



PanDA Update

Kaushik De

Univ. of Texas at Arlington

XRootD Workshop, UCSD

January 27, 2015



Overview



- At last XRootD workshop (OSG-AHM, April 2014) we showed first results from overflow jobs – looked promising but high failure rates caused us to switch off further testing
- Failover jobs were working very well
- In this talk I will focus on new results and future plans
- This work benefits from synergy with networking projects
 - While XRootD is an important component of US ATLAS core program, additional funding from DOE and NSF was crucial to success of this work
 - DOE-ASCR/HEP funded BigPanda project (BNL-UTA-ANL)
 - NSF funded CC-NIE:ANSE project (CalTech-UM-Vanderbilt-UTA)
 - Future work will be funded through Russian MegaPanDA grant



FAX and PanDA




- All WLCG ATLAS sites are becoming FAX enabled
 - But we are not replacing FTS with FAX – they are complementary
 - Gradually shift some PanDA workload to FAX
 - FAX allows more interesting use cases – focus on new capabilities in PanDA that are only possible through FAX
 - FAX deployment, robust operations and new software development must evolve together
- Many details already in other ATLAS talks at this meeting
 - ATLAS Status and Plans: Ilija Vukotic (UC)
 - ATLAS data formats and FAX: Doug Benjamin (Duke)
 - ATLAS experience: Rob Gardner (UC)
 - Event Service: Vakho Tsulaia (LBNL)
 - And others – I will try not to repeat their excellent content



PanDA Scale



dashboard



25M jobs every month at >150 sites
Processed more than one exabytes in 2013 and 2014

Case 1: Rescuing Failed Transfers



- PanDA pilot retries FTS transfer of input files a few times for production
- For analysis jobs, if local file cannot be read, jobs fail
- For sites with FAX enabled, why not try FAX to rescue job?
- This was implemented and working well – snapshot of past 6 hours

Summary:
[All Jobs](#) [CMS Jobs](#)
[My Jobs](#) [ATLAS Jobs](#)
Clouds' Summary:
[analysis](#) [production.test](#)
[jedi](#) [all](#)
Regions' Summary:
[analysis](#) [production.test](#)
[jedi](#) [all](#)
States:
[defined](#) [pending](#)
[waiting](#) [assigned](#)
[activated](#) [sent](#)
[starting](#) [running](#)
[holding](#) [transferring](#)
[finished](#) [failed](#)
[cancelled](#) [unassigned](#)

[Job Types](#)
[Jobs' Timing](#)
[Jobs' Metrics](#)
[Jedi Tasks](#)
[FAX](#)
[Quick Search](#)
[Errors](#) [Task Errors](#)
[Operations](#)
[Worker Nodes](#)
[Summaries](#)
[Tasks](#) [Search](#)
[Datasets](#)
[Search](#)
[Datasets](#)
[Distribution](#)

Site	Jobs	WithFAX [files]	WithoutFAX [files]	WithFAX [GB]	WithoutFAX [GB]
unknown: TOKYO_MCORE	56	539	65	225.98	44.46
unknown: INFN-NAPOLI-RECAS_MCORE	46	364	142	152.92	98.88
unknown: SWT2_CPB_MCORE	37	336	46	140.39	42.72
unknown: MWT2_MCORE	34	328	40	137.43	50.48
unknown: SLACXRD_MP8	31	255	29	105.89	26.76
unknown: INFN-NAPOLI-ATLAS_MCORE	27	198	93	74.46	57.76
unknown: BU_ATLAS_Tier2_MCORE	13	48	58	23.73	24.64
unknown: UKI-NORTHGRID-SHEF-HEP_MCORE	11	45	76	15.44	49.76
unknown: UKI-SCOTGRID-GLASGOW_MCORE	11	102	10	42.33	5.55
unknown: IFAE_MCORE	9	90	9	37.37	10.02
unknown: ANALY_BNL_LONG	8	8	98	28.72	327.33
unknown: AGLT2_MCORE	7	61	16	25.17	12.14
unknown: CA-MCGILL-CLUMEQ-T2_MCORE	7	54	23	23.00	11.65
unknown: CONNECT_MCORE	7	52	25	21.61	15.19
unknown: ANALY_IFIC	6	6	30	7.72	32.39
unknown: UKI-LT2-QMUL_MCORE	6	60	6	24.87	5.55
unknown: CPPM_MCORE	5	50	5	20.73	4.56
unknown: IFIC_MCORE	5	32	3	12.80	2.63
unknown: ANALY_GRIF-LAL	4	4	4	11.79	0.13
unknown: ANALY_SLAC	4	5	256	0.03	0.86
unknown: BNL_PROD	4	4	21	0.19	1.16
unknown: IN2P3-CC_MCORE	4	12	32	5.75	20.42
unknown: UAM-LCG2_MCORE	4	16	28	6.22	14.55
unknown: ANALY MANC SL6	3	3	16	4.00	16.24



