

# autopyfactory

Graeme Stewart  
2010-12-02

# Updates

- Feersum Endjinn (June):
  - Some bug fixes (passing site instead of siteid to pilot)
  - Better support for ANALY queue configuration
  - Suppress 'broken' queues
- Time out of Joint (October):
  - More minor bug fixes and improvements
  - Condor > 7.5 compatible
  - CREAM CE support returned

# Many potential contributions...

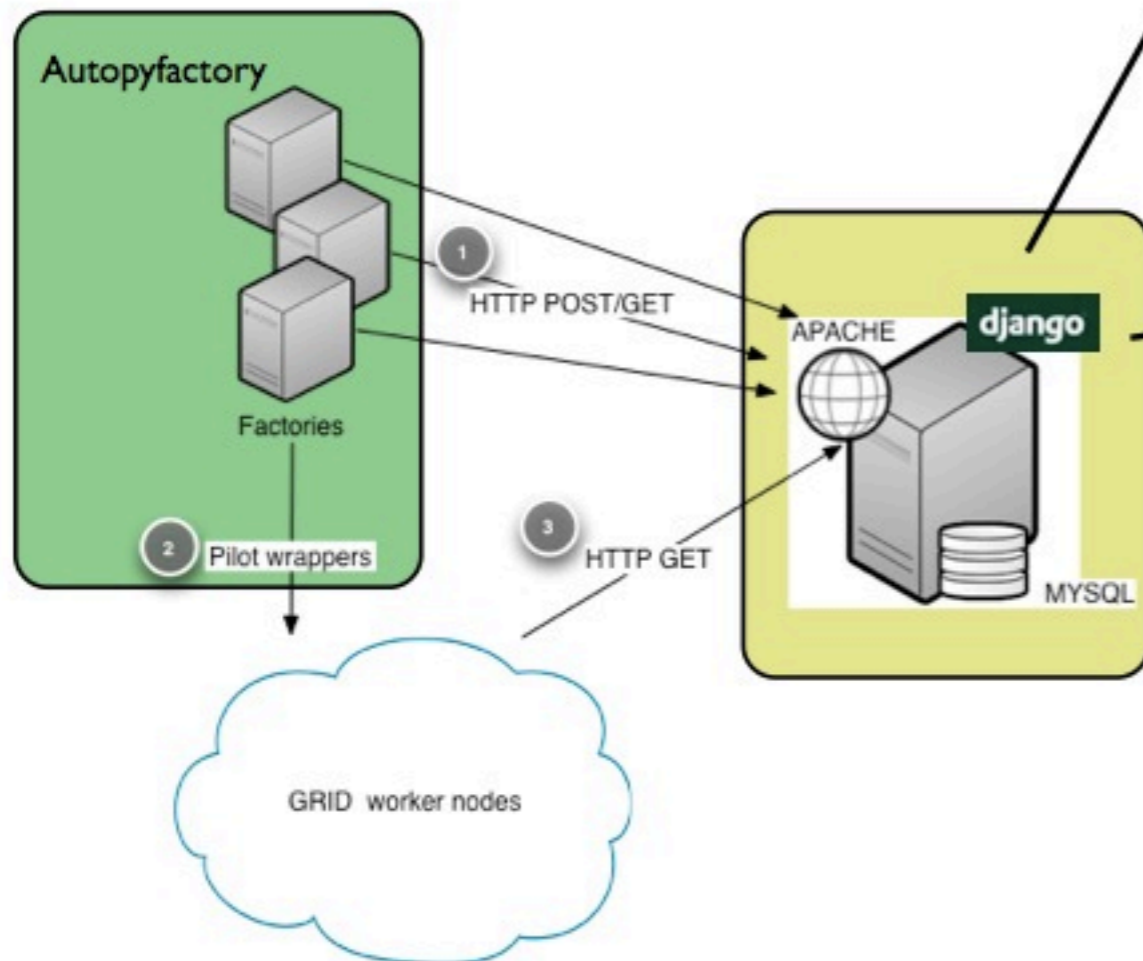
- VOMS renewals managed by factory (not cron)
- Wrapper slimming, submitting only a very thin script, then bootstrapping
- APF monitoring
- Support for local batch queues

