SKAF - Slovak Košice Analysis Facility

Martin Vaľa

Introduction to FAST ANALYSIS

- Fast Analysis
  - process same data many time with different cut parameters
  - fast processing on reasonable statistics
  - 1-10 milion of events
- Requriments
  - fast access to workers
  - fast reading of big input
  - splitting analysis to more workers
  - fast merging of output (usually histograms)
  - fast access to your final histograms
- AAF - ALICE Analysis Facitilites (CAF, SKAF, ...)

Analysis - GRID vs. PROOF

	GRID	PROOF
Mode	batch	interactive
Workers at same time	∼1000	∼100[1]
Start time	∼20 minutes[2]	∼1 minute
Short job (3 minute analysis)	∼30-60 minutes	∼5 minutes
Input reading	remotely	local disks
Network speed	1 Gbit/s	10 Gbit/s[3]
Split analysis	file (300 ev)	event
Output merging	manualy	automatic
Calibration and fast analysis	-	+
Final analysis (full statistics)	+	-[4]

[1] depends on AAF
[2] depends on queue
[3] for fast file staging on local disk
[4] depends on cache storage

History of AAF

- PROOF Analysis Facility History
  - CAF - first setup (since May 2006)
  - SKAF - new AF build from 5 AliEn workers (Jan 2010)
  - SKAF setup
    - new setup (different then CAF) (Jan 2010)
    - setup was ported to CAF in (May 2010)
    - AAF - ALICE Analysis Facilities
  - SKAF upgrade to 15x4 core machines (May 2010)
- Current AAF proof clusters
  - CAF (116/464)
  - SKAF (60/60)
  - KIAF (48/48)

Available PROOF clusters for all ALICE users

	Cluster					ROOT	Aggregated disk space			AF xrootd		xrootd
Name	Online	Status	Proof master	Workers	Users	Version	Total	Free	Used	Running	Latest	Version
1. CAF		Stable	alice-caf.cern.ch	116	5	v5-28-00d	162.4 TB	20.19 TB	142.2 TB	1.0.43	1.0.43	20100510-1509_dbg
2. KIAF		Stable	kiaf.sdfarm.kr	48	1	v5-28-00d	8.456 TB	6.926 TB	1.53 TB	1.0.43	1.0.43	20100510-1509_dbg
3. SKAF		Stable	skaf.saske.sk	60	0	v5-28-00d	53.72 TB	1.467 TB	52.26 TB	1.0.43	1.0.43	20100510-1509_dbg
Total				**224**	**6**		**224.6 TB**	**28.58 TB**	**196 TB**			

