

Monitoring the Flow

Lu Feng

The Pennsylvania State University

Ricardo Brito Da Rocha

Dr. Benjamin Gaidioz

CERN

Outline

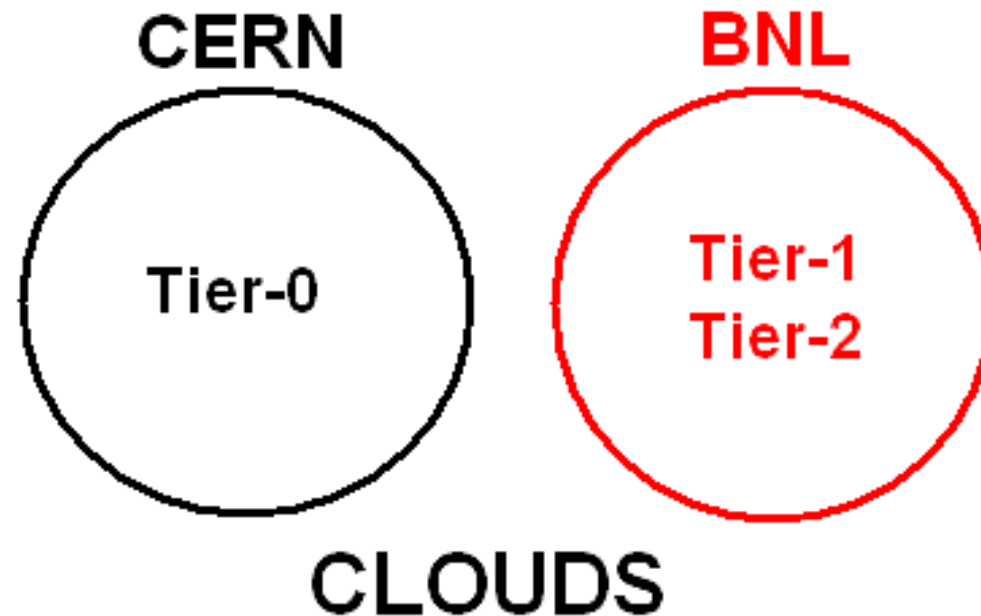
- The ATLAS Dashboard
- Projects
 - RSS View
 - Site Monitoring for Data Management
 - Administrative Monitoring for Production System
- Conclusion

The ATLAS Dashboard

- ARDA Group
- Enabling Grids for E-Science (EGEE)
- ATLAS Distributed Computing Monitoring
 - Data Management, e.g. data transfers
 - Production System, e.g. Grid jobs
 - Tier0 Export and Processing
 - Analysis Jobs
 - Panda Production

