

CMS Glidein Frontend Operations

JAMES LETTS
UCSD

Architecture

- CRAB Server: `glidein-2.t2.ucsd.edu`
 - Backup server: `submit-2.t2.ucsd.edu`
 - schedd with all CRAB user jobs
- Frontend: `glidein-collector.t2.ucsd.edu`
with `userid frontend`
 - Configuration file: `frontend.xml`

Operations

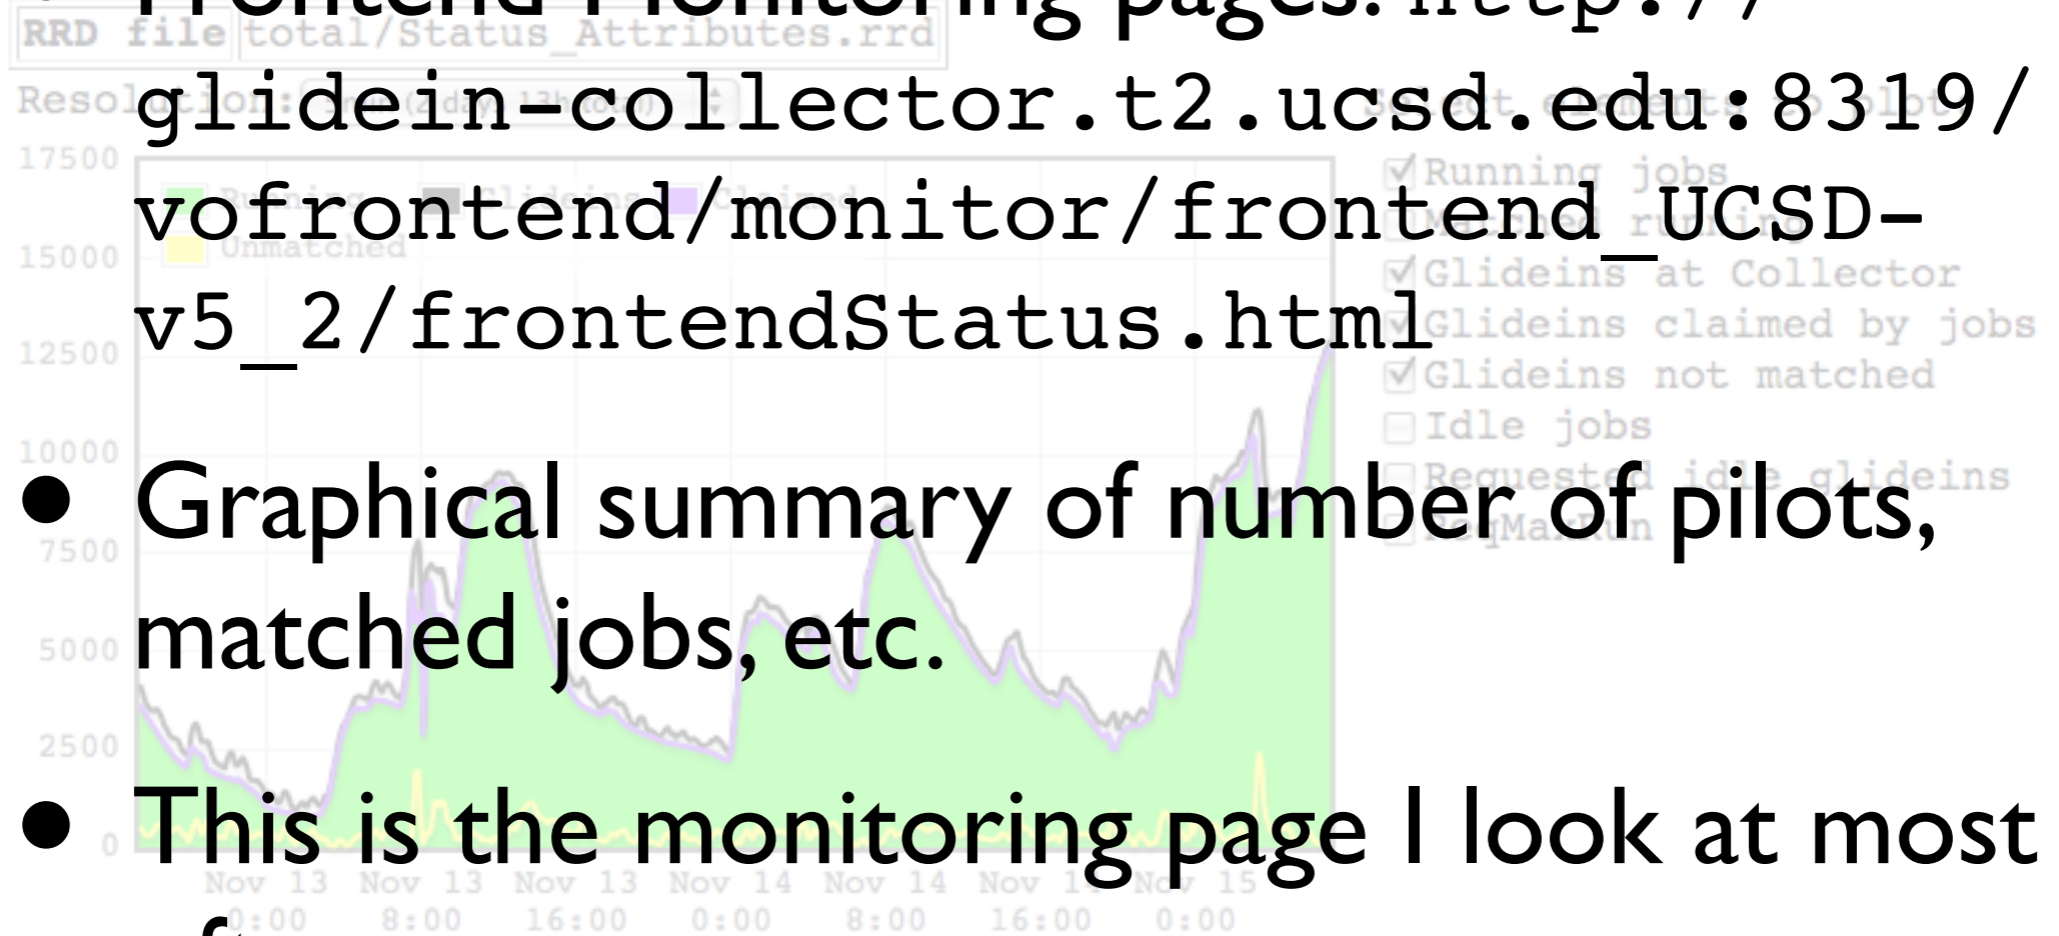
- Day-to-day operations usually trouble-free
 - Monitoring - keeping an eye on things
- Troubleshooting & Debugging Problems when they do arise
- Configuration changes as needed
- Development - validation

