
Subject: Re: [PATCH v7 09/10] IPC: message queue copy feature introduced
Posted by [ebiederm](#) on Thu, 18 Oct 2012 12:20:54 GMT

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

"Michael Kerrisk (man-pages)" <mtk.manpages@gmail.com> writes:

>>> Is there somewhere a *detailed* description of how this feature would
>>> be used? Lacking that, it's really hard to see how anything sane and
>>> reliable can be done with MSG_COPY.
>>>
>>
>> These patches are used by CRIU already.
>> So, you can have a look at the CRIU source code:
>>
>> [http://git.criu.org/?p=crtools.git](http://git.criu.org/?p=crtools.git;a=blob;f=ipc_ns.c;h=9e259fefcfc04ec0556bb722921545552e1c69f3;hb=HEAD)
>> ;a=blob;f=ipc_ns.c;h=9e259fefcfc04ec0556bb722921545552e1c69f 3;hb=HEAD
>>
>> Sanity and reliability on the level you are talking about can be achieved,
>> only if you'll freeze all message users before peeking.
>
> Okay -- that's the piece I was looking for. Thanks.
>
> In this scenario, how do you find all of the message users? Or do you
> simply ensure that everything is frozen beforehand?

The general design is a container is started with a fresh set of namespaces and then the entire is frozen using the process freezer control group. With all of the userspace process frozen checkpoint then happens.

Eric
