Subject: Re: [PATCH 05/11] sysfs: sysfs_chmod_file handle multiple superblocks
Posted by Tejun Heo on Tue, 24 Jun 2008 04:45:27 GMT
View Forum Message <> Reply to Message

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?

Thanks.

--
tejun

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers