



Lustre 2.0 Alpha Summary

Goal & Approach

The goal of the first Alpha release of Lustre 2.0 is to demonstrate basic stability of Lustre on a single platform/distro. For this release RHEL5/x86_64 was selected because it is the most often downloaded distribution from lustre.org. To achieve our goal, daily acc-sm runs will be executed on small test clusters and bug fixing priority will focus on fixing bugs which prevent clean acc-sm runs. Additionally, IOR and simul runs will be executed on a modest cluster (approx 75 clients) to verify basic stability. During this Alpha phase, only bug fixes advancing stability will be landed, all other HEAD landing requests will be held and deferred to a later landing date post Alpha.

Timeline

The start time for Alpha is somewhat arbitrarily set at March 1, 2009 and the end date for Alpha is set for April 16, 2009.

Fixes Landed

Below is a summary listing of all unique landings to HEAD during the Lustre 2.0 Alpha Milestone period. This is inclusive of all landings -- code fixes as well as test framework fixes.

1. 17954 some coverity fixes against b1_6 *Attachment #20756, Attachment #20757, Attachment #21931*
2. 11404 LustreError: 5482:0:(lu_object.c:693:lu_device_fini()) ASSERTION(atomic_read(&d->ld_ref) == 0) failed *Attachment #21925*
3. 17524 sanity.sh: FAIL: test_27n *Attachment #21130*
4. 18374 replay-single test_60: ASSERTION(rs->rs_service == svc) failed *Attachment #21489*
5. 18414 still busy with 2 active RPCs *Attachment #21851*
6. 17491 ptlrpc_at_add_timed() CPU hog *Attachment #21517*
7. 18380 lock ordering violation between &cli->cl_sem and _lprocfs_lock *Attachment #21519*
8. 3633 some open/close RPCs should not be replayed or resent *Attachment #21698*
9. 18471 badly use lprocfs in HEAD *Attachment #21853*
10. 17397 read-only cache data corruption with fsx *Attachment #22029*
11. 18609 kernel panic in llcd_get() on OST in replay-single 61a *Attachment #21984*
12. 18520 sanity test_151: @@@@ @@@@ FAIL: IN CACHE: before: 6, after: 9 *Attachment*

#21874

13. 18609 kernel panic in llcd_get() on OST in replay-single 61a *Attachment #21984*
14. 17682 Improve debug capability for production systems *Attachment #21709*
15. 17536 MDS create should not wait for statfs RPC while holding DLM lock *Attachment #21939*
16. 18154 conf-sanity test_47 hung at Stopping /mnt/mds *Attachment #22097*
17. 18558 package installation fails on Autoyast Changes committed in sunhpc-configuraiton-0.4.5
18. 17846 interop: sanity 102d, 102f fail, due to "trusted.lma" is restored. *Attachment #22012*
19. 18539 Do not allocate new files on degraded OSTs (SW RAID) *Attachment #21871*
20. 17644 blocking callback timeout results in eviction *Attachment #21772*
21. 12069 OST should not grant too much to client *Attachment #22209*
22. 17710 LDLM_TIMEOUT setting as an option to the MKFS.LUSTRE *Attachment #21616*
23. 18426 replay-single.sh: FAIL: test_68 2nd cp failed *1Attachment #22253*
24. 17701 liblustre: t55 incorrect striping information! *Attachment #22222*
25. 18676 recovery-small test_60: FAIL: Changelog not empty: 30005 entries *Attachment #22128*
26. 18696 recovery-random-scale test *Attachment #22099*
27. 18730 quota_save_version fix *Attachment #22163*
28. 18033 need to audit code for CONFIG_PREEMPT interaction *Attachment #22307*
29. 18514 sanity test_170: @@@@@@ FAIL: expected bad lines, but got *Attachment #22304*
30. 18761 recovery-mds-scale running issue *Attachment #22268*
31. 18623 Improve acc-sm to support ENABLE_QUOTA *Attachment #22282*
32. 18519 redesign lustre capability
33. 16277 O_CHECK_STALE should not conflict with VFS flags *Attachment #22070*
34. 17997 unchecked return values in lshowmount *Attachment #21839*
35. 18042 Oops: 0000 [1] SMP RIP: {:ldiskfs:move_entries+479} *Attachment #21618*
36. 3633 some open/close RPCs should not be replayed or resent *Attachment #22174*
37. 18711 kernel: LustreError: 987:0:(lu_object.c:1072:key_fini())
ASSERTION(atomic_read(&key->lct_used) > 1) failed *Attachment #22291*
38. 18589 replay-ost-single test_7: @@@@@@ FAIL: test_7 failed with 1 *Attachment #22106*
39. 18623 Improve acc-sm to support ENABLE_QUOTA *Attachment #22383*
40. 18811 sanity-quota test_22 defect *Attachment #22377*
41. 18804 test-framework cleanups *Attachment #22342*
42. 18662 client hung: RIP: {:osc:osc_grant_shrink_grant_cb+976} *Attachment #22280*
43. 18662 client hung: RIP: {:osc:osc_grant_shrink_grant_cb+976} *Attachment #22453*
44. 18676 recovery-small test_60: FAIL: Changelog not empty: 30005 entries *Attachment #22389*

