
Subject: Re: [PATCH 2/4] sysfs: Implement sysfs managed shadow directory support.

Posted by [ebiederm](#) on Tue, 31 Jul 2007 04:02:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

Tejun Heo <teheo@suse.de> writes:

> Hello,

>

> Eric W. Biederman wrote:

>> Ugh. I need to step back and carefully define what I'm seeing but it

>> looks like the current sysfs locking is wrong.

>>

>> I'm starting to find little inconsistencies all over the place

>> such as:

>>

>> Which lock actually protects sd->s_children?

>> - It isn't sysfs_mutex. (see sysfs_lookup)

>> - It isn't inode->i_mutex (we only get it if we happen to have the inode
>> in core)

>

> Yeah, I missed two places while converting to sysfs_mutex.

> sysfs_lookup() and rename(). I'm about to post patch to fix it.

Yes. Make certain to get the name change under sysfs_mutex while you are at it.

What do we use inode->i_mutex for? I think we might be able to kill that.

I'm starting to wonder if we can completely remove sysfs from grabbing inode->i_mutex.

>> At first glance sysfs_assoc_lock looks just as bad.

>

> I think sysfs_assoc_lock is okay. It's tricky tho. Why do you think

> it's bad?

I'm still looking. I just have a weird vibe so far. sysfs_get_dentry is really nasty with respect to locking.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

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