

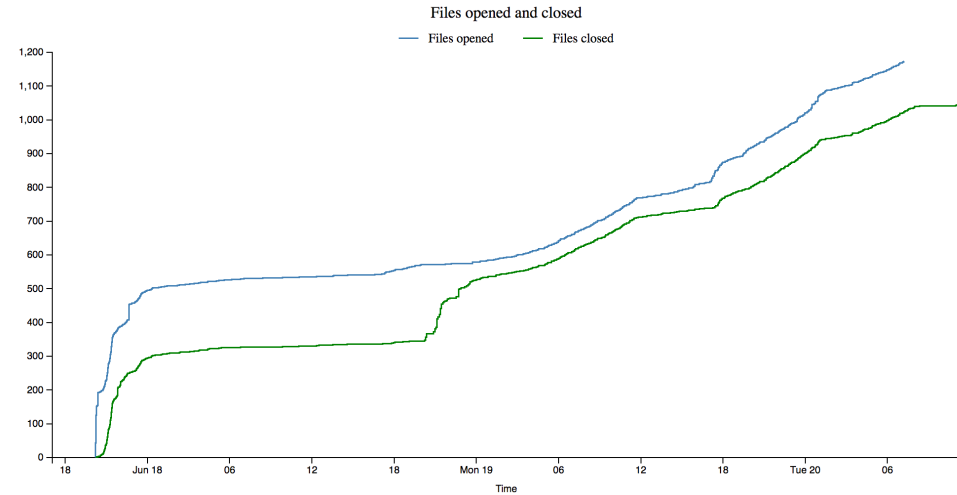

POMS in NOvA

Alex Himmel

FIFE Workshop

June 21st, 2017

Monitoring production jobs is complicated



Jobs: Fifemon

Files: Sam Station

- Need to handle two different “views”
 - “Jobs” which are the thing most tools are designed to monitor
 - “Files” which are the real measure of campaign status
- POMS has an ambitious goal: monitoring for both of these views together.
- What about submission through POMS?
 - Switching costs are prohibitive given an existing submission system.

POMS

Name	Campaign Stage ?	Active Jobs				Jobs in 7.0 days ending 2017-06-20 11:25				Stats	
		All	Total Idle ?	Running ?	Held ?	Total ?	Not Located ?	Completed Located ?	Removed ?	%Efficiency ?	Pending ?
NOvA Reconstruction Keepup fardet S16-12-07		915	0	0	1	914	458	456	0	90	-
NOvA Reconstruction Keepup neardet S16-12-07		141	0	2	2	137	45	92	0	93	-
fd_mcgen_cry_fhc_all		44248	0	0	0	44248	29738	14510	0	95	-
fd_mcgen_cry_fhc_all_draining		6731	0	0	0	6731	3421	3310	0	84	-
fd_data_period3		5250	0	0	0	5250	4558	692	0	76	-
prod_pid_caf_combo_R17-03-01-prod3reco.k_fd_num1_fhc_period5_v1_draining_bing		8	0	0	0	8	0	8	0	0	-
NOvA Raw2Root Keepup fardet S17-02-21		2478	0	4	0	2474	560	1914	0	25	-
NOvA Raw2Root Keepup neardet S17-02-21		138	0	0	0	138	46	92	0	23	-
prod_full_chain_fd_cry_fhc_all_draining		800	0	0	0	800	602	198	0	57	25
fd_data_period2		5000	1	203	2	4794	2630	2164	0	117	-
fd_data_period3_draining		250	0	133	32	85	24	61	0	6	0
prod_full_chain_fd_cry_fhc_all		30018	0	5	1	30012	7588	22424	0	81	-
fd_data_period1		5033	0	2140	2	2891	2671	220	0	19	-

