

# Lustre 2.2 Changelog

## version 2.2.0

### Support for networks:

o2iblnd - OFED 1.5.4

### Server support for kernels:

2.6.32-220.4.2.el6 (RHEL6)

### Client support for unpatched kernels:

2.6.18-274.18.1.el5 (RHEL5)

2.6.32-220.4.2.el6 (RHEL6)

2.6.32.36-0.5 (SLES11)

### Recommended e2fsprogs version:

1.41.90.wc4

### Known Issues in 2.2.0:

[LU-1188](#) Panics can occur running sanity test while running with kernel debug options turned on

[LU-1191](#) Due to a known NFS-related bug in the kernel - [https://bugzilla.kernel.org/show\\_bug.cgi?id=38572](https://bugzilla.kernel.org/show_bug.cgi?id=38572) - users wanting to re-export Lustre via NFS should apply the kernel patch for this issue

[LU-1244](#) MDS survey tool has a known problem of not deleting OST objects when run in unlink mode, which might lead to unexpected space and inode shortage on OSTs in case of multiple runs

### Issues fixes between 2.1.0 and 2.2.0

#### Bug

- [LU-78](#) - kiblnd\_check\_conns can deadlock
- [LU-81](#) - Some JBD2 journaling deadlock at BULL
- [LU-106](#) - unable to handle kernel paging request in lprocfs\_stats\_collect()
- [LU-110](#) - Allow the use of a 2.0 MDS on a filesystem with a 1.8 disk format
- [LU-176](#) - disable mb\_cache by default
- [LU-233](#) - "Russian doll" test log tarballs
- [LU-322](#) - Fix for ost pool test in test suite
- [LU-357](#) - port bug24451 to master(racer test defect)
- [LU-362](#) - lu\_kmem\_init() needs error handling
- [LU-412](#) - auster found /usr/sbin/lsfck instead of lfsck.sh
- [LU-417](#) - block usage is reported as zero by stat call for tens of seconds after creating a file
- [LU-424](#) - conf-sanity test 55, 56, 58 do not work with separate MGS and MDT
- [LU-427](#) - Test failure on test suite lfsck
- [LU-431](#) - retry interval for failed memory allocating in o2iblnd is too long
- [LU-443](#) - Only squawk when md->start is NULL on non-zero length
- [LU-462](#) - lut\_client\_alloc/free() shouldn't be called for self\_export
- [LU-488](#) - Lustre-Error: 4564:0:(connection.c:100:ptlrcp\_connection\_put()) ASSERTION(!cfs\_hlist\_unhashed(&conn->c\_hash)) failed
- [LU-499](#) - grant\_rate and cancel\_rate are static when an OST is idle
- [LU-508](#) - Kernel panic on ...BUILD/BUILD/lustre-ldiskfs-3.3.0/ldiskfs/extents.c:1920
- [LU-513](#) - cfs\_wait\_event\_interruptible\_exclusive is not exclusive
- [LU-515](#) - gpdd-survey.sh/sgpdd-survey doesn't play nice with devices specified by alias
- [LU-542](#) - Missing XATTR changelog record
- [LU-543](#) - Missing UNLINK record on overwritting rename
- [LU-574](#) - conf-sanity test\_41b: @@@@ FAIL: test\_41b failed with 1
- [LU-601](#) - kernel BUG at fs/jbd2/transaction.c:1030
- [LU-610](#) - port bug24512 to master(lfs find -s doesn't seem to work quite with >2GB args)
- [LU-613](#) - Lustre-Client dead-lock during binary exec() over Lustre FS
- [LU-617](#) - LBUG: (mdt\_recovery.c:787:mdt\_last\_rcvd\_update()) ASSERTION(req\_is\_replay(req)) failed
- [LU-620](#) - "Bad page state" reported after unlink
- [LU-625](#) - sanity test\_32a test failure
- [LU-629](#) - \_debug\_req() never prints out opcode flags/status
- [LU-633](#) - mds-survey script for MD echo client test
- [LU-635](#) - conf-sanity test\_29: @@@@ FAIL: test\_29 failed with 5
- [LU-639](#) - replay-dual test\_0b: @@@@ FAIL: mount1 fails
- [LU-645](#) - getcwd fails

