
Subject: Re: [Q] missing unused dentry in prune_dcache()?

Posted by [vaverin](#) on Thu, 26 Oct 2006 11:36:53 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello David

David Howells wrote:

> Vasily Averin <vvs@sw.ru> wrote:

>>>> The patch adds this dentry into tail of the dentry_unused list.

>>> I think that's reasonable. I wonder if we can avoid removing it from the
>>> list in the first place, but I suspect it's less optimal.

>> Could you please explain this place in details, I do not understand why tail
>> of the list is better than head. Also I do not understand why we should go
>> to out in this case. Why we cannot use next dentry in the list instead?

>

> I meant adding it back into the list is reasonable; I didn't actually consider
> where you were adding it back.

>

> So, given that the three ops that affect this are very unlikely to be happening
> at any one time, it's probably worth just slapping it at the head and ignoring
> it. The main thing is that it's reinserted somewhere.

I don't like to insert this dentry into tail of list because of it prevents
shrink_dcache_memory. It finds "skipped" dentry at the tail of the list, does
not free it and goes to out without any progress.

Therefore I've removed break of cycle and insert this dentry to head of the
list. Theoretically it can lead to the second using of the same dentry, however
I do not think that it is a big problem.

Signed-off-by: Vasily Averin <vvs@sw.ru>

--- linux-2.6.19-rc3/fs/dcache.c.prdch 2006-10-25 16:09:19.000000000 +0400

+++ linux-2.6.19-rc3/fs/dcache.c 2006-10-26 15:14:51.000000000 +0400

@@ -478,11 +478,9 @@ static void prune_dcache(int count, stru

up_read(s_umount);

}

spin_unlock(&dentry->d_lock);

- /* Cannot remove the first dentry, and it isn't appropriate

- * to move it to the head of the list, so give up, and try

- * later

- */

- break;

+ /* Inserting dentry to tail of the list leads to cycle */

+ list_add(&dentry->d_lru, &dentry_unused);

+ dentry_stat.nr_unused++;

}

spin_unlock(&dcache_lock);

}