Case 2: Faster User analysis



- Also first use case for network integration with PanDA
- Goal - reduce waiting time for user jobs
 - User analysis jobs normally go to sites with local input data
 - This can occasionally lead to long wait times (jobs are re-brokered if possible, or PD2P data caching will make more copies eventually to reduce congestion)
 - While nearby sites with good network access may be idle
- Brokerage uses concept of 'nearby' sites
 - Use cost metric generated with Hammercloud tests
 - Calculate weight based on usual brokerage criteria (availability of CPU resources, data location, release...) plus new network transfer cost
 - Jobs will be sent to the site with best overall weight
- Throttling is used to manage load on network

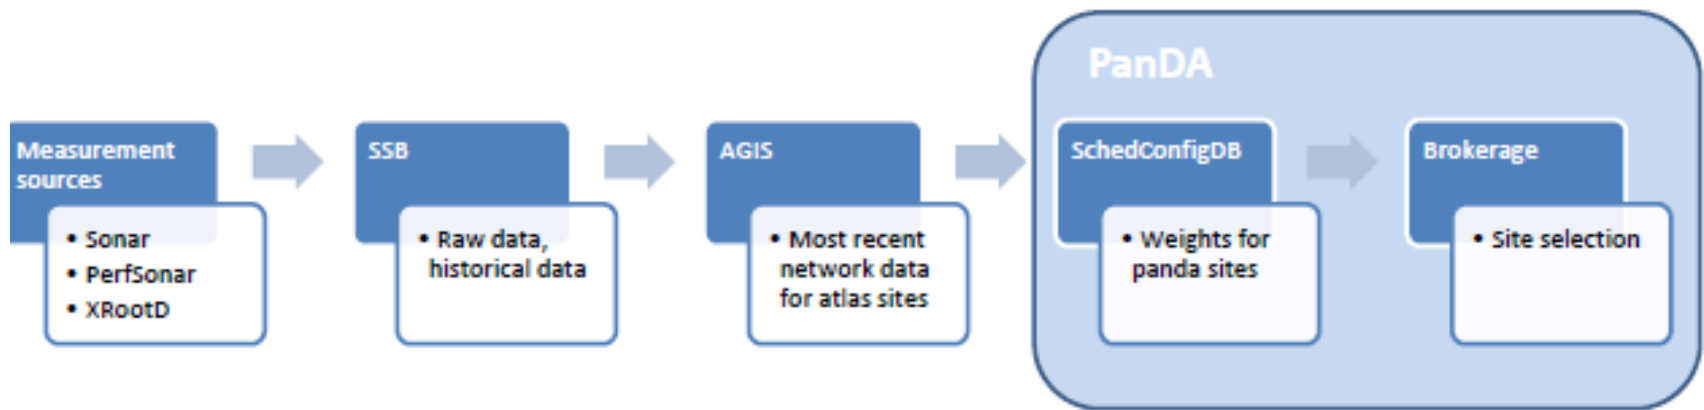


Network Data Repositories



- **Native data repositories**
 - Historical data stored from collectors
 - SSB – site status board for sonar and perfSonar data
 - FAX data is kept independently and uploaded
- **AGIS (ATLAS Grid Information System)**
 - Most recent / processed data only – updated periodically
 - Mixture of push/pull – depending on source of data
- **schedConfigDB**
 - Internal Oracle DB used by PanDA for fast access
 - Uses standard ATLAS collector

Dataflow



- Data is being transformed
 - Historical to most recent
 - Mb/sec to weights
 - Atlas sites to panda queues




Jobs using FAX for remote access



dashboard

Completed jobs

16 Weeks from Week 39 of 2014 to Week 04 of 2015




Maximum: 7,205,760, Minimum: 0.00, Average: 5,064,118, Current: 1,447,256

About 4-10% of jobs Access Data Remotely Through Federated XRootD

dashboard

Completed jobs (Sum: 4,067,722)