- LU-646 - clarification of lustre fsync behavior
- LU-649 - client receives reply to lock cancel yet server evicts it 100 seconds later
- LU-651 - Strange error message "is it ok to have flags 0x... and 0x... in the same brw?"
- LU-665 - Reduce unnecessary connection flags MDSs send to OSSs
- LU-666 - mdd depends on ldiskfs-only dynlocks feature
- LU-668 - Rename MDS\_FL to LUSTRE\_FL and update usage
- LU-669 - mea.c unconditionally uses ldiskfs hashes
- LU-672 - mdd had panic on getting version for non existent file.
- LU-685 - Wide busy lock in kiblnd\_pool\_alloc\_node
- LU-686 - conf-sanity test\_52: @@@@ FAIL: Unable to mount ost1 as ldiskfs
- LU-687 - Application is OOM-killed during page-fault resolution on its binary over Lustre when there is plenty of memory available
- LU-690 - port bug 24498 (conf-sanity test\_21a: @@@@ FAIL: can't put import for lustre-OST0000-osc(ost) into FULL state)
- LU-691 - sanity test\_133c defect
- LU-695 - port bz18750 attachment#32903(clean up for OBD\_FAIL\_CHECK\_ONCE)
- LU-699 - replay-dual test\_1 fails to remount mdt
- LU-700 - OSC BRW stats may be incorrect
- LU-707 - get\_version() in replay-vbr does not quote the FID string argument for do\_facet()
- LU-715 - Value read from qos\_threshold\_rr file does not match value written
- LU-717 - ext4 mmap block uses incorrect utsname field for nodename
- LU-719 - ldiskfs: MMP: kmmpd will print only one error message for any number of failed MMP block writes during its lifetime
- LU-721 - Parallel writes to same file results in a file of zeroes
- LU-728 - pdsh HOSTLIST expression support in test sutie
- LU-731 - export ldlm\_lock\_set\_data and implement it correctly
- LU-734 - improve recovery-\*scale tests to put debug logs into \$LOGDIR
- LU-743 - Test failure on test suite conf-sanity, subtest test\_46a
- LU-745 - ost-pools test 23 hung
- LU-746 - obdfilter-survey FAIL: test\_1b ost4: hndls expected > 8, have 2
- LU-748 - sanity test 220 must be placed in SLOW set
- LU-753 - kernel:LustreError: 4998.0:(filter\_io\_26.c:791:filter\_commitrw\_write()) LBUG
- LU-759 - mdc\_enqueue() should clear rq\_replay when mdc\_enter\_request() fails
- LU-760 - Test failure on test suite sanity, subtest test\_180a
- LU-761 - Return values of ll\_prep\_md\_op\_data() should be checked using IS\_ERR().
- LU-762 - Hyperion - mdtest failure
- LU-766 - Incorrect HAVE\_STATFS\_DENTRY\_PARAM result
- LU-769 - Test failure on test suite conf-sanity, subtest test\_23a
- LU-770 - sanity.sh fails in CLIENTONLY mode
- LU-773 - Test failure on test suite sanity, subtest test\_105b
- LU-774 - Test failure on test suite replay-ost-single, subtest test\_7
- LU-780 - improve parallel-scale to support hyperion run
- LU-791 - sanity 27q hang
- LU-795 - per-transaction commit callbacks (ORI-107 port)
- LU-797 - Test failure on test suite ost-pools, subtest test\_14, test\_18, test\_23
- LU-803 - Test failure on test suite obdfilter-survey, subtest test\_3a
- LU-805 - Ifs quota should show quota allocation rather than used
- LU-810 - Fix helpers for reading flags in HSM changelog records
- LU-817 - sgpd-d-survey is encountering r/w errors on arrays using 2TB drives
- LU-819 - "Ifs getstripe -M" broken
- LU-820 - Annoying "format at cl\_io.c:1551:cl\_req\_page\_add doesn't end in newline"
- LU-824 - fix corrupted ldiskfs after md rebuild
- LU-826 - Invalid memory free at mdd\_lov\_create\_finish.
- LU-831 - struct bit field type should be unsigned type
- LU-834 - fix page\_is\_vmlocked method for echo\_client
- LU-837 - missed semaphore release in gss\_do\_ctx\_init\_rpc
- LU-839 - conf\_sanity.sh test\_52 fails with loop devices
- LU-841 - sanity.sh 27q does not create a testing directory
- LU-843 - Several minor issues
- LU-844 - LBUG: ASSERTION(iobuf->dr\_npages < iobuf->dr\_max\_pages) when run obdfilter\_survey using rsz >= 2M
- LU-864 - Hostname name doesn't not equal NID and use facet\_mntpt instead of \$MOUNT%
- LU-867 - Fails to build with gss (lustre 2.1)
- LU-874 - Client eviction on lock callback timeout
- LU-878 - Test failure on test suite sanityn, subtest test\_40b
- LU-881 - shmem\_bump() notes "running" bumped without lock control
- LU-882 - Quota code causes run away ll\_ost\_io thread
- LU-887 - master (2.2) clients incompatible with 2.1 servers?
- LU-888 - lctl get\_params easily confused
- LU-889 - rework extent\_recovery\_timer()
- LU-898 - When copy data to 2.1 cluster via 1.8 client, all 2.1 OSS servers panic.
- LU-900 - Using ncli.sh as configuration file without SRUN installed will cause test failure
- LU-915 - MKFSOPTS fix for test suite
- LU-919 - Multiple wrong LBUGs checking cfs\_atomic\_t vars/fields with inaccurate poison value of 0x5a5a5a
- LU-924 - Test failure on test suite recovery-small, subtest test\_105
- LU-927 - cleanup ptlrpc\_thread::t\_flags processing
- LU-931 - fstress - LBUG (lustre\_idl.h:766:lu\_fid\_eq()) ASSERTION(fid\_is\_igif(f0) || fid\_ver(f0) == 0)
- LU-932 - cleanup ll\_inode\_info to reduce inode size
- LU-934 - external xattr inode support in e2fsprogs
- LU-935 - Crash lquota:dquot\_create\_oqaq+0x28f/0x510
- LU-948 - Client recovery hang
- LU-949 - sanity 42e failed
- LU-952 - Hung thread with HIGH OSS load