Data read last 6 months

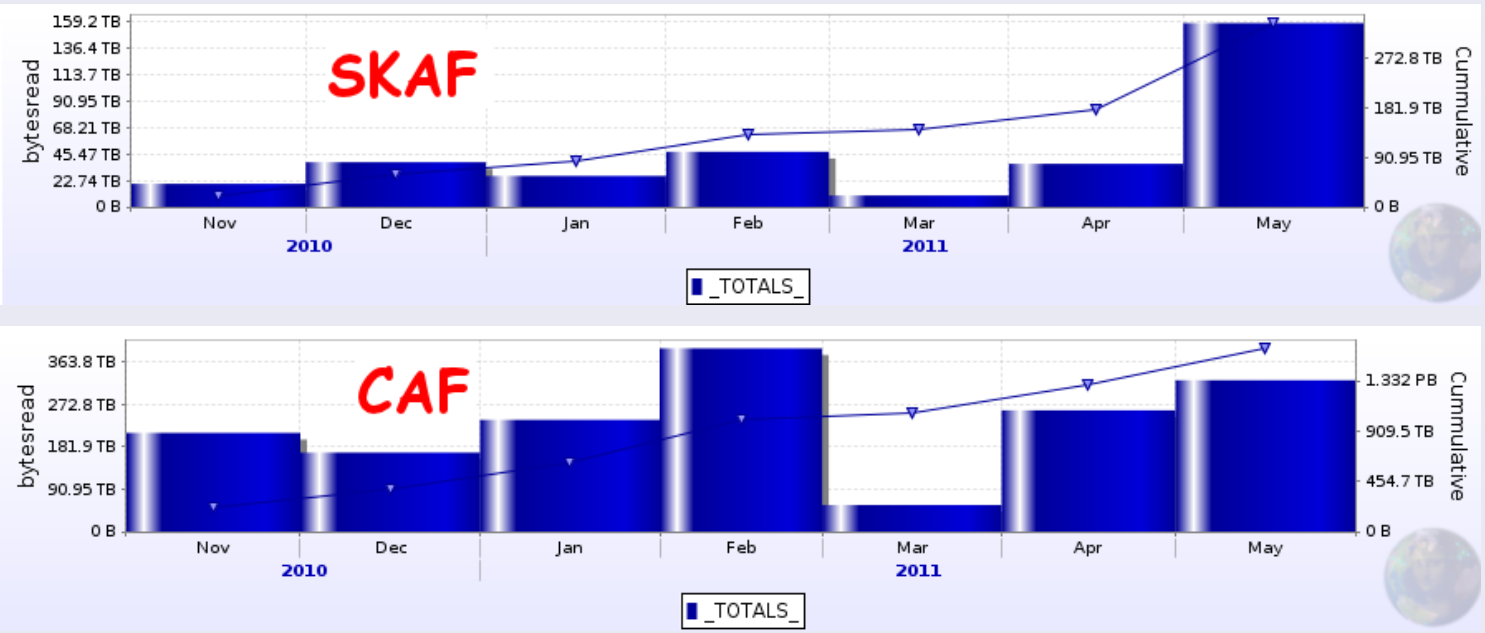

What is AAF made of?

