
Subject: Re: [PATCH 12/15] driver core: Implement tagged directory support for device classes.

Posted by [Tejun Heo](#) on Wed, 16 Jul 2008 06:48:39 GMT

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

Hello, Eric.

Eric W. Biederman wrote:

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

>

>> It's a bit scary tho. Working inode->i_dentry or dentry->d_alias

>> crosses multiple sb's. sysfs isn't too greedy about dcache/icache.

>> Only open files and directories hold them and only single copy of

>> sysfs_dirent is there for most nodes. Wouldn't it be better to stay on

>> the safer side and use separate inode hierarchy?

>

> To do that I believe we would need to ensure sysfs does not use

> the inode->i_mutex lock except to keep the VFS layer out. Allowing us

> to safely change the directory structure, without holding it.

I don't think sysfs is depending on i_mutex anymore but I need to go through the code to make sure.

> You raise a good point about inode->i_dentry and dentry->d_alias.

> Generally they are used by fat like filesystems but I am starting to

> see uses in generic pieces of code. I don't see any problems today

> but yes it would be good to do the refactoring to allow us to duplicate

> the inodes.

Yeah, I can't spot any place which can cause actual problem yet but it's still scary as we're breaking a vfs assumption and even if it's not a problem now, future seemingly unrelated changes can break things subtly.

Thanks.

--

tejun

Containers mailing list

Containers@lists.linux-foundation.org

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