Subject: Re: [PATCH 06/14] sysfs: Rewrite sysfs_get_dentry Posted by Tejun Heo on Wed, 01 Aug 2007 17:51:19 GMT

View Forum Message <> Reply to Message

Eric W. Biederman wrote:

> Tejun Heo <htejun@gmail.com> writes:

>

- >> Hmmm... I think we can live with a bit of complexity in
- >> sysfs get dentry(). It's very well localized and not even long. I have
- >> been trying hard to untangle sysfs internals from vfs and have a bit of
- >> resistance against going back. But, then again, if we can achieve
- >> something better and simpler, why not?

>

- > Right now my real challenge is untangling sysfs_move_dir. In the
- > case of multiple mounts of the sysfs tree.

OIC. Yeah, I guess moving a sd with multiple dentries attached could be pretty difficult.

- > My practical problem is that I need to hold a lock for the sysfs
- > dirents and while that lock is held I need to call sysfs_get_dentry
- > for the destination directory once for each superblock.

>

- > It might be that some kind of reader-writer lock strategy is what
- > I need to untangle this mess. Rather then making changing to i_mutex.
- > All I know is at the moment it is taking a lot of code reading and
- > brain storm to come of with something that is easy to maintain.

Just in case, sysfs used to have sysfs_rename_rwsem to protect move/rename against tree walking, which became unnecessary after i_mutex -> sysfs_mutex conversion. Move/rename can use stupid big fat locks if that helps.

Good luck.

--

tejun

Containers mailing list

Containers@lists.linux-foundation.org

https://lists.linux-foundation.org/mailman/listinfo/containers