



Lustre 2.0 Alpha-5 Summary

Summary

The focus of the fifth Alpha release (Alpha-5) of 2.0 was to continue improve stability of Lustre while landing fixes to HEAD and completing additional bug fixes. 155 total fixes were landed during this cycle.

This release was tested on RHEL5/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-5 bug fix period was from August 14, 2009 – October 8, 2009. The build on October 8 marks the Lustre 2.0 Alpha-5 release which was announced and available for download October 14, 2009.

Fixes Landed

Below is a summary listing of all unique landings to HEAD during the Lustre 2.0 Alpha-5 Milestone period. This is inclusive of all landings -- code fixes as well as test framework fixes.

1. Bug 13520: replay-single test_65: **** FAIL: No early reply
Attachment: 25647
2. Bug 14951: OST addition not being handled by client
Attachment: 25503
3. Bug 15284: write a simple utility to rename files in lost+found to object ID
Attachment: 25049 Attachment: 25050
4. Bug 15390: permission issue
Attachment: 25514
5. Bug 15504: acc-sm needs the possibility to start from the defined test
Attachment: 25775
6. Bug 15534: make HEAD safety run on BIG_ENDIAN machines.
Attachment: 25553 Attachment: 25577 Attachment: 25582
7. Bug 15902: make test for the 1.8->HEAD upgrade process
Attachment: 25653
8. Bug 16186: One down Lustre FS hangs ALL mounted Lustre filesystems
Attachment: 25045
9. Bug 16208: Land LLNL Patches for known and resolved bugzilla bugs
10. Bug 16403: test framework should use lustre_rmmod
Attachment: 25703
11. Bug 16890: small fixtures for quota
Attachment: 25405
12. Bug 17197: ll_read_ahead_pages() ASSERTION(page_idx > ria->ria_stoff) failed
Attachment: 25951
13. Bug 17517: LBUG in ldlm_lock_cancel() from ldlm_lock.c

- Attachment: 25845
14. Bug 17591: refactor acc-sm
 15. Bug 17873: not skipping tests due to env misconfiguration in YALA
Attachment: 25656 Attachment: 25657
 16. Bug 17951: conf-sanity test_46a: @@@@ @ FAIL: test_46a failed with 7
Attachment: 25772
 17. Bug 18054: build using in-kernel ldiskfs
Attachment: 25027
 18. Bug 18102: lnet_proc.c should be rewritten
Attachment: 25875
 19. Bug 18534: Unnecessary lookup RPC when doing non-open related creates
Attachment: 21800
 20. Bug 18548: Running of conf_param on 1.8 disk without --writeconf should return error
Attachment: 25284
 21. Bug 18623: Improve acc-sm to support ENABLE_QUOTA
Attachment: 25743
 22. Bug 18674: clients not allowed to reconnect because a glimpse callback cannot be sent
Attachment: 25038
 23. Bug 18775: cleanup lctl get_param interface
Attachment: 25698
 24. Bug 18866: sanity test_64b: FAIL: no OSTs are close to full
Attachment: 25851
 25. Bug 18881: recovery-double-scale test_2 failed, can not stop client
Attachment: 25197 Attachment: 25832
 26. Bug 18948: Snappy recovery
Attachment: 25372 Attachment: 25518 Attachment: xx
 27. Bug 18956: add write_disjoint to lustre tests
Attachment: 25271
 28. Bug 19027: lustre uuid specification should be more forgiving
Attachment: 25907
 29. Bug 19053: llapi_path2fid() for all object types
Attachment: 25370
 30. Bug 19124: t-f need to load modules from work dir on remote nodes.
Attachment: 25945
 31. Bug 19154: replay-single test_20b: @@@@ @ FAIL: after 2707848 > before 2701704
Attachment: 25741
 32. Bug 19205: lfs df for pools
Attachment: 25201
 33. Bug 19211: PIOS, LBUG, (mdc_locks.c:126:mdc_set_lock_data()) ASSERTION(lock != NULL) failed
Attachment: 25258
 34. Bug 19215: "MDT*" need be changed to "MDT0000" under CMD configuration
Attachment: 25519
 35. Bug 19312: add parallel_grouplock test to acc-sm
Attachment: 25655
 36. Bug 19380: LBUG hit while running racer
Attachment: 24688
 37. Bug 19418: port cmd3-38 test to acc-sm
Attachment: 25629
 38. Bug 19442: kernel panic when setting parameter mdt.quota_type
Attachment: 25394
 39. Bug 19525: MGS doesn't export RPC stats in /proc
Attachment: 23586
 40. Bug 19537: fsx failed

