
Subject: Re: namespaces compatibility list

Posted by [ebiederm](#) on Tue, 06 Nov 2007 17:00:27 GMT

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

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

> right. I think we can address Ulrich concerns first because we have
> a solution for it (which looks like unsharing all namespaces at once,
> here comes back the container object story :)

It doesn't work because we can't create a fresh mount namespace.

We need to create all new mounts (and deny access to the old ones)
if we want to prevent all possibility of user space goof ups.

While that is easy enough to build an application to do we can't
easily enforce that in the kernel. Currently this is all
CAP_SYS_ADMIN so only root can do this anyway. So we can easily
say don't do that then.

Clone flag consistency checking should only be used to enforce
cases where the kernel side cannot support correctly. Currently
the kernel has no problems with the current mix and match possibilities
short of implementation deficiencies. So I do not see us
addressing Ulrich's concerns with clone flags.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

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