Subject: Re: [PATCH 2.6.19-rc3] VFS: per-sb dentry Iru list Posted by Neil Brown on Tue, 14 Nov 2006 05:44:46 GMT

View Forum Message <> Reply to Message

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

>

> Currently we have 3 type of functions that works with dentry_unused list:

>

- > 1) prune_dcache(NULL) -- called from shrink_dcache_memory, frees the memory and
- > requires global LRU. works well in current implementation.
- > 2) prune_dcache(sb) -- called from shrink_dcache_parent(), frees subtree, LRU
- > is not need here. Current implementation uses global LRU for these purposes, it
- > is ineffective, and patch from Neil Brown fixes this issue.
- > 3) shrink_dcache_sb() -- called when we need to free the unused dentries for
- > given super block. Current implementation is not effective too, and per-sb LRU
- > would be the best solution here. On the other hand patch from Neil Brown is much
- > better than current implementation.

>

- > In general I think that we should approve Neil Brown's patch. We (I and Kirill
- > Korotaev) are ready to acknowledge it when the following remarks fill be fixed:

>

- > it seems for me list_splice() is not required inside
- > prune_dcache(),

Yes, the list should be empty when we finish so you are right.

- > DCACHE_REFERENCED dentries should not be removed from private list to
- > dentry_unused list, this flag should be ignored if the private list is used,

Agreed.

- > count argument should be ignored in this case too, we want to free all the
- > dentries in private list,

Agreed.

- > when we shrink the whole super block we should free per-sb anonymous dentries
- > too (please see Kirill Korotaev's letter)

>

Yes. Unfortunately I don't think it is as easy as it sounds. I'll have a closer look.

- > Then I'm going to prepare new patch that will enhance the shrink dcache sb()
- > performance:

- > we can add new list head into struct superblock and use it in
- > shrink_dcache_sb() instead of temporal private list. We will check is it empty
- > in dput() and add the new unused dentries to per-sb list instead of
- > dentry_unused list.

I think that makes sense. It means that you end up doing less work in select_parent, because the work has already been done in dput.

How is the patch going?

NeilBrown