# Monitoring

- Of these, monitoring is in advanced beta testing, thanks to Peter Love
- Well developed and in 'current' branch
- Next slides stolen from Peter's ADC Dev talk: <http://indico.cern.ch/conferenceDisplay.py?confId=112990>

# Architecture

Django app running on standard LAMP stack, minimal changes to autopyfactory code



- Post submitted jobs to monitor in main loop
- Condor monitoring crons on factory machine
- Hooks in pilot wrapper to post start/finish events

# MySQL tables

```
mysql> describe mon_factory;
```

Field	Type	Null	Key	Default	Extra
id	int(11)	NO	PRI	NULL	auto_increment
name	varchar(64)	NO		NULL	
ip	char(15)	NO		NULL	
url	varchar(200)	NO		NULL	

```
4 rows in set (0.04 sec) mysql> describe mon_job;
```

Field	Type	Null	Key	Default	Extra
id	int(11)	NO	PRI	NULL	auto_increment
created	datetime	NO		NULL	
cid	varchar(16)	NO	MUL	NULL	
fid_id	int(11)	NO	MUL	NULL	
last_modified	datetime	NO		NULL	
state_id	int(11)	NO	MUL	NULL	
pandaq_id	int(11)	NO	MUL	NULL	
label_id	int(11)	NO	MUL	NULL	
result	smallint(6)	NO		NULL	
flag	tinyint(1)	NO		NULL	

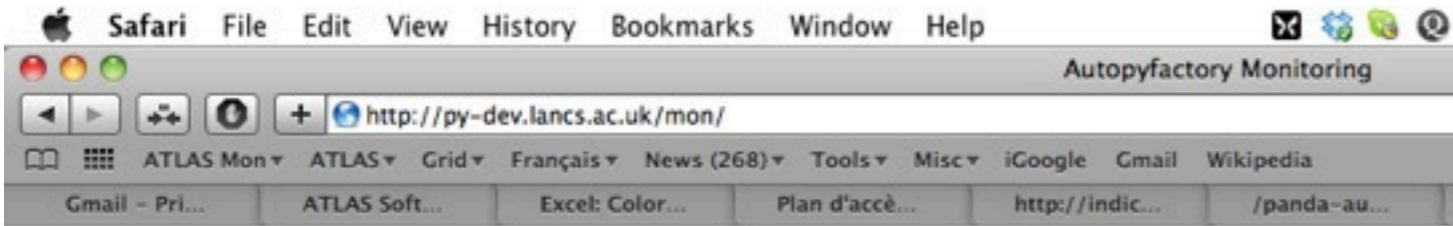
```
10 rows in set (0.00 sec) mysql> describe mon_label;
```

Field	Type	Null	Key	Default	Extra
id	int(11)	NO	PRI	NULL	auto_increment
name	varchar(64)	NO		NULL	
fid_id	int(11)	NO	MUL	NULL	
pandaq_id	int(11)	NO	MUL	NULL	
msg	varchar(140)	NO		NULL	
last_modified	datetime	NO		NULL	

```
6 rows in set (0.00 sec)
```

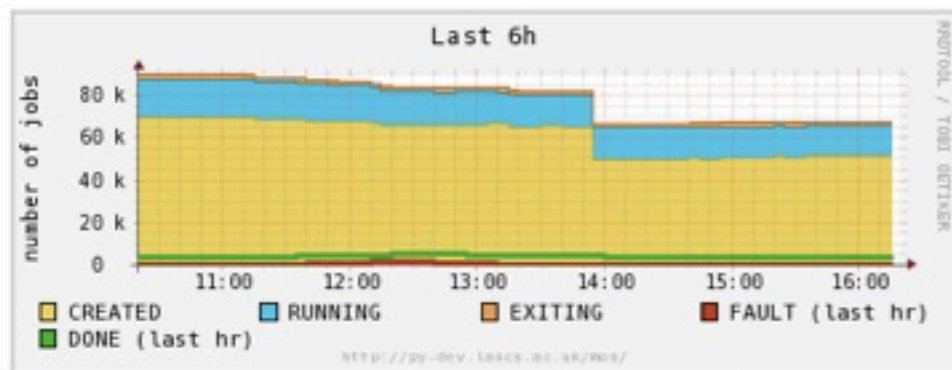
+ others modelling the site topology

# In the box...



[Home](#) | [Admin](#)