45. 17336 sanity.sh: FAIL: test_127 sum is wrong: 2101248 Attachment #22354
46. 18804 test-framework cleanups Attachment #22370
47. 18741 sanity-quota test_18b: @@@@@@ FAIL: resetquota failed Attachment #22536
48. 18656 build system changes to get xen domU packages out of lbuild Attachment #22018
49. 18829 osc_lock.c:1408:osc_lock_cancel()) ASSERTION(dlmlock->l_readers == 0 && dlmlock->l_writers == 0) failed Attachment #22450
50. 18850 racer, BUG: soft lockup - CPU#0 stuck for 10s! [ll_imp_inval:32389] Attachment #22466
51. 18901 ASSERTION(m->mo_ops->moo_version_get) failed Attachment #22538
52. 18867 sanity test 77c was skipped on HEAD Attachment #22491
53. 17195 create VBR large scale tests Attachment #22557
54. 18874 conf-sanity test 48: cannot open file `/etc/groups' for reading Attachment #22498
55. 18854 (mdd_lov.c:424:mdd_lov_create()) ASSERTION(eadata != NULL) failed Attachment #22497
56. 18656 build system changes to get xen domU packages out of lbuild Simple fix directly committed without an attachment
57. 18656 build system changes to get xen domU packages out of lbuild Attachment #22019
58. 18794 conf-sanity test_30: Oops: 0010 [1] SMP ll_ping Tainted Attachment #22501
59. 18903 llmount.sh: BUG: spinlock wrong owner on CPU#0, mount.lustre/8180 Attachment #22567
60. 18487 Failed to mount over MX Attachment #21837
61. 18972 sanity.sh test 24i contains a typo Attachment #22647
62. 18861 HEAD with intermixed platforms can not show the owner for normal users Attachment #22541
63. 18554 sanity-quota : FAIL: unable to write to /mnt/lustre/d0_runas_test as UID 60000. Attachment #22674 (Change reverted a day later)
64. 18973 LustreError: 27229:0:(lov_lock.c:1043:lov_closure_get()) ASSERTION(list_empty(&closure->clc_list)) Attachment #22648
65. 17710 LDLM_TIMEOUT setting as an option to the MKFS.LUSTRE Attachment #22452
66. 18675 conf-sanity test 41: failed to start MDT with "-o nomgs" Attachment #22136
67. 18796 replay-single test_50: __ptlrpc_req_finished()) LBUG Attachment #22340
68. 18711 kernel: LustreError: 987:0:(lu_object.c:1072:key_fini()) ASSERTION(atomic_read(&key->lct_used) > 1) failed Attachment #22743
69. 18895 replay-single test_61d: LustreError: 17392:0: (lprocfs_status.c:859:lprocfs_obd_cleanup()) ASSERTION(obd->obd_proc_exports_entry->subdir == NULL) failed Attachment #22664
70. 18810 ondisk corruption on OST Attachment #22642
71. 18994 REPLAY_VBR is missing on HEAD Attachment #22689
72. 18404 improve and document write_append_truncate test Attachment #21958
73. 18806 fix the debug_daemon commands to wait for completion Attachment #22742
74. 17973Fix cmd3-11 cmd3-17 defects which do not fit the Hendrix requirements Attachment #22492

