



Rebrokerage in PanDA

Alden Stradling
6 April 2011

ATLAS Software Week

Unrefined

- This is not something I have spent very long poring over, yet – I am very interested in useful critique and questions
- Plots need some work as well
- ~290,000 jobs rebrokered over the last 5 months



Oracle

- VERY difficult to extract this from Oracle. Not a shock – huge table, parsing text blobs
- Too large to handle neatly on a laptop with Python without trimming
- Pickled dictionaries work far worse than pickled lists (for those who are interested)



Rebrokerage

- Job is brokered to a particular site
- Site delays or downtime, dataset problems:
 - Rebroker. Cancel previous job (taskbuffererrorcode = 105)
 - Create new (build, run) jobs on new site if possible, new PandaID and etc. Job runs.




Things to Examine

- Delay before rebrokering
- Average wait for rebrokered job:
 - By site
 - By cloud
- Check destination sites for anything particularly popular



Site Proportions

ANALY_LONG_BNL_ATLAS	32535
ANALY_RAL	26128
ANALY_FZK	13752
ANALY_BNL_ATLAS_1	13271
ANALY_AGLT2	11821
ANALY_SLAC	11026
ANALY_LANCS	9702
ANALY_SARA	7955
ANALY_MWT2	7156
ANALY_SWT2_CPB	7043
ANALY_CERN_XROOTD	6708
ANALY_GLASGOW	6604
ANALY_QMUL	6550
ANALY_INFN-T1	6470




Top Sites (rebrokerage)

ANALY_LONG_BNL_ATLAS	32535
ANALY_RAL	26128
ANALY_FZK	13752
ANALY_BNL_ATLAS_1	13271
ANALY_AGLT2	11821
ANALY_SLAC	11026
ANALY_LANCS	9702
ANALY_SARA	7955
ANALY_MWT2	7156
ANALY_SWT2_CPB	7043
ANALY_CERN_XROOTD	6708
ANALY_GLASGOW	6604
ANALY_QMUL	6550
ANALY_INFEN-T1	6470
ANALY_TR-10-ULAKBIM	6018
ANALY_TOKYO	5338
ANALY_TW-FTT	5087
ANALY_GRIF-LPNHE	4892
ANALY_GRIF-IRFU	4864
ANALY_PIC	4776
ANALY_LONG_LYON-T2	4407
ANALY_RHUL	4313
ANALY_IFAE	4231
ANALY_NIKHEF-ELPROD	4208
ANALY_DESY-HH	4107
ANALY_CPPM	3817
ANALY_CERN	3810

ANALY_wuppertalprod	3694
ANALY_MANC	3589
ANALY_INFEN-MILANO-ATLASC	3563
ANALY_TRIUMF	2819
ANALY_FREIBURG	2804
ANALY_LRZ	2784
ANALY_ARC	2638
ANALY_IllinoisHEP	2490
ANALY_WEIZMANN	2474
ANALY_TAIWAN	2363
ANALY_LYON_DCACHE	2303
ANALY_CAM	2198
ANALY_SCINET	2162
ANALY_LYON-T2	1899
ANALY_TECHNION-HEP	1820
ANALY_UAM	1660
ANALY_LONG_LYON_DCACHE	1540
ANALY_SLAC_LMEM	1500
ANALY_IHEP	1480
ANALY_IFIC	1467
ANALY_OX	1365
ANALY_LIV	1275
ANALY_CYF	1120
ANALY_LIP-Coimbra	1091
ANALY_IL-TAU-HEP	1055
ANALY_AUSTRALIA	1052
ANALY_BEIJING	1039

Major sites seem to have had really big days, but the largest spikes correspond to a busy Grid day

