



Lustre 2.0 Alpha-3 Summary

Summary

The focus of the third Alpha release (Alpha-3) of 2.0 was to continue improve stability of Lustre while landing queued fixes to HEAD and completing additional bug fixes. 95 total fixes were landed during this cycle. Notable check-ins during the period:

- Final patches for Size on MDS (SOM) preview were checked-in (Bug 1028)
- OFED 1.4.1 support was landed

This release was tested on RHEL5/x86_64 and SLES 10/x86_64 with both IB and TCP connectivity. Known failures are documented in the HEAD Daily Testing Document available from lustre.org. Additionally, the backlog of HEAD patches that were queued for landing once HEAD was stabilized has been worked through and patches have been landed.

Timeline

The Alpha-3 bug fix period was from May 8, 2009 – June 4, 2009. The build on June 4 marks the Lustre 2.0 Alpha-3 release which was announced and available for download June 10, 2009.

Fixes Landed

Below is a summary listing of all unique landings to HEAD during the Lustre 2.0 Alpha-3 Milestone period. This is inclusive of all landings -- code fixes as well as test framework fixes.

1. Bug 1028: cache file size on MDS Attachment #23578; Attachment #23774; Attachment #23890
2. Bug 3633: some open/close RPCs should not be replayed or resend Attachment #23487
3. Bug 11505: find_task_by_pid() export is going away in 2.6.27 Attachment #23527
4. Bug 12069: OST should not grant too much to client Attachment #23069
5. Bug 13277: Null pointer deref in ptlrpc_set_destroy Attachment #20898
6. Bug 14109: Making new e2fsprogs releases Attachment #17803
7. Bug 15059: Over-writes are not tracked properly Attachment #23828
8. Bug 15266: port cmd3 performance tests to acc-sm t-f Attachment #22573
9. Bug 15870: error building rpms for v1_6_5_RC4 on Cray XT3 Attachment #23309
10. Bug 16261: MOUNTOPT="user_xattr,acl,flock" would not work due to a test script issue in cfg/local.sh Attachment #23449

11. Bug 16312: build the lustre kernel the way the vendor does
Attachment #23213
12. Bug 16855: Filesystem Replication / Backup
Attachment #23759
13. Bug 17086: LSI Fusion MPT driver hacks to improve performance on Thor and Riverwalk
Attachment #23704
14. Bug 17584: integrate kDMU with lustre build system
Attachment #22293
15. Bug 17605: use mpi compiler wrapper for building mpitests
Attachment #23356
16. Bug 17614: LBUG on PPC64 Client :(ldlm_lock.c:159:ldlm_lock_put())
ASSERTION(list_empty(&lock->l_pending_chain)) failed
Attachment #20413
17. Bug 17671: o2iblnd with OFED1.4 support
Attachment #22815; Attachment #22950; Attachment #23285
18. Bug 18015: ldlm_exit() ASSERTION(rc == 0) slab error in kmem_cache_destroy():
cache `ldlm_locks': Can't free all objects
Attachment #22890
19. Bug 18031: LBUG: check_for_next_transno()
ASSERTION(lustre_msg_get_transno(req->rq_reqmsg) >= next_transno) failed
Attachment #22506
20. Bug 18075: LNet selftest hit ASSERTION(ev->ev_fired)
Attachment #23415
21. Bug 18204: MDS Oops'd at kernel groups_search+76
Attachment #22165
22. Bug 18382: replay-single test_48: @@@@ @@@@ FAIL: test_48 failed with 2;
input/output error
Attachment #23528
23. Bug 18399: OSS Deadlock
Attachment #22920
24. Bug 18450: better error reporting, and verify correct local device
Attachment #23361
25. Bug 18498: trace_get_console_buffer() calculates index incorrectly
Attachment #23789
26. Bug 18511: recovery-mds-scale tar script encountered the space issue
Attachment #23800
27. Bug 18525: FID2path on lustre clients
Attachment #22739
28. Bug 18556: replay-dual test_19 timed out
Attachment #22680
29. Bug 18561: lu_ref dump
Attachment #21937
30. Bug 18645: Too many small read RPCs on BGP
Attachment #23585
31. Bug 18654: MXLND: eliminate hosts file, use arp for peer nic_id
Attachment #23381
32. Bug 18668: SLES11 support
Attachment #23204; Attachment #23439
33. Bug 18674: clients not allowed to reconnect because a glimpse callback cannot be sent
Attachment #23339
34. Bug 18688: allow tuning service thread count via /proc
Attachment #23090
35. Bug 18696: recovery-random-scale test
Attachment #23655; Attachment #22589

