
Subject: Re: [PATCH] namespaces: introduce sys_hijack (v7)

Posted by [serue](#) on Wed, 31 Oct 2007 23:13:59 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Serge E. Hallyn (serue@us.ibm.com):

> This is just a first stab at doing hijack by cgroup files. I force
> using the 'tasks' file just so that I can (a) predict and check
> the name of the file, (b) make sure it's a cgroup file, and then
> (c) trust that taking `__d_cont(dentry parent)` gives me a legitimate
> container.
>
> Seems to work at least.
>
> Paul, does this look reasonable? `task_from_cgroup_fd()` in particular.

The patch I sent does 'hijack a cgroup' by taking a task in the cgroup and hijacking it.

Paul would like to be able to 'enter a cgroup', even if it is empty. Hijack takes more than just the nsproxy from the hijacked task, so this would result in different behavior between hijacking a populated cgroup and an empty cgroup. So we might want to introduce a third type of hijacking, so we have `HIJACK_PID`, `HIJACK_CGROUP`, and `HIJACK_EMPTY_CGROUP`.

It also then acts like the nsproxy cgroup patchset I sent out months ago for simply entering namespaces. In fact this would need to be restricted to ns cgroups, and ns cgroups would need to grab a reference to the nsproxy.

So do we want to allow hijacking/entering an empty cgroup?

-serge

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