- US-MWT2 - 21.23% (863,502)
- US-WT2 - 11.59% (471,587)
- FR-CCIN2P3 - 5.06% (205,760)
- CA-WEST-T2 - 2.84% (115,482)
- NL-T1 - 0.63% (25,590)
- US-T1-BNL - 16.33% (664,241)
- US-NET2 - 8.22% (334,262)
- IT-INFN-CNAF - 4.34% (176,532)
- DE-MCAT - 2.35% (95,475)
- CA-TRIUMF - 0.25% (10,243)
- US-AGLT2 - 15.46% (628,828)
- DE-DESY-ATLAS-T2 - 6.46% (262,972)
- US-SWT2 - 3.82% (155,385)
- UK-London-Tier2 - 1.23% (50,189)
- CA-EAST-T2 - 0.19% (7,676)

Remote site Is selected Based on Network Performance


Data since October 2014



Early Example from Oct, 2014




How do we measure success?
Let's look at example task.






Jobs from task 4199733 on oct 3





Job wait times for example task





Other Tasks and other Metrics?



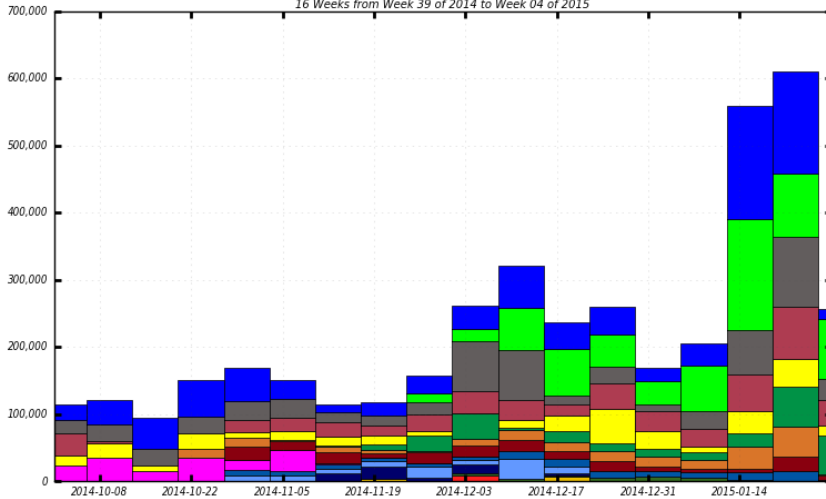
- **There are three metrics to quantify performance**
 - Average wait times with/without FAX
 - Execution time of jobs – averages and tails
 - Success rate of jobs with/without FAX
- **We do not yet have comprehensive analytics**
 - Ilija is exporting and studying PanDA job data in Hadoop
- **Future optimizations**
 - Algorithm has many tunable parameters
 - Work is underway (for example, site black-listing)
 - Complex optimization problem over three dimensions
- **We have moved from development to operations**
 - Which is always harder
 - ATLAS scale is challenging - >150 sites, millions jobs/day

Measure of Success



Completed jobs

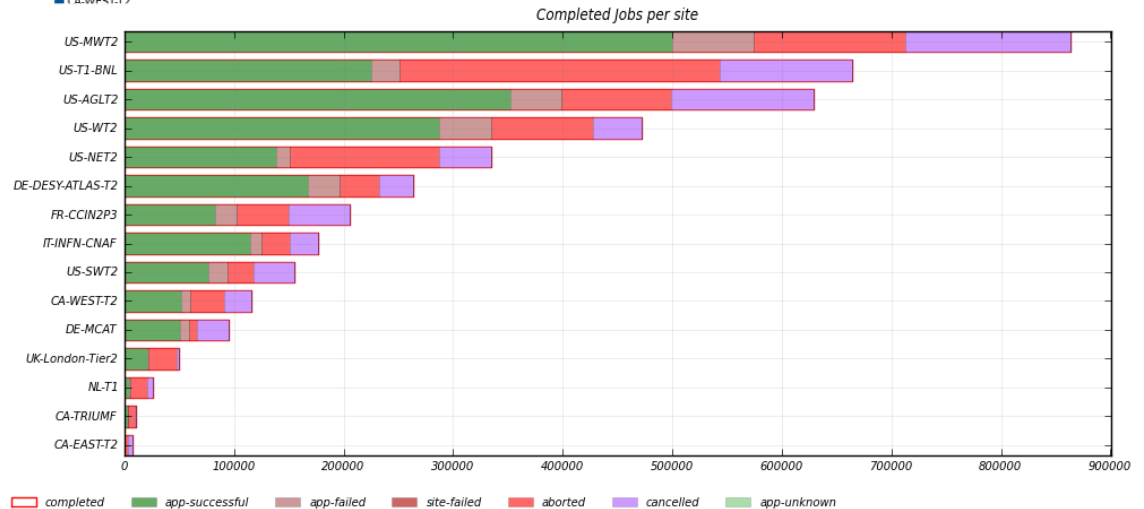
16 Weeks from Week 39 of 2014 to Week 04 of 2015