## Global view - all factories



Factory	created	running	exiting	done (last hour)	fault (last hour)
<a href="#">peter-UK-devel</a>	25	36	0	3	0
<a href="#">peter-voatlas60</a>	29634	3608	676	808	0
<a href="#">graeme-UK-apf</a>	18849	9916	1314	2762	0

## Notes

- CREATED: condor\_submit executed, condor\_id returned
- RUNNING: pilot wrapper has started on Worker Node
- EXITING: pilot wrapper is finishing up on Worker Node
- DONE: condor jobstate is Completed (or Removed)
- FAULT: condor jobstate indicates a fault
- Webservice apache [usage](#)

- Overview of factory activity

# Drill down

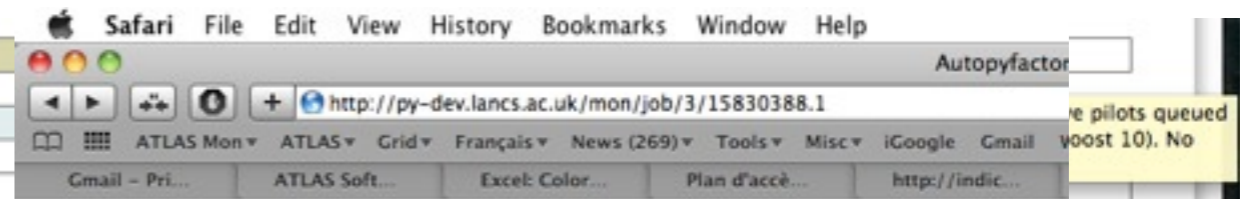
<a href="#">ANALY_GLASGOW-svr26-atlanaly-cream</a>	1697	286	243	0	0	2 minutes ago
<a href="#">ANALY_LANCS-ce44</a>	2	87	32	129	0	749 activated jobs, 200 inactive pilots queued (>= queue depth 20 * depth boost 10). No extra pilots needed. 2 minutes ago
<a href="#">ANALY_LIV</a>	44	94	20	40	0	2 minutes ago

- Highlight queues in trouble
- Drill down to logs and status



# Drill down

<a href="#">ANALY_GLASGOW-svr26-atlanaly-cream</a>	1697	286	243
<a href="#">ANALY_LANCS-ce44</a>	2	87	32
<a href="#">ANALY_LIV</a>	44	94	20



Home | Admin

```

FACTORY      graeme-UK-apf
CONDOR_ID    15830388.1 (atdout, atderr, atdlog, only last seven days)
QUEUE        ANALY_GLASGOW
CREATED      2010-12-01 17:16:21
MODIFIED     2010-12-01 17:16:21
STATE        CREATED
RESULT       0
FLAG         0
    
```

2010-12-01 17:16:21 130.209.239.17 CREATED


globusstate (gs)	Reason
1	PENDING
2	ACTIVE
4	FAILED
8	DONE
16	SUSPENDED
32	UNSUBMITTED
64	STAGE_IN
128	STAGE_OUT

JobState (js)	Reason
0	Unexpanded
1	Idle
2	Running
3	Removed
4	Completed
5	Held
6	Submission_err

- Highlight queues in trouble
- Drill down to logs and status

# Drill down

<a href="#">ANALY_GLASGOW-svr26-atlanaly-cream</a>	<a href="#">1697</a>	<a href="#">286</a>	<a href="#">243</a>
<a href="#">ANALY_LANCS-ce44</a>	<a href="#">2</a>	<a href="#">87</a>	<a href="#">32</a>
<a href="#">ANALY_LIV</a>	<a href="#">44</a>	<a href="#">94</a>	<a href="#">20</a>



```
Home | Admin

FACTORY      graeme-UK-apf
CONDOR_ID    15830388.1 (atdout, atderr, atdlog, only last seven days)
QUEUE        ANALY_GLASGOW
CREATED      2010-12-01 17:16:21
MODIFIED     2010-12-01 17:16:21
STATE        CREATED
RESULT       0
FLAG         0

2010-12-01 17:16:21 130.209.239.17 CREATED
```

globusstate (gs)	Reason
1	PENDING
2	ACTIVE
4	FAILED
8	DONE
16	SUSPENDED
32	UNSUBMITTED
64	STAGE_IN
128	STAGE_OUT