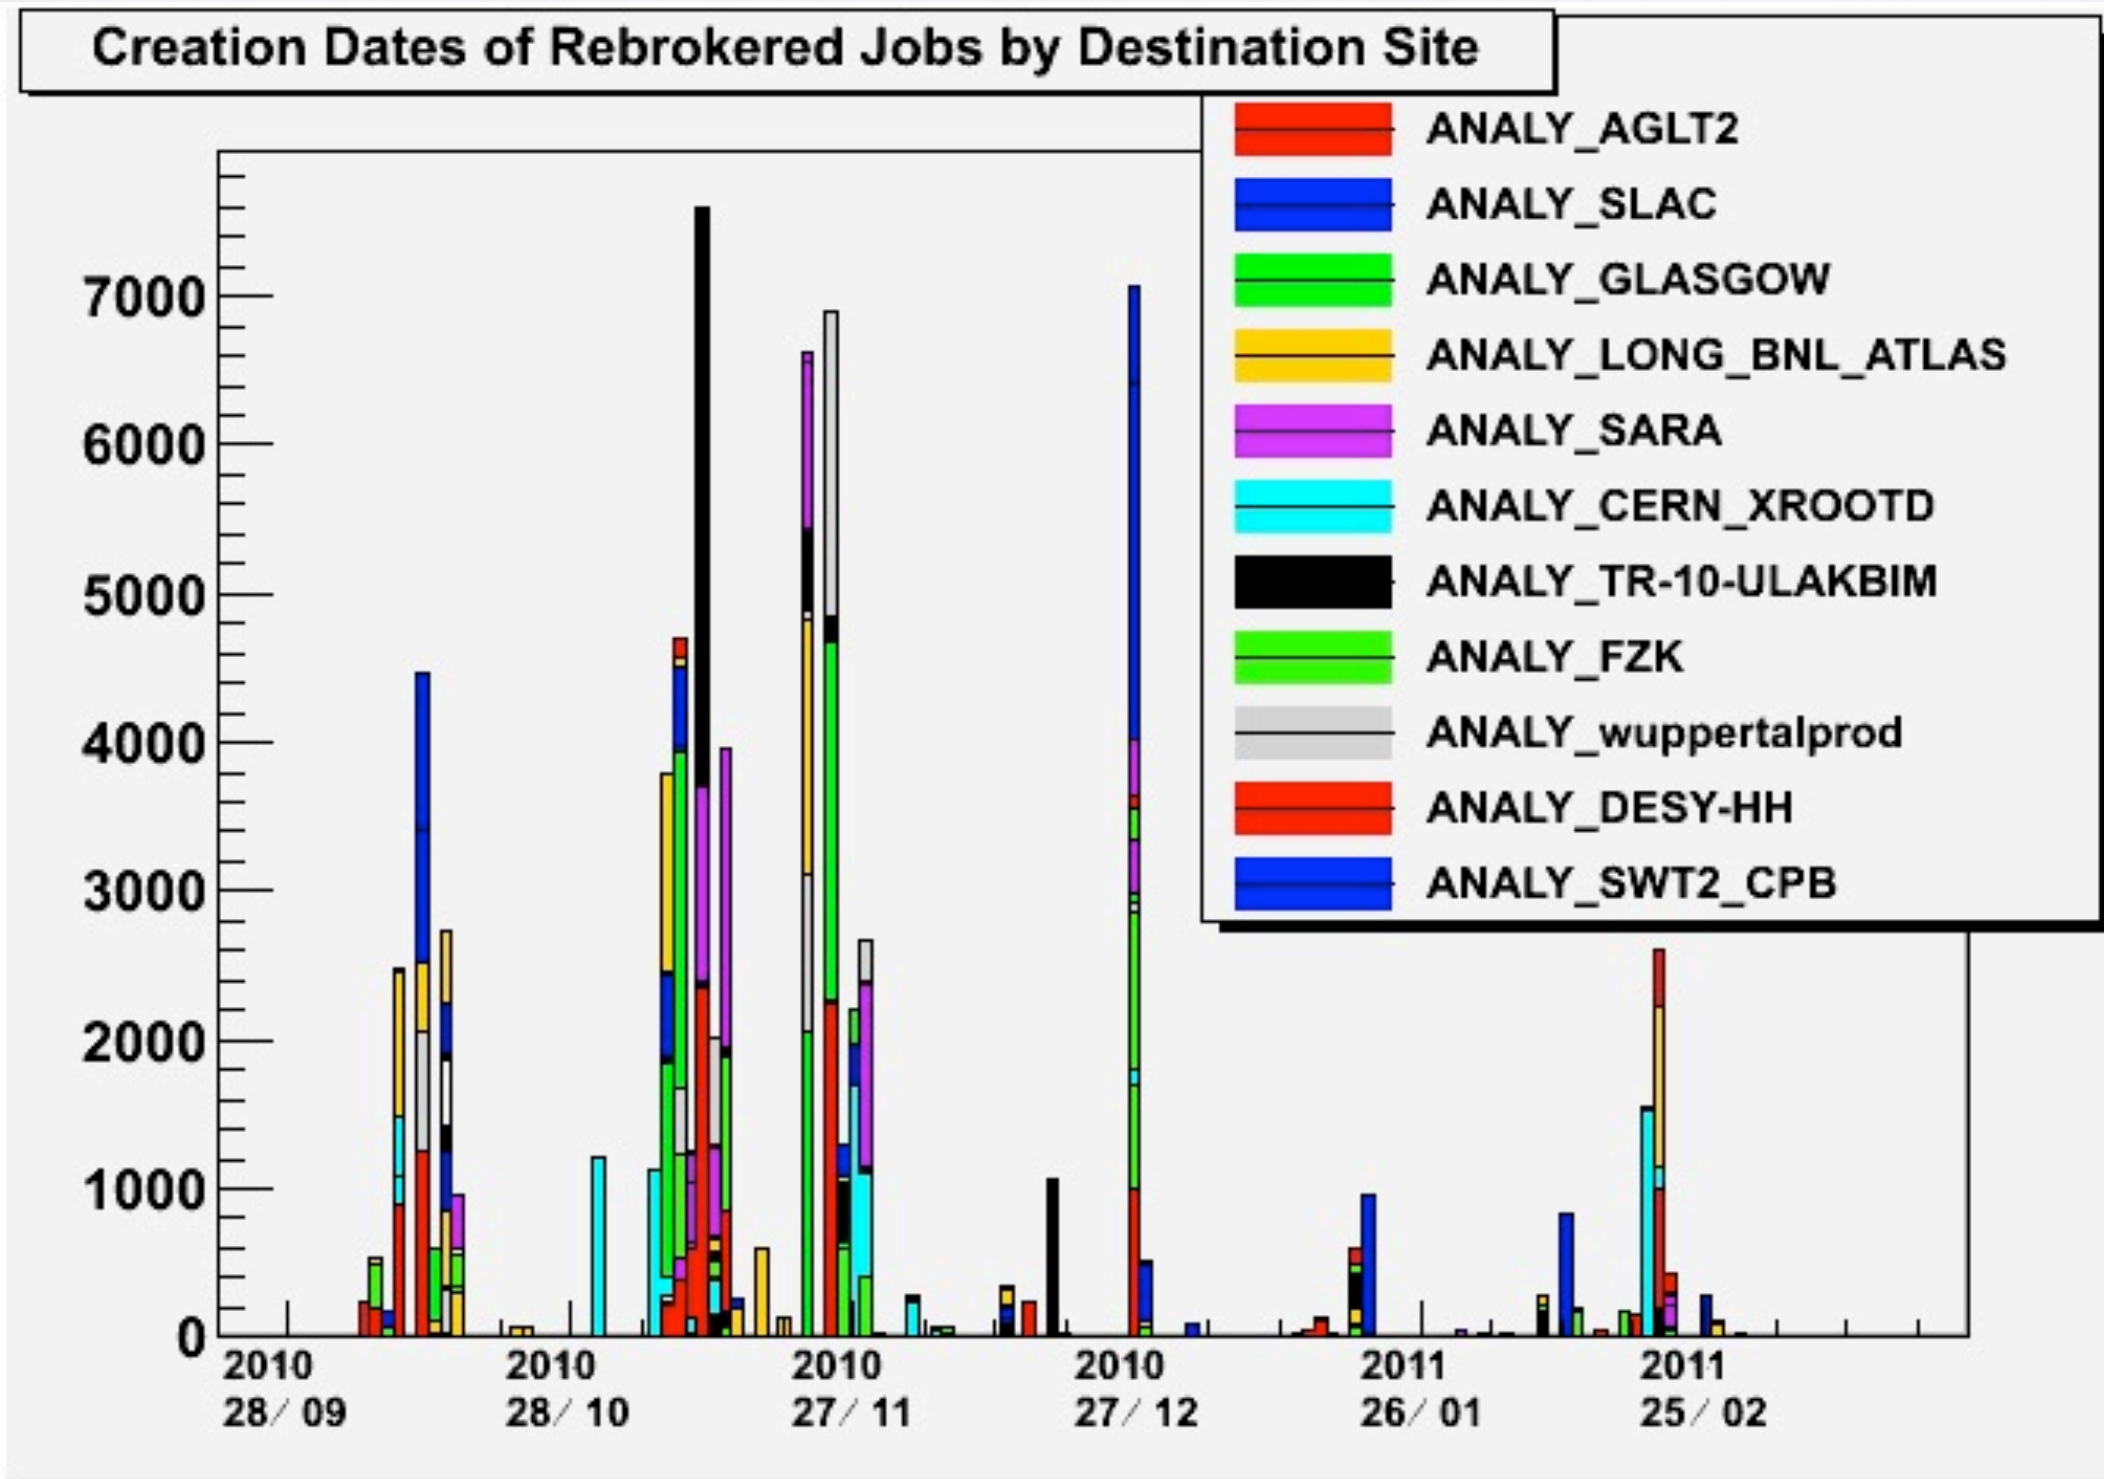


Sites

- BNL takes top honors, especially the LONG queue – unsurprising, given the routine overload that BNL faces
- RAL, FZK, etc – big T1 and T2 sites. No real surprise
- The drop is pretty sharp – rebrokerage really benefits overloaded sites, less impact on less loaded systems

