Subject: Re: [netns] sysfs: issues porting shadow directories on top of 2.6.21-mm2
Posted by ebiederm on Thu, 21 Jun 2007 20:48:10 GMT
View Forum Message <> Reply to Message

Dave Hansen <hansendc@us.ibm.com> writes:

> On Mon, 2007-06-18 at 19:57 +0200, Benjamin Thery wrote:
>> /sys/class/net/ ("real" net class)
>> /sys/class/net-shadow1/
>> /sys/class/net-shadow2/
>
> This seems like a nice "quick fix", but do we really want to be hacking
> sysfs around like this?
>
> We have the backing sysfs_* entries that are separate from the vfs
> entities already. Why can't we simply have different /sys vfsmounts
> with different views of the backing sysfs_* entries?

So far this has been easier. sysfs is so tightly coupled to the kobject tree decoupling them is seriously non-trivial.

When I can I do prefer separate mounts because that fixes a lot of cache issues. For /proc we may be able to go as far as breaking up all of it's pieces into separate filesystems, and do magic mounts and umounts behind the users back. We can't even contemplate that with sysfs as it has symlinks all over the place. A remount of sysfs that displays just the information from the current processes namespaces would be possible. If you can figure how to make those mods to sysfs I will be happy to use it.

We actually must modify /proc to do multiple mounts because a follow link method on the root directory entry is not honored. With a magic follow link method we can use the dcache as intended except for the handful of magic directories.

So since I have working code today (baring one or two bugs I'm in the process of fixing) and that greg has given his basic nod of acceptance I figure I will pursue this route and we can update things later if we need to.

Eric

_____

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