



Lustre 2.0 Alpha-4 Summary

Summary

The focus of the fourth Alpha release (Alpha-4) of 2.0 was to continue improve stability of Lustre while landing fixes to HEAD and completing additional bug fixes. 120 total fixes were landed during this cycle.

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.

Timeline

The Alpha-3 bug fix period was from May 20, 2009 – June 26, 2009. The build on June 29 marks the Lustre 2.0 Alpha-4 release which was announced and available for download July 1, 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 11505: find_task_by_pid() export is going away in 2.6.27
Attachment #23527
3. Bug 12069: OST should not grant too much to client
Attachment #23096
4. Bug 13277: Null pointer deref in ptrpc_set_destroy
Attachment #20898
5. Bug 15059: Over-writes are not tracked properly
Attachment #23828
6. Bug 15266: port cmd3 performance tests to acc-sm t-f
Attachment #22573
7. Bug 15870: error building rpms for v1_6_5_RC4 on Cray XT3
Attachment #23309
8. Bug 16317: exports in /proc are broken for 1.6.5
Attachment #18040
9. Bug 16424: Examine for shorten the build version string
Attachment #23839; Attachment #24051
10. Bug 16823: poor Lustre level Shared File performance on Jaguar/ DDN 9900
Attachment #21598
11. Bug 16855: Filesystem Replication / Backup
Attachment #23759

12. Bug 16893: port ldiskfs to ext4
Attachment #23641; Attachment #23281; Attachment #23771; Attachment #23280
13. Bug 17086: LSI Fusion MPT driver hacks to improve performance on Thor and Riverwalk
Attachment #23704
14. Bug 17546: An ost failover test failed when the osts did not recover all the clients.
Attachment #23840; Attachment #23844
15. Bug 17560: readdir : Page-wide hash collision error on client
Attachment #23837
16. Bug 17670: find -type f performance degradation with 2.0 servers
Attachment #23932; Attachment #23664
17. Bug 17671: o2iblnd with OFED1.4 support
Attachment #22815; Attachment #22950; Attachment #23285
18. Bug 17986: Automatically tune RAID stripe cache size in mount.lustre
Attachment #23791
19. Bug 18015: ldlm_exit() ASSERTION(rc == 0) slab error in kmem_cache_destroy(): cache `ldlm_locks': Can't free all objects
Attachment #22890; Attachment #yyyyy; Attachment #zzzzz
20. Bug 18021: test-framework: CLIENTONLY should not depend on *_HOST settings
Attachment #24214
21. Bug 18031: LBUG: check_for_next_transno()
ASSERTION(lustre_msg_get_transno(req->rq_reqmsg) >= next_transno) failed
Attachment #22506
22. Bug 18075: LNet selftest hit ASSERTION(ev->ev_fired)
Attachment #23415
23. Bug 18204: MDS Oops'd at kernel groups_search+76
Attachment #22165
24. Bug 18357: Down OSTs block all MDT threads resulting in client hangs
Attachment #24225
25. Bug 18364: cxfscp w/ lfs setstripe causes target (lustre) filesystem to lockup
Attachment #22593
26. Bug 18382: replay-single test_48: @@@@ @@@@ FAIL: test_48 failed with 2;
Input/output error
Attachment #23528
27. Bug 18399: OSS Deadlock
Attachment #22920
28. Bug 18450: better error reporting, and verify correct local device
Attachment #23361
29. Bug 18498: trace_get_console_buffer() calculates index incorrectly
Attachment #23789
30. Bug 18511: recovery-mds-scale tar script encountered the space issue
Attachment #23800
31. Bug 18574: Quotas causes ~15% slowdown in metadata rates
Attachment #23957; Attachment #24303
32. Bug 18616: optimize the quota code
Attachment #24261
33. Bug 18618: All MDT Threads Blocked
Attachment #23872
34. Bug 18619: applications not exiting for extended amounts of time
Attachment #22487
35. Bug 18624: Unable to run several mkfs.lustre on loop devices at the same time
Attachment #22002
36. Bug 18638: (ldlm_lockd.c:2157:ldlm_exit()) ASSERTION(rc == 0) failed:couldn't free
ldlm lock slab
Attachment #22626
37. Bug 18645: Too many small read RPCs on BGP