- Attachment: 25757 Attachment: 25970
41. Bug 19551: continuing LustreError "mds adjust qunit failed!"
Attachment: 25934
 42. Bug 19557: BUG: soft lockup - CPU#1 stuck for 10s! [ll_mdt_07:4523]
 43. Bug 19566: successful umounts although umount failed
Attachment: 25467
 44. Bug 19580: BUG: soft lockup - CPU#7 stuck for 10s! [ll_imp_inval:18451]
Attachment: 25829
 45. Bug 19657: pios: Segmentation fault
Attachment: 26011
 46. Bug 19674: problems seen building 1.8.1 prerelease on sles 11
Attachment: 25733
 47. Bug 19686: SLES11-ext4: replay-single test 5: kernel BUG at fs/jbd2/transaction:1017!
Attachment: 23854
 48. Bug 19690: provide "make cscope" target, like "make tags"
Attachment: 23733
 49. Bug 19720: various -Werror illicited errors
Attachment: 25382
 50. Bug 19721: debian/ubuntu .deb build bits
Attachment: 25383
 51. Bug 19743: 1.8.1 VBR Testing
Attachment: 25310
 52. Bug 19778: MDS refuses reconnection due to still active RPCs
Attachment: 25699
 53. Bug 19780: Tests for the OST Pools feature
Attachment: 25584 Attachment: 25937
 54. Bug 19781: Tests for the OSS Read Only Cache feature
Attachment: 25841
 55. Bug 19791: ASSERTION(atomic_read(&de->d_count) > 0) failed
Attachment: 25126
 56. Bug 19808: 2.6.29-fc11 patchless support
Attachment: 25636
 57. Bug 19835: sanity-quota test_11 : l_getgroups: no such user
Attachment: 25411
 58. Bug 19844: Timer expired but recovery doesn't end ever
Attachment: 25210 Attachment: 25222
 59. Bug 19934: some files created in lustre 1.8 are not valid under lustre 1.6
Attachment: 25312
 60. Bug 19955: add a reverse comm path from MDT to MDC for copytool reqs
Attachment: 25104
 61. Bug 19994: replay-single test_0c failed with 2
Attachment: 25433
 62. Bug 20008: quota related jbd/ldiskfs/quota BUG
Attachment: 25168
 63. Bug 20044: Assertion failure in journal_start() at fs/jbd/transaction.c:270: "handle->h_transaction->t_journal == journal"
Attachment: 25093
 64. Bug 20082: sanity quota : test 15 failed
Attachment: 25777
 65. Bug 20126: replay-single FAIL: test_53b failed with 2
Attachment: 25091
 66. Bug 20131: obdfilter-survey: LBUG (echo_client.c:939:cl_echo_object_put())
ASSERTION(&eco->eo_hdr == luh2coh(loh)) failed
Attachment: 25606
 67. Bug 20183: kernel oops during parallel_scale: (lov_pack.c:240:lov_packmd())
ASSERTION(loi->loi_id) failed

