

---

Subject: Re: [PATCH 05/11] sysfs: sysfs\_chmod\_file handle multiple superblocks  
Posted by [Daniel Lezcano](#) on Tue, 24 Jun 2008 10:39:46 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Tejun Heo wrote:

> Hello,

>

> Daniel Lezcano wrote:

>>> I think it would be great if sysfs\_chmod\_file can do all-or-nothing

>>> instead of failing half way through but given the interface of

>>> notify\_change(), it could be difficult to implement. Any ideas?

>> Is it acceptable to queue the notifications in a list until we are in

>> the loop and loop again to notify when exiting the first loop without

>> error ?

>

> Can you please take a look at the following patch?

>

> <http://article.gmane.org/gmane.linux.file-systems/24484>

>

> Which replaces notify\_change() call to two calls to sysfs\_setattr() and

> fsnotify\_change(). The latter never fails and the former should always

> succeed if inode\_change\_ok() succeeds (inode\_setattr() never fails

> unless the size is changing), so I think the correct thing to do is...

>

> \* Separate out sysfs\_do\_setattr() which doesn't do inode\_change\_ok() and

> just sets the attributes. Making it a void function which triggers

> WARN\_ON() when inode\_setattr() fails would be a good idea.

>

> \* Implement sysfs\_chmod\_file() in similar way rename/move are

> implemented - allocate all resources and check conditions and then iff

> everything looks okay commit the operation by calling sysfs\_do\_setattr().

>

> How does that sound?

Much better than my first proposition :)

I will do a separate patchset.

Thanks.

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

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

---