
Subject: Re: Mapping PIDs from parent->child namespaces

Posted by [Mike Heffner](#) on Tue, 04 Jan 2011 20:17:48 GMT

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

On 01/04/2011 11:44 AM, Cedric Le Goater wrote:

> On 01/04/2011 05:04 PM, Daniel Lezcano wrote:

>> On 01/04/2011 12:02 AM, Mike Heffner wrote:

>>> Hi,

>>>

>>> Is it possible for a process running in a parent PID namespace to map
>>> the PID of a process running in a child's namespace from the
>>> parent->child's namespace? For example, if I spawn the process "myproc"
>>> with CLONE_NEWPID then a call to getpid() inside myproc will return "1"
>>> whereas in the parent's namespace that process could actually be PID
>>> "23495". I'd like to be able to know that 23495 maps to 1 in the new NS.
>>> Obviously, just mapping the first PID is straightforward since I can
>>> just look at the result of clone(). However, mapping the PIDs of
>>> processes subsequently forked from "myproc" -- in this example -- I
>>> haven't been able to figure out.

>>

>> AFAIK, it is not possible.

>>

>> That would be very nice to show the pid<-> vpid association.

>>

>> The procfs is a good candidate to show these informations.

>>

>> That would makes sense to show the content of /proc/<pid>/status with
>> the pid relatively to the namespace.

>>

>> Let me give an example:

>>

>> Assuming the process '1234' creates a new pid namespace, and the child
>> which is '1' in the new namespace has the real pid '4321'. This one
>> mounts its /proc.

>>

>> If the process '1234' looks at /proc/4321/root/proc/1/status, it sees:

>>

>> ...

>> Tgid: 1

>> Pid: 1

>> PPid: 0

>> ...

>>

>>

>> It could be:

>>

>> ...

>> Tgid: 4321

>> Pid: 4321
>> PPid: 1234
>> ...
>>
>> as the file is inspected from the parent namespace. Of course, if the
>> file is looked from the child namespace context, we will see '1', '1'
>> and '0'.
>>
>> I suppose the patch in the kernel should very small also.
>>
>> Thoughts ?
>
> we use the following patch to get the pid of a task as seen from its
> pid namespace. It can be useful to identify tasks writing pids in files.
>

Great, I'll try it out. Has there been any interest in getting this into the mainline? Are there negatives to advertising child vpid's?

Cheers,

Mike

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