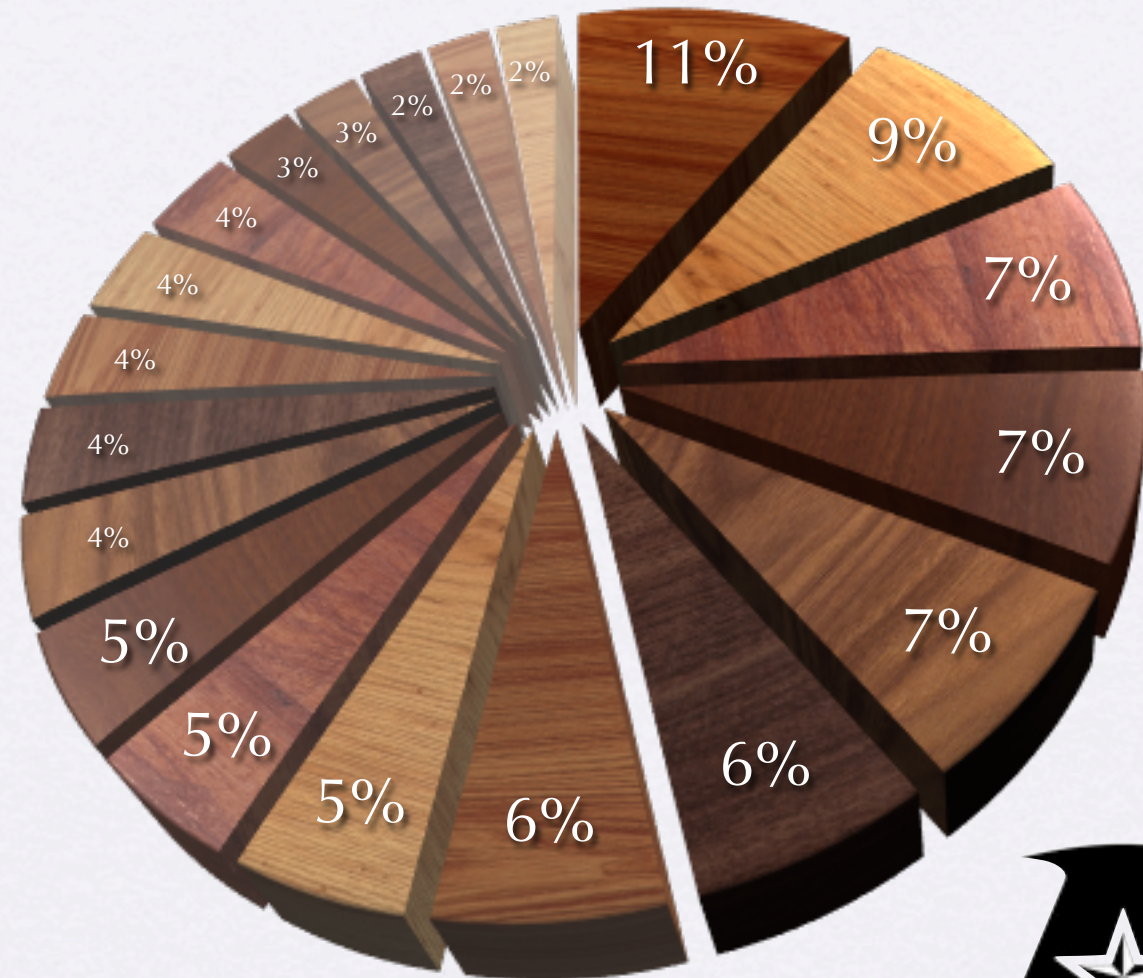
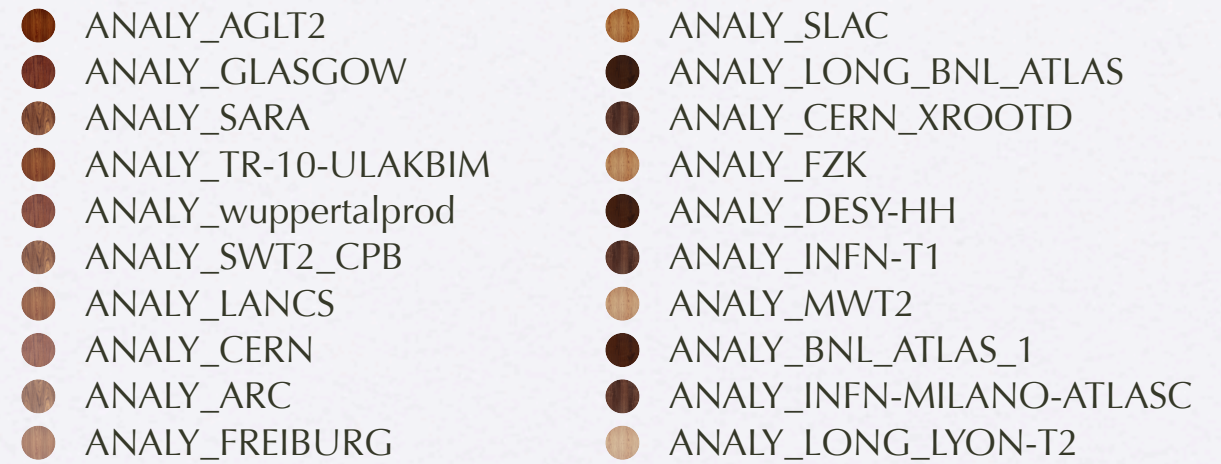


