
Subject: Re: [PATCH -mm 1/2] user namespace : add unshare

Posted by [akpm](#) on Fri, 08 Jun 2007 19:22:15 GMT

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

On Fri, 08 Jun 2007 17:14:07 +0200

Cedric Le Goater <clg@fr.ibm.com> wrote:

> Basically, it will allow a process to unshare its user_struct table, resetting
> at the same time its own user_struct and all the associated accounting.
>
> A new root user (uid == 0) is added to the user namespace upon creation. Such
> root users have full privileges and it seems that theses privileges should be
> controlled through some means (process capabilities ?)

This second paragraph is distressingly indecisive. How much thought has gone into this??

For a start, it seems wrong for the kernel to hardwire knowledge about UID 0 in this fashion.

I'd have thought that a better model for user-namespace unsharing would be to do a copy-by-value of the entire namespace, then permit a suitably-privileged application to go through and kill off any unwanted users from the now-unshared user namespace.

Or maybe just remove that "Insert new root user" altogether? What would then go wrong?

Containers mailing list

Containers@lists.linux-foundation.org

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