Monitoring

- Factory Monitoring pages
- Frontend Monitoring pages
- Frontend Logs

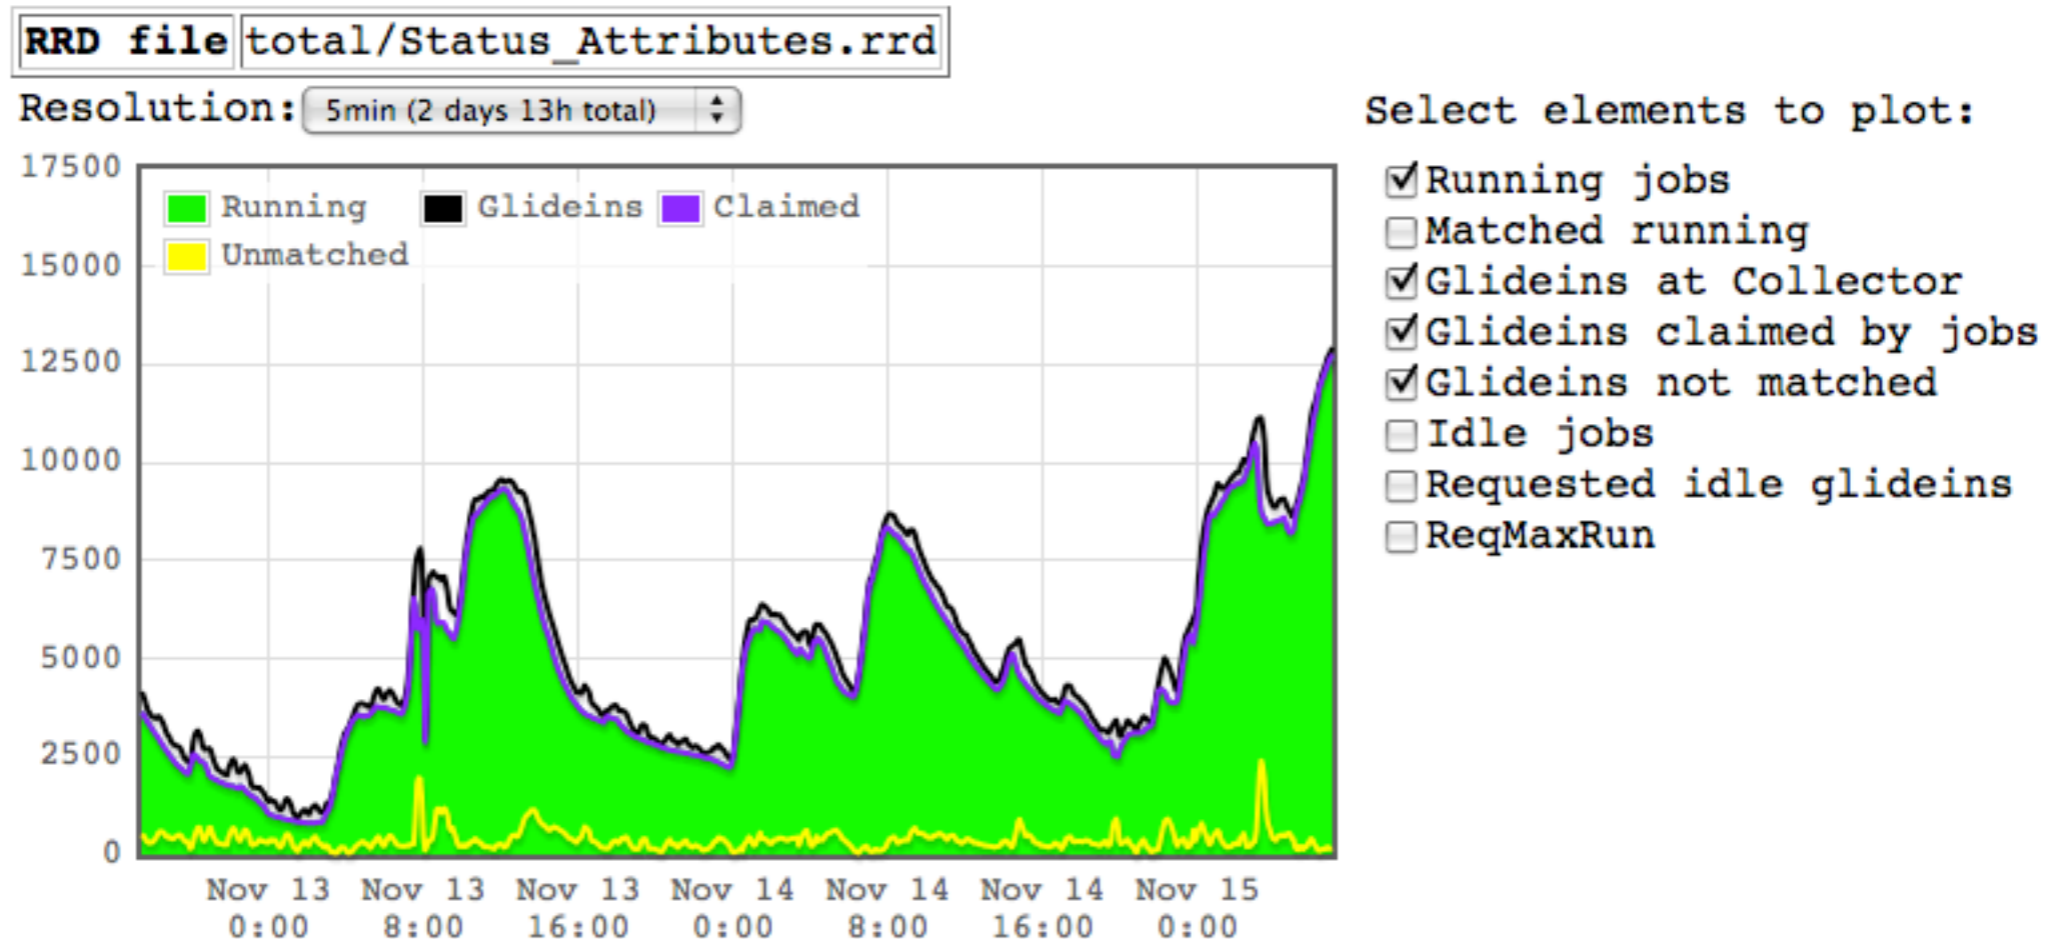
Monitoring

- Frontend Monitoring pages: http://glidein-collector.t2.ucsd.edu:8319/vofrontend/monitor/frontend_UCSD-v5_2/frontendStatus.html
- Graphical summary of number of pilots, matched jobs, etc.
- This is the monitoring page I look at most often.



