
Subject: [PATCH 18/25] elevate write count for do_sys_utime() and touch_atime()
Posted by [Dave Hansen](#) on Mon, 11 Dec 2006 22:30:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

Signed-off-by: Dave Hansen <haveblue@us.ibm.com>

lxc-dave/fs/inode.c | 20 ++++++++-----
lxc-dave/fs/utimes.c | 15 ++++++++-----
2 files changed, 23 insertions(+), 12 deletions(-)

diff -puN fs/inode.c~17-24-elevate-write-count-for-do-sys-utime-and-touch-atime fs/inode.c
--- lxc/fs/inode.c~17-24-elevate-write-count-for-do-sys-utime-and-touch-atime 2006-12-11

14:22:08.000000000 -0800

+++ lxc-dave/fs/inode.c 2006-12-11 14:22:08.000000000 -0800

@@ -1167,22 +1167,23 @@ void touch_atime(struct vfsmount *mnt, s

```
    if (IS_RDONLY(inode))
        return;
-   if (inode->i_flags & S_NOATIME)
+   if (mnt && mnt->mnt_want_write(mnt))
        return;
+   if (inode->i_flags & S_NOATIME)
+   goto out;
    if (inode->i_sb->s_flags & MS_NOATIME)
-   return;
+   goto out;
    if ((inode->i_sb->s_flags & MS_NODIRATIME) && S_ISDIR(inode->i_mode))
-   return;
+   goto out;
```

```
/*
 * We may have a NULL vfsmount when coming from NFS
 */
if (mnt) {
    if (mnt->mnt_flags & MNT_NOATIME)
-   return;
+   goto out;
    if ((mnt->mnt_flags & MNT_NODIRATIME) && S_ISDIR(inode->i_mode))
-   return;
-
+   goto out;
    if (mnt->mnt_flags & MNT_RELATIME) {
/*
 * With relative atime, only update atime if the
@@ -1193,16 +1194,19 @@ void touch_atime(struct vfsmount *mnt, s
    &inode->i_atime) < 0 &&
    timespec_compare(&inode->i_ctime,
```

```

        &inode->i_atime) < 0)
-   return;
+   goto out;
    }
}

    now = current_fs_time(inode->i_sb);
    if (timespec_equal(&inode->i_atime, &now))
-   return;
+   goto out;

    inode->i_atime = now;
    mark_inode_dirty_sync(inode);
+out:
+ if (mnt)
+   mnt_drop_write(mnt);
}
EXPORT_SYMBOL(touch_atime);

```

diff -puN fs/utimes.c~17-24-elevate-write-count-for-do-sys-utime-and-touch-atime fs/utimes.c
--- lxc/fs/utimes.c~17-24-elevate-write-count-for-do-sys-utime-and-touch-atime 2006-12-11

```

14:22:08.000000000 -0800
+++ lxc-dave/fs/utimes.c 2006-12-11 14:22:08.000000000 -0800
@@ -3,6 +3,7 @@
#include <linux/linkage.h>
#include <linux/namei.h>
#include <linux/sched.h>
+#include <linux/mount.h>
#include <linux/utime.h>
#include <asm/uaccess.h>
#include <asm/unistd.h>
@@ -32,6 +33,10 @@ asmlinkage long sys_utime(char __user *
    goto out;
    inode = nd.dentry->d_inode;

+   error = mnt_want_write(nd.mnt);
+   if (error)
+   goto mnt_drop_write_and_out;
+
    error = -EROFS;
    if (IS_RDONLY(inode))
        goto dput_and_out;
@@ -41,7 +46,7 @@ asmlinkage long sys_utime(char __user *
    if (times) {
        error = -EPERM;
        if (IS_APPEND(inode) || IS_IMMUTABLE(inode))
-   goto dput_and_out;
+   goto mnt_drop_write_and_out;

```

```

    error = get_user(newattrs.ia_atime.tv_sec, &times->actime);
    newattrs.ia_atime.tv_nsec = 0;
@@ -49,21 +54,23 @@ asmlinkage long sys_utime(char __user *
    error = get_user(newattrs.ia_mtime.tv_sec, &times->modtime);
    newattrs.ia_mtime.tv_nsec = 0;
    if (error)
-   goto dput_and_out;
+   goto mnt_drop_write_and_out;

    newattrs.ia_valid |= ATTR_ATIME_SET | ATTR_MTIME_SET;
} else {
    error = -EACCES;
    if (IS_IMMUTABLE(inode))
-   goto dput_and_out;
+   goto mnt_drop_write_and_out;

    if (current->fsuid != inode->i_uid &&
        (error = vfs_permission(&nd, MAY_WRITE)) != 0)
-   goto dput_and_out;
+   goto mnt_drop_write_and_out;
}
mutex_lock(&inode->i_mutex);
error = notify_change(nd.dentry, &newattrs);
mutex_unlock(&inode->i_mutex);
+mnt_drop_write_and_out:
+ mnt_drop_write(nd.mnt);
dput_and_out:
path_release(&nd);
out:
-

```

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