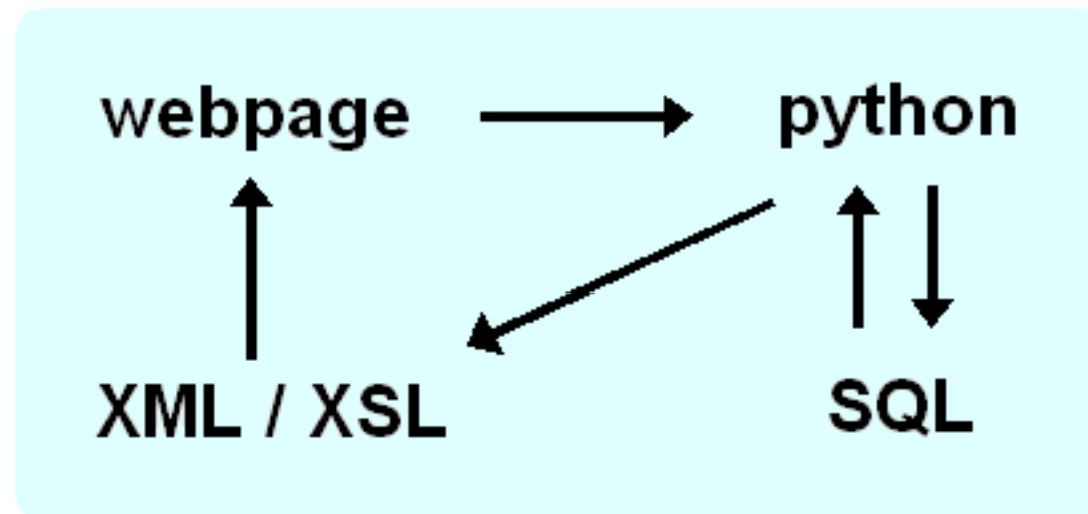
The ATLAS Dashboard

- Sites: Tier-0, Tier-1, Tier-2



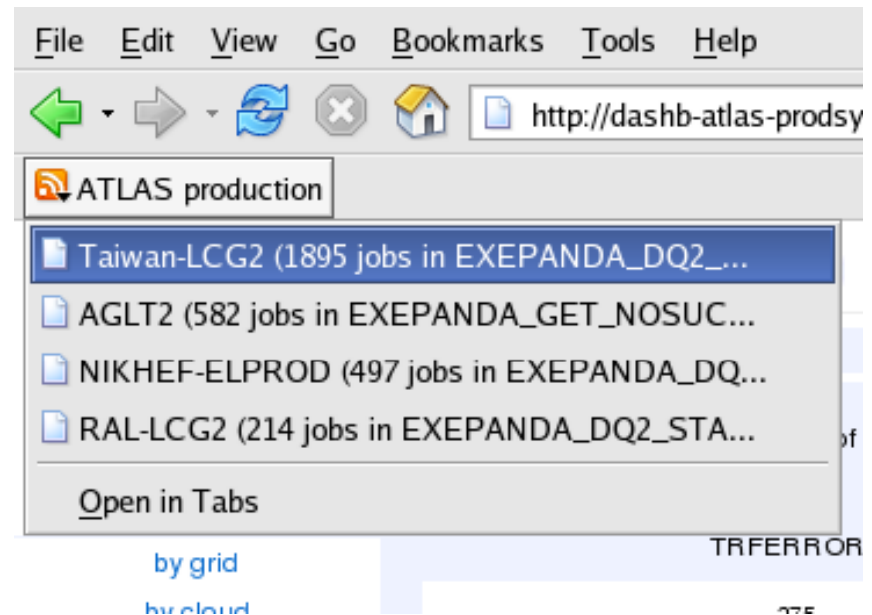
The Programming

Python Action \longrightarrow Database



RSS View

- generic output format
- RSS Feed
- implemented for Shifters of Production System
- documentation



- Folders
- bgaidioz@cern.ch
 - inbox (189)
 - Drafts
 - Sent
 - Trash (164)
 - ATLAS DA (267)
 - Cern Spam (5)
 - dashboard_notifications (90)
 - dashboard_stuff
 - ganga-dev (12)
 - grooveberries
 - operations
 - prod-log
 - shifts (163)
 - wasabi
 - RSS
 - Trash
 - ATLAS prod_sites (3)
 - bgaidioz@gmail.com
 - inbox (17)
 - Drafts
 - Sent
 - Trash (4)
 - (Gmail)
 - All Mail (55)
 - Drafts
 - Sent Mail
 - Spam
 - Stared
 - Trash
 - Local Folders
 - Unsent
 - Drafts
 - Sent
 - Trash
 - gerhild
 - raquel
 - stuff (3)
 - wasabi
 - CERN news

View: All

Subject or Sender

| Subject | Sender | Date |
|---|-----------|----------|
| Taiwan-LCG2 (1895 jobs in EXEPANDA_DQ2_STAGEIN) | Dashboard | 10:55 AM |
| AGLT2 (342 jobs in EXEPANDA_GET_NOSUCHFILE) | Dashboard | 10:55 AM |
| NIKHEF-ELPROD (497 jobs in EXEPANDA_DQ2GET_INFILE) | Dashboard | 10:55 AM |
| RAL-LCG2 (217 jobs in EXEPANDA_DQ2_STAGEIN) | Dashboard | 10:55 AM |

Subject: NIKHEF-ELPROD (497 jobs in EXEPANDA_DQ2GET_INFILE)
 From: Dashboard <dashboard.support@cern.ch>
 Date: 10:55 AM
 Website: http://dashb-atlas-prod.sys.cern.ch/dashboard/request.py/errors-detail?groupingsite&site=NIKHEF-ELPROD&status=failure&error=EXEPANDA_DQ2_STAGEIN&...