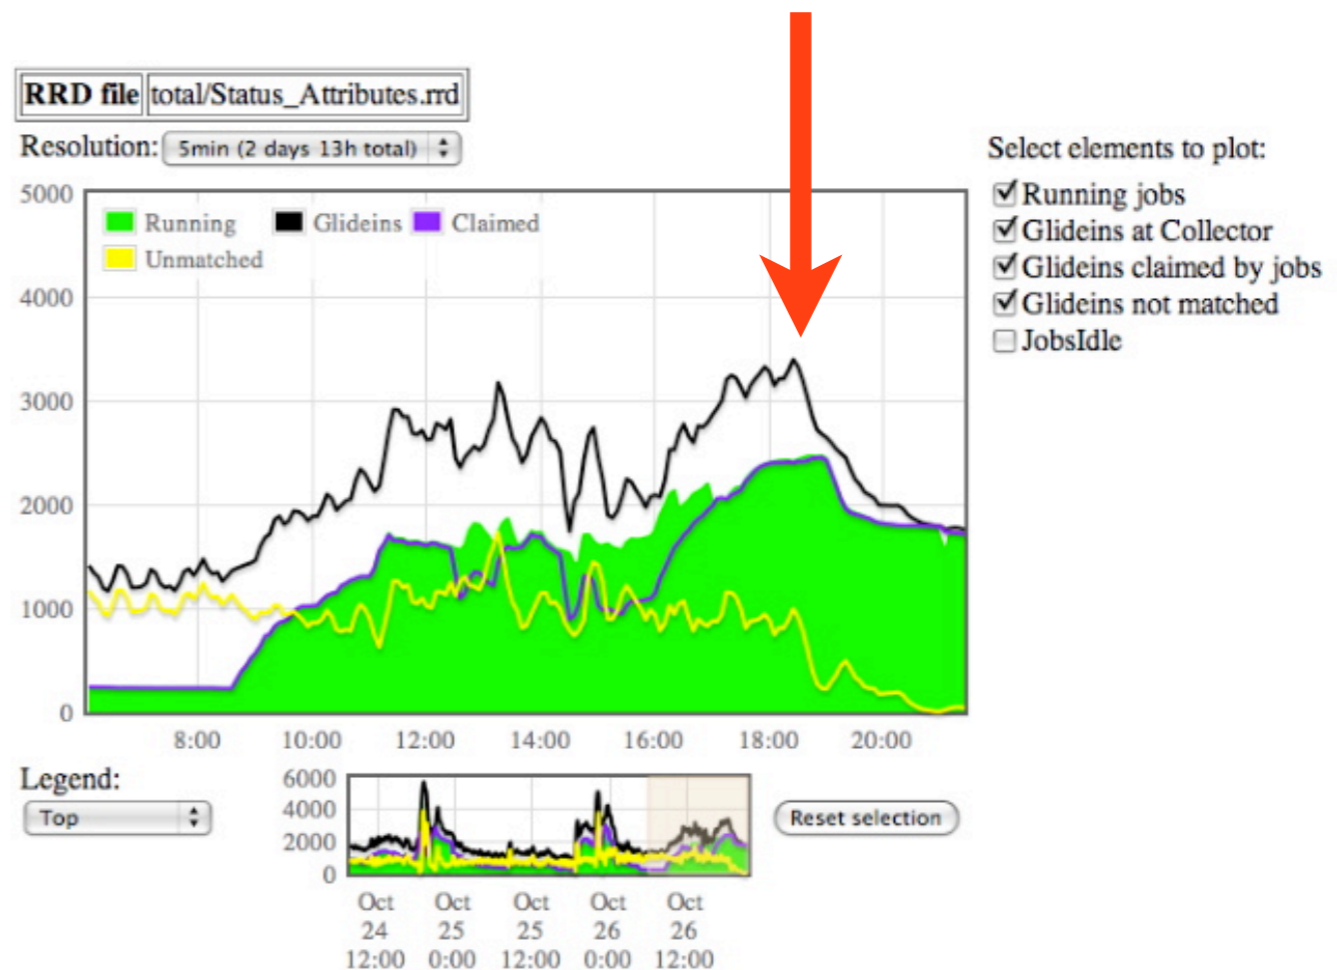
Monitoring

- Problems typically manifest themselves as many unmatched pilots, e.g.



Problem Day

- Many jobs failing the ImageSize < 2GB limit
- Symptom: many unmatched pilots
- Solution: remove the problem jobs from the queue ... they would never run anyway.




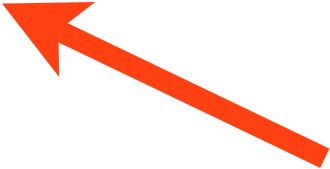
Factory Daily Analysis

From: gfactory@glidein-1.t2.ucsd.edu
Subject: [Cms-wms-support] UCSD Factory analyze_entries report for last 24h on 11-11-2011
Date: November 11, 2011 7:01:21 AM PST
To: UCSD glidein support
Reply-To: osg-gfactory-support@physics.ucsd.edu

Glidein log analysis for frontend frontend_UCSDCMS_cmnpilot - 11-11-2011_07:01:21

Past 24.0 hours:

Total Glideins: 5320
Total Jobs: 5306 (Average jobs/glidein: 1.00) 

Total time: 62.6M (17.4K hours - 725.0 slots)
Total time used: 53.7M (14.9K hours - 621.5 slots - 85%)
Total time validating: 378.6K (105.2 hours - 4.4 slots - 0%)
Total time idle: 8.5M (2.4K hours - 98.5 slots - 13%)
Total time wasted: 8.9M (2.5K hours - 103.5 slots - 14%)