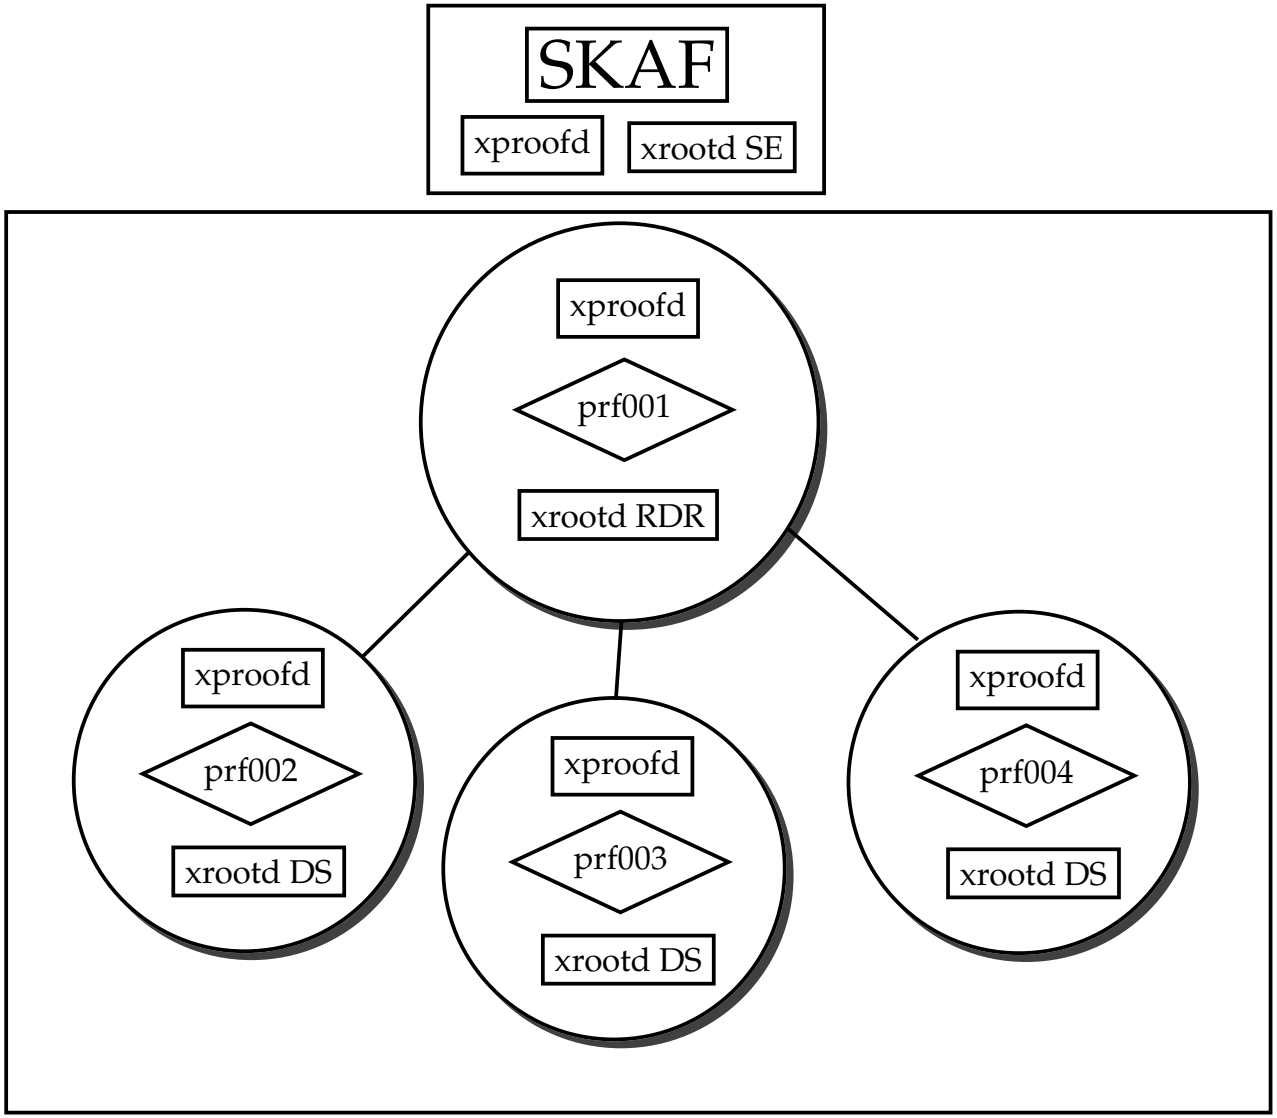

AAF performance (Download speed)

SKAF

	Machine status						CPU				Networking		Disk space			
Machine	Online	xproofd	xrootd	cmsd	load1	# proof	usr	sys	nice	idle	IN	OUT	Total	Free	Used	%
1. prf000-iep-grid					1.15	0	2.572	1.72	0	84.84	5.26 Mbps	68.25 Kbps	-	-	-	-
2. prf001-iep-grid					2.89	0	3.744	2.389	0	68.02	4.506 Mbps	69.07 Kbps	3.582 TB	2.013 TB	1.569 TB	43.81
3. prf002-iep-grid					2	0	4.134	2.723	0	66.56	5.459 Mbps	79.9 Kbps	3.582 TB	2.021 TB	1.56 TB	43.57
4. prf003-iep-grid					2.49	0	4.36	2.767	0	71.78	6.122 Mbps	86.56 Kbps	3.582 TB	2.012 TB	1.569 TB	43.81
5. prf004-iep-grid					1.7	0	3.275	2.227	0	75.8	5.171 Mbps	74.93 Kbps	3.582 TB	2.02 TB	1.562 TB	43.61
6. prf005-iep-grid					1.58	0	3.596	2.363	0	70.02	4.471 Mbps	65.45 Kbps	3.582 TB	2.028 TB	1.554 TB	43.38
7. prf006-iep-grid					1.32	0	3.586	2.531	0	70.74	6.082 Mbps	86.89 Kbps	3.582 TB	2.01 TB	1.572 TB	43.88
8. prf007-iep-grid					2.03	0	4.275	2.762	0	71.63	4.471 Mbps	65.76 Kbps	3.582 TB	2.031 TB	1.551 TB	43.3
9. prf008-iep-grid					2.47	0	4.652	3.02	0	64.13	4.547 Mbps	70.08 Kbps	3.582 TB	2.194 TB	1.388 TB	38.75
10. prf009-iep-grid					1.3	0	4.047	2.737	0	73.5	4.674 Mbps	69.68 Kbps	3.582 TB	2.036 TB	1.546 TB	43.16
11. prf010-iep-grid					2.39	0	3.2	2.141	0	69.4	3.329 Mbps	52.33 Kbps	3.582 TB	2.018 TB	1.564 TB	43.67
12. prf011-iep-grid					1.76	0	3.989	2.853	0	75.5	5.966 Mbps	90.12 Kbps	3.582 TB	2.026 TB	1.555 TB	43.43
13. prf012-iep-grid					1.92	0	3.815	2.554	0	74.42	3.511 Mbps	56.94 Kbps	3.582 TB	2.036 TB	1.546 TB	43.16
14. prf013-iep-grid					2	0	3.963	2.809	0	66.91	4.769 Mbps	72.68 Kbps	3.582 TB	2.025 TB	1.557 TB	43.46
15. prf014-iep-grid					1.47	0	3.961	2.559	0	70.06	4.407 Mbps	67.19 Kbps	3.582 TB	2.055 TB	1.527 TB	42.63
16. prf015-iep-grid					1.22	0	3.894	2.647	0	75.89	4.979 Mbps	74.28 Kbps	3.582 TB	1.884 TB	1.697 TB	47.39
Total	**16**	**16**	**16**	**16**	**29.69**	**0**					**77.72 Mbps**	**1.123 Mbps**	**53.72 TB**	**30.41 TB**	**23.32 TB**	
Average	**1**	**1**	**1**	**1**	**1.856**	**0**	**3.816**	**2.55**	**0**	**71.83**	**4.858 Mbps**	**71.88 Kbps**	**3.582 TB**	**2.027 TB**	**1.554 TB**	

