Subject: Re: [PATCH 2.6.19-rc3] VFS: per-sb dentry Iru list Posted by Neil Brown on Tue, 14 Nov 2006 04:51:08 GMT

View Forum Message <> Reply to Message

On Wednesday November 1, dev@sw.ru wrote:

> <<< AFAICS, doing so you introduced a leak of anonymous dentries.

>

- > d_alloc_anon() calls d_alloc() with parent == NULL, i.e. dentries have no parent
- > and are not linked to the sb->s root...

Yep, thanks.

- > BTW, looking at it, I found that s_anon field on super block is not
- > used any more.

I don't know what you mean by that. It is still used...

> we can add BUG_ON(!hlist_empty(&sb->s_anon)) in generic_shutdown_super to avoid such issues like this.

>

- > maybe we can fix it adding something like:
- > while (!list empty(&sb->s anon)))
- > prune_dcache(MAX_INT, &sb->s_anon);

It seems that anon dentries can now have children (I think someone explained that too me - shrink_dcache_for_umount certainly suggests it).

Also, in this context we cannot be sure that all dentries can be freed. This is being called a remount time remember, and some stuff might still be in use.

We probably need to move the s_anon list to a temporary list and repeatedly:

move the top entry back to s_anon and call shrink_dcache_parent on it.

Needs more thought.

NeilBrown