- LU-955 - Lustre-modules 2.1 client rpm wrongly requires lustre-backend-fs
- LU-958 - tests: debug\_mb set incorrectly for smp or vm
- LU-960 - Bad stripe count reported by lfs getstripe when default stripe count set to -1
- LU-966 - post-fsck MDS LBUG during recovery due to missing FID
- LU-969 - 2.1 client stack overruns
- LU-971 - Re-sent open requests cause mod\_count leakages
- LU-972 - server shouldn't log a message to console each time a client connects
- LU-976 - Test failure for mmap on ZFS filesystems
- LU-980 - incompatible error handling
- LU-983 - On a 1.8.6 client/host, tar is performing 10K reads, and the RPCs are typically one page in size for small files==tar is slow for small files using 1.8.6 vs 1.8.5
- LU-985 - buffer overflow resulting in stack corruption
- LU-987 - Fail to create ldisk rpms RHEL5
- LU-993 - 1.8->2.1.54 Test failure on test suite sanity :osd\_attr\_set() ASSERTION((oh) ->ot\_declare\_attr\_set > 0) failed
- LU-997 - port the fix of 'max nested link count' from upstream e2fsprogs to master-lustre branch
- LU-998 - Test failure on test suite sanity, subtest test\_103
- LU-1014 - MGS with sys.timeout is ignored and if sys.timeout is changed its not synced across the file system.
- LU-1017 - MDS oops when running racer test
- LU-1023 - strange numbers shown in obdfilter-survey
- LU-1028 - Bus error (core dumped) during fsx test
- LU-1029 - EL6 kernel 2.6.32-220.el6 fails to build with OFA OFED 1.5.4
- LU-1039 - data corruption in check\_set
- LU-1042 - 1.8 clients show wrong dates with 2.1 servers
- LU-1046 - fsstress - Watchdog detected hard LOCKUP on cpu 0
- LU-1048 - Test failure on test suite conf-sanity
- LU-1049 - lc\_net does not parse unexpected output from pdsh well
- LU-1050 - configure broken on centos6.2 with OFED-1.5.4
- LU-1051 - removing file failed due to no space left on device
- LU-1053 - git commits fail when deleting files in the repo
- LU-1055 - recovery-mds osd\_handler.c:1966:osd\_declare\_object\_destroy() ASSERTION(!lu\_object\_is\_dying(dt->do\_lu.lo\_header)) failed
- LU-1059 - vvp\_page\_unmap() ASSERTION(PageLocked(vmpage))
- LU-1060 - Test failure on test suite replay-vbr, subtest test\_7c
- LU-1061 - LBUG on cl\_locks\_prune()
- LU-1067 - Crash on mds-survey. osd\_child\_dentry\_get+0x3c/0xb0 osd\_ldisks
- LU-1069 - 2.1 <-> 2.1.55 Test failure on test suite sanity, subtest test\_103
- LU-1070 - Inconsistent CLIO lock state maybe cause cl\_lock cannot be cached
- LU-1073 - mandatory (LOCK\_MAND) flock requests cause unknown fcntl lock type: 32 Lustre error messages
- LU-1074 - servers are umounted while the client is still doing job when running racer.sh by auster
- LU-1075 - cannot run racer.sh by auster since there is a dir with the same name racer
- LU-1080 - mds-survey crash
- LU-1082 - add mds-survey to lustre-tests
- LU-1084 - Console messages from sec stuff MUST GO
- LU-1091 - tag matching not quite working
- LU-1098 - Object doesn't exist!
- LU-1109 - NFS server not responding when running parallel-scale test\_iorsff
- LU-1122 - MDS got LBUG after upgrade from 1.8.7-RHEL5 to 2.1.55-RHEL6 and then downgrade to 2.1.0-RHEL6
- LU-1125 - crash in lc\_watchdog\_del\_pending during tgt\_recov
- LU-1128 - Two ll\_ost\_io threads apparently stuck in shrinker code
- LU-1133 - sanity-quota 33 fails
- LU-1163 - Never try to invalidate a dirty page
- LU-1206 - Executing unwriteable, recently touched file returns ETXTBSY (Text file busy)
- LU-1212 - On MDS startup upon client connection mdt\_xx threads Consume All Available CPU
- LU-1213 - Duplicate setting of mdtest\_nFiles - parallel-scale.sh
- LU-1222 - MDS crashes while running simul