- Attachment: 25226
68. Bug 20200: Allow lfs quota command to use numeric user/group id.
Attachment: 25678 Attachment: 25744
69. Bug 20219: system degrades after a while during 1.8.1 test
Attachment: 24902
70. Bug 20231: How to map objid to file name.
71. Bug 20237: acc-sm: recovery-* -scale tests gather logs fn
Attachment: 25785
72. Bug 20298: Merge head ldiskfs and b1_8 ldiskfs
Attachment: 25567
73. Bug 20305: racer:LBUG:osc_io.c:637:osc_req_attr_set()) ASSERTION(lock != NULL) failed
Attachment: 25825
74. Bug 20339: HEAD does not compile for ppc
Attachment: 25482 Attachment: 25176 Attachment: 25178
75. Bug 20346: performance-sanity test_8: @@@@ @@@@ FAIL: test_8 failed with 1
Attachment: 25290
76. Bug 20347: replay-vbr: rm: cannot remove `/mnt/lustre/f0n': Operation not permitted
Attachment: 25410
77. Bug 20368: fsck fixes for fid in dirent
78. Bug 20392: Oops: RIP: ptlrpc:ptlrpc_at_check_timed+0x290/0x1480
Attachment: 25272
79. Bug 20400: Kernel update for RHEL5.3 2.6.18-128.7.1.el5.
Attachment: 25245
80. Bug 20411: Modify changelog to use LustreNetLink instead of seq_file
Attachment: 26023
81. Bug 20433: mdsrate fails to write after 1.3+M files opened Cray XT3
Attachment: 25483
82. Bug 20456: ASSERTION(atomic_read(&imp->imp_inflight) == 0) failed
Attachment: 25622
83. Bug 20459: make LST test and client RPC timeouts tunable
Attachment: 25371
84. Bug 20463: getstripe doesn't report stripe count = -1 sanely
Attachment: 25520
85. Bug 20464: resolve possible race between update and write lov objid file
Attachment: 25307
86. Bug 20478: version_tag.pl needs to support no-kernel builds
Attachment: 25314
87. Bug 20482: All compute nodes panic after boot.
Attachment: 25682 Attachment: 25684
88. Bug 20500: Doxygen on HEAD
Attachment: 25817
89. Bug 20521: ost-pools test_3d: @@@@ @@@@ FAIL: pool_add failed: pool1; lustre-OST[0-2/1]
Attachment: 25473
90. Bug 20533: Lustre not doing large IOs
Attachment: 25754 Attachment: 25755
91. Bug 20550: missing tree_status on 1.8.1 RPM build
Attachment: 25840
92. Bug 20560: File access errors and checksum failures with OST read cache on
Attachment: 25833 Attachment: 25896 Attachment: 25912
93. Bug 20595: Support OFED 1.4.2
Attachment: 25605
94. Bug 20605: e2label Segfault
Attachment: 25665
95. Bug 20620: ASSERTION(obd->obd_stopping) failed

- Attachment: 25761
96. Bug 20626: Move "struct obd_stats" definition to "lustre_user.h".
Attachment: 25671
 97. Bug 20653: ost-pools test_1 failed due to no mds_HOST defined
Attachment: 25842
 98. Bug 20695: configure: error: You have got no 64-bit kernel quota support
Attachment: 25831
 99. Bug 20698: integrate PIOS into acc-sm
Attachment: 25931
 100. Bug 20703: trivial code cleanups on HEAD
Attachment: 25794 Attachment: 25925
 101. Bug 20705: typo in _debug_req()
Attachment: 25795
 102. Bug 20719: replay-single test_80a: @@@@ @@@@ FAIL: needs >= 2 MDTs
Attachment: 25857
 103. Bug 20734: ior test should enable striping
Attachment: 25858
 104. Bug 20748: sanity test_55 failed
Attachment: 25892
 105. Bug 20771: vulnerable sprintf usage in ptlrpcd_start()
Attachment: 25928
 106. Bug 20784: sanity test_215: @@@@ @@@@ FAIL: /proc/sys/lnet/peers
misformatted
Attachment: 25974
 107. Bug 20814: enable frame pointers in our kernels
Attachment: 26013
 108. Bug 3633: some open/close RPCs should not be replayed or resent
Attachment #25758
 109. Bug 17153: remove 2.4 compatibility from b1_8
Attachment #26208
 110. Bug 17670: find -type f performance degradation with 2.0 servers
Attachment #24900; Attachment #25374
 111. Bug 17873: not skipping tests due to env misconfiguration in YALA
Attachment #25776
 112. Bug 18505: OOM on MDS when lru resize is enabled (i686)
Attachment #26113
 113. Bug 18629: sanityN test_32a: @@@@ @@@@ FAIL: wrong file size
Attachment #26316
 114. Bug 18674: clients not allowed to reconnect because a glimpse callback cannot
be sent
Attachment #26039
 115. Bug 18698: incorrect debug message in request_in_callback()
Attachment #25793
 116. Bug 18857: namei.c:732:ll_create_node()) ASSERTION(list_empty(&inode->i_dentry)) failed
Attachment #24626; reverted 2 days later.
 117. Bug 19113: umounting MGS hangs
Attachment #26066
 118. Bug 19235: LBUG: (recov_thread.c:436:llog_recov_thread_stop()) Busy llcfs
found (0) on lcm ffff8100c4ae16e0
Attachment #25830
 119. Bug 19242: replay-single, test_44a : lctl failed with: Try interactive use without
arguments or use one of
Attachment #26032; second landing "minor fix" 40 minutes later, no patch attached.
 120. Bug 19317: allow multiple instances of echo OBD
Attachment #23152

