



PROOF Cluster Management in ALICE

Jan Fiete Grosse-Oetringhaus, CERN PH/ALICE

CAF / PROOF Workshop, 29.11.07




CAF Test Setup




- Test setup since May 2006
 - 40 machines, 2 CPUs each, 200 GB disk
 - Normal lxbatch machines (taken out of ALICE's LSF share)
 - 5 machines as development partition
 - 35 machines as production partition
- Machines are a xrootd disk pool
 - xrootd redirector and PROOF master on head node
 - Other nodes are xrootd disk servers and PROOF slaves

Installation




- ROOT (including xrootd) installed as RPM (deployed by IT, QUATTOR)
 - tar files for quick updates
 - ROOT version including sources also installed on AFS
 - As reference
 - For users running on lxplus
- AliRoot built against ROOT version on CAF, installed to AFS

Configuration



- For each partition (development, production)
 - xrootd/PROOF configuration (3 files) on AFS (public, no token needed!)
 - Staging script + alien api commands (runs on all slaves) on AFS
 - AliEn user certificate deployed to each machine (to read files from AliEn)
- Starting/ Stopping of services
 - Using wassh tool

Monitoring




- Usual lemon monitoring by IT
- ALICE monitoring powered by MonaLISA
- Machines
 - CPU, Memory, Network
 - Alerts
- Queries → CPU quotas
- Staging → Disk quotas
- Detailed query monitoring (to be enabled in query)
 - CPU, disk used by single queries
 - Network traffic correlation

Machine	Machine status					
	Online	xrootd	olbd	Load	Proof users	Disk
1. lxb6041.cern.ch	Green	Green	Green	0	0	Red
2. lxb6042.cern.ch	Green	Green	Green	0.08	0	White
3. lxb6043.cern.ch	Green	Green	Green	0.02	0	White
4. lxb6044.cern.ch	Green	Green	Green	0.33	0	White
5. lxb6045.cern.ch	Green	Red	Green	0	0	White
6. lxb6046.cern.ch	Green	Green	Green	0.11	2	Blue
7. lxb6047.cern.ch	Green	Green	Green	0.02	1	Red
8. lxb6048.cern.ch	Green	Green	Green	0.09	1	Red

Traffic between the cluster machines (MB/sec) (last 0.5h average)																							
Machine	6047	6048	6049	6050	6052	6053	6054	6055	6056	6057	6058	6059	6060	6061	6062	6063	6064	6065	6066	6067	6068	6069	607
1. 6047	0	-	-	-	-	-	2.927	2.018	-	-	1.094	-	-	-	1.908	4.112	-	-	0.974	0.614	0	0	-
2. 6048	-	9.406	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
3. 6049	-	-	8.678	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
4. 6050	-	-	-	6.692	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
5. 6052	-	-	-	-	3.913	-	1.454	-	-	-	-	3.084	-	0.317	0	0	-	0	-	-	0.985	4.447	-
6. 6053	-	-	-	-	-	6.603	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7. 6054	0	-	1.363	-	-	-	6.195	-	-	-	0	-	-	-	0	-	-	-	-	-	0	-	1.68
8. 6055	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
9. 6056	-	-	-	-	-	-	-	-	4.962	-	2.442	0.525	-	-	-	-	-	-	-	-	-	-	-
10. 6057	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
11. 6058	1.164	-	-	-	0	-	-	2.531	-	0	0	-	-	-	-	1.103	-	-	0	-	-	-	-
12. 6059	3.755	-	0.622	-	-	-	-	-	-	-	-	11.76	1.955	0	0.677	1.848	0	-	-	-	2.812	-	0.78
13. 6060	-	-	-	-	-	-	-	2.068	-	-	-	11.59	-	-	-	1.06	-	-	-	-	-	-	2.02
14. 6061	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
15. 6062	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
16. 6063	-	-	-	-	1.655	0.27	2.416	-	-	-	-	-	0	-	6.38	-	0	-	0	-	-	-	-
17. 6064	-	-	-	-	-	1.123	-	2.822	-	-	-	1.621	-	0	-	3.117	-	0	0	-	-	-	0.58
18. 6065	0	-	-	-	3.52	3.165	-	0	-	-	0	-	-	-	-	3.034	0	1.579	-	0	-	-	-
19. 6066	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Experience



- xrootd/olbd configuration sometimes requires a bit of trying, calm reflection and thinking outside of the box
→ in the end everything is logical 😊
- Certain turn-around time for RPMs, therefore "manual" deployment is advisable at the beginning
- Restarting kills all user sessions
 - Makes updates difficult in a production system
- Conclusion: Administration and user support always takes more time than one thinks...