

First look at PoD using AliEn

Martin Vala

November 18, 2010

Static

- Same set of machines
- Fast response
- Data can be saved locally

Dynamic on GRID (AliEn) using PoD

- Different set of machines (possible to get more workers)
- Slow response
- Data are processes from AliEn
- Maybe they can be cached for short time

- Pod - Multi PROOF master (one per user)
- CPU dynamically available, large pot of GRID cpus, every job can potentially become a PROOF node
- fits well with a 'diskless cluster' concept, i.e. the sotrage is independent

- Master at CERN
- Create workers at CERN, by sending jobs to CERN site
- requested 100 workers
 - in 5 minutes 10 workers
 - in 20 minutes all workers ready
- Process same dataset from CAF

Static Proof cluster - CAF (CAF is SE)

PROOF Query Progress: mvala@alice-caf.cern.ch

Executing on PROOF cluster "alice-caf.cern.ch" with 208 parallel workers:
Selector: AliAnalysisSelector
7284 files, number of events 34224567, starting event 0

100%

Initialization time: 13.0 secs
Processing time: 20 min 5 sec
Processed: 34224567 events (848.67 GB)
Processing rate: 28391.2 evts/sec (720.9 MB/sec)

Close dialog when processing is complete Smooth speedometer update

Show Logs Performance plot Memory Plot

Run in background Stop Cancel Close



The speedometer displays the processing rate in events per second (Ev/s). The needle is positioned at approximately 120.5 on a scale from 0 to 150. The scale is marked with 0, 37, 75, 112, and 150. The unit is $\times 10^3$. Below the needle, the text 'Proc Time [m s]' is displayed with the value '1205' and another $\times 10^3$ multiplier.

First tests PoD with AliEn (CAF is SE)

PROOF Query Progress: alicaf@lbsq1410.cern.ch

Executing on PROOF cluster "lbsq1410.cern.ch" with 61 parallel workers:
Selector: AllAnalysisSelector
7284 files, number of events 34224567, starting event 0

100%

Initialization time: 3.3 secs
Processing time: 21 min 0 sec
Processed: 34224567 events (841.38 GB)
Processing rate: 27160.5 evts/sec (683.7 MB/sec)

