

---

# JEDI overflow jobs

performance, tuning, and issues

---

Ilija Vukotic • University of Chicago

Computing Integration and Operations meeting  
September 3, 2014

---

# Introduction

---

- Overflow was turned back on for the US cloud on September 28th.
  - Here a simple overview of the first steps in checking and tuning the overflow access.
-

# Bugs found

---

- Immediately got struck with another rogue path lookup that killed all the cmsd services. Patch produced and deployed almost everywhere quite quickly.
  - Discovered that schedconfig FAX cost table contained Panda queue names and not queue PandaIDs, so jobs were sent only to sites where these two match (only MWT2). Artem fixed it so jobs started flowing everywhere.
  - Found that in some situations jobs at BNL were still getting “old style” gLFNs. Pilot got fixed and should be in production today. In a meanwhile BNL was turned off as a “sink” site.
  - Endpoints at WT2, OU\_OCHEP\_SWT2, and SWT2\_CPB, failed several times. Fixes are in testing.
-

# Overflow jobs

---

Opening 5 sites to overflow released the flood of jobs.

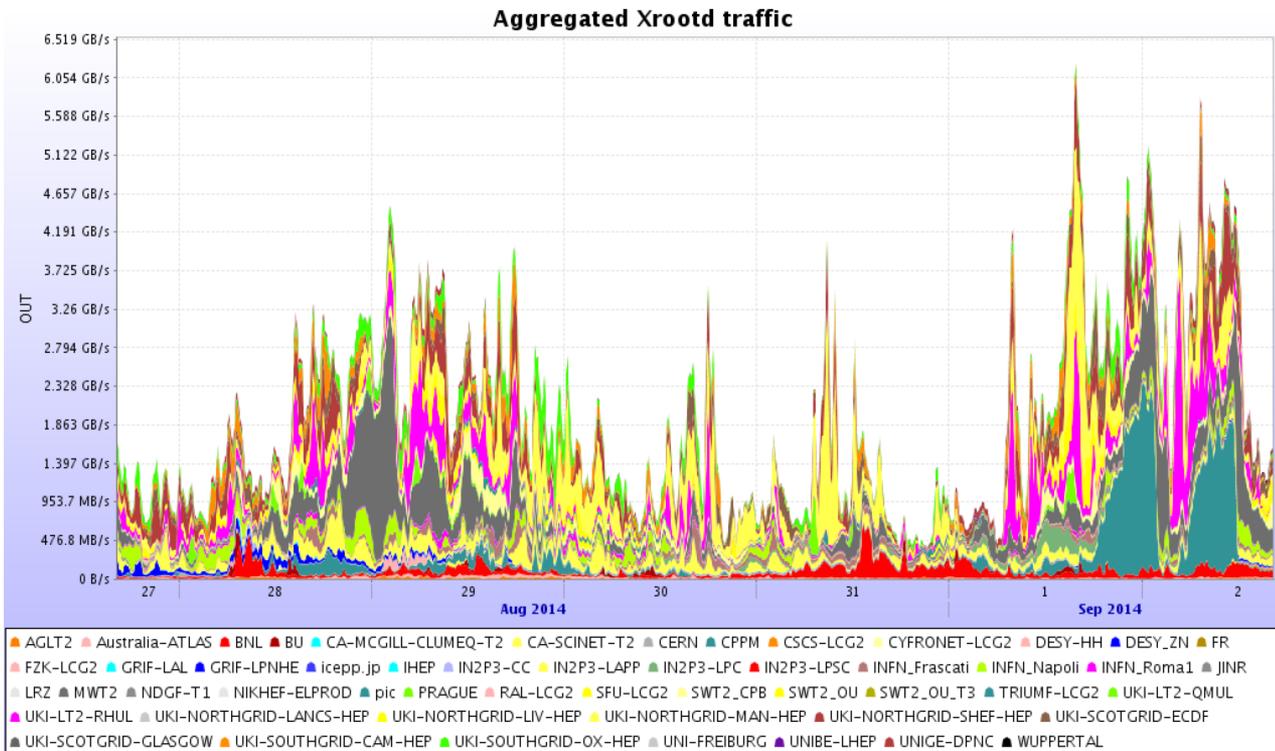
Trying to reduce number of concurrently running overflow jobs by reducing WANsinkLimits to 3Gbps. ( some queues have seen >4k )

With a lot of jobs in “running” quite some delay until effects of tuning are seen.

---

# Overflow jobs

Not much stress  
on WAN links.

