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

Posted by Tejun Heo on Mon, 30 Jul 2007 07:09:06 GMT

View Forum Message <> Reply to Message

Hello, Eric.

Tejun Heo wrote:

- > Eric W. Biederman wrote:
- >> Further while there are a few little nits I think mostly Tejun is
- >> mostly objecting to the fundamental complexity of the problem rather
- >> then to things that can be fixed by a cleaner implementation.

>

> Oh well, I don't think so but I might be wrong.

And I'm wrong. Mine didn't turn out to be much cleaner than yours. What I did was (still broken)...

- * No shadower/shadowee. Each dentry is tagged.
- * dentries of tagged sd's are taken out of dcache and always go through
- ->lookup() where the correct sd is looked up considering the current tag.

Tagging and adding new entries could be done rather cleanly but shooting down existing dentries on rename/move turned out to be a mess. Things will be much simpler if no sysfs dentry is hashed on dcache and always go through ->lookup() but that will hurt big machines.

The basic problem here is that dcache layer doesn't allow different views and sysfs shadow is trying to work behind its back. I don't think this is a viable approach. Both implementations bend too many rules and are too fragile. It will be a genuine pain in the ass to maintain.

Sorry that I can't come up with an alternative but NACK.

tejun

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