



Lustre 2.0 Beta1 Summary

Summary

The focus of the first Beta release (Beta1) of 2.0 was to continue improve stability of Lustre while landing fixes to HEAD and completing additional bug fixes. 116 total fixes were landed during this cycle.

This release was tested on RHEL5/x86_64 and OEL5/x86_64 with both IB and TCP connectivity. Known failures are documented in the HEAD Daily Testing Document available from lustre.org.

Timeline

The Beta1 bug fix period was from February 6, 2010 – April 9, 2010. The build on April 9 marks the Lustre 2.0 Beta1 release which was announced and available for download on April 10, 2010.

Fixes Landed

Below is a summary listing of all unique landings to HEAD during the Lustre 2.0 Beta1 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: 27764
2. Bug 11063: file mtime changes after a rename
Attachment: 28376 Attachment: 29304
NOTE: These attachments appear to be correct but no landing flag has actually been set for either.
3. Bug 11680: make panic on lbug configurable
Attachment: 29134
NOTE: There were 2 landings in the GIT log but I only found one in the bug report.
4. Bug 15057: files ownership in lustre-modules
Attachment: 20753
5. Bug 15962: LustreError: 14169:0:(namei.c:773:ll_create_node())
ASSERTION(list_empty(&inode->i_dentry)) failed
Attachment: 28310
6. Bug 15977: conf-sanity test_39: @@@@ @@@@ FAIL: memory leak detected
Attachment: 27959
7. Bug 16195: ASSERTION(list_empty(&pool->prp_req_list) ll size == pool->prp_rq_size)
Attachment: 27763
8. Bug 16526: No permission to write in /mnt/lustre/dir**/file** with intermixed distributions on Lustre client and servers.
Attachment: 28384
9. Bug 16680: BUG in ldiskfs/mballoc.c:3636 ldiskfs_mb_release_inode_pa()
Attachment: 28170

10. Bug 16721: 4*racer.sh: osd_object_ref_del()) ASSERTION(inode->i_nlink > 0)
Attachment: 28819
11. Bug 16909: Console message cleanup
Attachment: 28032
12. Bug 17197: ll_read_ahead_pages()) ASSERTION(page_idx > ria->ria_stoff) failed
Attachment: 27949
13. Bug 17591: refactor acc-sm
Attachment: 26140
14. Bug 17660: possible deadlock between &obd->obd_processing_task_lock and
&export->exp_lock
Attachment: 28474
15. Bug 17670: find -type f performance degradation with 2.0 servers
Attachment: 28761
16. Bug 17952: --with-linux requires a FQPN but does not enforce it
Attachment: 28571
17. Bug 18030: replay-single.sh test_33a hang during mds_client_add
Attachment: 28142
18. Bug 19069: ASSERTION(strlen(rm_entry->name) == rm_entry->namelen) failure on
umount of OST
Attachment: 28280
19. Bug 19208: sanity.sh: FAIL: test_52a lsattr
Attachment: 28845 Attachment: 29350
20. Bug 19505: Changelogs: SATTR not raised if CREATE not acknowledged
Attachment: 27961
21. Bug 19944: error writing to file /proc/sys/lnet/debug_mb
Attachment: 29003
22. Bug 20057: Autovetting and test-framework enhancements
Attachment: 28393 Attachment: 28585
23. Bug 20101: sanity.sh 27w does not work as expected, but still does not fail
Attachment: 28111
24. Bug 20315: Use libexecdir for scripts
Attachment: 27264
25. Bug 20339: HEAD does not compile for ppc
Attachment: 25178
26. Bug 20383: configure fails with syntax error
Attachment: 26864
27. Bug 20433: mdsrate fails to write after 1.3+M files opened Cray XT3
Attachment: 28782
28. Bug 20492: allow llverdev/llverfs to handle media errors better
Attachment: 28685
29. Bug 20562: increase LUSTRE_SEQ_META_WIDTH to keep FLD compact
Attachment: 26675
30. Bug 20647: replay-single FAIL: test_53f failed with 2
Attachment: 28808 Attachment: 29001
NOTE: There were 2 landings referenced in the bug report but only 1 landing in the
GIT log. They were probably just lumped into a single landing but this is an
assumption.
31. Bug 20680: replay-vbr test_0r: @@@@ @@@@ FAIL: time not changed
Attachment: 28997
32. Bug 20731: 1.8<->2.0 interop: error: llapi_obd_statfs(): Inappropriate ioctl for device (-
25)
Attachment: 28450
33. Bug 20738: FC11 support in Yala is needed
Attachment: 28626
34. Bug 20814: enable frame pointers in our kernels
Attachment: 26013

