
Subject: Re: Re: containers and cgroups mini-summit @ Linux Plumbers
Posted by [Tejun Heo](#) on Thu, 26 Jul 2012 18:15:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello, Eric.

On Thu, Jul 26, 2012 at 03:42:50AM -0700, Eric W. Biederman wrote:

- > - Create a new mount namespace.
- > - Create fresh mounts of all of the control groups like I would do at
- > boot, with no consideration to any other control group state.
- > - Start forking processes.
- >
- > The expected semantics would be something like chroot for control
- > groups, where all of the control groups that are created by fresh mounts
- > are relative to whatever state the process of being in a control group
- > that the process that mounted them was in.

No, any attempt to build namespace support into cgroup core code will be nacked with strong prejudice. I still think it was a mistake to add that to sysfs. Thankfully, procfs is going the FUSE way and I hope in time we could convert sysfs to a similar mechanism and deprecate the in-kernel support.

So, no, no, no, no, no, no, no, no, no, no, no, no, no. :P

--
tejun
