
Subject: Re: Userspace checkpoint/restart hack: cryo

Posted by [serue](#) on Mon, 09 Jun 2008 15:23:36 GMT

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

Quoting Nadia Derby (Nadia.Derbey@bull.net):

> Serge E. Hallyn wrote:

>> Quoting Nadia Derby (Nadia.Derbey@bull.net):

>>

>>> Cedric Le Goater wrote:

>>>

>>>>

>>>>

>>>>

>>>>> Looks like it has worked for me (msg1 creates 1000 msg queues, sleeps

>>>>> for a while and then removes the msg queues).

>>>>

>>>>

>>>> cool. which kernel are you using ?

>>>>

>>>> C.

>>>>

>>>>

>>>

>>> 2.6.25-mm1 (sorry for the late answer - plenty of days off in France

>>> in May).

>>>

>>> Regards,

>>> Nadia

>>

>>

>> I'm playing with features in cryo, and keeping a git tree at:

>>

>> [git://git.sr71.net/~hallyn/cryodev.git](http://git.sr71.net/~hallyn/cryodev.git)

>>

>> It's meant to exploit the extras which are in the -lxc kernel at

>> lxc.sf.net. Current latest kernel patch is at

>> <http://lxc.sourceforge.net/patches/2.6.26/2.6.26-rc2-mm1-lxc4/> This -lxc

>> tree contains, for instance, Nadia's next_id patches, exploitation for

>> setting ids for sysvipc and for tasks at fork, and updated ipc_setall

>> patches (also using next_id). The version of cryo in my git tree

>> exploits these. If you're root when you restart a task, it will clone a

>> new set of namespaces and recreate your sysvipc objects, and it will

>> reset your pids (even if you're not root if the pids are available).

>

> Serge,

>

> I noticed that the sys_hijack() has disappeared from the lxc dev tree:

> would you mind putting it back. I think it might be useful if we want to

> start a task in a newly defined cgroup, checkpoint it and then try to
> restart it. We will need to 'join' the restarted container to check if
> everything has correctly be restarted. Or may be is there another way to
> do that?

Well it *could* be done similar to how cryo itself works, by ptracing the destination task and making it fork a task which `execve()`s a process to do the querying. But yuck.

Kathy, I'm sorry, I know I asked you to take `sys_hijack()` out. Could you please put it back in? Preferably at the end of the queue, as I don't want other patches having to be ported on top of it since its future is very suspect... Let me know if you have trouble porting it, but it should be pretty simple.

thanks,
-serge

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