



Lustre Kernel Patch

所有描述基于 Lustre 2.8 2015.10

概述

- 为什么需要 kernel patch
- 两大部分，kernel 和 ldiskfs
- kernel patch 的类型
 - 修复bug (来自upstream或者其它途径)
 - 增加功能
 - 提高性能
- 尽力减少patch的数量
- Patchless Client早已实现。 Patchless Server

历史

- 两个主要的发行版本
 - RHEL
 - RHEL5, RHEL6, RHEL7
 - SLES
 - SLES9, SLES10, SLES11, SLES12
- 组织方式和工具
- 多版本共存

RHEL6 patch

目前应用最广泛的版本

Kernel patch (rhel6)

9个patch

- 服务于性能5个
- 服务于功能的2个
- bug fix 2个

2.6-rhel6.series

mpt-fusion-max-sge-rhel6.patch
raid5-mmp-unplug-dev-rhel6.patch
dev_read_only-2.6.32-rhel6.patch
blkdev_tunables-2.6-rhel6.patch
bh_lru_size_increase.patch
quota-replace-dqptr-sem.patch
quota-avoid-dqget-calls.patch
jbd2-log_wait_for_space-2.6-rhel6.patch
module-load-deadlock-rhel6.patch

Ldiskfs patch (rhel6)

47个patch

- 大约8700多行
- 功能 23 个
- 性能 7 个
- Bug fix 17 个

Ldiskfs patch (rhel6)

第一部分：功能

- data-in-directory-entry
- journal-callback
- large-eas
- disable-mbcache
- preallocation

ldiskfs-2.6-rhel6.7.series

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

ext4-kill-dx_root.patch

ext4-data-in-dirent.patch

ext4-journal-path-opt.patch

ext4-large-eas.patch

ext4-prealloc.patch

ext4-journal-callback.patch

ext4-disable-mb-cache.patch

Ldiskfs patch (rhel6)

第二部分：性能

- extent-walk-space
- large-directory
- parallel-directory-operation

ldiskfs-2.6-rhel6.7.series

ext4-ext_generation.patch
ext4-ext-walk-space.patch
ext4-store-tree-generation-at-find.patch
ext4-map_inode_page-2.6.18.patch
ext4-large-dir.patch
ext4-pdirop.patch
ext4-pdir-fix.patch

Ldiskfs patch (rhel6)

第三部分：其它

- Bug fix
- 导出变量或函数
- 清理和改进

ext4-remove-truncate-warning.patch
ext4-inode_info_reorganize.patch
ext4-print-inum-in-htree-warning.patch
ext4-nocmtime-2.6.patch
ext4-use-correct-inode.patch
ext4-fix-mbgroups-access.patch
ext4-quota-dont-update-cmtime.patch
ext4-quota-first-class.patch
ext4-nlink-2.6.patch
ext4-extra-isize.patch
ext4-fix-ext4_mb_add_n_trim.patch
ext4-max-dir-size.patch

RHEL 7 patch

即将使用的版本

Kernel patch (rhel7)

3个patch

- 得益于patch并入upstream
- dev_read_only 是唯一必需的
- 基于dm-flakey的替代方案
- 稳定性仍旧有待改进
- 在生产环境可实现patchless server

3.10-rhel7.series

raid5-mmp-unplug-dev-3.7.patch
dev_read_only-3.7.patch
blkdev_tunables-3.7.patch

Ldiskfs patch (rhel7)

19个patch

- patch减少到5900多行
- 尝试合并更多的patch到upstream
- 仍旧需要广泛的测试

ldiskfs-3.10-rhel7.series

ext4-print-inum-in-htree-warning.patch
ext4-large-eas.patch
ext4-disable-mb-cache.patch
ext4-large-dir.patch
ext4-kill-dx-root.patch
ext4-pdirop.patch
ext4-prealloc.patch
ext4-remove-truncate-warning.patch