CAF

56. lxfssl3401					2.73	2	22.68	1.528	0	69.76	17.2 MB/s	3.483 MB/s	3.635 TB	192 GB	3.447 TB	94.84
57. lxfssl3402					2.46	2	24.67	1.833	0	70.58	6.057 MB/s	2.715 MB/s	3.635 TB	188 GB	3.451 TB	94.95
58. lxfssl3403					2.23	2	23.52	1.066	0	72.28	5.853 MB/s	3.124 MB/s	3.635 TB	191.1 GB	3.448 TB	94.87
59. lxfssl3404					2.47	2	23.51	1.103	0	72.02	4.562 MB/s	3.544 MB/s	3.635 TB	187.9 GB	3.451 TB	94.95
60. lxfssl3405					2.53	2	24.49	1.879	0	69.43	12.84 MB/s	3.394 MB/s	3.635 TB	185.5 GB	3.454 TB	95.02
Total	**59**	**59**	**59**	**59**	**162.5**	**117**					**1.254 GB/s**	**92.43 MB/s**	**162.4 TB**	**35.45 TB**	**126.9 TB**	
Average	**0.983**	**0.983**	**0.983**	**0.983**	**2.754**	**1.983**	**18.84**	**2.044**	**0**	**72.06**	**21.77 MB/s**	**1.567 MB/s**	**2.799 TB**	**625.8 GB**	**2.188 TB**	

AAF performance (Download speed) [2]

