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

Posted by ebiederm on Mon, 30 Jul 2007 15:51:28 GMT

View Forum Message <> Reply to Message

Tejun Heo <teheo@suse.de> writes:

- > Kirill Korotaev wrote:
- >> Tejun Heo wrote:
- >>> I thought something like supermount plus some twists or fuse based sysfs
- >>> proxy would fit better. Dunno whether or how uevent and polling stuff
- >>> can work that way tho. Note that sysfs no longer keeps dentries and
- >>> inodes pinned. It might make the shared dentry stuff harder.

>>

- >> We simply don't share sysfs dentries/inodes between containers.
- >> It's not that frequently used time critical fs to be super-optimized...:)

>

- > OIC, dentries and inodes are not shared. Good then. Agreed that sysfs
- > doesn't need to be super-optimized as long as big machines aren't
- > penalized too much (both memory and cpu cycle wise).

>

- >> I don't like the idea with fuse, since sysfs exports kernel-related stuff,
- >> so doing it via user-space would be pain.

>

- > Yeah, it would be cumbersome to setup but it's also fast and easy to toy
- > with for prototypes at least.

How close are we to the point where we can get mount sysfs multiple times and get multiple dentry trees with different super blocks?

That really does sound like the right way to go. Especially as it simplifies the monitoring of containers. If you want to watch what the view looks like in some container your bind mount his sysfs and look at that.

If we can do that the dcache side at least will be beautiful. And with a little care we may be able to reduce the work to a special case in lookup, some extra handling to mark directories as belonging only to a certain mount of sysfs.

If we can find something that is stupid and simple I'm all for that.

To reach the no-kobj utopia we may also need a special device_migrate that is a super set of device_rename (because sometimes we need to rename devices when we move them between namespaces).

So are we close to having a sysfs that we can have multiple super blocks for?

I'm on a sysctl tangent, but I should be able to look at that in just a little bit.

Eric

Containers mailing list

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