
Subject: Re: [PATCH 11/11] sysfs: user namespaces: fix bug with
clone(CLONE_NEWUSER) with fairsched

Posted by [Tejun Heo](#) on Mon, 23 Jun 2008 02:18:58 GMT

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

Benjamin Thery wrote:

- > Mark the /sys/kernel/uids directory to be tagged so that processes in
- > different user namespaces can remount /sys and see their own uid
- > listings.
- >
- > Without this patch, having CONFIG_FAIR_SCHED=y makes user namespaces
- > unusable, because when you
- > clone(CLONE_NEWUSER)
- > it will auto-create the root userid and try to create
- > /sys/kernel/uids/0. Since that already exists from the parent user
- > namespace, the create fails, and the clone misleadingly ends up
- > returning -ENOMEM.
- >
- > This patch fixes the issue by allowing each user namespace to remount
- > /sys, and having /sys filter the /sys/kernel/uid/ entries by user
- > namespace.
- >
- > Signed-off-by: Serge Hallyn <serue@us.ibm.com>
- > Signed-off-by: Benjamin Thery <benjamin.thery@bull.net>

Ditto as patch #10.

Thanks.

--

tejun

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