- Attachment #23585
- 38. Bug 18668: SLES11 support
Attachment #23204; Attachment #23439
- 39. Bug 18674: clients not allowed to reconnect because a glimpse callback cannot be sent
Attachment #23339
- 40. Bug 18688: allow tuning service thread count via /proc
Attachment #23090
- 41. Bug 18696: recovery-random-scale test
Attachment #22655; Attachment #22589
- 42. Bug 18721: liblustre failed
Attachment #23633
- 43. Bug 18734: parallel-scale binaries env.
Attachment #23796
- 44. Bug 18787: OSS memory usage before and after we mount the clients
Attachment #23701
- 45. Bug 18788: Serialized client eviction processing causes overall I/O performance degradation
Attachment #22580
- 46. Bug 18881: recovery-double-scale test_2 failed, can not stop client
Attachment #23891
- 47. Bug 18914: recovery-mds-scale was interrupted after running for 8 hours
Attachment #22625
- 48. Bug 18918: server_fill_super logs server start when using mount.lustre --nosvc
Attachment #22650
- 49. Bug 19024: Kernel update for RHEL5.3 2.6.18-128.1.6.el5.
Attachment #23859
- 50. Bug 19044: HEAD conf-sanity test_46 :(layout.c:1623:req_capsule_shrink()) LBUG
Attachment #23599
- 51. Bug 19086: replay-single FAIL: test_18 failed with 4
Attachment #24316
- 52. Bug 19149: sanity-quota test_6: @@@@ @@@@ FAIL: write filea failure, expect success
Attachment #24189
- 53. Bug 19151: performance-sanity test_3: @@@@ @@@@ FAIL: test_3 failed with 1
Attachment #23795
- 54. Bug 19154: replay-single test_20b: @@@@ @@@@ FAIL: after 2707848 > before 2701704
Attachment #23606; Attachment #23988; Attachment #23990; Attachment #24188
- 55. Bug 19156: Test node can't be in multiple groups of a session
Attachment #23671
- 56. Bug 19188: racer: LBUG: (cl_lock.c:287:cl_lock_put()) ASSERTION(list_empty(&lock->cl_linkage)) failed
Attachment #24041; Attachment #24172; Attachment #24254
- 57. Bug 19194: performance-sanity test failures
Attachment #23358
- 58. Bug 19195: Oops: RIP: 0010:[<ffffff88a1fb4f>] [<ffffff88a1fb4f>]
:obdfilter:filter_map_remote_to_local+0x1df/0x320
Attachment #23374
- 59. Bug 19212: Kernel update for SLES10 SP2 2.6.16.60-0.37.
Attachment #23585
- 60. Bug 19259: killed app doesn't exit
Attachment #23923
- 61. Bug 19263: Sanity test_170: soft lockup - CPU#0 stuck for 10s! [ktracefiled:4769]
Attachment #23631
- 62. Bug 19264: sanity-quota: multiple failures: FAIL: System is error when query quota for block (u:quota_usr)

- Attachment #23865
63. Bug 19280: conf-sanity test_50e: error : glob /proc/{fs,sys}/{lnet,lustre}/osc/lustre-OST0000-osc/ost_server_uuid: No such process
Attachment #24226
 64. Bug 19293: move AT tunables to proc
Attachment #23475
 65. Bug 19321: llmountcleanup.sh failed service didn't stop after 318 seconds
Attachment #24021
 66. Bug 19385: LBUG: (mdt_recovery.c:897:mdt_txn_stop_cb()) ASSERTION(req != NULL && req->rq_repmsg != NULL) failed
Attachment #23580
 67. Bug 19386: Oops: RIP :ptlrpc:lustre_msg_get_opc+0x8/0xe0
Attachment #24082
 68. Bug 19389: llapi_fid2path() introduces dependencies to lustre internal libs
Attachment #23472
 69. Bug 19398: BUG: soft lockup - CPU#13 stuck for 10s! [umount:16789]
Attachment #23272
 70. Bug 19407: use do_nodes() instead of looping on all nodes
Attachment #23284
 71. Bug 19413: sanity-gss: Error checking null of all2all: expect 20, actual 8
Attachment #24269
 72. Bug 19414: port cmd3-33 to acc-sm
Attachment #23435
 73. Bug 19435: quotacheck turn on quota failed: Device or resource busy
Attachment #24297
 74. 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
 75. Bug 19451: IOR load for recovery scale tests
Attachment #23362
 76. Bug 19479: obdfilter-survey causes oops:
:obdecho:echo_device_alloc+0x1368/0x1760
Attachment #23413
 77. Bug 19481: acc-sm: Initail test setup fails to configure quotas correctly
Attachment #23518
 78. Bug 19484: conf-sanity test_41 Hang running on Cray XT3
Attachment #23812
 79. Bug 19528: replay-single: test_70b ASSERTION(list_empty(&exp->exp_mds_data.med_open_head)) failed
Attachment #23648
 80. Bug 19536: turning mmp off leaves filesystem with incorrect block bitmap/summary
Attachment #23999
 81. Bug 19544: lustre_{read,write}_dqot() use variable on-stack allocations
Attachment #23989
 82. Bug 19550: MDS hung at stopping mds: Stopping /mnt/mds1
Attachment #24252; Attachment #24317; Attachment #24318
 83. Bug 19577: sanity-quota failed in acc-sm test for HEAD
Attachment #24174
 84. Bug 19590: v1_8_0_110 rhel5 i686 patchless build failure
Attachment #23608
 85. Bug 19593: memory leak in ptlrpc_init()
Attachment #23581
 86. Bug 19598: sanity-quota test 11 hung
Attachment #23680
 87. Bug 19612: incorrect lctl quota commands causes kernel general protection fault
Attachment #23620; Attachment #23638