75. 18696 recovery-random-scale test(revert mistakenly added patch from 18696) Attachment #22099
76. 19010 Yala doesn't check whether lustre mount succeeds Attachment #22688
77. 18994 REPLAY_VBR is missing on HEAD Attachment #22711
78. 17973 Fix cmd3-11 cmd3-17 defects which do not fit the Hendrix requirements Attachment #22531
79. 18855 Compile error building v1_9_164 on Cray Attachment #22543
80. 18798 add import status info Attachment #22572
81. 17378 LBUG in lov_free_memmd() Attachment #22465
82. 17378 LBUG in lov_free_memmd() Attachment #22763
83. 18554 sanity-quota : FAIL: unable to write to /mnt/lustre/d0_runas_test as UID 60000. Attachment #22782
84. 18988 kernel BUG at mm/filemap.c:2370! Attachment #22692
85. 18414 still busy with 2 active RPCs Attachment #22175
86. 19087 Inet memory optimizations Attachment #22780
87. 18902 racer test failed to kill `ls -R /mnt/lustre/racer` Attachment #22603
88. 18882 ldiskfs_free_inode: bit already cleared for inode Attachment #22643
89. 18556 replay-dual test_19 timed out Attachment #22509
90. 17491 ptlrpc_at_add_timed() CPU hog Attachment #22515
91. 16839 ll_sync thread stay in waiting mds<ost recovery finished Attachment #22802
92. 16417 Lustre doesn't delete files Attachment #22407
93. 17926 iozone failed Attachment #22776
94. 18849 racer, umount /mnt/mds1 failed, LBUG ASSERTION(imp->imp_sec == NULL) failed Direct commit, no attachment
95. 17560 readdir : Page-wide hash collision error on client Attachment #22840
96. 18293 sanity-quota.sh: FAIL: test_23 (1) write failure, but expect success: 24576 Attachment #22406
97. 19016 ASSERTION(ergo(cl_lock_nesting(lock) == CNL_TOP, cl_lock_nr_mutexed(env) == 1)) failed Attachment #22781
98. 18080 sanityN test 34 hangs running on Cray XT3 Attachment #22786
99. 19129 syntax error on replay_vbr Attachment #22847
100. 19139 minor typo in END_RUN_FILE name Attachment #22852
101. 18978 MDS LBUG: LustreError: 3838:0:(cmm_object.c:979:cmr_version_set()) Attachment #22761
102. 18556 replay-dual test_19 timed out Attachment #22676, Attachment #22677
103. 19025 recovery-small OOPS: [<ffffff88034387>]
:jbd:journal_commit_transaction+0xc44/0x134d Attachment #22846
104. 18792 lcommon_cl.c:1178:cl_inode_fini() ASSERTION(atomic_read(&head->loh_ref) == 1) failed Attachment #22853
105. 18863 HEAD: sanity tests failed due to "Out of memory" in clinet Attachment #22790
106. 18271 conf-sanity test_23a hang Attachment #22574

Results

Please review the Lustre 2.0 Alpha Test Plan document and Head Daily Testing Results spreadsheet (both available on lustre.org) for detailed plans and daily testing results.

The Lustre 2.0 Alpha (build 16) has the following characteristics:

- Testing completed for RHEL5/x86_64 and SLES10x85_64 using IB and TCP
- Passing consistently: ior, simul, runtests, sanity, dbench, bonnie, iozone, fsx, recovery small, replay_dual, sanity_quota, sanity_sec, replay_vbr, replay_ost_single
- Known Failures: sanityn (bug 16417), liblustre (19104 intermittent), replay single (19085), insanity (19150), performance_sanity (19151), conf_sanity(17931)
- IOR and simul passed on a 74 client cluster
- CMD, Quotas, Recovery, and 1.8->2.0 upgrades not yet tested

Lustre 2.0 Alpha was delivered on April 16, 2009 and is available from:

http://downloads.lustre.org/public/lustre/Lustre_2.0_Alpha/