Maximum: 611,268 , Minimum: 0.00 , Average: 214,090 , Current: 256,851

Many sites used for overflow

Failure rate is manageable




Maybe not so Rosy?



dashboard

Running jobs

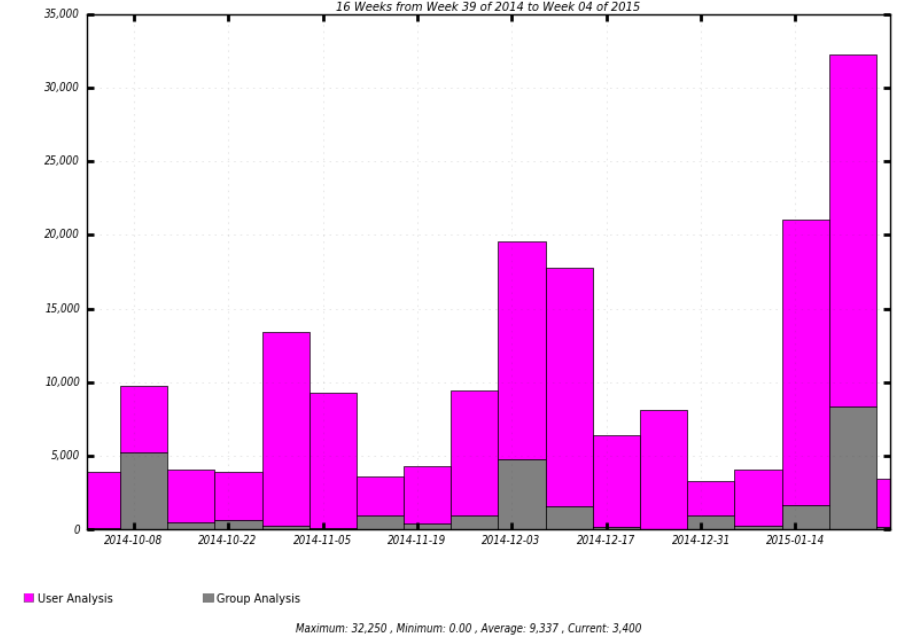
16 Weeks from Week 39 of 2014 to Week 04 of 2015



dashboard

Pending jobs

16 Weeks from Week 39 of 2014 to Week 04 of 2015





Conclusions for Case 2



- **Network data collection working well**
 - Additional algorithms to combine network data will be tried
 - HC tests working well – PS data not robust yet
- **PanDA brokerage worked well**
 - Achieved goal of reducing wait time – though anecdotally
 - Well balanced local vs remote access
 - Need fine tuning – we have a lot of data now (4M jobs)
- **We have overflow working for FAX**
- **We need deeper study to optimize**



Case 3: FAX for Data Transfer



- **Production jobs move data**
 - PanDA transfers input data for MC production, reprocessing, group production (DPD making), and other privileged processing
 - Transfer is done asynchronously, if needed (preference given to sites which already have data by brokerage)
 - Once transfer is completed, PanDA receives callback, which activates processing
- **Work started to send callback to PanDA from FAX**
 - Once callback is available, we will instrument a few sites



Other Possibilities



- **FAX for data caching**
 - PanDA will send jobs to sites which may have data (not guaranteed)
- **Workload aware data access**
 - PanDA can choose FAX direct access for IO intensive workloads, and copy data for CPU intensive workloads
- **Optimized ROOT access**
 - pRoot?
- **FAX for writing output files**
 - So far, all use cases involve read access
 - PanDA can use the flexibility of writing via FAX
- **FAX for event service**
 - Event service needs object store – FAX can provide this



Conclusion



- Finally, good results from 2 out of our 3 use cases discussed at previous FAX meetings
- Lot of work to move from development to operations
 - Ongoing work and future work
 - We need to focus on operations and analytics
- Many new possibilities
- FAX is now integral part of PanDA WMS