
Subject: Re: [PATCH 4/4] The control group itself
Posted by [Paul Menage](#) on Tue, 15 Jan 2008 16:13:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Jan 15, 2008 6:44 AM, Serge E. Hallyn <serue@us.ibm.com> wrote:
>
> I don't think so... Wouldn't really make sense for the cgroup
> infrastructure to presume to know what to enforce, and I don't see any
> checks around the `_write` functions in `cgroup.c`, and no `capable()` calls
> at all.

The cgroup filesystem can provide simple unix-level permissions on any given file. Am I right in thinking that having an entry in the mapper doesn't automatically give privileges for a device to the members of the cgroup, but they also have to have sufficient privilege in their own right? If so, that might be sufficient.

One other thought - should the parse/print routines themselves do a translation based on the device mappings for the writer/reader's cgroup? That way you could safely give a VE full permission to write to its children's device maps, but it would only be able to add/remap device targets that it could address itself.

Paul

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