Time used/time wasted: 6.0
Time efficiency: 0.86 

Monitoring

● Factory Monitoring pages: http://glidein-1.t2.ucsd.edu:8319/glidefactory/monitor/glidein_Production_v4_0/factoryStatusNow.html

XML last update: Tue Aug 30 11:36:40 2011

Entry Name	Requested:										Client Monitor:						
	Running	Idle	Waiting	Pending	Staging in	Staging out	Unknown	Held	Max run	Idle	Claimed	User run here	User running	Unmatched	User idle	Registered	Info age
CDF_DE_ce4_cef1	0	0	0	27	0	0	0	0	625	0	0	1	543	0	0	1	81
CDF_DE_KIT_cream_cdi...	0	0	0	0	0	0	0	151	625	0	0	0	543	0	0	0	27
CDF_ES_PIC_ce05_glong	0	0	0	0	0	0	0	0	625	0	0	0	543	0	0	0	21
CDF_ES_PIC_ce06_glong	0	0	0	0	0	0	0	0	625	0	0	0	543	0	0	0	13
CDF_FR_CCIN2P3_cclcgceli01_long	32	9	0	9	0	0	0	0	474	0	30	30	412	0	0	31	71
CDF_FR_CCIN2P3_cclcgceli03_...	0	2	0	2	0	0	0	0	734	1	0	0	635	0	9	0	15
CDF_IT_Bari_ce01	0	14	0	14	0	0	0	0	734	1	3	3	635	0	9	3	12
CDF_IT_ce2	0	0	0	0	0	0	0	0	900	1	0	0	635	0	0	0	6

● Information about number of pilots and user jobs running on individual sites.

