
Subject: Re: [Announce] Checkpoint-restore tool v0.1

Posted by [cyrill](#) on Tue, 31 Jul 2012 10:30:07 GMT

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

On Tue, Jul 31, 2012 at 12:21:58PM +0200, richard -rw- weinberger wrote:

> On Tue, Jul 31, 2012 at 12:16 PM, Cyrill Gorcunov <gorcunov@openvz.org> wrote:

> > On Tue, Jul 31, 2012 at 12:08:22PM +0200, richard -rw- weinberger wrote:

> > > On Tue, Jul 31, 2012 at 11:54 AM, Pavel Emelyanov <xemul@parallels.com> wrote:

> > > > Yeah, but I fear it's not that easy.

> > > > We'd have to change crtools to work without ptrace().

> > > >

> > > > Well, this is hard. Using ptrace saved us from having many special-purpose

> > > APIs for dumping various stuff (there will be an article about it). Thus I

> > > don't know which way is simpler -- stop using ptrace or teach ptrace to allow

> > > several tracers to attach to one task %)

> > >

> > > Allowing multiple tracers in a safe way is IMHO even more harder.

> > >

> > > BTW: While reading prctl_set_mm() I noticed two things.

> > > 1. Why isn't the return value of find_vma() verified?

> >

> > prctl_set_mm

> > vma = find_vma(mm, addr);

> > ...

> > if (!vma) {

> > error = -EFAULT;

> > goto out;

> > }

> >

> > these values are used in procfs statistics only. So I don't get

> > which verify you mean here.

>

> If I do PR_SET_MM_START_BRK the if(!vma) will never be executed because

> there a break in case PR_SET_MM_START_BRK.

Yes, and this is done by purpose, since we need to setup _completely_
new memory map on restore procedure.

There is a minimal check for value being sane

```
if (addr >= TASK_SIZE || addr < mmap_min_addr)
    return -EINVAL;
```

and the address belongs to mm::start_data|end_data area. But sure,
better to add checks that at least code/data areas do exist, otherwise
the proc output will not reflect the real state of memory maps.

Cyrill
