
Subject: Re: [PATCH 06/14] sysfs: Rewrite sysfs_get_dentry
Posted by [Tejun Heo](#) on Wed, 01 Aug 2007 17:20:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

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

>

>> Hello, Eric.

>>

>> Eric W. Biederman wrote:

>>> I will look a little more and see. But right now it looks like the

>>> real problem with locking is that we use sysfs_mutex to lock the

>>> sysfs_dirent s_children list.

>>>

>>> Instead it really looks like we should use i_mutex from the appropriate

>>> inode. Or is there a real performance problem with forcing the directory

>>> inodes in core when we modify the directories?

>> I don't think there is any performance problem. Problems with using

>> i_mutex were...

>>

>> * It was messy. I don't remember all the details now but IIRC symlink

>> walk code was pretty complex.

>>

>> * And more importantly, inodes are reclaimable and might or might not be

>> there.

>

> Yes. But we can always force inodes into the cache when we need them.

> When I complete it I will have to show you a patch using the inode lock

> for locking directory modifications. From what I can tell so far it allows

> me to fix the weird lock order problems and generally simplify the locking.

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?

--

tejun

Containers mailing list

Containers@lists.linux-foundation.org

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