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

Posted by [ebiederm](#) on Sun, 22 Jul 2007 22:19:42 GMT

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

Greg KH <gregkh@suse.de> writes:

> On Mon, Jul 23, 2007 at 04:47:21AM +0900, Tejun Heo wrote:
>> Sorry but I don't think the current approach is the correct one. It's
>> too painful and too much complexity is scattered all over the place.
>> I'm afraid this implementation is going to be a maintenance nightmare.
>
> In looking over this again, and due to the fact that there are no
> in-kernel users of this code, I'm going to drop the other patches in
> this series and only keep the first one that removes the current
> implementation.
>
> Eric, is that ok? It gives you time to revisit these changes.

Greg a big part of the reason I don't have internal kernel users is keeping up with changes of sysfs has meant I haven't had time to clean up and submit the patches for the users. With a little luck I will have users by 2.6.24. I am in the process of cleaning up those patches right now. So I can send the off to Dave Miller.

So please can we try and at least keep these sysfs patches in a development tree so that people will see them.

Further while there are a few little nits I think mostly Tejun is mostly objecting to the fundamental complexity of the problem rather than to things that can be fixed by a cleaner implementation.

If it didn't take me a week every time I had to update this code after Tejun changes the locking rules in fs/sysfs/dir.c or if there was someone I could delegate the work of maintaining this code to I probably would not mind dropping the patches for a little bit. As it stands I am having horrible nightmares about how the internals of sysfs will be completely different if you drop the last 3 patches by the time I come back and I will need to spend several more weeks just catching up.

Eric

Containers mailing list