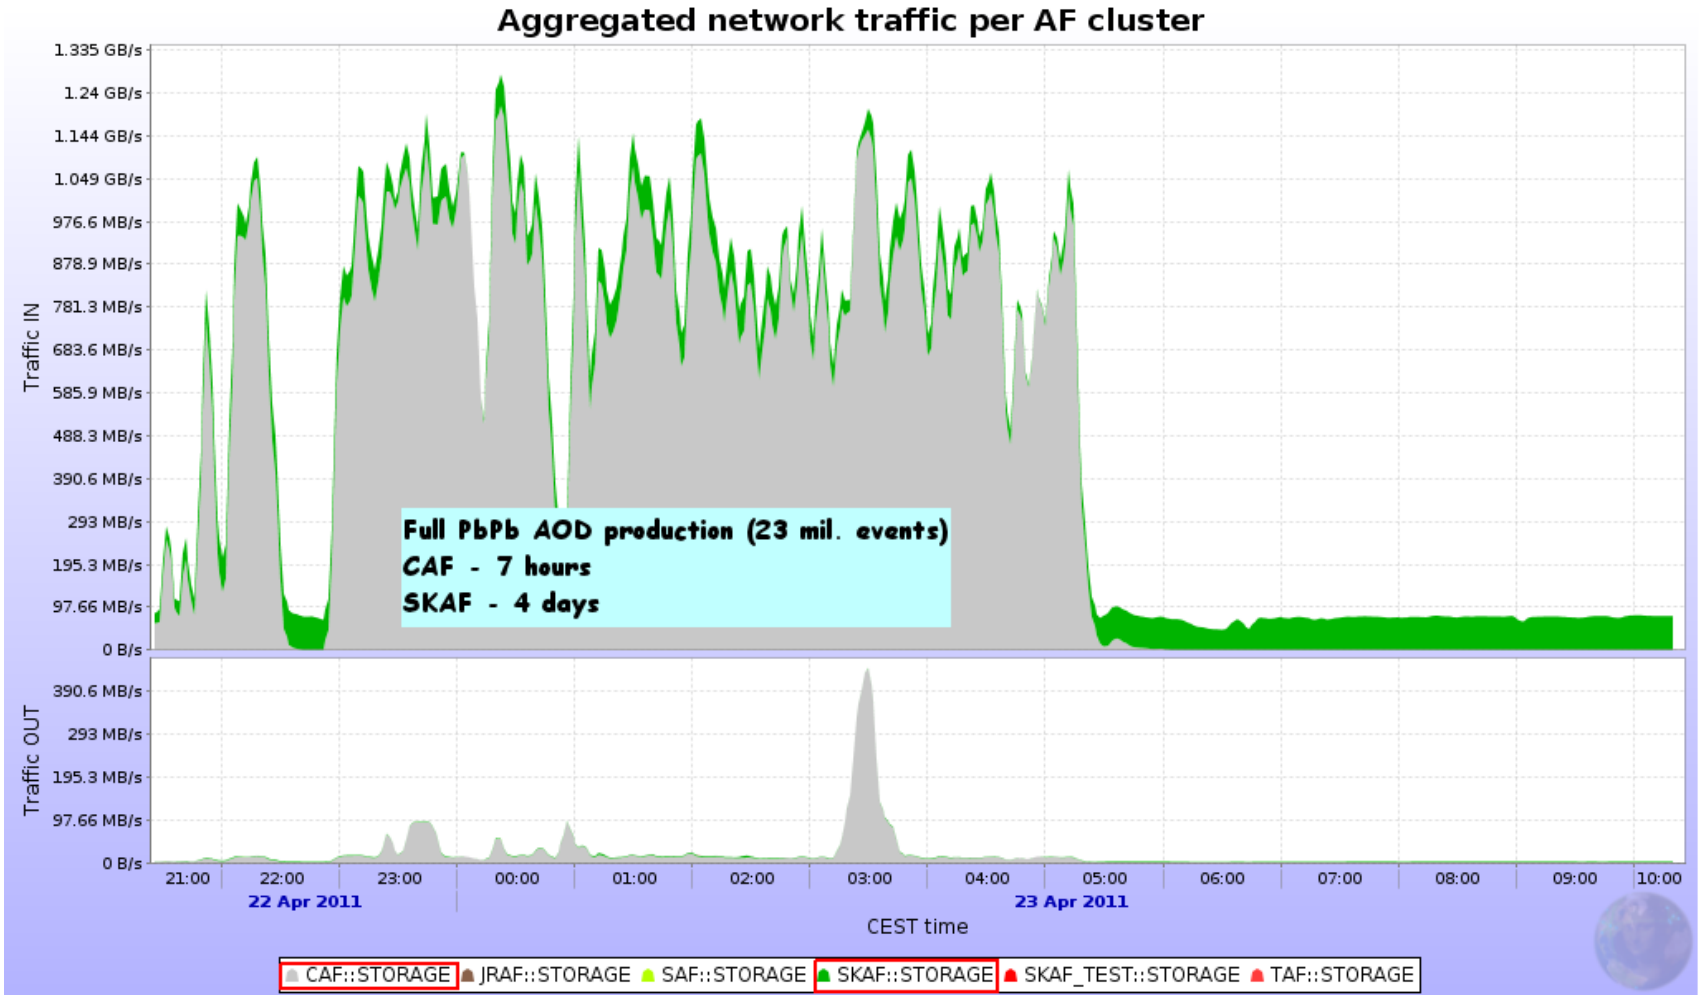

AAF performance (Processing speed)