Destination Sites



Top Destinations

ANALY_AGLT2	6609
ANALY_SLAC	5306
ANALY_GLASGOW	4466
ANALY_LONG_BNL_ATLAS	4356
ANALY_SARA	4334
ANALY_CERN_XROOTD	3936
ANALY_TR-10-ULAKBIM	3898
ANALY_FZK	3206
ANALY_wuppertalprod	3063
ANALY_DESY-HH	2995
ANALY_SWT2_CPB	2585
ANALY_INFN-T1	2448
ANALY_LANCS	2264
ANALY_MWT2	2255
ANALY_CERN	2205
ANALY_BNL_ATLAS_1	1706
ANALY_ARC	1588
ANALY_INFN-MILANO-ATLASC	1504
ANALY_FREIBURG	1493
ANALY_LONG_LYON-T2	1464



How many tasks does this represent?

- 1886 job definitions were rebrokered
 - Average subjobs: 155, mean:12, mode:1
 - 916 of 1886 (49%) contain less than 10 subjobs
- No time distribution plot on this – I can't imagine that it would be clear/useful.




Top JobIDs

(for what it's worth)


JobID	Subjobs
Job 576	4920
Job 850	3575
Job 624	3545
Job 768	2766
Job 128	2640
Job 828	2640
Job 766	2573
Job 132	2494
Job 750	2494
Job 644	2373
Job 620	2356
Job 854	2356
Job 328	2276
Job 856	2222
Job 1596	2149
Job 762	2140
Job 866	1989
Job 642	1978
Job 586	1946

- Job 576
- Job 128
- Job 750
- Job 328
- Job 866
- Job 850
- Job 828
- Job 644
- Job 856
- Job 642
- Job 624
- Job 766
- Job 620
- Job 1596
- Job 586
- Job 768
- Job 132
- Job 854
- Job 762




Cloud Proportions

US	88047
UK	63501
FR	32881
DE	30951
NL	26037
ES	13377
IT	10700
CERN	10518
TW	7450
CA	6904
ND	2638



UK spikes correlate to the huge RAL spike we saw in the Sites plot. US seems to consistently do the lion's share of rebrokering



