

Herbert Poetzl <herbert@13thfloor.at> writes:

>
> again, we basically support 3 different guest models
> (regarding init) which probably can be best explained
> with an example ...
>
> 1) blend through/fake init (from the host system)
>
> USER PID %CPU %MEM VSZ RSS TTY STAT START TIME COMMAND
> root 1 6.0 1.9 2036 1096 ? S 14:24 0:06 init
> root 38 0.7 0.8 2832 448 ? S 14:26 0:00 sleep 1000
> root 43 50.0 1.2 2536 676 ? R 14:26 0:00 ps auxwww
>
> 2) a real init process (running inside the guest with pid=1)
>
> USER PID %CPU %MEM VSZ RSS TTY STAT START TIME COMMAND
> root 1 1.6 0.7 2832 444 ? S 14:26 0:00 sleep 1000
> root 44 0.0 1.2 2536 676 ? R 14:26 0:00 ps auxwww
>
> 3) no init process (inside a guest)
>
> USER PID %CPU %MEM VSZ RSS TTY STAT START TIME COMMAND
> root 42 0.4 0.7 2828 444 ? S 14:26 0:00 sleep 1000
> root 45 38.0 1.2 2536 676 ? R 14:26 0:00 ps auxwww
>
>
> in cases 1) and 3) the 'first' process is in no
> way special for the Guest, and must not be treated
> special .. it can also go away anytime without
> affecting the other guest processes ...
>
> case 2) could in theory handle the pid=1 process
> (which might not be the first process, but usually
> is a special init process) special, and it would
> be acceptable to zap the context when this process
> dies off ...
>
> note that the cases 1) and 2) are the most commonly
> used cases as most init processes do not handle case
> 3) yet. still case 3) is important for application
> isolation too (which doesn't need any init)

>From a maintenance standpoint options like this can be horrible.

The practical question is this. For application isolation what problems have you encountered with running an application as pid == 1?

Why do you need the no init process inside a guest case?

If you can answer this question when it comes time to optimize things it will give us incentive to solve these cases.

Eric

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