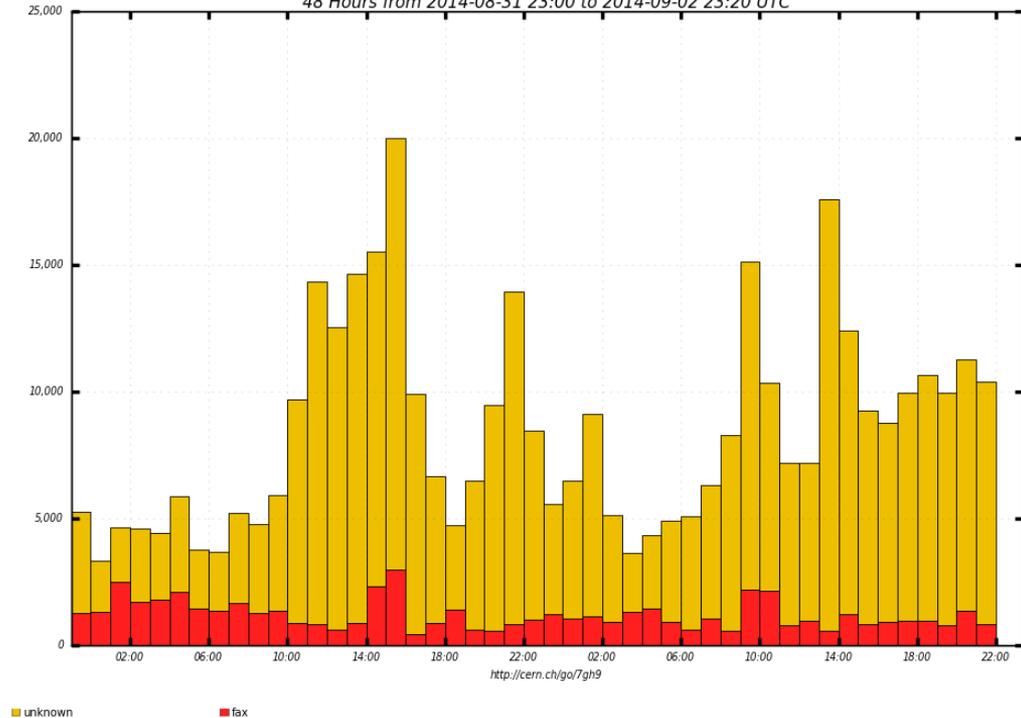


# Comparing “overflow” and “regular”

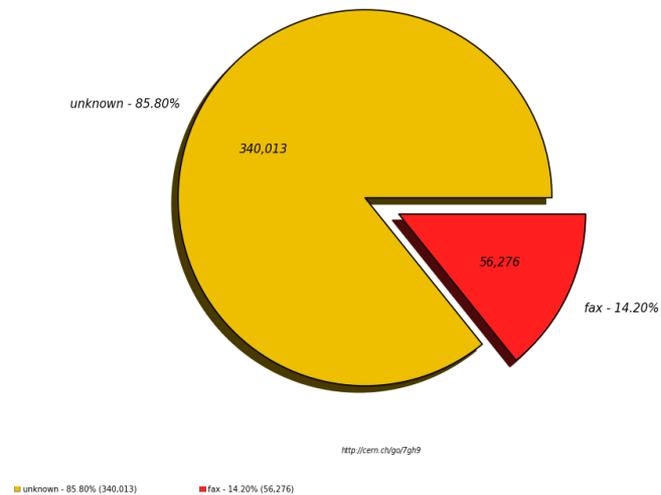


Completed jobs

48 Hours from 2014-08-31 23:00 to 2014-09-02 23:20 UTC



Completed jobs Pie (Sum: 396,289)



