




# Toward patchless ldiskfs kernels

**Andreas Dilger, High Performance Data Division**

# Patch changes by release



# Current RHEL 7 patch series

ext4-print-inum-in-htree-warning.patch

ext4-remove-truncate-warning.patch

ext4-hash-indexed-dir-dotdot-update.patch

ext4-mballoc-extra-checks.patch

ext4-mballoc-pa-free-mismatch.patch

ext4-disable-mb-cache.patch

ext4-inode-version.patch

ext4-lookup-dotdot.patch

ext4-osd-iop-common.patch

ext4-kill-dx-root.patch

ext4-data-in-dirent.patch

ext4-large-eas.patch

ext4-nocmtime.patch

ext4-pdirop.patch

ext4-max-dir-size.patch

ext4-prealloc.patch

ext4-misc.patch

14 files changed, 2428 insertions(+), 249 deletions(-)

# Could land with update to latest kernel

`ext4-print-inum-in-htree-warning.patch`

`ext4-remove-truncate-warning.patch`

`ext4-mballoc-extra-checks.patch`

`ext4-mballoc-pa-free-mismatch.patch`

`ext4-kill-dx-root.patch`

`ext4-nocmtime.patch`

`ext4-disable-mb-cache.patch`

`ext4-max-dir-size.patch`

## Maybe drop

`ext4-prealloc.patch` - mballoc tunables that are likely no longer needed

`ext4-lookup-dotdot.patch` - Need to determine if still needed after upstream fixes

14 files changed, 430 insertions(+), 149 deletions(-)

# Could land with rework/enhancement

`ext4-large-eas.patch` - medium effort

- Needs support for 64KB direct-addressed xattrs (LU-908)

`ext4-inode-version.patch` - medium effort

- Could rework `i_version` increment to `ext4_inc_iversion()` wrapper
- `ldiskfs` would only need to patch `ext4_inc_iversion()`

`ext4-pdirop.patch` - already done

- Being split into two patches (<http://review.whamcloud.com/14264>)
- 3-level htree + >2GB dir size likely to be accepted upstream
- `ext4` parallel dir locking needs VFS-level parallel dir locking
  - Major amount of effort, won't land upstream without it

14 files changed, 2213 insertions(+), 156 deletions(-)

# Unlikely to land upstream

[ext4-osd-iop-common.patch](#)

- EXPORT\_SYMBOL()
- Two isolated functions, could move to osd-lldiskfs?

[ext4-data-in-dirent.patch](#)

- Patch largely ready, need to find upstream use case for it first

[ext4-misc.patch](#)

- Mostly EXPORT\_SYMBOL()

EXPORT\_SYMBOL() is easily patched or worked around if needed

# Conclusion

Unlikely to become completely patchless in near future

Reducing patch count is good anyway, makes updates easier

LU-6030 cleanup for ldiskfs patches

- <https://jira.hpdd.intel.com/browse/LU-6030>

LU-6220: push ext4/ldiskfs patches upstream if possible

- <https://jira.hpdd.intel.com/browse/LU-6220>