Subject: Re: [patch 1/2] [RFC] Simple tamper-proof device filesystem.
Posted by serue on Tue, 18 Dec 2007 02:09:33 GMT

View Forum Message <> Reply to Message

Quoting Oren Laadan (orenl@cs.columbia.edu):

>
> I hate to bring this again, but what if the admin in the container
> mounts an external file system (eg. nfs, usb, loop mount from a file,
> or via fuse), and that file system already has a device that we would
> like to ban inside that container ?

Miklos' user mount patches enforced that if !capable(CAP_MKNOD),
then mnt->mnt_flags |= MNT_NODEV. So that's no problem.

But that's been pulled out of -mm! ? Crap.

> Since anyway we will have to keep a white- (or black-) list of devices
> that are permitted in a container, and that list may change even change
> per container -- why not enforce the access control at the VFS layer ?
> It's safer in the long run.

By that you mean more along the lines of Pavel's patch than my whitelist
LSM, or you actually mean Tetsuo's filesystem (i assume you don't mean that
by 'vfs layer' :), or something different entirely?

thanks,
-serge

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