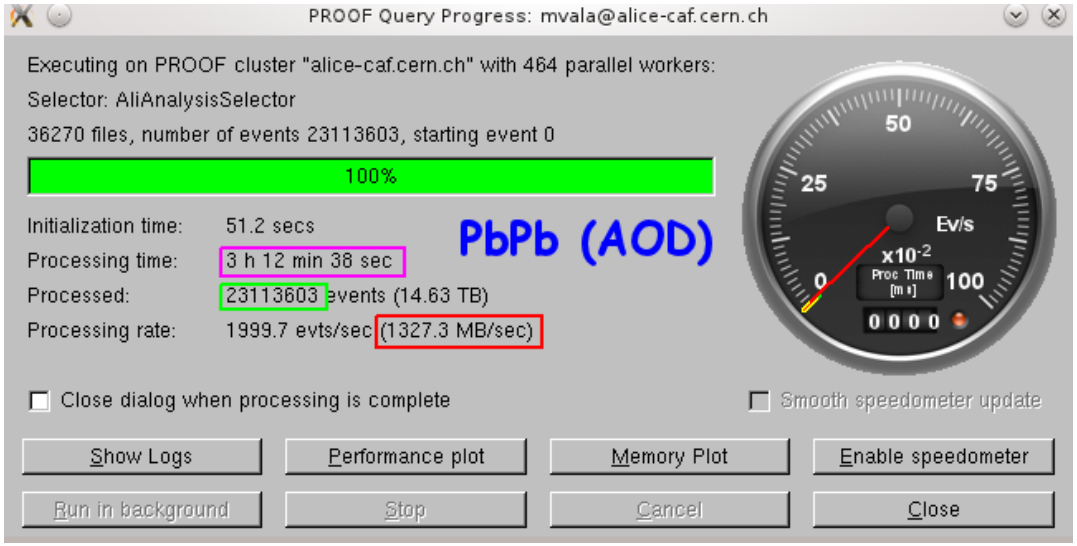

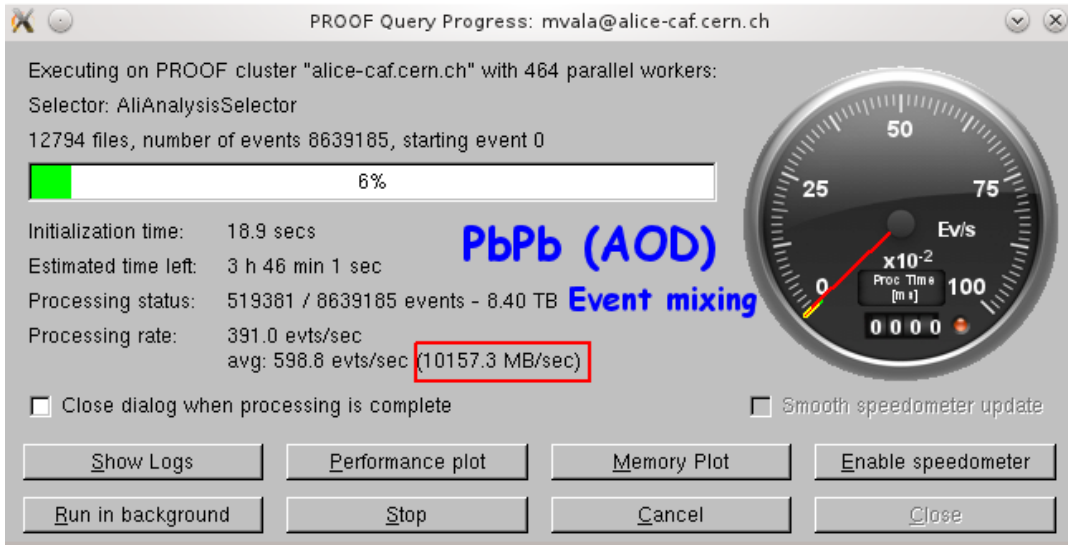

SKAF upgrade needed

- 8 core machines [4 core]
- 3 GB/core RAM (24 GB per machine) [2 GB/core]
- 4 disk solution (1 x system + 3 x data) [2 data disks]
  - 3 x 2TB = 6TB per machine [2 x 2TB = 4TB]
- 10 Gbit/s network connection [1 GBit/s]
- PROOF worker is 1.5x more expensive then GRID worker
- more financial support needed
  - PROOF cluster [100 000 EUR]
  - Maintenance [10 000 EUR per year]

Backup slides

SKAF overloaded + not enough RAM memory

	Machine status						CPU				Networking		Disk space			
Machine	Online	xproofd	xrootd	cmsd	load1	# proof	usr	sys	nice	idle	IN	OUT	Total	Free	Used	%
1. prf000-iep-grid					0.18	2	3.035	0.338	0	96.51	21.48 Kbps	530.7 Kbps	-	-	-	-
2. prf001-iep-grid					4.74	5	97.69	1.48	0	0.04	43.69 Kbps	7.221 Kbps	3.582 TB	97.88 GB	3.486 TB	97.33
3. prf002-iep-grid					4.29	5	96.87	1.232	0	0.361	16.39 Kbps	0.834 Kbps	3.582 TB	97.55 GB	3.486 TB	97.34
4. prf003-iep-grid					4.75	5	98.24	0.757	0	0.177	44.09 Kbps	1.61 Kbps	3.582 TB	105.9 GB	3.478 TB	97.11
5. prf004-iep-grid					4.06	5	98.41	0.299	0	0.203	16.19 Kbps	0.745 Kbps	3.582 TB	98.76 GB	3.485 TB	97.31
6. prf005-iep-grid					4.26	5	97.67	1.024	0	0.294	7.459 Kbps	0.472 Kbps	3.582 TB	100 GB	3.484 TB	97.27
7. prf006-iep-grid					4.3	5	98.24	0.888	0	0.084	43.11 Kbps	1.594 Kbps	3.582 TB	97.26 GB	3.487 TB	97.35
8. prf007-iep-grid					5.65	5	99.24	0.696	0	0	42.77 Kbps	1.577 Kbps	3.582 TB	97.62 GB	3.486 TB	97.34
9. prf008-iep-grid					5.9	5	97.85	1.284	0	0.104	31.88 Kbps	1.295 Kbps	3.582 TB	100.5 GB	3.483 TB	97.26
10. prf009-iep-grid					5.66	5	99.19	0.709	0	0	26.25 Kbps	1.327 Kbps	3.582 TB	97.08 GB	3.487 TB	97.35
11. prf010-iep-grid					5.11	5	99.7	0.224	0	0.002	33.25 Kbps	1.208 Kbps	3.582 TB	99.43 GB	3.484 TB	97.29
12. prf011-iep-grid					4.2	5	99	0.357	0	0.015	41.52 Kbps	1.412 Kbps	3.582 TB	97.49 GB	3.486 TB	97.34
13. prf012-iep-grid					4.9	5	97.66	1.167	0	0.315	42.27 Kbps	1.552 Kbps	3.582 TB	100.8 GB	3.483 TB	97.25
14. prf013-iep-grid					4.38	5	98.67	0.523	0	0.093	42.15 Kbps	1.459 Kbps	3.582 TB	97.56 GB	3.486 TB	97.34
15. prf014-iep-grid					4.71	5	97.04	1.785	0	0.279	42.53 Kbps	1.493 Kbps	3.582 TB	97.91 GB	3.486 TB	97.33
16. prf015-iep-grid					4.1	5	97.91	1.051	0	0.316	0.858 Kbps	0.424 Kbps	3.582 TB	97.55 GB	3.486 TB	97.34
Total	**16**	**16**	**16**	**16**	**71.19**	**77**					**495.9 Kbps**	**554.9 Kbps**	**53.72 TB**	**1.449 TB**	**52.28 TB**	
Average	**1**	**1**	**1**	**1**	**4.449**	**4.813**	**92.28**	**0.863**	**0**	**6.175**	**30.99 Kbps**	**34.68 Kbps**	**3.582 TB**	**98.89 GB**	**3.485 TB**	

