|
| by grid (not avail.) | 8495531 | 5855147 | 10169 | misal1_mc12.005808.JF35_pythia_loosejet_filter.digit.v12000604_tid010169_39813.job | WRAPLCG_STAGEOUT_LCGCR | FAILEDSRMSIN srm://sc.cern.ch |
| by cloud (not avail.) | 8495026 | 5885763 | 10483 | ideal0_csc11.005014.J5_pythia_jet_filter.digit.v12000604_tid010483_01312.job | WRAPLCG_STAGEOUT_LCGCR | OKSRMSIN srm://srm.triumf.ca |
| by executortype | 8515677 | 5786239 | 10174 | misal1_mc12.005807.JF35_pythia_jet_filter.digit.v12000604_tid010174_23493.job | WRAPLCG_STAGEOUT_LCGCR | OKSRMSIN srm://ccsrm.in2p3.fr |
| by executor | 8514339 | 5786662 | 10174 | misal1_mc12.005807.JF35_pythia_jet_filter.digit.v12000604_tid010174_23916.job | WRAPLCG_STAGEOUT_LCGCR | OKSRMSIN srm://ccsrm.in2p3.fr |
| by cluster | 8520509 | 5787350 | 10174 | misal1_mc12.005807.JF35_pythia_jet_filter.digit.v12000604_tid010174_24604.job | WRAPLCG_STAGEOUT_LCGCR | FAILEDSRMSIN srm://srm-disk |
| select executor | 8496914 | 5784669 | 10174 | misal1_mc12.005807.JF35_pythia_jet_filter.digit.v12000604_tid010174_21923.job | WRAPLCG_STAGEOUT_LCGCR | FAILEDSRMSIN srm://srm-disk |
| | | | | error text: FAILEDERMEIN srm://srm-disk.pic.es OKSRMSIN srm://srm.grid.sara.nl FAILEDERMBOUT srm://srm-disk.pic.es srm://grid002.ft.uam.es srm://ifaesrm.pic.es srm://castorsrm.ific.uv.es FAILEDERBOUT srm://grid002.ft.uam.es OKERMBOUT srm://srm.cern.ch FAILEDERMBOUT srm://srm-disk.pic.es srm://grid002.ft.uam.es srm://ifaesrm.pic.es srm://castorsrm.ific.uv.es srm://srm.cern.ch FAILEDERBOUT srm://grid002.ft.uam.es | | |
| | | | | 8496914 supervisor: david_bulkMM#2122801152 infoexecutor: https://wms003.cnaf.infn.it:9000/3El69zLvDmvqUHKJnMRNwg creationtime: 2007-06-22 18:04:01 attemptnr: 1 errorcode: 410403 partnr: 21923 endtime: 2007-06-24 17:33:59 modificationtime: 2007-06-24 20:20:02 nevents: starttime: 2007-06-24 11:18:45 execcluster: lcgce01.triumf.ca facilityid: https://wms003.cnaf.infn.it:9000/VdnUwLmib7vcFakewEAxWA | | |
| | | | | 8509627 5784638 10174 misal1_mc12.005807.JF35_pythia_jet_filter.digit.v12000604_tid010174_21892.job | WRAPLCG_STAGEOUT_LCGCR | OKSRMSIN srm://ccsrm.in2p3.fr |
| | | | | 8508730 5784701 10174 misal1_mc12.005807.JF35_pythia_jet_filter.digit.v12000604_tid010174_21955.job | WRAPLCG_STAGEOUT_LCGCR | OKSRMSIN srm://ccsrm.in2p3.fr |

click a job to
get more
details


www.eu-egee.org


- **A python API**
 - For automatization instead of browsing a web page:

```
from dashboard.api.production import ProductionQuery
query = ProductionQuery()
errors = query.errors(site="IN2P3-LPC", grouping='clusters')
for error in errors:
    # etc.
```

- **Integration in the tutorials:**

- ATLAS was organizing tutorials in sep 2007,
- Dashboard part of it (as “next monitoring tool, please send feedback”),
 - Demo based on the instructions found in tutorial slides,
 - Howto in the user's guide.
- Lots of feedback (most from managers and expert shifters),
- However the existing tools were still the ones to use (sep 2007).

- **Developments then (sep 2007 to dec 2007):**

- Integration of middleware information systems (LCG BDII, OSG, Nordugrid ARC) to recognize grids, sites, from the cluster name reported by the production system,
- Manager's display (select dates, plots of success rate, share between the three grids, etc.),
- Integration with the monitoring of submission tools (links to log files for example).