ATLAS dashboard

Data: Task | Data Production | Jobs Production | Jobs Analysis | Panda Production

Tasks | Grid jobs | Summaries | **Errors** | User Guide | Feedback

OVERVIEW

view

- by grid
- by cloud
- by dest_cloud
- by executortype
- by executor
- by site
- by cluster
- by task

select site

NIKHEF-ELPROD

restrict to: EXEPANDA_DQ2GET_INFILE (497), EXEPANDA_GET_FAILEDTOGETLFCREPLICA (7), EXEPANDA_DQ2PUT_FILECOPYERROR (2).

| site | 1 | 2 | 3 | others |
|-------|---|---|---|--------|
| count | 1 | 7 | 2 | 0 |

6.4.1.4. Generic RSS View

The RSS view can be used to create RSS feeds for any page as long as it is properly configured.




The RSS view takes flagged objects from the *ActionContext*. To flag an object in *ActionContext*, add `flagged=True`:

```
actionCtx.put("exListofDictionaries", exListofDictionaries, flagged=True)
```

The following action mapping file will show how to make use of the RSS view:

```
<action path="example.rss" method="GET" type="dashboard.http.actions.production.exampleAction"
  description="a description of the action">
  <output format="*/*" contentType="application/rss+xml" type="dashboard.http.views.GenericRSSView"
    title="the title of the feed"
    link="http://the_main_website_hosting_the_feed"
    managingEditor="your.name@cern.ch (Your Name)"
    ttl="120"
    description="a description of the feed"
    itemTitle="#field2 (#field1)"
    itemLink="http://link_to_each_feed_item/for_example_#field1"
    itemDescription="field2 contains the value #field2, and similarly for field1 (#field1)"
    itemAuthor="your.name@cern.ch (Your Name)"
  />
</action>
```

The above action mapping file contains all of the RSS elements accepted by the view. Note that the elements with "item" in the naming is an RSS *item* element, e.g. `itemTitle`.

 `ttl` (time to live) is the element specifying the number of minutes before the feed refreshes. Some browsers (e.g. Firefox) might have a refresh time (e.g. 60 minutes) that would overwrite your specification if it is short (e.g. 3 minutes).

You will notice that in the above example, there are keys in the *item* elements with a special format "#key_name." This is where you can specify the items from the *ActionContext*. Here, the example makes use of the XML file from the Generic XML View section. What happens is that the any occurrence of `#field2` will be replaced by `elem1-value2` for the first feed item, and by `elem2-value2` for the second feed item. Similarly, `#field1` will be replaced by `elem1-value1` and `elem2-value1` for respective feed items. If, instead of `#field2` or `#field1`, you have `#field3`, which does not exist in the *ActionContext*, `#field3` will appear as it is in the feed.

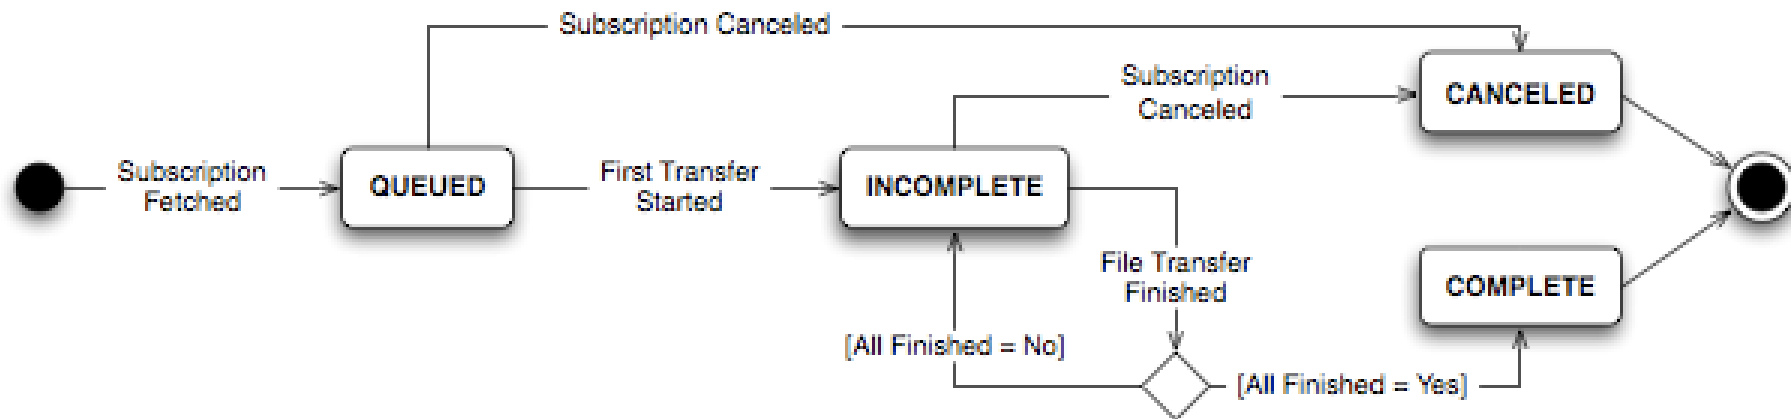
The special keys are *only* available for *item* elements.

