Subject: Re: [RFC][PATCH 4/4] PID: use the target ID specified in procfs
Posted by Oren Laadan on Thu, 13 Mar 2008 23:24:10 GMT
View Forum Message <> Reply to Message

Eric W. Biederman wrote:
> Oren Laadan <orenl@cs.columbia.edu> writes:
>
>>>> I'm sorry but I'm pretty new in this domain, so I don't see what are the
>>>> namespaces where setting (or pre-setting) the id would be a problem?
>>> pids to some extent as people use them in all kinds of files. Being
>>> able to force the pid of another process could make a hard to trigger
>>> security hole with file permissions absolutely trivial to hit.
>> Since the intent of this mechanism is to allow ckpt/restart, it makes
>> sense to only allow this operation during restart. For example, in zap,
>> containers have a state, e.g. running, stopped, ckpt, restart, and this
>> is only possible in restart state; Furthermore, a container can only be
>> put in restart state at creation time, and only by root. Of course, you
>> should only trust that as much as you trust the root :O
>
> Yes and thanks.
>
> The notion of the state of a container makes a lot of sense (even if
> we never implement explicit state bits).

I found it extremely helpful in managing containers (pods) in zap. There are three more states, actually: stopping, reviving and dead. It is like extending the notion of process state into their collective representation which is the container. In fact, restricting certain operations to specific states was instrumental in eliminating a myriad of potential races in the implementation.

I believe this belongs to the ever-pending ckpt/restart discussion :)

Oren.

>
> Eric

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