Repeated Rebrokeraage

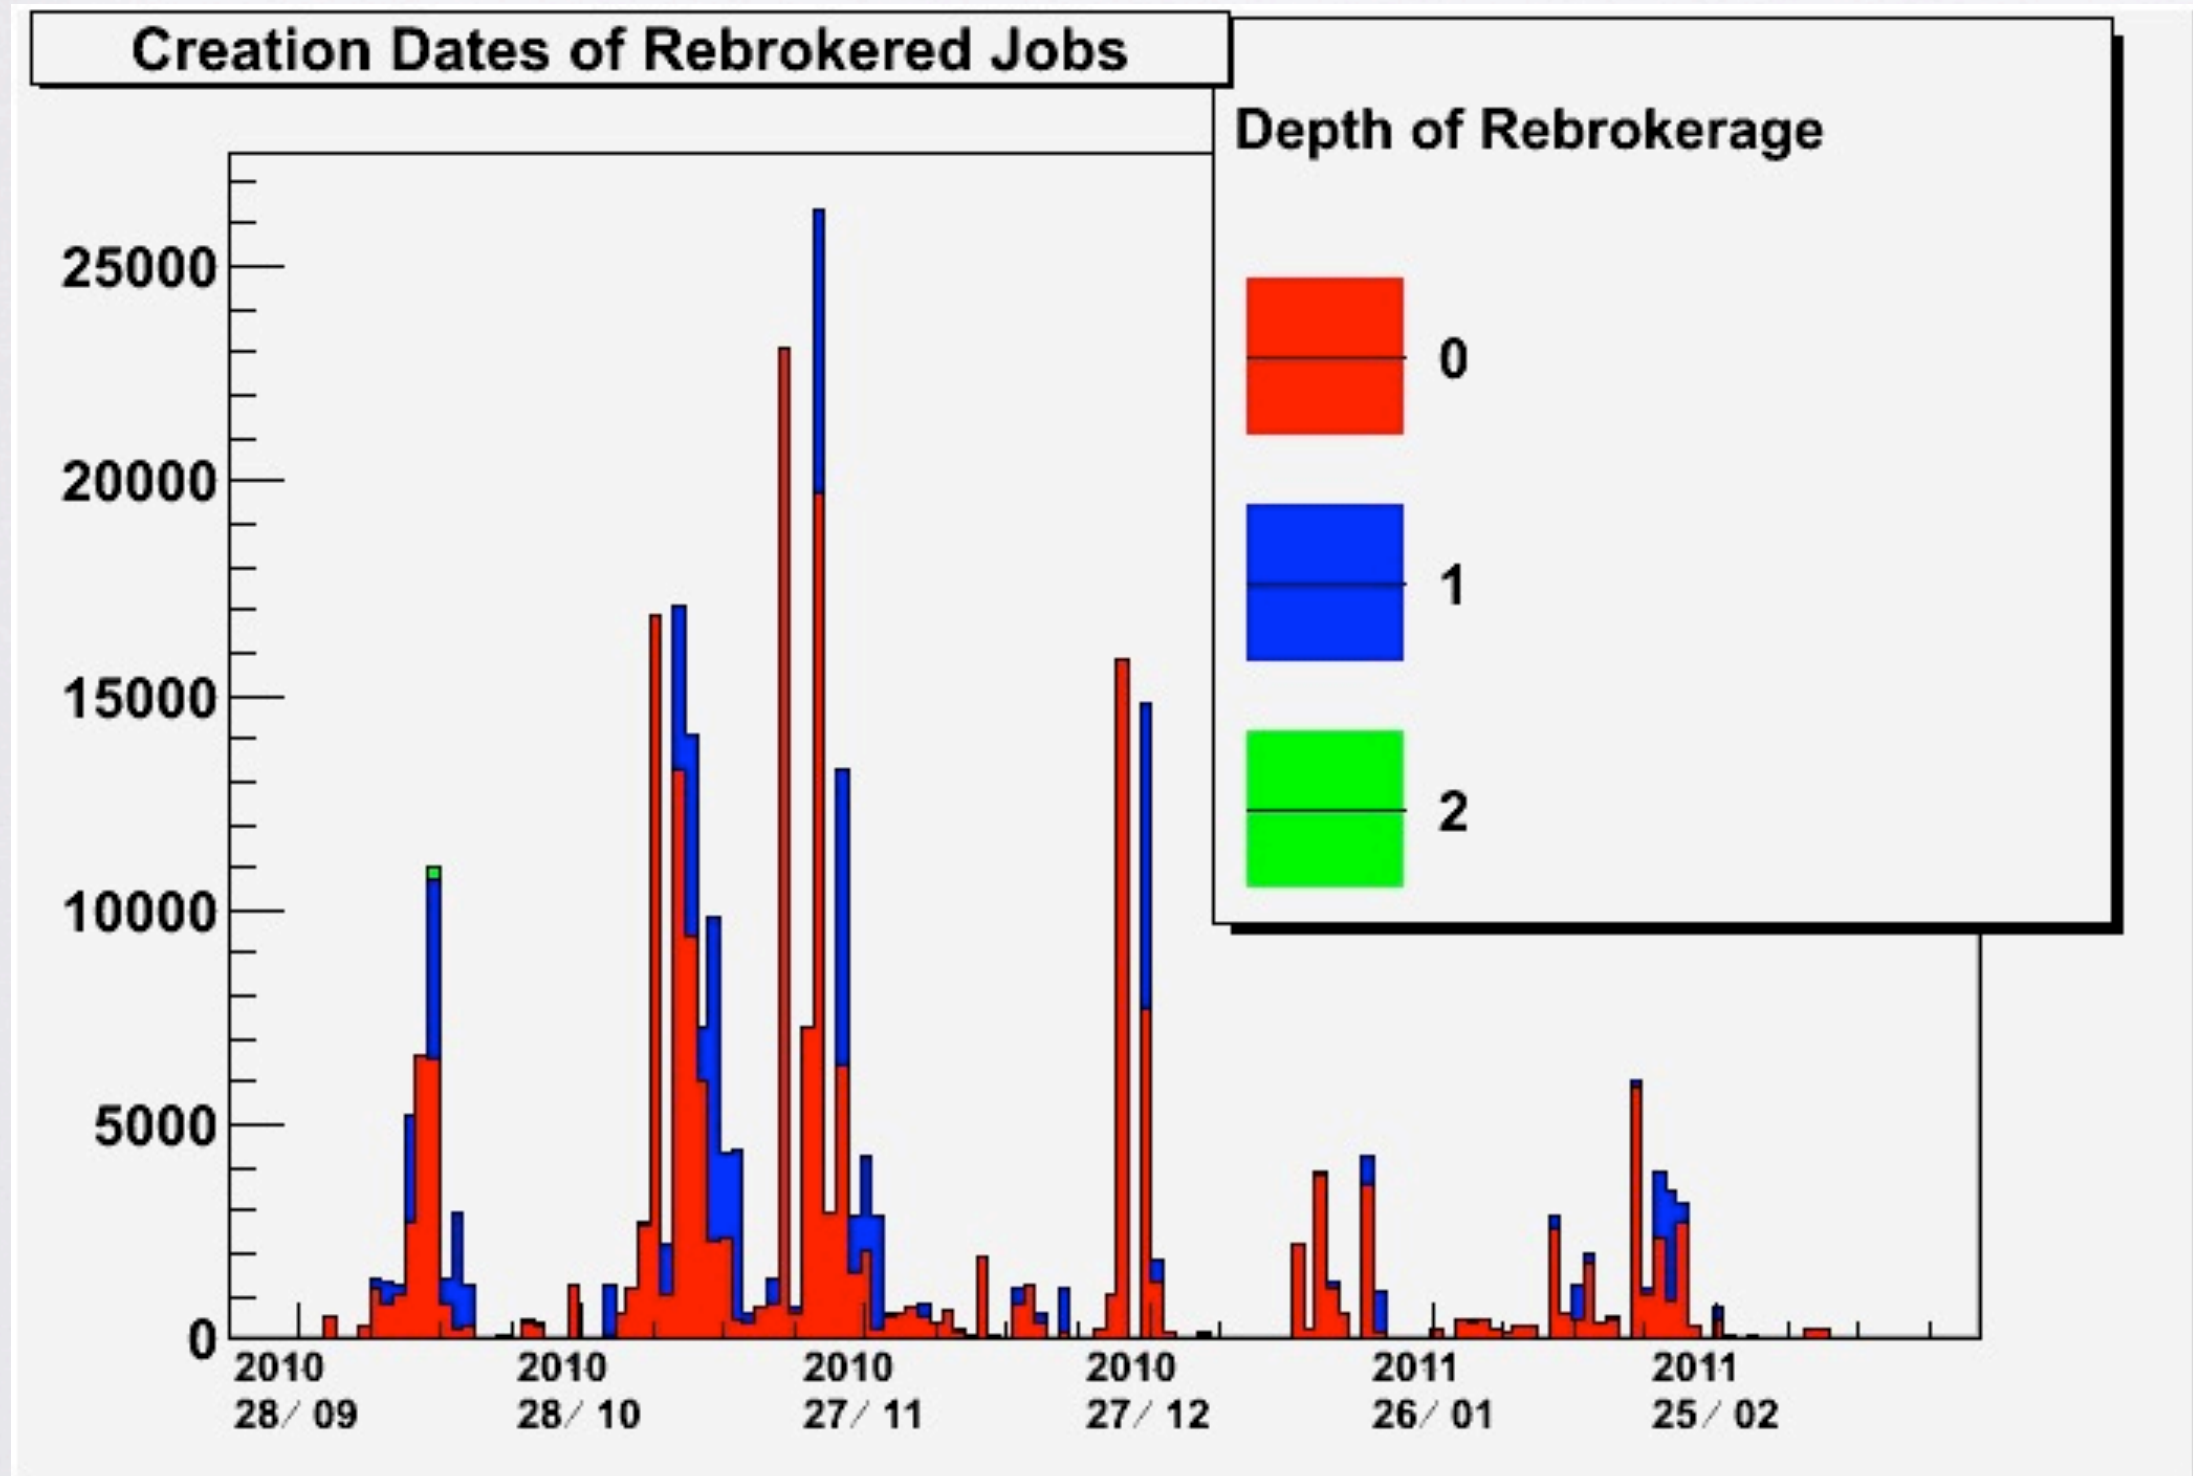
0	215851
1	76826
2	327

● One Hop ● Two Hops
● Three Hops




Pretty clear that re-rebrokeraage is common
Three hops is the max seen so far, and is rare