● Typical problem: DESIRED_SE unknown to factory (new site, changed SE name ...)

Monitoring

- Frontend Logs: `~/frontlogs/frontend_UCSD-v5_2/group_main`
- Contains summary information:

```
[2011-11-14T10:51:19-07:00 2202]           Jobs in schedd queues           |           Glideins           |           Request           |
[2011-11-14T10:51:19-07:00 2202] Idle (match  eff  old  uniq )  Run ( here  max ) | Total Idle  Run | Idle MaxRun Down Factory
[2011-11-14T10:51:19-07:00 2202] 7951( 120k  7929  4647    0) 397k(12674 7260k) | 12864  163 12688 | 2223  608k Up   Sum of useful factories
[2011-11-14T10:51:19-07:00 2202]  513(16853  513  513    0) 70234(    0  720k) |    0    0    0 |  276  106k Down Sum of down factories
[2011-11-14T10:51:19-07:00 2202]    1(    1    1    0    1)    0(    0    0) |    0    0    0 |    0    0 Down Unmatched
[2011-11-14T10:51:19-07:00 2202] Advertizing 266 requests
[2011-11-14T10:52:49-07:00 2202] Done advertizing
[2011-11-14T10:52:49-07:00 2202] Advertising 266 glideresource classads to the user pool
[2011-11-14T10:52:49-07:00 2202] Done advertising glideresource classads
[2011-11-14T10:52:49-07:00 2202] Writing stats
```

Troubleshooting

- Typical Problems:
 - Configuration change breaks job matching
- Specialized Condor Queries - why are jobs not matching?
 - Involves looking at user jobs on the CRAB Server schedd in detail and comparing with matching requirements.
 - Nice glideinWMS python API available:

```
[1030] frontend@glidein-collector ~$ ./scripts/test_monitor.py
Number of Pilots Running          8174
Number of Claimed Pilots          7964
Number of Unclaimed Pilots        200
Efficiency of claimed pilots      97.4 %
```

Configuration Changes

- frontend.xml is the main configuration file
- Changes introduced in copy instance_v5_2.cfg, then reconfigure FE:

```
./frontend_startup reconfig ../instance_i5_2a.cfg/frontend.xml
```

