
Subject: Re: restarting tests/sleep

Posted by [Dave Hansen](#) on Mon, 09 Jun 2008 16:32:28 GMT

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

On Mon, 2008-06-09 at 11:23 -0500, Serge E. Hallyn wrote:

>
>
> Yeah that's insurmountable - notice the stack in the process which was
> fork()ed to be the restarted process topped at bfcab000, while the
> checkpointed stack topped at bfdae000. You're not allowed to write
> above the stack. So the only things to do are
>
> 1. keep trying the restart in the hopes you get a task with
> stack topping at or above bfdae000
> 2. if the checkpointed stack is too high to be likely to be
> restartable, generate a new checkpoint image and you
> should get a lower stack top.
>
> (Dave, maybe you had other ideas I haven't considered)

Have you tried turning of stack randomization? It should make the stack more dependable at exec. We could also provide a hint somewhere on exec() where to place the stack.

-- Dave

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: restarting tests/sleep

Posted by [serue](#) on Mon, 09 Jun 2008 18:19:12 GMT

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

Quoting Dave Hansen (dave@linux.vnet.ibm.com):

> On Mon, 2008-06-09 at 11:23 -0500, Serge E. Hallyn wrote:
> >
> >
> > Yeah that's insurmountable - notice the stack in the process which was
> > fork()ed to be the restarted process topped at bfcab000, while the
> > checkpointed stack topped at bfdae000. You're not allowed to write
> > above the stack. So the only things to do are
> >
> > 1. keep trying the restart in the hopes you get a task with
> > stack topping at or above bfdae000
> > 2. if the checkpointed stack is too high to be likely to be

> > restartable, generate a new checkpoint image and you
> > should get a lower stack top.
> >
> > (Dave, maybe you had other ideas I haven't considered)
>
> Have you tried turning of stack randomization? It should make the stack
> more dependable at exec. We could also provide a hint somewhere on
> exec() where to place the stack.

Yup, that should prevent the spurious -EFAULTS. Suka, you were getting these too last week, so I recommend doing

```
echo 0 > /proc/sys/kernel/randomize_va_space
```

on startup.

-serge

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