Subject: Re: containers - bug
Posted by Paul Menage on Sat, 01 Sep 2007 00:44:41 GMT
View Forum Message <> Reply to Message

On 8/31/07, Daniel Lezcano <dlezcano@fr.ibm.com> wrote:
> Hi Paul,
>
> I was playing with the container filesystem (very nice) and I fall
> inside a kbug.
>
> I did the following:
>
>   mkdir /dev/container
>   mount -t container -o cpuset cpuset /dev/container
>   cd /dev/container/
>   mkdir Charlie
>   cd Charlie
>   echo $$ > tasks

FYI, this bit didn't have any effect, since the cpuset has no
mems/cpus by default.

>   bash
>   cd ..
>   rmdir Charlie
>   exit
>   ls => bang !

The basic problem appears to be that the reference count on a
containerfs directory is one too low. The root cause is a change that
I made when adapting the cpuset filesystem to create the container
filesystem - in order to implement container_clone() I rearranged the
way that a dentry was passed down to the directory creation code, and
managed to lose a call to lookup_one_len() (since I was getting the
dentry directly from a container structure).

Normally this didn't seem to cause a problem, since dput() doesn't
appear to care if you dput() on something with a refcount of 0. (It
should probably BUG() in that case, I suspect).

But in your case, by making the dead directory some process' cwd, when
you tried to do an ls, the dget(current->fs->pwd) blew up since
current->fs->pwd had a d_count of 0.

I have a fix for this that I'm testing and should be able to send out soon.

Paul

_____

Subject: Re: containers - bug
Posted by Paul Menage on Mon, 10 Sep 2007 21:55:42 GMT

View Forum Message <> Reply to Message

On 8/31/07, Paul Menage <menage@google.com> wrote:
>
> I have a fix for this that I'm testing and should be able to send out soon.
>

Sorry for the delay on this - I was away all last week at the kernel
summit. I've sent out a patch to the list and to andrew, as a separate
thread.

Paul

_____