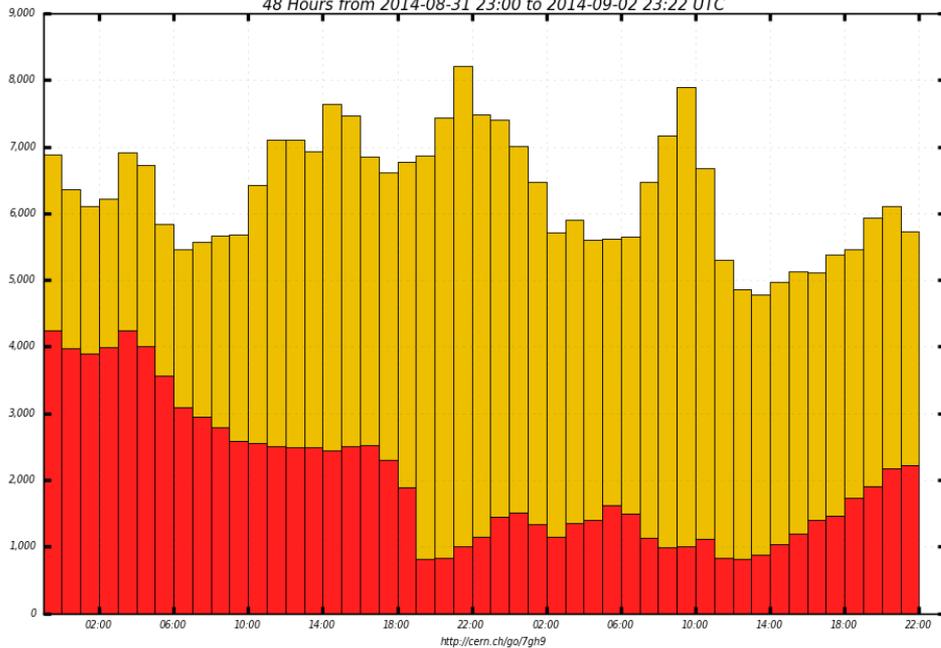
<http://cern.ch/ga/7gh9>

# Comparing “overflow” and “regular”



Running jobs

48 Hours from 2014-08-31 23:00 to 2014-09-02 23:22 UTC



<http://cern.ch/go/7gh9>

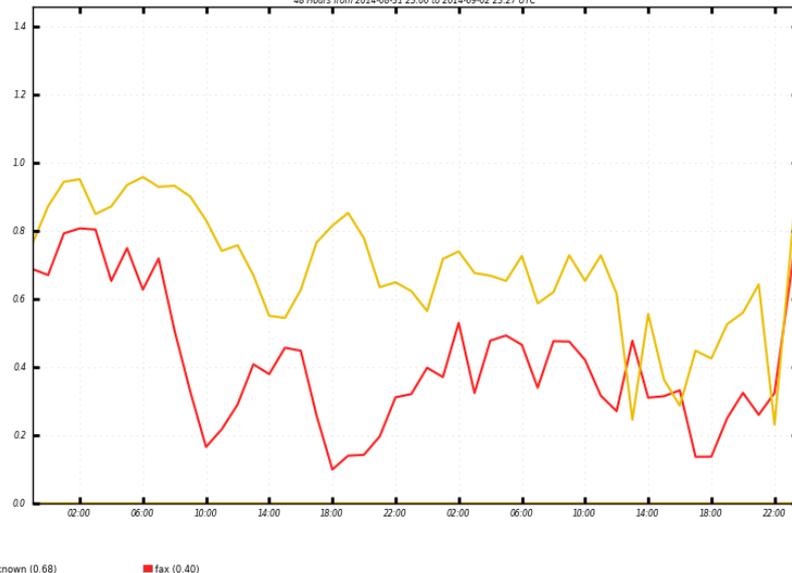
unknown fax

Maximum: 8,215 , Minimum: 4,777 , Average: 6,313 , Current: 5,737



Efficiency Good Jobs

48 Hours from 2014-08-31 23:00 to 2014-09-02 23:27 UTC



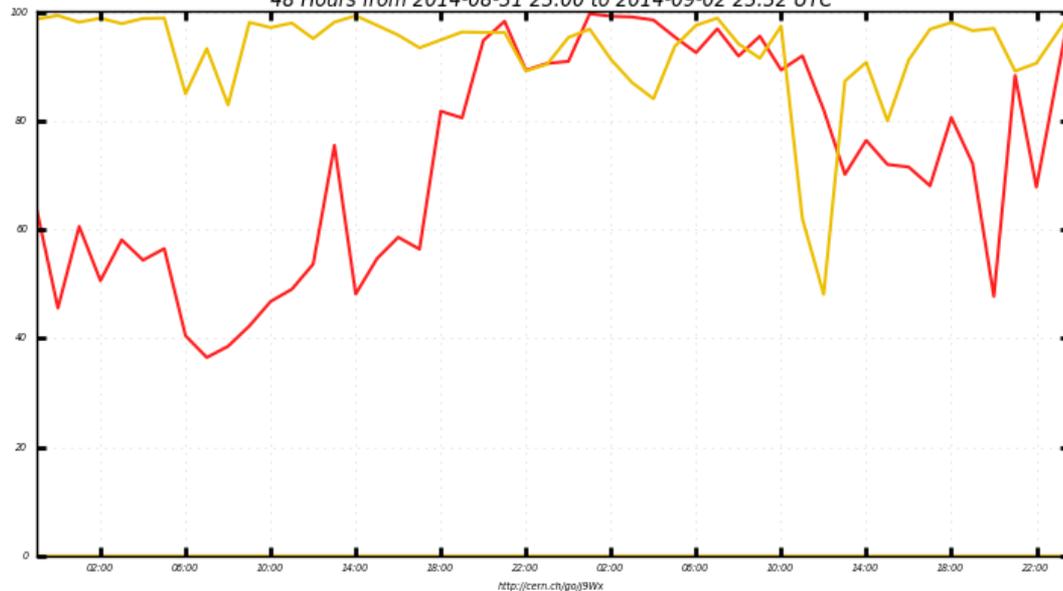
unknown (0.68) fax (0.40)

