
Subject: Re: [patch -mm] update mq_notify to use a struct pid
Posted by [ebiederm](#) on Tue, 12 Sep 2006 01:22:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

> Eric W. Biederman wrote:

>

>> Cedric you mentioned a couple of other patches that are in flight.

>> In the future could you please Cc: the containers list so independent

>> efforts are less likely to duplicate work, and we are more likely

>> to review each others patches instead?

>

> yes sure, i was relying on the openvz wiki to avoid duplicated efforts on

> this topic but i guess email is just the one and only tool for this kind of

> development :)

Sure. Especially when it comes to helping review each others code :)

Not duplicating work is not really my goal, not submitting a patch
after a patch has been reviewed and accepted is.

Plus we need patch review.

Several people working on a patch in parallel if it is difficult
can frequently find a solution that a single person would miss.

>>>> Filling in a struct pid through sysctl is extremely ugly at the

>>>> moment, plus cad_pid needs some locking.

>>> Which distros use /proc/sys/kernel/cad_pid and why ? I can image the need

>>> but i didn't find much on the topic.

>>

>> I'm not at all certain, and I'm not even certain I care. The concept

>> is there in the code so it needs to be dealt with.

>

> OK. It would be nice to make sure this is still in use before trying to

> deal with /proc/sys/kernel/cad_pid.

>

>> Although if I we extend the cad_pid concept it may make a difference.

>

> what do you mean by extending cad_pid ? kill_init() ?

My meaning was every time we are sending a signal to init. It is quite
possible we should be using cad_pid instead.

>>> is that about updating the siginfos in collect_signal() to hold the right

>>> pid value depending on the pid namespace they are being received ?

>>

>> Yes in send_signal, and in collect signal. To make it work easily I needed

>> to add a struct pid to struct sigqueue. So in send_signal I generate
>> the struct pid from the pid_t value and in collect signal I regenerate
>> the numeric value.
>
> OK. That's what i imagined also but we need a bit more of the pid namespace
> to regenerate the numerical value. So, how will you convert this 'struct
> pid*' in a pid value using the current pid namespace ?

By calling pid_nr :) The question I guess is how will pid_nr be implemented.

> thinking aloud :
>
> * if the pid namespace of the sending struct pid and current match,
> use nr.
> * if they don't,
> if the sending pid namespace is the ancestor of the current pid
> namespace
> use 0
> else
> it's a bug.
>
> struct pid* needs a pid namespace attribute and pid namespace needs to know
> its parent.

Yes, that sounds correct.

There is also the case that should not come up with signals where
we have a pid from a child namespace, that we should also be able to
compute the pid for.

In essence I intend to have a list of pid_namespace, pid_t pairs connected
to a struct pid that we can look through to find the appropriate pid.

Eric

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