OVERVIEW

view
by grid
by cloud
by dest_cloud
by executortype
by executor
by site
by cluster

select grid

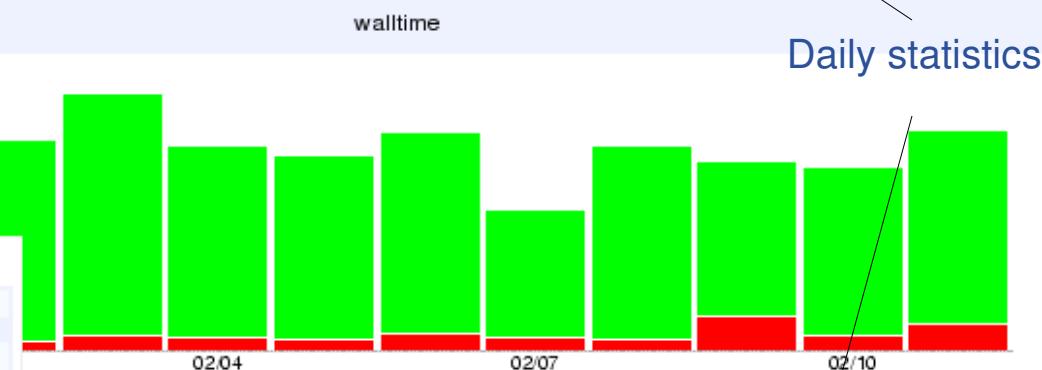
LCG
OSG
None
Nordugrid

Activity in ...

Start Date / Time

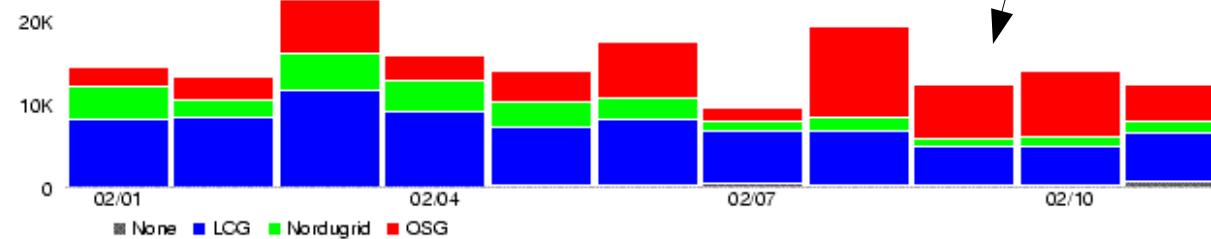
| February - 2008 | | | | | | |
|-----------------|----|----|-------|----|----|----|
| S | M | T | W | T | F | S |
| 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 10 | 11 | 12 | 13 | 14 | 15 | 16 |
| 17 | 18 | 19 | 20 | 21 | 22 | 23 |
| 24 | 25 | 26 | 27 | 28 | 29 | |
| 2008-02-12 | | | 07:00 | | | |

walltime



Daily statistics

jobs (successful jobs)