Total: 0.56 , Average Rate: 0.00 /s

# Comparing “overflow” and “regular”



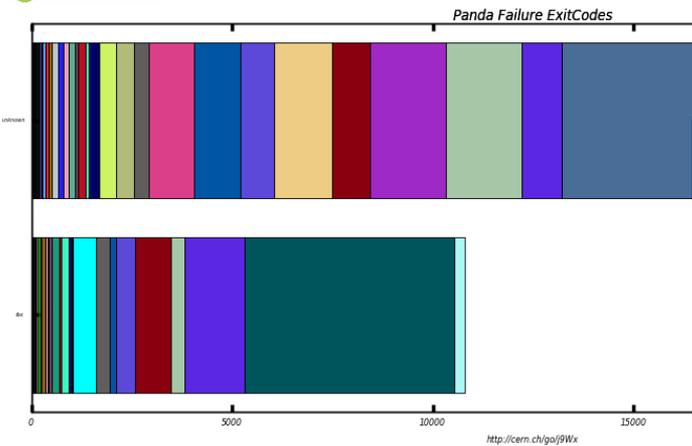
Efficiency over time based on success/all accomplished jobs  
48 Hours from 2014-08-31 23:00 to 2014-09-02 23:32 UTC



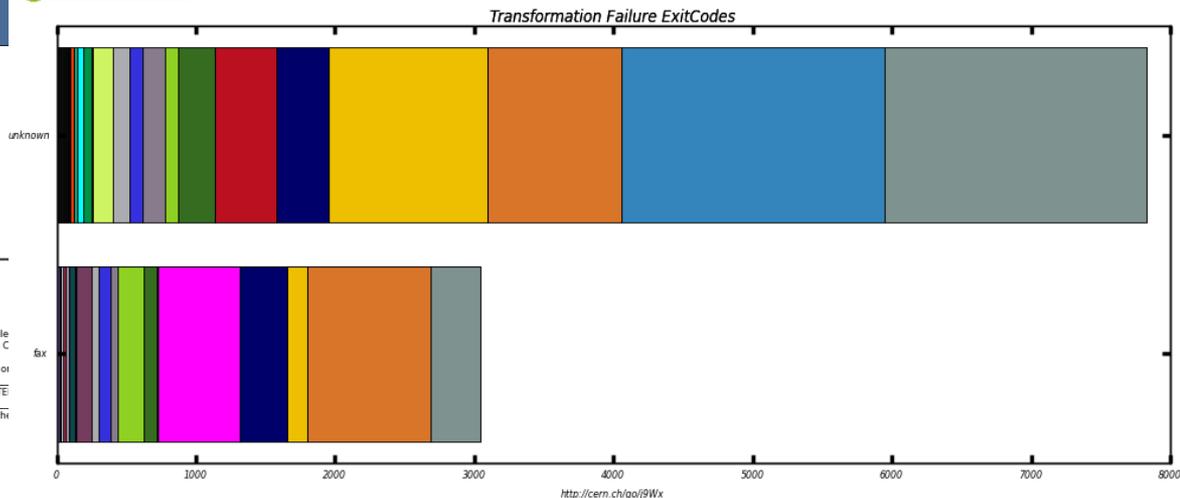
■ unknown (92.52) ■ fax (72.12)

