Subject: Re: [PATCH -mm] remove node_ prefix_from ns subsystem
Posted by Paul Menage on Thu, 22 May 2008 08:23:35 GMT
View Forum Message <> Reply to Message

Hi Andrew,

On Thu, May 22, 2008 at 1:10 AM, Andrew Morton
<akpm@linux-foundation.org> wrote:
>
> What the change is, why it is being made, what the user-visible
> presentation is, what the impact upon users is, why we think it won't
> be a problem, etc? The stuff which should have been right there from
> day one, before the code change was even made?

The change is that previously when cgroup_clone() was called
(currently only from the unshare path in ns_proxy cgroup, you'd get a
new group named "node_$pid" whereas now you'll get a group named after
just your pid.)

The only users who would notice it are those who are using the
ns_proxy cgroup subsystem to auto-create cgroups when namespaces are
unshared - something of an experimental feature, which I think really
needs more complete container/namespace support in order to be useful.
I suspect the only users are Cedric and Serge, or maybe a few others
on containers@lists.linux-foundation.org. And in fact it would only be
noticed by the users who make the assumption about how the name is
generated, rather than getting it from the /proc/<pid>/cgroups file
for the process in question.

Whether the change is actually needed or not I'm fairly agnostic on,
but I guess it is more elegant to just use the pid as the new group
name rather than adding a fairly arbitrary "node_" prefix on the
front.

Paul

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