JobState (js)	Reason
0	Unexpanded
1	Idle
2	
3	
4	
5	
6	



```
000 (15830366.016.000) 12/01 17:11:55 Job submitted from host: <130.209.239.17:38172>
...
012 (15830366.016.000) 12/01 17:12:04 Job was held.
Failed to create proxy delegation
Code 0 Subcode 0
...
```

- Highlight queues in trouble
- Drill down to logs and status

# Queue View

Safari File Edit View History Bookmarks Window Help

Autopyfactory Monitoring

http://py-dev.lancs.ac.uk/mon/q/31/

ATLAS Mon ATLAS Grid Français News (269) Tools Misc iGoogle Gmail Wikipedia

Gmail - Pri... ATLAS Soft... Excel: Color... Plan d'accè... http://indic... /panda-au... Top Sites Autopyfact... BBC News ... ADC Devel...

[Home](#) | [Admin](#)

## Panda Queue view - ANALY\_GLASGOW

[panda](#) | [admin](#)

Factory label	Factory	created	running	exiting	done (last hour)	fault (last hour)	Latest action
<a href="#">ANALY_GLASGOW-svr014-cream</a>	<a href="#">graeme-UK-apf</a>	<a href="#">59</a>	<a href="#">764</a>	<a href="#">526</a>	<a href="#">383</a>	<a href="#">0</a>	1 minute ago
703 activated jobs, 20 inactive pilots queued (< queue depth 20 * depth boost 10). Will submit full pilot load.							
<a href="#">ANALY_GLASGOW-svr21-atlanaly</a>	<a href="#">graeme-UK-apf</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	1 minute ago
Site ANALY_GLASGOW is offline - will not submit pilots.							
<a href="#">ANALY_GLASGOW-svr26-atlanaly</a>	<a href="#">graeme-UK-apf</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	1 week ago
Site ANALY_GLASGOW is offline - will not submit pilots.							
<a href="#">ANALY_GLASGOW-svr26-atlanaly-cream</a>	<a href="#">graeme-UK-apf</a>	<a href="#">1697</a>	<a href="#">286</a>	<a href="#">243</a>	<a href="#">0</a>	<a href="#">0</a>	1 minute ago
703 activated jobs, 200 inactive pilots queued (>= queue depth 20 * depth boost 10). No extra pilots needed.							
<a href="#">ANALY_GLASGOW_SVR021</a>	<a href="#">peter-voatlas60</a>	<a href="#">2700</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	7 hours, 53 minutes ago
3 activated jobs, 40 inactive pilots queued (>= queue depth 20 * depth boost 10). No extra pilots needed.							
<a href="#">ANALY_GLASGOW_SVR026</a>	<a href="#">peter-voatlas60</a>	<a href="#">3063</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	<a href="#">0</a>	7 hours, 53 minutes ago
3 activated jobs, 20 inactive pilots queued (>= queue depth 20 * depth boost 10). No extra pilots needed.							

Factory	Job	State	Result	Created	Last modified
graeme-UK-apf	<a href="#">15830379.8</a>	DONE	0	26 minutes	0 minutes ago
graeme-UK-apf	<a href="#">15830426.0</a>	DONE	0	17 minutes	0 minutes ago
graeme-UK-apf	<a href="#">15830379.4</a>	DONE	0	26 minutes	0 minutes ago
graeme-UK-apf	<a href="#">15830414.9</a>	DONE	0	19 minutes	0 minutes ago
graeme-UK-apf	<a href="#">15829946.0</a>	DONE	0	1 hour, 54 minutes	0 minutes ago
graeme-UK-apf	<a href="#">15830426.1</a>	DONE	0	17 minutes	0 minutes ago

# Performance

- Peter reported some server scaling problems 3 weeks ago
- Now much improved, even though still running on underpowered VM
- Ready to add further factories
- Still problems with hovering up stale jobs (awol, held), possibly we can share the condor\_q query with the main factory?

# Monitor: Next Steps

- Make a release
- Get a production APF monitor at CERN
  - Independent and probably better that way...
- Cross link from panda monitor
- Alarm page - concentrate on broken sites

# APF: Prioritised Work List?

- High: Local batch support - seems like its needed by the T3s, not so much work
- High: Deliver and improve monitoring
- Medium: More setup 'out of the box'
- Medium: Thin pilot wrapper
- Low: Factory foreman (bigger project)