To add the feed onto the browser, add the following to your web XSL template:

```
<link rel="alternate" type="application/rss+xml" title="Title of the Feed" href="/dashboard/request.py/example.rss?otherparameters=otherparameters/">
```


Site Monitoring

- Distributed Data Management (DDM) Monitoring



Site Monitoring

- additional site statistics

| Clouds | Transfers | | | Registrations | | Movement | | | | | | Errors | | Services | |
|--------|-----------|---------|---------|---------------|------------|-------------|--------------|-------------|---------------|--------------------|--------------------|---------|----------|----------|------|
| | Eff. | Rate | Success | Sets Done | Files Done | Sets Queued | Sets Transf. | Queue Stuck | Transf. Stuck | Avg. Comp. Time(s) | Avg. MvSt. Time(s) | Transf. | Register | Services | Grid |
| ASGC | 91% | 14 MB/s | 867 | 652 | 1738 | 0 | 0 | 0 | 0 | 951 | 0 | 87 | 0 | 0 | |
| BNL | 85% | 21 MB/s | 1645 | 797 | 14228 | 0 | 0 | 0 | 0 | 5435 | 0 | 282 | 0 | 0 | |

| Movement | | | | | |
|-------------|--------------|-------------|---------------|--------------------|--------------------|
| Sets Queued | Sets Transf. | Queue Stuck | Transf. Stuck | Avg. Comp. Time(s) | Avg. MvSt. Time(s) |
| 0 | 0 | 0 | 0 | 951 | 0 |
| 0 | 0 | 0 | 0 | 5435 | 0 |

- new site statistics page

TAIWAN-LCG2_DATADISK

[Dataset List](#)

[Site Statistics](#)

[Most Recent Events](#)

[Most Recent Errors](#)

[All Cloud Sites](#)

[ASGCDISK](#)

[ASGCDISK_V2](#)

[ASGCTAPE](#)

[ASGCTAPE_V2](#)

[TAIWAN-LCG2_DATADISK](#)

[TAIWAN-LCG2_DATATAPE](#)

[TAIWAN-LCG2_MCDISK](#)

[TAIWAN-LCG2_MCTAPE](#)

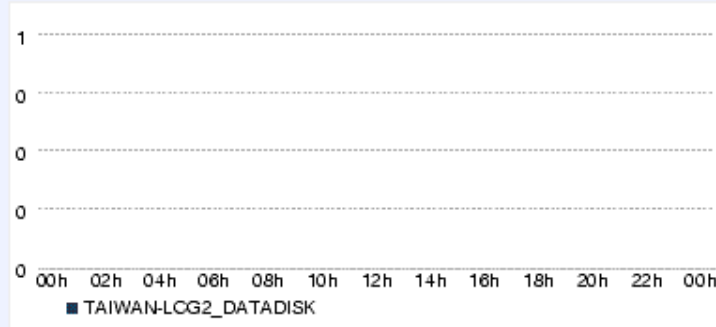
[AU-ATLAS](#)

[TW-FTT](#)