36. Bug 18721: liblustre failed
Attachment #23484
Attachment #23633
37. Bug 18734: kernel BUG in ptrlpc_commit_replies() during liblustre unmount
Attachment #23796
38. Bug 18788: Serialized client eviction processing causes overall I/O performance degradation
Attachment #22580
39. Bug 18849: racer, umount /mnt/mds1 failed, LBUG ASSERTION(imp->imp_sec == NULL) failed
Attachment #23279; Attachment #23282
40. Bug 18868: lov_add_target is leak lov_lock.
Attachment #22472
41. Bug 18881: recovery-double-scale test_2 failed, can not stop client
Attachment #23410
42. Bug 18884: implement GROUP locks for CLIO code
Attachment #22940
43. Bug 18948: Snappy recovery
Attachment #23424
44. Bug 18949: e2fsck does not atomically read mmp block
Attachment #23417
45. Bug 19024: Kernel update for RHEL5.3 2.6.18-128.1.6.el5.
Attachment #23859
46. Bug 19026: improve recovery-double-scale output
Attachment #23257
47. Bug 19044: HEAD conf-sanity test_46 :(layout.c:1623:req_capsule_shrink()) LBUG
Attachment #23599
48. Bug 19080: ASSERTION(sublock->cl_state != CLS_FREEING) failed
Attachment #23422
49. Bug 19151: performance-sanity test_3: @@@@ @@@@ FAIL: test_3 failed with 1
Attachment #23795
50. Bug 19154: replay-single test_20b: @@@@ @@@@ FAIL: after 2707848 > before 2701704
Attachment #23606
51. Bug 19156: Test node can't be in multiple groups of a session
Attachment #23671
52. Bug 19163: yala didn't install RHEL5 kernel
Attachment #23332; Attachment #23246
53. Bug 19194: performance-sanity test failures
Attachment #23358
54. Bug 19195: Oops: RIP: 0010:[<ffffff88a1fb4f>] [<ffffff88a1fb4f>]
:obdfilter:filter_map_remote_to_local+0x1df/0x320
Attachment #23374
55. Bug 19212: Kernel update for SLES10 SP2 2.6.16.60-0.37.
Attachment #23858
56. Bug 19263: Sanity test_170: soft lockup - CPU#0 stuck for 10s! [ktracefiled:4769]
Attachment #23631
57. Bug 19264: sanity-quota: multiple failures: FAIL: System is error when query quota for block (u:quota_usr)
Attachment #23865
58. Bug 19293: move AT tunables to proc
Attachment #23475
59. Bug 19327: Cleanup of o2iblnd protocol V2
Attachment #23164
60. Bug 19352: LBUG: (cl_page.c:1253:cl_page_completion()) !(pg->cp_flags & CPF_READ_COMPLETED) ASSERTION(0) failed