What is this about?

Memory status

	Memory					Swap		
Machine	Total	Used	Cached	Buffers	Free	Total	Used	Free
1. prf000-iep-grid	22.02 GB	1.635 GB	12.95 GB	925.8 MB	20.38 GB	29.76 GB	0.18 MB	29.76 GB
2. prf001-iep-grid	7.307 GB	6.171 GB	797.5 MB	9.41 MB	1.136 GB	16 GB	10.83 GB	5.171 GB
3. prf002-iep-grid	7.307 GB	5.614 GB	1.641 GB	11.21 MB	1.693 GB	16 GB	10.75 GB	5.253 GB
4. prf003-iep-grid	7.307 GB	6.284 GB	969.9 MB	10.3 MB	1.023 GB	16 GB	10.97 GB	5.032 GB
5. prf004-iep-grid	7.307 GB	6.113 GB	1.149 GB	7.313 MB	1.194 GB	16 GB	10.87 GB	5.13 GB
6. prf005-iep-grid	7.307 GB	5.774 GB	1.481 GB	9.535 MB	1.533 GB	16 GB	10.87 GB	5.13 GB
7. prf006-iep-grid	7.307 GB	6.182 GB	1.025 GB	8.695 MB	1.125 GB	16 GB	11.56 GB	4.446 GB
8. prf007-iep-grid	7.05 GB	5.674 GB	1.309 GB	11.21 MB	1.375 GB	16 GB	6.283 GB	9.72 GB
9. prf008-iep-grid	7.05 GB	4.895 GB	2.039 GB	26.86 MB	2.155 GB	16 GB	7.168 GB	8.834 GB
10. prf009-iep-grid	7.05 GB	5.119 GB	1.834 GB	22.42 MB	1.931 GB	16 GB	6.779 GB	9.224 GB
11. prf010-iep-grid	7.05 GB	5.27 GB	1.718 GB	22.36 MB	1.78 GB	16 GB	6.905 GB	9.097 GB
12. prf011-iep-grid	7.05 GB	5.237 GB	1.705 GB	20.21 MB	1.813 GB	16 GB	9.035 GB	6.968 GB
13. prf012-iep-grid	7.05 GB	5.399 GB	1.323 GB	10.71 MB	1.651 GB	16 GB	9.459 GB	6.544 GB
14. prf013-iep-grid	7.05 GB	5.789 GB	1.172 GB	9.758 MB	1.261 GB	16 GB	8.879 GB	7.124 GB
15. prf014-iep-grid	7.05 GB	5.726 GB	1.275 GB	9.477 MB	1.324 GB	16 GB	8.79 GB	7.212 GB
16. prf015-iep-grid	7.05 GB	5.184 GB	1.807 GB	19.04 MB	1.866 GB	16 GB	9.011 GB	6.991 GB
Total	**129.3 GB**	**86.06 GB**	**34.15 GB**	**1.108 GB**	**43.25 GB**	**269.8 GB**	**138.2 GB**	**131.6 GB**
Average	**8.082 GB**	**5.379 GB**	**2.134 GB**	**70.89 MB**	**2.703 GB**	**16.86 GB**	**8.635 GB**	**8.227 GB**