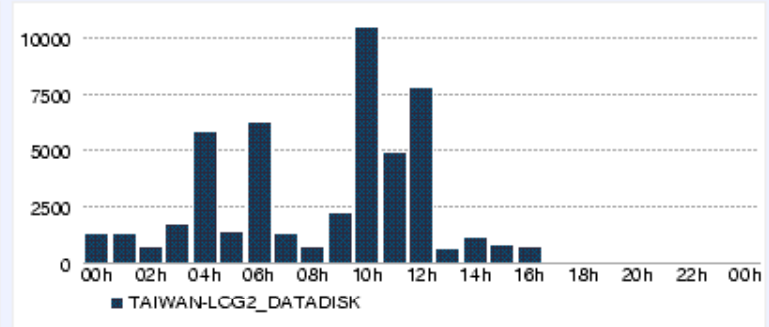
[TW-FTT_DATADISK](#)

[TW-FTT_MCDISK](#)


Testing Queued




Testing Average Completion Time




Testing In-Transfer




Testing Average Move Start Time




Testing Frozen-Transfer



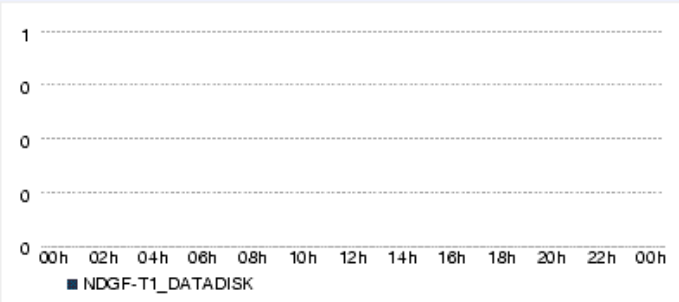
Testing Frozen-Queue



Datasets Stuck In Transfer



Datasets Stuck In Queue



Destination Site Activity Summary ('2008-07-17 00:00' to '2008-07-18 00:00')

Click on the site name to go to the site page.

| Destination Site | Transfers | | | Registrations | | Movement | | | | | | Errors | | Services | |
|----------------------------------|-----------|---------|---------|---------------|------------|-------------|--------------|-------------|----------------|--------------------|--------------------|---------|----------|----------|------|
| | Eff. | Rate | Success | Sets Done | Files Done | Sets Queued | Sets Transf. | Queue Stuck | Transfer Stuck | Avg. Comp. Time(s) | Avg. MvSt. Time(s) | Transf. | Register | Services | Grid |
| NDGF-T1_DATADISK | 73% | 13 MB/s | 802 | 639 | 1686 | 0 | 0 | 0 | 0 | 1096 | 0 | 301 | 0 | 0 | ok |

Source Sites Activity Summary

| Source Sites | Transfers | | | Errors | | |
|------------------------|------------|------------|-----------|----------|--------------|----------|
| | Efficiency | Throughput | Successes | Transfer | Registration | Services |
| BNL-OSG2_DATADISK | 100% | 4 MB/s | 135 | 0 | 0 | |
| CERN-PROD_DATATAPE | 100% | 2 MB/s | 396 | 0 | 0 | |
| FZK-LCG2_DATADISK | 68% | 1 MB/s | 39 | 18 | 0 | |
| IN2P3-CC_DATADISK | 21% | 2 MB/s | 61 | 223 | 0 | |
| INFN-T1_DATADISK | 100% | 1 MB/s | 27 | 0 | 0 | |
| NIKHEF-ELPROD_DATADISK | 73% | 2 MB/s | 57 | 21 | 0 | |
| PIC_DATADISK | 100% | 0 MB/s | 12 | 0 | 0 | |
| RAL-LCG2_DATADISK | 50% | 1 MB/s | 39 | 39 | 0 | |
| TAIWAN-LCG2_DATADISK | 100% | 0 MB/s | 18 | 0 | 0 | |
| TRIUMF-LCG2_DATADISK | 100% | 1 MB/s | 18 | 0 | 0 | |

CRITICAL
WARNING
NORMAL
GOOD
NO_ACTIVITY
SCHED DOWNTIME

Administrative Monitoring

- Production System Monitoring
- new administrator page
 - search function

Admin

Please enter your site or cloud name:

LYON

Search

Administrative Monitoring

- Production System Monitoring
- new administrator page
 - search function
 - link to statistics summary page

[Admin](#)

[Click here for the cloud statistics of the last 30 days](#)