35. Bug 20837: incomplete test output for ost-pools
Attachment: 28541
36. Bug 20878: yala test start timed out
Attachment: 28690
37. Bug 20984: sanity test_27o: @@@@@@ FAIL: mdt-s : inodes count OLD 236 NEW 236
Attachment: 28518
38. Bug 20989: (lov_merge.c:74:lov_merge_lvb()) ASSERTION(spin_is_locked(&ism->ism_lock)) failed (SR 71691004)
Attachment: 28751
39. Bug 20997: Cannot send after transport endpoint shutdown
Attachment: 28414
40. Bug 21049: sanity test_27p: @@@@@@ FAIL
NOTE: Robert noted that this was a missed change from a previous patch but the attachment is not referenced, so I'm not sure which one this relates to.
41. Bug 21066: ost-pools test_14: @@@@@@ FAIL: /mnt/lustre/d0.ost-pools/d14/dir1/file46 not allocated from OSTs 0.
Attachment: 28540
42. Bug 21097: metadata-updates : @@@@@@ FAIL: md5sum verification failed
Attachment: 28381
43. Bug 21106: client softlock up in rm command during recovery-mds-scale
Attachment: 28261
44. Bug 21147: mdsrate-create-large test_4: @@@@@@ FAIL: quotacheck has failed: EBUSY
Attachment: 28099
45. Bug 21255: parallel-scale: short timeout
Attachment: 28661
46. Bug 21278: ASSERTION(fid_is_igif(dst) || fid_ver(dst) == 0) failed: [0x6d6d69747465640a:0x7f000000:0x181ffff] running fs_test
Attachment: 28034 Attachment: 28363
47. Bug 21380: make dist seems to exclude the "darwin" bits
Attachment: 28635
48. Bug 21396: LNET soft lockups in socknal_cd thread
Attachment: 28397
49. Bug 21406: Applications stuck in jbd2_log_wait_commit during exit
Attachment: 28090
50. Bug 21411: recovery-small test_17 hang
Attachment: 28141
51. Bug 21454: cl_io_prepare_write()) ASSERTION(cl_page_in_io(page, io)) failed
Attachment: 28955
52. Bug 21476: Unable to handle kernel NULL pointer dereference at 0000000000000000
RIP: [<ffffff88a7aca4>] :mds: __mds_lov_synchronize+0x444/0x2740
Attachment: 29166
53. Bug 21485: kernel oops in mdt_handle_common, replay-single test 70b
Attachment: 28706 Attachment: 28750
54. Bug 21486: ASSERTION(ev->md.start == req->rq_rebuf) failed
Attachment: 28398
55. Bug 21500: 2.6.31-fc12 patchless client support
Attachment: 27315 Attachment: 27313 Attachment: 27311
56. Bug 21501: posix lock does not failed at eviction
Attachment: 28709 Attachment: 28709 reverted Attachment: 28709 (again)
57. Bug 21502: symlink compatibility between 1.6 and 2.0
Attachment: 29314
58. Bug 21527: use RPMs %configure macro in our lustre.spec file
Attachment: 28250 Attachment: 29179 Attachment: 29402
NOTE: There is a 4th landing referenced in the GIT log but no specific entry for it in the

bug.

59. Bug 21551: racer, (cl_page.c:233:cl_page_gang_lookup()) slice != NULL
Attachment: 28178
60. Bug 21564: e2fsck should warn when MMP update interval is extended
Attachment: 27897
61. Bug 21636: performance-sanity.sh: FAIL: test_8 test_8 failed with 1; client evicted due to a lock blocking callback time out
Attachment: 29219
62. Bug 21666: Building lustre fails on SUSE10 SP2 with Ofed v1.4.1
Attachment: 27932
63. Bug 21669: RHEL4, RHEL5 & OEL5 ia64 builds fail for both patched and patchless
Attachment: 27713
64. Bug 21718: Unable to handle kernel NULL pointer dereference at 0000000000000010
RIP: [<ffffff88742699>] :obdclass:lu_context_key_get+0x49/0x60
Attachment: 28162
65. Bug 21738: conf_sanity.sh: error: conf_param: No such device
Attachment: 27829
66. Bug 21754: suse releases rpms which have a different name than their version
Attachment: 27838
67. Bug 21829: llobdstat is broken on Lustre 1.8.1
Attachment: 28152
68. Bug 21835: performance_sanity: No space left on device
Attachment: 28440
69. Bug 21853: Build problem with cfs_spin_lock_nested when CONFIG_DEBUG_LOCK_ALLOC=y
Attachment: 28165
70. Bug 21877: ASSERTION(list_empty(&req->rq_timed_list)) failed
Attachment: 29131
71. Bug 21881: mdt_num_threads tuning
Attachment: 28465
NOTE: An additional landing was done but no bugzilla reference for it. It appears to be for the same patch.
72. Bug 21882: lru resize SLV can get stuck
Attachment: 28243 Attachment: 28382
73. Bug 21888: simul: timeout
Attachment: 29099
74. Bug 21900: ost-pools test_25: @@@@ @@@@ FAIL: /mnt/lustre/d0.ost-pools/d25/file1 not allocated from OSTs 0.
Attachment: 28206 Attachment: 28791
75. Bug 21911: kiblnd_check_sends() should always be called with a ref on conn (was: double release of ibc_lock in o2iblnd)
Attachment: 28360
76. Bug 21919: quota -v hangs when OST's are failing over
Attachment: 28735
77. Bug 21927: interop: replay-vbr test_8e: @@@@ @@@@ FAIL: Restart of mds failed
Attachment: 28902
78. Bug 21932: parallel-scale test_ior: @@@@ @@@@ FAIL: ior failed! 1, no space left on device
Attachment: 28375
79. Bug 21948: NFSv3 parallel-scale test_parallel_grouplock subtests failed: Inappropriate ioctl for device.
Attachment: 28240
80. Bug 21957: debug_mb not initialized on Lustre 2.0 client node (kernel 2.6.31)
Attachment: 28945
81. Bug 21966: can not mount mdt after umount
Attachment: 28290

