
Subject: Re: [PATCH 2.6.19-rc3] VFS: per-sb dentry lru list
Posted by [Eric Dumazet](#) on Mon, 30 Oct 2006 16:09:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Monday 30 October 2006 16:34, Kirill Korotaev wrote:

> > Quick search maybe, but your patch adds 2 pointers to each dentry in the
> > system... That's pretty expensive, as dentries are already using a *lot*
> > of ram.

>

> I don't see much problems with it... it is cache and it can be pruned if
> needed. Some time ago, for example, my patch introducing the same list for
> inodes was committed.

The ratio of dentries/PAGE is higher than ration of inodes/PAGE.

(On a x86_64 machine, 19 dentries per PAGE, 5 ext3 inodes per PAGE=

Therefore I suspect that the number of pages locked in dentry_cache because of
one inuse dentry is higher.

>

> > Maybe an alternative would be to not have anymore a global dentry_unused,
> > but only per-sb unused dentries lists ?

>

> I don't know global LRU implementation based on per-sb lists, do you?

> If someone suggest the algorithm for more or less fair global LRU

> based on non-global list we will implement it. However, so far,

> AFAICS there were problems with it.

Using 32 bytes per dentry for unused LRUs sounds too much, maybe we should
revert LRUS handling to timestamps or things like that.