Administrative Monitoring

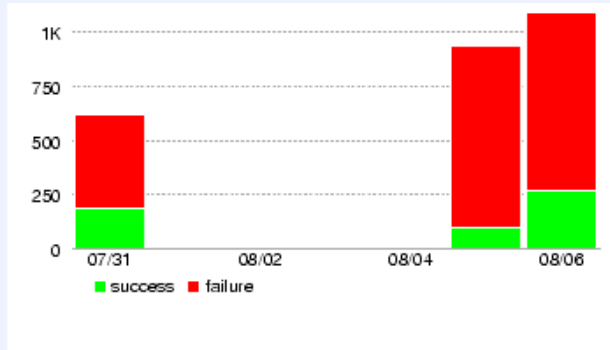
- Production System Monitoring
- new administrator page
 - search function
 - link to statistics summary page
 - additional plots



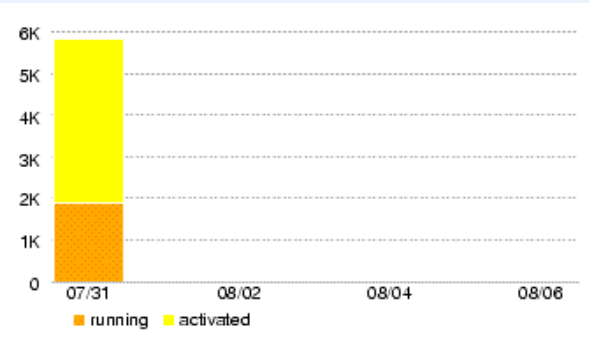
Admin

[Click here for the cloud statistics of the last 30 days](#)

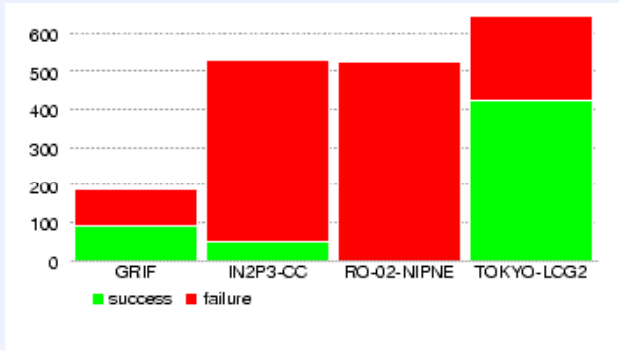
jobs last seven days



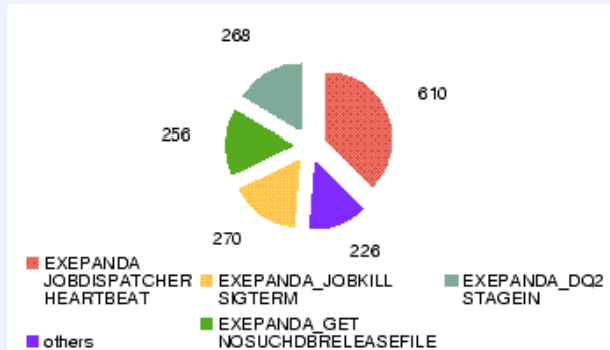
queued jobs last seven days



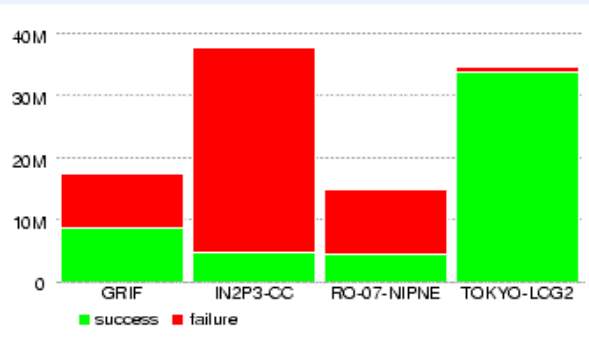
jobs (last 24h)



