Subject: Re: [Q] missing unused dentry in prune_dcache()?
Posted by vaverin on Fri, 27 Oct 2006 06:50:45 GMT
View Forum Message <> Reply to Message

Vasily Averin wrote:
> David Howells wrote:
>> Vasily Averin <vvs@sw.ru> wrote:
>>
>>> I've noticed one more minor issue in your patch: in
>>> shrink_dcache_for_umount_subtree() function you decrement
>>> dentry_stat.nr_dentry without dcache_lock.
>> How about the attached patch?
> I'm sorry, but your patch is wrong:
> you have mixed calculation of 2 variables:
> dentry_stat.nr_unused -- were correct, it was decremented under dcache_lock.
> dentry_stat.nr_dentry -- were incorrect, it was decremented without dcache_lock.
>
> You should correct dentry_stat.nr_dentry, but instead you broke calculation of
> dentry_stat.nr_unused.
>
> I've fixed this issue by following patch.
corrected version, extra space were removed

Thank you,
 Vasily Averin

VFS: Fix an error in dentry_stat.nr_dentry counting

From: Vasily Averin <vvs@sw.ru>

Fix an error in dentry_stat.nr_dentry counting in
shrink_dcache_for_umount_subtree() in which the count is modified without the
dcache_lock held.

Signed-Off-By: Vasily Averin <vvs@sw.ru>

```
--- linux-2.6.19-rc3/fs/dcache.c.nrdntr 2006-10-26 15:14:51.000000000 +0400
+++ linux-2.6.19-rc3/fs/dcache.c 2006-10-27 10:45:11.000000000 +0400
@@ -554,6 +554,7 @@ repeat:
 static void shrink_dcache_for_umount_subtree(struct dentry *dentry)
 {
  struct dentry *parent;
+ unsigned detached = 0;

  BUG_ON(!IS_ROOT(dentry));

@@ -618,7 +619,7 @@ static void shrink_dcache_for_umount_sub
```

```
     atomic_dec(&parent->d_count);

    list_del(&dentry->d_u.d_child);
-   dentry_stat.nr_dentry--; /* For d_free, below */
+    detached++;

    inode = dentry->d_inode;
    if (inode) {
@@ -636,7 +637,7 @@ static void shrink_dcache_for_umount_sub
     * otherwise we ascend to the parent and move to the
     * next sibling if there is one */
    if (!parent)
-    return;
+     goto out;

    dentry = parent;

@@ -645,6 +646,11 @@ static void shrink_dcache_for_umount_sub
   dentry = list_entry(dentry->d_subdirs.next,
         struct dentry, d_u.d_child);
  }
+out:
+ /* several dentries were freed, need to correct nr_dentry */
+ spin_lock(&dcache_lock);
+ dentry_stat.nr_dentry -= detached;
+ spin_unlock(&dcache_lock);
 }

 /*
```