Job Matching

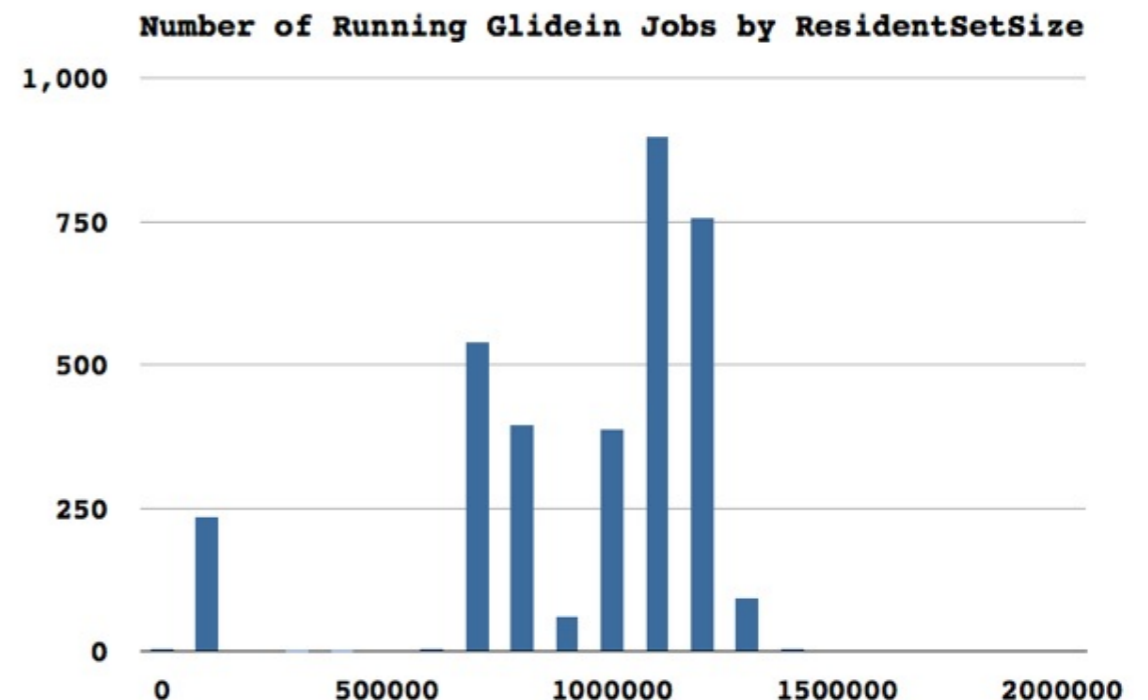
- Main sections of frontend.xml
 - Job Matching Requirements: `<match>`
 - Validation scripts for pilots: `<files>`
- Configuration not particularly reader-friendly: easy to break the matching if you are not careful with changes.

Job Matching

- CMS matches:
 - DESIRED_SE with Glidein_SEs
 - NormMaxWallTimeMins with time left for pilot to run, up to 8 or 16h.
 - If job is rescheduled, ImageSize < 2GB
 - This sometimes causes issues with unmatchable pilots if there are too many such jobs.