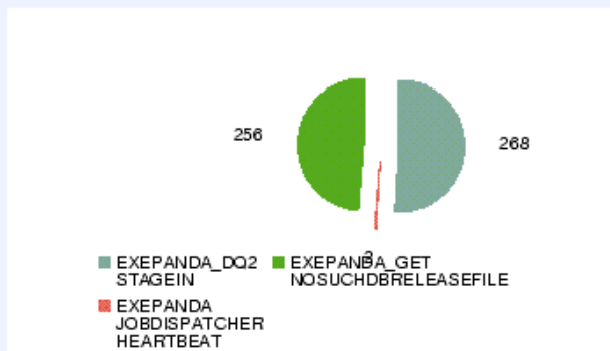
errors (jobs, last 24h)




walltime (seconds, last 24h)



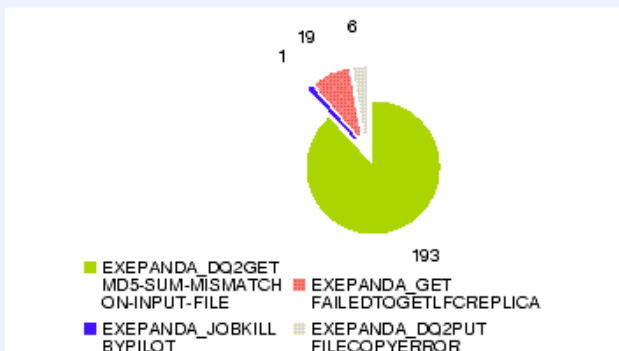
errors in RO-02-NIPNE (jobs, last 24h)



errors in IN2P3-CC (jobs, last 24h)



errors in TOKYO-LCG2 (jobs, last 24h)



Administrative Monitoring

- Production System Monitoring
- new administrator page
 - search function
 - link to statistics summary page
 - additional plots
 - embedded error messages: JavaScript

A blue rounded rectangular button with the text "Admin" in a light blue font, positioned to the right of the list items.

| | | | | | | | | |
|---|---|---|---|---|---|------|------|-------|
| UKI-LT2-RHUL + | 0 | 0 | 0 | 0 | 1 | 1413 | 1542 | 47.8% |
| UKI-LT2-QMUL + | 0 | 0 | 0 | 0 | 0 | 819 | 2250 | 26.7% |
| Get error: (...) (...)get(...)replica(...)) Log put error: Copy command returned error code (...) and output: /opt/grid/glite/3.1.13/lcg/bin/lcg-cr lcg_util-1.6.11 GFAL(...): (283) Get error: (...)l(...)d (...)t(...) get LFC replica(...): -1 (lfc_getreplica(...) failed for guid: (...): (242) (...)ut error: Copy command returned error code 256 and output: /opt/grid/glite/3.1.13/lcg/bin/lcg-cr lcg_util-1.6.11 GFAL-client-1.10.11 send2nsd: NS002 - (...): (39) (...)ut error: Copy command returned error code 256 and output: /opt/grid/glite/3.1.13/lcg/bin/lcg-cr lcg_util-1.6.11 GFAL-client-1.10.11 (...) 0 bytes 0.00 KB/sec(...): (4) | | | | | | | | |
| UKI-NORTHGRID-MAN-HEP + | 0 | 0 | 0 | 0 | 1 | 0 | 14 | 0% |

| | | | | | | | | |
|-------------------------------|----|-----|-----|---|---|----|----|-------|
| BEIJING-LCG2 + | 21 | 0 | 0 | 0 | 0 | 0 | 0 | - |
| No error messages to display. | | | | | | | | |
| RO-07-NIPNE + | 0 | 426 | 143 | 0 | 1 | 16 | 41 | 28.1% |

| | |
|-------------------------------|----|
| BEIJING-LCG2 + | 21 |
| No error messages to display. | |

| | |
|--|---|
| UKI-LT2-QMUL + | 0 |
| Get error: (...) (...)get(...)replica(...)) Log put error: Copy comm Get error: (...)l(...)d (...)t(...) get LFC replica(...): -1 (lfc_getrepli (...)ut error: Copy command returned error code 256 and outq (...)ut error: Copy command returned error code 256 and outq | |

Conclusion

- The ATLAS Dashboard – monitoring
- Projects – added new features to existing web interface to improve monitoring ability

Acknowledgements

- Benjamin Gaidioz, Ricardo Brito Da Rocha
- Myron Campbell, Steve Goldfarb, Jeremy Herr, Jean Krisch, Homer Neal
- CERN, Ford, NSF, University of Michigan

Photos