Close dialog when processing is complete Smooth speedometer update

Show Logs Performance plot Memory Plot

Run in background Stop Cancel Close



First tests PoD with AliEn

Machines status

Machine	Machine status						CPU				Networking		Disk space			
	Online	xproof	xrootd	cmsd	load1	# proof	usr	sys	nice	idle	IN	OUT	Total	Free	Used	%
1. lxbsq1409					2.5	2	9.769	0.846	0	89.16	6.517 MB/s	383.7 KB/s	-	-	-	-
2. lxbsq1410					-	-	-	-	-	-	-	-	-	-	-	-
3. lxbsq1411					2.5	3	17.45	1.376	0	69.73	202 KB/s	37.32 MB/s	2.498 TB	1.134 TB	1.364 TB	54.6
4. lxbsq1412					2.54	3	16.95	1.337	0	70.49	289.7 KB/s	34.52 MB/s	2.498 TB	1.1 TB	1.398 TB	55.97
5. lxbsq1413					1.55	3	13.31	1.292	0	77.75	1.677 MB/s	28.24 MB/s	2.498 TB	1.013 TB	1.485 TB	59.45
6. lxbsq1414					1.24	3	13.25	0.866	0	85.53	1.567 MB/s	274.9 KB/s	2.498 TB	1.186 TB	1.312 TB	52.51
7. lxbsq1415					-	-	-	-	-	-	-	-	-	-	-	-
8. lxbsq1416					2.26	3	12.48	1.293	0	78	1.374 MB/s	33.17 MB/s	2.498 TB	1.074 TB	1.424 TB	56.99
9. lxbsq1417					1.85	3	15.98	1.249	0	73.4	218.6 KB/s	38.42 MB/s	2.498 TB	1.065 TB	1.433 TB	57.36
10. lxbsq1418					2.5	3	15.39	1.317	0	70.11	274.2 KB/s	41.18 MB/s	2.498 TB	997 GB	1.524 TB	61.02
11. lxbsq1419					2.25	3	15.89	1.215	0	74.06	232 KB/s	35.53 MB/s	2.498 TB	1.05 TB	1.448 TB	57.98
12. lxbsq1420					2.3	3	14.78	1.217	0	74.32	133.1 KB/s	36.68 MB/s	2.498 TB	1.002 TB	1.496 TB	59.88
13. lxbsq1421					2.23	3	14.62	1.134	0	73.79	111 KB/s	31.5 MB/s	2.498 TB	1.031 TB	1.467 TB	58.73
14. lxbsq1422					2.46	3	14.21	1.241	0	72.11	233.2 KB/s	38.71 MB/s	2.498 TB	1.011 TB	1.486 TB	59.51
15. lxbsq1423					2.06	3	12.97	1.2	0	74.05	269.6 KB/s	32.75 MB/s	2.498 TB	1.084 TB	1.414 TB	56.61
16. lxfss13309					2.2	3	19.6	1.56	0	70.63	177.4 KB/s	38.38 MB/s	3.635 TB	2.228 TB	1.407 TB	38.7
17. lxfss13310					2.6	3	18.31	1.653	0	68.99	385.7 KB/s	43.1 MB/s	3.635 TB	2.278 TB	1.357 TB	37.34
18. lxfss13311					2.34	3	18.63	1.533	0	72.08	146.7 KB/s	35.85 MB/s	3.635 TB	2.256 TB	1.379 TB	37.93
19. lxfss13312					2.27	3	17.81	1.606	0	71.6	295.5 KB/s	39.66 MB/s	3.635 TB	2.22 TB	1.415 TB	38.93
20. lxfss13313					2.26	3	15.85	1.71	0	69.86	212.6 KB/s	44.47 MB/s	3.635 TB	2.223 TB	1.412 TB	38.83
21. lxfss13314					2.8	3	16.97	1.611	0	70.17	426.3 KB/s	45.62 MB/s	3.635 TB	2.187 TB	1.448 TB	39.84
22. lxfss13315					-	-	-	-	-	-	-	-	-	-	-	-
23. lxfss13316					1.98	3	17.65	1.512	0	72.23	223.3 KB/s	37.26 MB/s	3.635 TB	2.289 TB	1.346 TB	37.03
24. lxfss13317					1.89	2	16.79	1.379	0	74.31	204.3 KB/s	28.64 MB/s	3.635 TB	2.298 TB	1.337 TB	36.78
25. lxfss13318					2.44	2	17.55	1.367	0	73.89	237 KB/s	30.6 MB/s	3.635 TB	2.261 TB	1.374 TB	37.81
26. lxfss13401					2.17	2	16.29	1.503	0	73.52	161.2 KB/s	37.81 MB/s	3.635 TB	2.271 TB	1.364 TB	37.52
27. lxfss13402					2.42	2	16.24	1.546	0	72.12	253.4 KB/s	37.16 MB/s	3.635 TB	2.264 TB	1.371 TB	37.72
28. lxfss13403					2.07	2	15.08	1.422	0	74.98	306.4 KB/s	37.82 MB/s	3.635 TB	2.28 TB	1.355 TB	37.27
29. lxfss13404					2.17	2	14.07	1.116	0	77.38	177 KB/s	28.04 MB/s	3.635 TB	2.249 TB	1.386 TB	38.13
30. lxfss13405					2.41	2	14	1.368	0	74.4	364.9 KB/s	41.96 MB/s	3.635 TB	2.207 TB	1.428 TB	39.28
Total	27	27	27	27	58.76	73					16.54 MB/s	915 MB/s	80.86 TB	44.23 TB	36.63 TB	
Average	0.9	0.9	0.9	0.9	2.176	2.704	15.63	1.351	0	74.02	627.3 KB/s	33.89 MB/s	3.11 TB	1.701 TB	1.409 TB	

- User can create its own PoD PROOF cluster, by submitting job
- User will have
 - ALICE software by default
 - afdutils.C available (ds management)
- PROOF cluster reserved for 24 hours (idle 2 hours = PROOF killed)
- XROOTD for storage
- User can analyse file
 - remotely
 - ESD to AOD filter once
 - process its own AOD after that many times
 - 20 GB of space available

ALICE borrow me a Ferrari

- ALICE will create 1000 workers
- Filter ESD to AOD
- Save it in RAM (/dev/shm/)
- User will submit task (send it to ALICE Ferrari store)
 - Verify driver lesson (task will be tested on small subset of data with proof lite)
 - User can drive Ferrari (runs user's task)
 - Output will be send to SE and user will get list of files (for merging)