Total: 158.34 , Average Rate: 0.00 /s

# Comparing “overflow” and “regular”

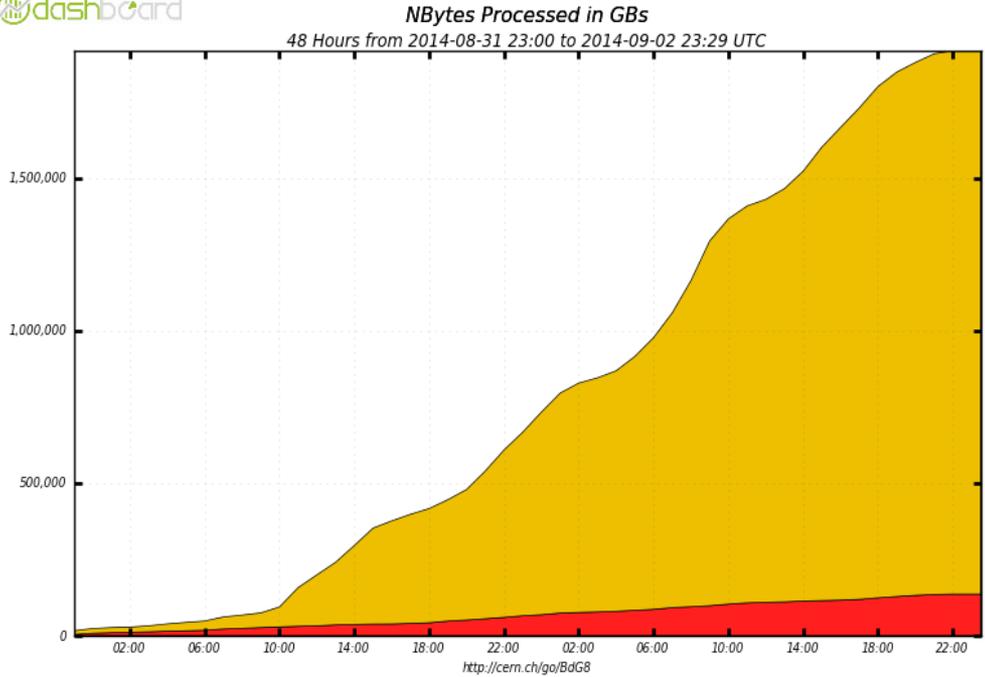


- Put error: Local output file missing
- Get error: Replica not found
- Unspecified error: consult log file
- Athena crash - consult log file
- Undocumented TaskBuffer Error Code : taskbuffer113
- Expired three days after submission
- Error found with transform input file\_\_\_Get error: Input file staging timed out
- Reached batch system time limit
- No events to be processed in the input file - skipEvents is higher than actual event number\_\_\_Looping job killed by pilot
- Get error: Input file staging timed out
- Looping job killed by pilot
- Get error: Staging input file failed
- Transformation not installed in C
- lost heartbeat
- Athena core dump or timeout, or Athena crash - consult log file
- Transform received signal SIGTERM
- Athena crash - consult log file
- plus 35 more



- Unspecified error, consult log file
- Athena crash - consult log file
- Athena crash - consult log file\_\_\_Looping job killed by pilot
- Transform received signal SIGTERM\_\_\_Reached batch system time limit
- No events to be processed in the input file - skipEvents is higher than actual event number\_\_\_Looping job killed by pilot
- analysis output merge crash - consult log file
- Athena core dump or timeout, or conddb DB connect exception\_\_\_Athena ran out of memory
- FAX error: site local redirector problems due to FAX functional test failures
- Unspecified error, consult log file\_\_\_Looping job killed by pilot
- Transformation not installed in CE
- Athena core dump or timeout, or conddb DB connect exception
- Error found with transform input file\_\_\_Get error: Input file staging timed out
- Athena crash - consult log file\_\_\_Athena ran out of memory
- No events to be processed in the input file - skipEvents is higher than actual event number\_\_\_cancelled
- Transform received signal SIGTERM\_\_\_Reached batch system time limit
- Transform received signal SIGTERM\_\_\_Looping job killed by pilot
- General failure in transform substep executor
- Not documented. Exitcode: 103
- plus 32 more

# Comparing “overflow” and “regular”



■ unknown (1,780,968)    ■ fax (136,229)

Total: 1,917,198 , Average Rate: 10.98 /s

# Comparison table

---

	<b>regular</b>	<b>overflow</b>
<b>jobs/day</b>	170k / 85.8%	28k / 14.2%
<b>job efficiency</b>	92%	72%
<b>cpu efficiency</b>	68%	40%

---

# Conclusions

---

- There is a need for it - enough jobs overflow even with very stringent limit on per file rate - still at 10MB/s.
  - Still not a serious load on the WAN links.
  - Some endpoints stability issues found and are being addressed.
  - Job's CPU efficiency will benefit a lot from switch to xAODs.
  - Two hours delay in getting cost matrix data to JEDI. This increases cost of and endpoint going down.
  - Needs debugging and tuning.
-