82. Bug 21968: conf_sanity mount.lustre: mount /dev/sdb1 at /mnt/mds1 failed: Operation already in progress
Attachment: 28318
83. Bug 21982: long file truncate problem
Attachment: 28326
84. Bug 21986: racer: MDS crashed at mdt_reint_open()) ASSERTION(!created) failed
Attachment: 28732
85. Bug 21988: test 23: client hang when close a file after mds crash
Attachment: 28473
86. Bug 22030: "link_ea add failed -28" error messages during testing
Attachment: 28718
87. Bug 22033: IOR data check error
NOTE: landing to make ext3 default but no reference in the bugzilla report.
88. Bug 22069: ost-pools test_23: FAIL: dd did not fail with EDQUOT
Attachment: 29057
89. Bug 22070: conf-sanity test 20 hung at remounting 'ro' lustre on /mnt/lustre
Attachment: 28627
90. Bug 22074: Incorrect triggering of synchronous IO by OSC
Attachment: 20603
91. Bug 22079: Moving definition of cfs_trace_buf_type_t conflicts with Darwin/Winnt definition
Attachment: 28807
92. Bug 22097: mount.lustre fails to pass some options to mount()
Attachment: 28644
93. Bug 22126: ost-pools: multiple tests failed
Attachment: 28849
94. Bug 22127: lustre 1.8.2 lfs permissions
Attachment: 28780
95. Bug 22147: Fix a typo in echo client which causes echo client always sends INTENT dlm lock
Attachment: 28639
96. Bug 22169: recovery_mds_scale : failed
Attachment: 28756 Attachment: 28757
97. Bug 22181: interval_erase does not update maxhigh properly
Attachment: 28681
98. Bug 22187: NULL dereference in ll_setxattr() if value = NULL
Attachment: 29196
99. Bug 22190: Oops: RIP: obdfilter:filter_map_remote_to_local+0x1df/0x320
Attachment: 29362
100. Bug 22225: cleanup test banners
Attachment: 28729 Attachment: 29138
101. Bug 22235: llapi_uuid_match() prints bogus error message on upgraded filesystem
Attachment: 29151 Attachment: 29153
102. Bug 22258: remove "old school" build
Attachment: 28855
103. Bug 22261: sanity test_34f: @@@@ @ @ FAIL: /tmp/f34f not 48000 bytes
Attachment: 28876
104. Bug 22272: recovery_random_scale failed on OSS for ASSERTION(led->led_lcd)
Attachment: 28851
105. Bug 22273: recovery_random_scale failed on OSS for ASSERTION(led->led_lcd)
Attachment: 28893
106. Bug 22299: "lctl --device %lustre-MDT0000 readonly" does not work properly
Attachment: 29169

107. Bug 22312: write_bytes are missing in obdfilter stats
Attachment: 29331 Attachment: 29427
NOTE: Another landing is referenced in the GIT log which appears related to attachment 29331.
108. Bug 22319:statahead in parallel-scale-nfsv3 failed
Attachment: 28986
109. Bug 22327: "lfs df" does not print stats for all mountpoints
Attachment: 29333
110. Bug 22334: swap the declarations for OBD_FAIL_MDS_READLINK_EPROTO and OBD_FAIL_MDS_OPEN_WAIT_CREATE
Attachment: 29137
111. Bug 22342: do_statahead_enter() ASSERTION(dchild == *dentryp) failed
Attachment: 29011
112. Bug 22359: LBUG: (filter.c:1942:filter_iobuf_get()) ASSERTION(thread_id < filter->fo_iobuf_count) failed
Attachment: 29220
113. Bug 22403: Patchless client build fails for sles11 and fc11
Attachment: 29183
114. Bug 22470: Add header files to doxygen groups.
Attachment: 29249
115. Bug 22505: liblustre test_1: @@@@ @ FAIL: test_1 failed with 134
Attachment: 29403
116. Bug 22526: num_possible_cpus() should be cfs_num_possible_cpus()
Attachment: 29376

Results

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

The Lustre 2.0 Beta1 exhibits the following characteristics:

- Testing completed for RHEL5/x86_64 and OEL5/x86_64 using IB and TCP
- Passing consistently: all daily tests are passing consistently on small clusters
- Known Failures: NFS, Interoperability

Lustre 2.0 Beta1 was delivered on April 10, 2010 and is available from:

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

or

http://downloads.lustre.org/public/lustre/v2.0/beta/Lustre_2.0_Beta1