OVERVIEW

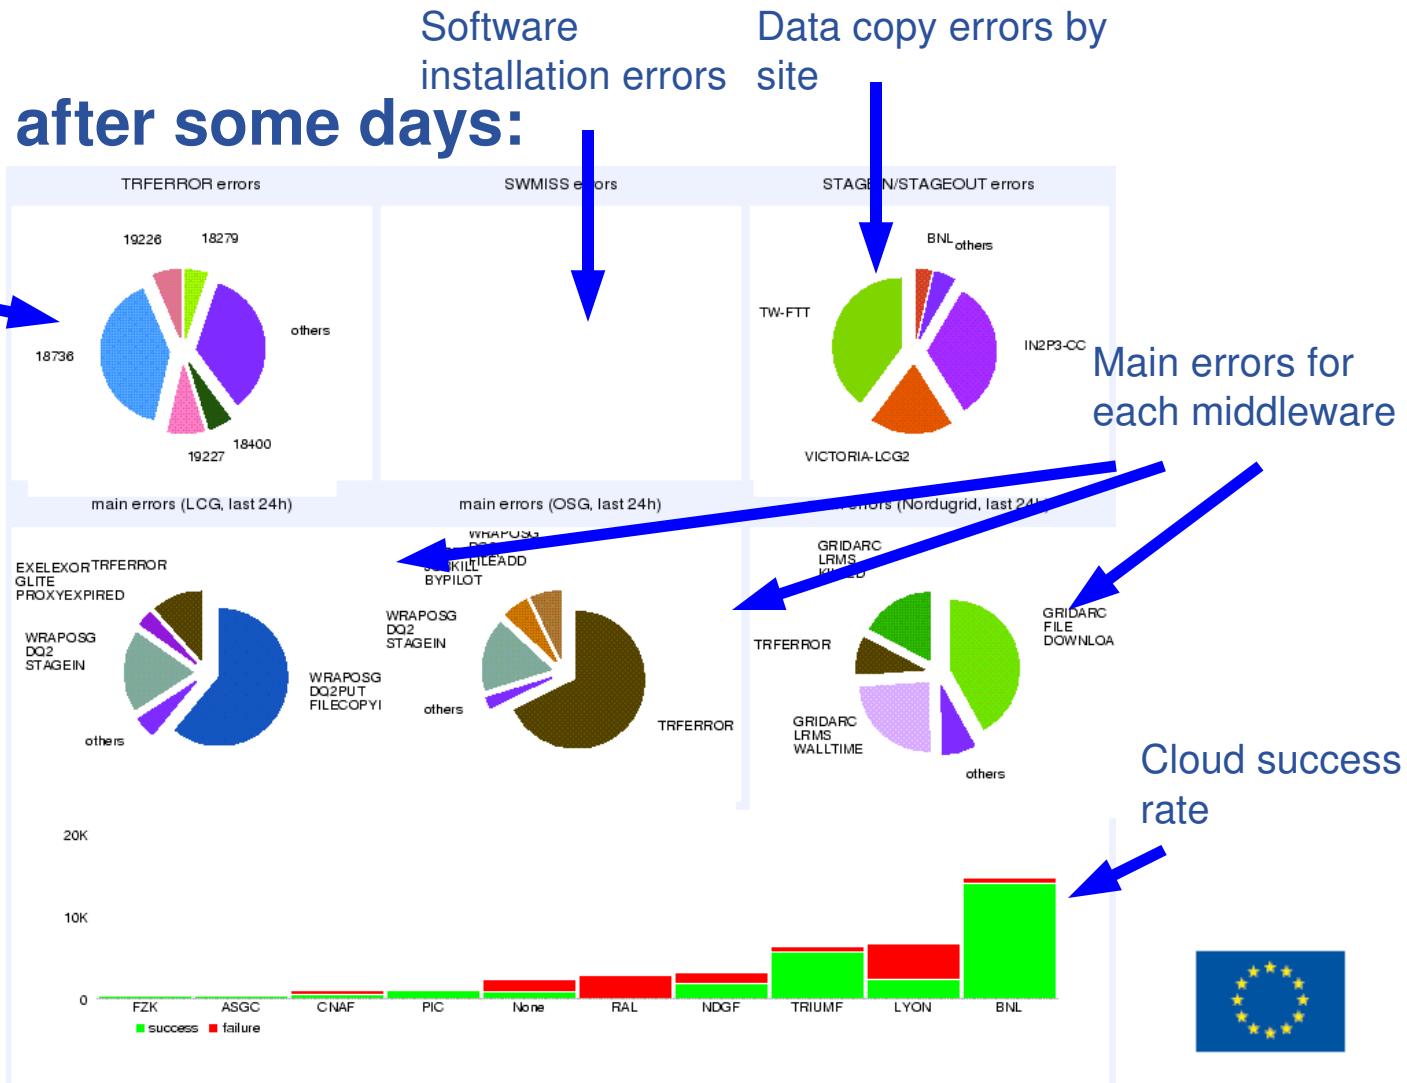


- **Official switch to the dashboard (jan 2008):**
 - Big warning in the former tool,
 - Tutorial slides based on the dashboard,
 - Shifter's workbook referring to the dashboard,
 - New feedback from shifters (non experts),
- **The developer did a shift (using the dashboard),**
 - Suggested by the developer of the former ATLAS tool,
 - B. Gaidioz did two days of shifts following the workbook, using the dashboard.

- **Outcome after two days of shifts:**
 - Three clicks each time is too much!
 - Lots of cut & paste from dashboard to GGUS, etc.
 - One needs **an immediate display matching the shifter's workbook:**
 - “for error X, submit a savannah ticket, write into it this and that”,
 - “for error Y, submit a GGUS ticket with these details”,
 - “look at clouds (tier1+ its tier2) with high failures and do that”.
- **All the information was in the dashboard**
 - But the interface is practical enough.

- Prototype after some days:

ATLAS software errors by task



- **Feedback for the shifter's display:**
 - Very positive feedback from everyone.
- **Other improvements to come:**
 - Simple ones (but were never expressed before):
 - Simple information missing (need other tool or three more clicks),
 - Need links to some tools (logbook, etc.) with URL params ready,
 - Need a pre-filled template for GGUS ticket / savannah bug,
 - More sophisticated improvements:
 - Error codes are not always sufficient, need to check some correlations before taking actions



Enabling Grids for E-sciencE

Summary & conclusion

www.eu-egee.org



- ## Summary

- Job monitoring installed in VL-eMed,
- User task monitoring implemented with job monitoring,
- The dashboard framework getting stronger and integrated with Lemon, Nagios, SAM.
- A specific monitoring application for ATLAS production
 - Now an operation tool.

• Conclusion

- Requirements came naturally after two days of operation,
 - Dashboard: “monitoring the grid from the VO point of view”,
 - (Obviously, running the VO systems teaches you a lot).
- Need flexibility in web interface:
 - Because we will need very often custom displays,
 - XSL web pages are hard to design as “components”
- Push for using the python APIs:
 - Permits to get very specific information in zero click,
 - Perhaps integrate it in the VO system.

- **More info:**

- <http://dashboard.cern.ch>
- Experience with the VL-eMed installation (poster: the experiment dashboard for medical applications).