88. Bug 19624: sanity test_118b: @@@@ @@@@ FAIL: Must return error due to dropped pages, rc=0
Attachment #23659
89. Bug 19669: Store HSM flags on MDT and release
Attachment #24197
90. Bug 19671: store_for_reuse in old-school build does not account for OFED pre-releases
Attachment #23699
91. Bug 19672: MDS lockup during inode quota test on Cray XT3
Attachment #23955
92. Bug 19685: include the ChangeLogs in the lustre-modules RPM
Attachment #23725
93. Bug 19690: provide "make cscope" target, like "make tags"
Attachment #23733
94. Bug 19693: lbuild's reuse caching is racing with cleanup scripts
Attachment #23748
95. Bug 19694: LustreError: 8032:0:(md_object.h:636:mo_xattr_list()) ASSERTION(m->mo_ops->mo_xattr_list) failed
Attachment #23752
96. Bug 19695: build with -Werror fails
Attachment #23754
97. Bug 19702: runtests : @@@@ @@@@ FAIL: old and new files are different 11
Attachment #23866
98. Bug 19725: Oops: 0000 [1] SMP RIP:fsfilt_ldiskfs_quotacheck+0x2a5/0x26c0 [fsfilt_ldiskfs]
Attachment #23867
99. Bug 19730: patchless client configure should not consider the kernel's support for 64-bit quotas
Attachment #23867
100. Bug 19731: build ofed 1.4.1 with mlx4_en (Mellanox ConnectX drivers in 10GbE mode) enabled
Attachment #23816
101. Bug 19734: parallel-scale binaries env.
Attachment #23819
102. Bug 19741: HEAD build failed
Attachment #23825
103. Bug 19756: recovery-small test 27 multiple failures.
Attachment #23860; Attachment #24016
104. Bug 19762: need to redirect stdout in build_kernel_with_srpm
Attachment #23874
105. Bug 19769: update "make checkstack" target
Attachment #23878
106. Bug 19773: FAIL: mdt-s : inodes count OLD (41) NEW (42)
Attachment #23889
107. Bug 19784: remove lu_env_init in mdt_fid2path
Attachment #23917
108. Bug 19802: sanity test_132: @@@@ @@@@ FAIL: som is still disabled
Attachment #23981
109. Bug 19820: recovery-small fails to start after running conf-sanity
Attachment #23971
110. Bug 19826: failed to build v1_9_0_200 with GSS enabled
Attachment #23997
111. Bug 19850: shouldn't specify both --with{,out}-srp-target-mod
Attachment #24053
112. Bug 19852: Need to upgrade e2fsprogs to 1.41.6 version
Attachment #24063

113. Bug 19853: /cache/build/RPMS/i386/mbuild_subject.tag: No such file or directory
Attachment #24052
114. Bug 19856: HSM copytool API
Attachment #24077
115. Bug 19867: quotacheck: invalid option -- 3
Attachment #24100
116. Bug 19885: Evict clients after upgrade to 2.0 MDS from incompatible version
Attachment #24224
117. Bug 19913: sanity test_151: @@@@ @@@@ FAIL: NOT IN CACHE: before: , after:
Attachment #24187; Attachment #24210
118. Bug 19920: LIBCFS_ALLOC_GFP should not print 'kmallocced' message at allocation failures
Attachment #24206
119. Bug 19927: unable to open the tree_status file
Attachment #24276
120. Bug 19933: sanity test_103: @@@@ @@@@ FAIL:
Attachment #24272

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-4 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, performance_sanity
- Known Failures: replay single (19994, 20021), sanity_quota (19553), racer (19698),

Lustre 2.0 Alpha-3 was delivered on July 1, 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_Alpha4/