CAF performance

- Processing data 10 Gbytes/s (event mixing)
- Staging new data 750 MBytes/s

40. lxfssi3310					9.98	8	94.53	3.169	0	0.311	20.73 MB/s	352.9 KB/s	2.554 TB	1.424 TB	1.13 TB	44.25
41. lxfssi3311					9.72	8	95.04	2.843	0	0.341	16.36 MB/s	202.8 KB/s	2.554 TB	1.427 TB	1.127 TB	44.14
42. lxfssi3401					10.39	9	87.03	11.84	0	0.073	11 MB/s	190.6 KB/s	2.554 TB	1.424 TB	1.13 TB	44.25
43. lxfssi3402					9.94	8	95.42	2.538	0	0.314	15.69 MB/s	289.5 KB/s	2.554 TB	1.328 TB	1.226 TB	47.99
44. lxfssi3403					9.19	8	95.78	2.617	0	0.125	11.04 MB/s	209.7 KB/s	2.554 TB	1.429 TB	1.125 TB	44.06
45. lxfssi3404					10.29	8	84.72	11.58	0	0.025	15.8 MB/s	234.3 KB/s	2.554 TB	1.423 TB	1.131 TB	44.29
46. lxfssl3309					9.87	8	89.09	9.808	0	0.095	4.735 MB/s	77.63 KB/s	3.635 TB	190 GB	3.449 TB	94.89
47. lxfssl3310					9.27	8	96.26	2.316	0	0.047	12.59 MB/s	144.9 KB/s	3.635 TB	200.2 GB	3.439 TB	94.62
48. lxfssl3311					9.62	8	93.05	2.924	0	1.283	28.86 MB/s	768 KB/s	3.635 TB	203.5 GB	3.436 TB	94.53
49. lxfssl3312					9.04	8	97.44	1.519	0	0.056	2.161 MB/s	46.7 KB/s	3.635 TB	186.9 GB	3.452 TB	94.98
50. lxfssl3313					9.19	8	96.09	2.68	0	0.061	11.14 MB/s	131.6 KB/s	3.635 TB	190.3 GB	3.449 TB	94.89
51. lxfssl3314					9.21	8	96.45	2.393	0	0.108	8.391 MB/s	124.3 KB/s	3.635 TB	199 GB	3.441 TB	94.65
52. lxfssl3315					9.73	8	94.02	4.508	0	0.047	12.04 MB/s	298.7 KB/s	3.523 TB	1.682 TB	1.84 TB	52.24
53. lxfssl3316					9.15	8	96.81	1.981	0	0.021	4.909 MB/s	78.22 KB/s	3.635 TB	188.2 GB	3.451 TB	94.94
54. lxfssl3317					9.34	8	95.21	2.031	0	0.767	9.835 MB/s	111.2 KB/s	3.635 TB	201.3 GB	3.438 TB	94.59
55. lxfssl3318					9.6	8	91.69	2.753	0	2.022	26.1 MB/s	314.1 KB/s	3.635 TB	198.3 GB	3.441 TB	94.67
56. lxfssl3401					9.43	8	95.21	2.31	0	0.391	16.94 MB/s	381.3 KB/s	3.635 TB	198.8 GB	3.441 TB	94.66
57. lxfssl3402					8.79	8	95.83	2.151	0	0.157	7.058 MB/s	130.6 KB/s	3.635 TB	190.8 GB	3.449 TB	94.87
58. lxfssl3403					9.32	8	96.61	2.049	0	0.067	8.589 MB/s	110.3 KB/s	3.635 TB	193.9 GB	3.446 TB	94.79
59. lxfssl3404					9.36	8	96.13	2.152	0	0.155	14.33 MB/s	243.3 KB/s	3.635 TB	194.9 GB	3.445 TB	94.76
60. lxfssl3405					9.18	8	94.87	2.414	0	0.592	8.695 MB/s	160.5 KB/s	3.635 TB	199.1 GB	3.441 TB	94.65
Total	**59**	**59**	**59**	**59**	**566.3**	**472**					**746.4 MB/s**	**10.19 MB/s**	**162.4 TB**	**33.47 TB**	**128.9 TB**	
Average	**0.983**	**0.983**	**0.983**	**0.983**	**9.598**	**8**	**91.84**	**4.367**	**0**	**2.057**	**12.65 MB/s**	**176.8 KB/s**	**2.799 TB**	**590.8 GB**	**2.223 TB**	