- Attachment #23335
61. Bug 19385: LBUG: (mdt_recovery.c:897:mdt_txn_stop_cb()) ASSERTION(req != NULL && req->rq_repmsg != NULL) failed
Attachment #23580
 62. Bug 19389: llapi_fid2path() introduces dependencies to lustre internal libs
Attachment #23472
 63. Bug 19405: client.c:1744:__ptlrpc_free_req()) LBUG
Attachment #23283
 64. Bug 19414: port cmd3-33 to acc-sm
Attachment #23435
 65. Bug 19426: fid_pack incorrectly packing "short" FIDs
Attachment #23355
 66. Bug 19431: sanity-sec test 5 and 6 were always skipped
Attachment #23466
 67. Bug 19447: LBUG: (quota_context.c:1430:lqs_exit()) ASSERTION(atomic_read(&q->lqs_refcount) == 0) failed: Busy lqs ffff8100af1df180 with 2 refs
Attachment #23558
 68. Bug 19448: file handle leak in liblustre_init_random
Attachment #23345
 69. Bug 19450: sanity-quota test_18c hung
Attachment #23465
 70. Bug 19478: sanity test 65l very slow
Attachment #23456
 71. Bug 19479: obdfilter-survey causes oops:
:obdecho:echo_device_alloc+0x1368/0x1760
Attachment #23413
 72. Bug 19481: acc-sm: Initail test setup fails to configure quotas correctly
Attachment #23518
 73. Bug 19484: conf-sanity test_41 Hang running on Cray XT3
Attachment #23812
 74. Bug 19497: tune2fs should warn if needs_recovery is set
Attachment #23419
 75. Bug 19511: questionable EIO handling in e2fsck
Attachment #23442
 76. Bug 19520: LBUG in ofed 1.4.1
Attachment #23498
 77. Bug 19523: sanity test 100 : failure message.
Attachment #23504; Attachment #23575
 78. Bug 19527: HEAD build is broken
Attachment #23457
 79. Bug 19528: replay-single: test_70b ASSERTION(list_empty(&exp->exp_mds_data.med_open_head)) failed
Attachment #23648; Attachment #23751
 80. Bug 19577: sanity-quota failed in acc-sm test for HEAD
Attachment #23835
 81. Bug 19578: wrong mode in replay-vbr test 3b
Attachment #23551
 82. Bug 19590: v1_8_0_110 rhel5 i686 patchless build failure
Attachment #23608
 83. Bug 19598: sanity-quota test 11 hung
Attachment #23680
 84. Bug 19612: incorrect lctl quota commands causes kernel general protection fault
Attachment #23620; Attachment #23638
 85. Bug 19624: sanity test_118b: @@@@ @@@@ FAIL: Must return error due to dropped pages, rc=0
Attachment #23659

86. Bug 19671: store_for_reuse in old-school build does not account for OFED pre-releases
Attachment #23699
87. Bug 19685: include the ChangeLogs in the lustre-modules RPM
Attachment #23725
88. Bug 19690: provide "make cscope" target, like "make tags"
Attachment #23733
89. Bug 19693: lbuild's reuse caching is racing with cleanup scripts
Attachment #23748
90. Bug 19695: build with -Werror fails
Attachment #23754
91. Bug 19730: patchless client configure should not consider the kernel's support for 64-bit quotas
Attachment #23876
92. Bug 19731: build ofed 1.4.1 with mlx4_en (Mellanox ConnectX drivers in 10GbE mode) enabled
Attachment #23816
93. Bug 19734: parallel-scale binaries env.
Attachment #23819
94. Bug 19741: HEAD build failed
Attachment #23825
95. Bug 19756: recovery-small test 27 multiple failures.
Attachment #23860

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-3 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, sanityn, liblustre, recovery small, replay_dual, insanity, sanity_sec, replay_vbr, replay_ost_single, conf_sanity
- Known Failures: replay single (19154), sanity_quota (19577), performance_sanity (19588), racer (19188), connectathon (18087)
- CMD, Quotas, Recovery, and 1.8->2.0 upgrades not yet tested

Lustre 2.0 Alpha-3 was delivered on June 10, 2009 and is available from:

<http://downloads.lustre.org/public/lustre/v2.0/latest/>

or

http://downloads.lustre.org/public/lustre/v2.0/alpha/Lustre_2.0_Alpha3/