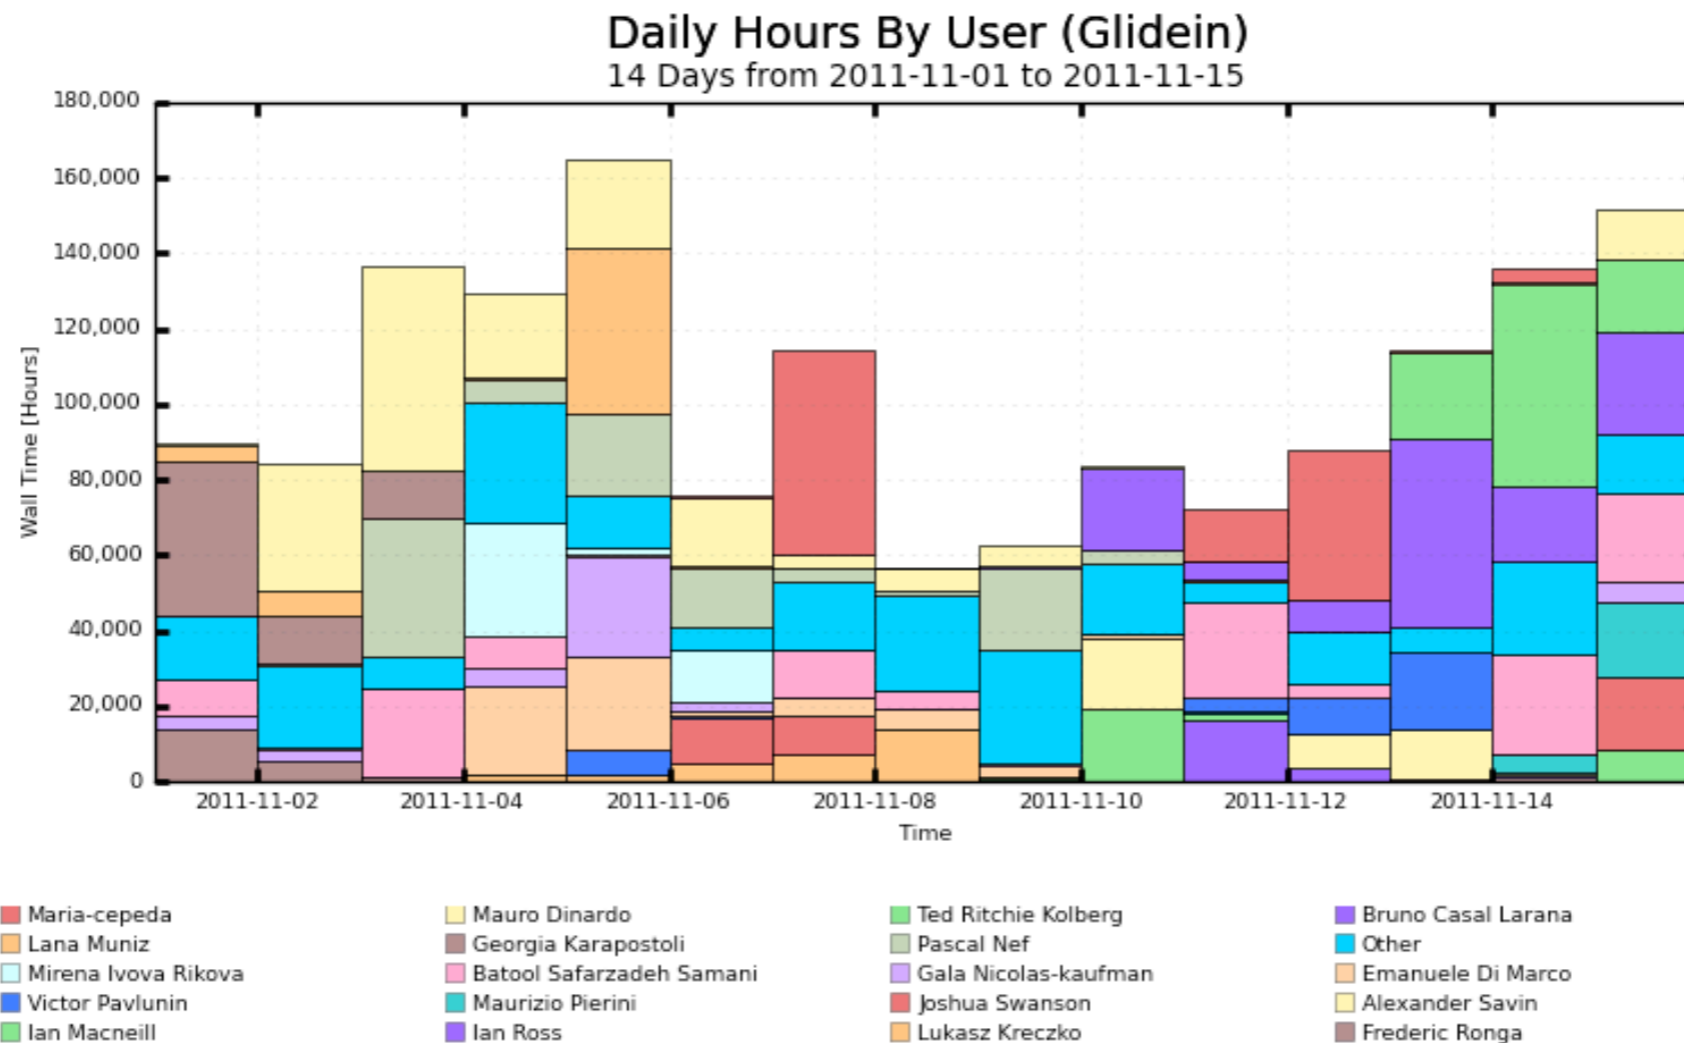
Other Operations

- Updating pilot certificates (once/year/cert)
- Developing CMS-specific validation scripts: basic SAM tests
 - Required CMSSW version
 - /store/temp/user writable?
 - SQUID/Frontier running?
- Ongoing discussion about ImageSize or RSS limits. Currently left to sites. Site-dependent memory limits?



Gratia Probe

http://gratiaweb.grid.iu.edu/gratia/xml/glidein_dn_hours_bar?probe=condor:glidein-2.t2.ucsd.edu



Maximum: 165,200 Hours, Minimum: 56,750 Hours, Average: 104,064 Hours, Current: 151,582 Hours