Repeated Rebrokage



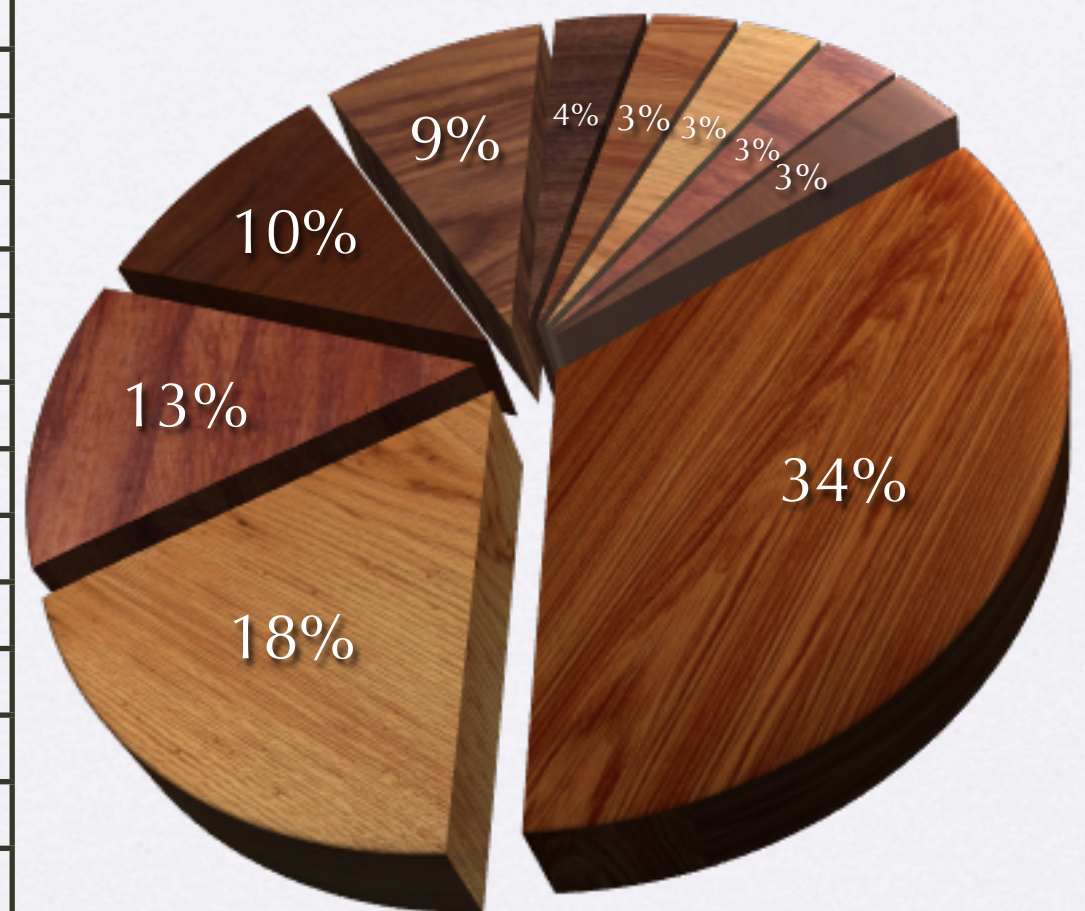
Rebrokerage by Tool



By User

Yousuke Kataoka	60363	Frederic Derue	2872
Evgeny Khramov	32359	Albert Olariu	2784
Maarten Boonekamp	22732	Matthew Relich	2778
Enrique Kajomovitz	18312	Nuno de Castro	2650
Wen Guan	15990	German Carrillo Montoya	2626
Nuno dos Anjos	6452	Tommaso Lari	2588
Iro Koletsou	6078	christina potter	2413
Kunihiro Akiyama	5812	Jean-Francois Marchand	2376
Matthew Beckingham	5581	Trevor Vickey	2350
Laurent Vacavant	5296	Matthew Graham Chapman	2332
Haifeng Li	4688	Lidia DellAsta	2280
Helena dos Santos	4394	Yuto Komori	2216
Michela Biglietti	4154	Fabrizio Parodi	2190
Andrea Ferretto Parodi	4040	Hovhannes Khandanyan	1836
Tomonori Nishiyama	3913	Takayasu Hayashi	1594
Louise Oakes	3906	Mikhail Demichev	1496
Peter Loscutoff	3770	Jean-Baptiste De Regie	1493
Moritz Bunse	3628	Nicole Manuela Ruckstuhl	1421
Robert Calkins	3255	Elizabeth Jue Hines	1408
Alexander Vincent Penson	3082	Andrea Knue	1354

- Yousuke Kataoka
- Evgeny Khramov
- Maarten Boonekamp
- Enrique Kajomovitz
- Wen Guan
- Nuno dos Anjos
- Iro Koletsou
- Kunihiro Akiyama
- Matthew Beckingham
- Laurent Vacavant



User Variations

- Not sure why this is the case – what is it about the three primary users that gave them so many rebrokered jobs?
 - Blind submission of massive jobs?
 - Production role, doing group submission on a specific site?



Conclusions

- Fair amount of mining yet to do on this dataset
- Rebrokering seems to be doing its job, and does not seem to have any obvious pathologies
- Possible expansion of its role – active aggregation/
matchmaking of jobs to cached files on sites
 - Small set of files, probably – but very visible and
critical path