- 121. Bug 19526: conf-sanity test_46a: stat: cannot stat `/mnt/lustre/widestripe':
Cannot send after transport endpoint shutdown
Attachment #26072
- 122. Bug 19657: pios: Segmentation fault
Attachment #26011; Attachment #26275
- 123. Bug 19662: Remove set_info(KEY_UNLINKED) from MDS/OSC
Attachment #26009
- 124. Bug 19688: move duplicate FID-to-inode functions to lclient
Attachment #26186
- 125. Bug 19763: quotacheck performance/scaling issues
Attachment #24175
- 126. Bug 19917: repeated atomic allocation failures in kptllnd_rx_alloc() causing OST
to be unreachable
Attachment # 26214
- 127. Bug 19990: additional software raid patch needed to avoid counter overflow
Attachment #26225
- 128. Bug 20098: mod use after free
Attachment #25823; reverted 5 days later
- 129. Bug 20205: Performance degradation with O_DIRECT between 1.6 & 1.8.1
b190
Attachment #26012; Attachment #25458
- 130. Bug 20317: sanity.sh test_13: lockless truncate does not work
Attachment #25969
- 131. Bug 20349: don't access to invalid data.
Attachment #26248
- 132. Bug 20482: All compute nodes panic after boot - accessing deactivated OST
Attachment #26156
- 133. Bug 20501: performance_sanity : test 8: getattr large files : timed out
Attachment #26293
- 134. Bug 20505: sanity test_163: @@@@ @@@@ FAIL: Inl send failed - Cray XT3
Attachment #26114
- 135. Bug 20529: Patch to fix the problem in vvp_mmap_locks
Attachment #25681
- 136. Bug 20539: Lustre support for OEL
Attachment #25734; Attachment #26209
- 137. Bug 20575: pios: Watchdog; ll_ost_io_119 S
Attachment #26321
- 138. Bug 20603: recovery-small test_60: @@@@ @@@@ FAIL: Recorded 0 unlinks out
of 15000
Attachment #26297
- 139. Bug 20650: memory leak issue in Lustre 2.0
Attachment #26416
- 140. Bug 20672: sanity-benchmark with ext4, ldiskfs_mb_free_blocks: IO failure
Attachment #26135; Attachment #26136
- 141. Bug 20689: Server downgrade testing for the OST Pools feature
Attachment #25939; Attachment # ; Attachment #
- 142. Bug 20698: integrate PIOS into acc-sm
Attachment #26139
- 143. Bug 20703: trivial code cleanups on HEAD
Attachment #26134
- 144. Bug 20719: replay-single test_80a: @@@@ @@@@ FAIL: needs >= 2 MDTs
Reported as landed but no attachment in this bug was actually landed.
Possibly it was landed with bug 20784.
- 145. Bug 20736: ptlrpc: Unknown symbol cpu_online_map
Attachment #25861
- 146. Bug 20756: filter_cancel_cookies_cb: no valid context for group 0

- Attachment #26237
147. Bug 20762: sanity.sh: test_156 IN CACHE: before: 2072, after: 2075
Attachment #25998
148. Bug 20771: vulnerable sprintf usage in ptrpcd_start()
Attachment #25928
149. Bug 20784: sanity test_215: @@@@ @@@@ FAIL: /proc/sys/lnet/peers
misformatted
Attachment #25974; Attachment #26361
150. Bug 20877: sanity-quota.sh test_12 doesn't scale well
Attachment #26167
151. Bug 20878: yala test start timed out
Attachment #26335
152. Bug 20902: conf-sanity test_50g hang
Attachment #26314
153. Bug 20903: sanity-quota hang while stopping mds
Attachment #26355
154. Bug 20913: e2fsck waits long time for MMP despite clean unmount
Attachment #26218
155. Bug 20929: target_request_copy_put()) ASSERTION(list_empty(&req-
>rq_replay_list))
Attachment #26249

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-5 has the following characteristics:

- Testing completed for RHEL5/x86_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

Lustre 2.0 Alpha-5 was delivered on October 14, 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_Alpha5/