POMS


Jobs

Jobs x Files

Name	Campaign Stage ?	All	Active Jobs			Jobs in 7.0 days ending 2017-06-20 11:25				Stats	
			Total ?	Running ?	Held ?	Total ?	Completed	Removed ?	%Efficiency ?		Pending ?
			Idle ?			Not Located ?	Located ?				
NOvA Reconstruction Keepup fardet S16-12-07		915	0	0	1	914	458	456	0	90	-
NOvA Reconstruction Keepup neardet S16-12-07		141	0	2	2	137	45	92	0	93	-
fd_mcgen_cry_fhc_all		44248	0	0	0	44248	29738	14510	0	95	-
fd_mcgen_cry_fhc_all_draining		6731	0	0	0	6731	3421	3310	0	84	-
fd_data_period3		5250	0	0	0	5250	4558	692	0	76	-
prod_pid_caf_combo_R17-03-01-prod3reco.k_fd_num1_fhc_period5_v1_draining_bing		8	0	0	0	8	0	8	0	0	-
NOvA Raw2Root Keepup fardet S17-02-21		2478	0	4	0	2474	560	1914	0	25	-
NOvA Raw2Root Keepup neardet S17-02-21		138	0	0	0	138	46	92	0	23	-
prod_full_chain_fd_cry_fhc_all_draining		800	0	0	0	800	602	198	0	57	25
fd_data_period2		5000	1	203	2	4794	2630	2164	0	117	-
fd_data_period3_draining		250	0	133	32	85	24	61	0	6	0
prod_full_chain_fd_cry_fhc_all		30018	0	5	1	30012	7588	22424	0	81	-
fd_data_period1		5033	0	2140	2	2891	2671	220	0	19	-

Benefits

- Gives a nice 30,000 ft overview of what is happening.
 - See previous slide.
- Allows you to start at the top level and drill all the way down to issues with individual jobs.
- Best tool I've found so far for studying job efficiency.
 - I just discovered this yesterday, actually.



Note: jobs without cpu time/wallclock time in Condor stats are listed as "Unk."

Issues

Proxy Error

The proxy server received an invalid response from an upstream server.
The proxy server could not handle the request [GET /poms/show_campaigns](#).

Reason: **Error reading from remote server**

Apache/2.2.15 (Scientific Linux) Server at pomsgpvm01.fnal.gov Port 443

- Performance
 - I get the error above constantly when interacting with POMS.
 - Even when working, pages can be very slow to load.
- The issue is likely scale...
 - previous page showed ~100,000 jobs.
- ...but this is the scale we operate at, and the tool needs to meet it.

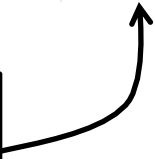
- Inconsistencies in numbers can make it difficult to trust what is reported.
- Problems with reporting from OSG
 - Sometimes jobs get evicted with no information and then re-started.
 - System needs to be robust enough to handle this situation and report it correctly.
- Misunderstanding column headings
 - I think some progress has happened on this recently
 - However, let me give an example that confused me for a while when preparing for this talk

Campaign Stage ?						Jobs in 7.0 days ending 2017-06-20 11:25				
Name	All	Total	Active Jobs			Completed			%Efficiency	Stats
		Idle ?	Running ?	Held ?	Total ?	Not Located ?	Located ?	Removed ?		
fd_mcgen_cry_fhc_all	44248	0	0	0	44248	29738	14510	0	95	-

“2/3 of the files were never located?!
What is happening with these jobs?”

Name	Campaign Stage ?	Jobs in 7.0 days ending 2017-06-20 11:25									
		All	Total	Active Jobs			Completed			Stats	
			Idle ?	Running ?	Held ?	Total ?	Not Located ?	Located ?	Removed ?	%Efficiency ?	Pending ?
fd_mcgen_cry_fhc_all		44248	0	0	0	44248	29738	14510	0	95	-

“2/3 of the files were never located?!
What is happening with these jobs?”



- I checked on these files directly in SAM.
- First, the input dataset only has 30k files
 - So 14k should end up “Not Located” even if everything is working correctly.
 - Extra jobs might be a mistake by our producer, or be restarts when jobs get evicted.
- Next, of those 30k input files, only 310 are missing outputs.
 - So, this sample is totally complete, but it doesn’t show up that way.
 - The key information of files, independent of jobs, is not there.

POMS is still Promising

- It is part of the way to something really great.
- We automatically include POMS hooks in our submission scripts, so its available for every job.
 - However, it is not really in regular use by our production team.
- Critical issues required for broader adoption:
 - Performance needs to be addressed.
 - If it is faster to look at both the SAM Station and Fifemon, that's what will happen.
 - We need to report the “right numbers” at the top level.
 - Again, if the most pertinent information isn't right up front, people will go to where it is.
- I have lots of cosmetic suggestions, too, but we should fix the fundamentals first!