## Improvement

- LU-32 - always keep buffer\_head reference on root node of iam\_container
- LU-143 - hash function of ldlm\_namespace and lu\_site is not good enough
- LU-241 - support crc32c with hardware accelerated instruction as one of lustre checksums
- LU-250 - Router random shuffling support.
- LU-258 - obdfilter-survey needs cleanup
- LU-278 - Tag check is broken in configure script
- LU-335 - SuSE 11.1 and SuSE server pack 1 support for Lustre 1.8 and 2.X
- LU-433 - remove jbd2-jcberr patch from kernel
- LU-448 - add lstat --count
- LU-459 - silence excess 2.x error messages
- LU-527 - Update RHEL6 kernel to 2.6.32-131.6.1 security update
- LU-532 - 2.1 does not filter extended attributes list based on permissions, showing entire list at all times
- LU-537 - LNET fails to start with aliased interface in configuration
- LU-553 - Add checkpatch.pl to validate patches for Lustre coding style
- LU-554 - add gnlnd awareness to LNet
- LU-555 - ll\_have\_md\_lock() code optimization
- LU-570 - find connection uid by peer nid
- LU-571 - Remove Parallel AST limits
- LU-572 - Make ptrpc\_recover\_import async
- LU-575 - lfs find fixes
- LU-579 - add CLOSE event into changelog
- LU-611 - Add lfs find --stripe\_count --stripe\_index --stripe\_count

- LU-624 - Kernel update [RHEL6.1 2.6.32-131.17.1.el6](#)
- LU-641 - Use cfs\_rwlock\_t for obd\_dev\_lock instead of spinlock
- LU-663 - Some architectures do not have NUMA features anymore
- LU-680 - superblock <-> mount point mapping
- LU-723 - External ldiskfs package support
- LU-735 - Remove fixme() macro
- LU-737 - can not attach device name starting with a digit
- LU-778 - RHEL 6.2: rdma\_create\_id() takes four arguments
- LU-781 - Kernel update [RHEL6.2 2.6.32-220.el6](#)
- LU-792 - lbuild-rhel5 should use redhat's SRPM repo
- LU-794 - allow larger O\_DIRECT requests with large allocations
- LU-814 - automate NFSv3/v4 over Lustre Testing
- LU-835 - add files to .gitignore
- LU-846 - add Ifs\_migrate.1 man page to list of installed pages
- LU-847 - Merge some quota patches from Orion to 2.2
- LU-848 - clio page writeback support
- LU-871 - code cleanups for clang static analysis
- LU-879 - Detail rename operation in stats file
- LU-883 - commit-msg disallows common reference tags
- LU-905 - keep BH reference on root node of IAM container
- LU-928 - Update or improvement to a code comment
- LU-930 - Update or improvement to a Lustre man page
- LU-954 - lbuild should produce lustre-iokit rpms
- LU-962 - Add a feature to run a callback in ptlrpcd context
- LU-963 - update lustre/ChangeLog to reflect new e2fsprogs
- LU-991 - cleanup LMV for DNE
- LU-1025 - Lustre crc32c implementation not use final bit inversion
- LU-1054 - Kernel update [RHEL6.2 2.6.32-220.4.1.el6](#)
- LU-1077 - build/autoconf/lustre-build-ldiskfs.m4 missing in tarball
- LU-1081 - LUSTRE\_BFL\_FID has a conflicting namespace
- LU-1095 - Console message cleanup
- LU-1120 - Kernel update [RHEL6.2 2.6.32-220.4.2.el6](#)
- LU-1146 - update copyright for all WC modified files in master before 2.2 release
- LU-1195 - ll\_sai\_put() ASSERTION(thread\_is\_stopped(&sai->sai\_agl\_thread)) failed

## New Feature

- LU-50 - pdirops patch for ldiskfs
- LU-80 - Wide striping support
- LU-389 - Statahead
- LU-580 - Imperative Recovery
- LU-593 - echo client for MDS
- LU-822 - allow multiple Object Index files to be created
- LU-884 - Client In-Memory Data Checksum
- LU-909 - changes to OSD API
- LU-925 - land Async Glimpse Lock (AGL) to master to pre-fetch size/block attributes from OSTs for statahead
- LU-1019 - general ptlrpcd pool support to process kinds of async RPC efficiently on SMP client