
Subject: RE: trying to build simple checkpoint/restart recipes
Posted by [Rob Landley](#) on Wed, 08 Dec 2010 21:10:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

> > The restoration of the mounts is not scriptable however. It involves
> > parsing the mountinfo file and coordinating the mounts with those done by
> > lxc itself during lxc-restart. I honestly haven't looked at that closely
>
> I'd be fine with requiring some bit of hand-parsing. But right, even
> once we get a list of the mounts to be restored, I don't know of any
> good way to get those mounts re-created at the right time.

Mount code is one of my old stomping grounds from back when I wrote the busybox mount and switch_root commands and had to learn more implementation details about it than I ever wanted to know. :)

I never could find a proper mount spec, and kept meaning to write one, but I blathered about some of the less obvious details here:

<http://www.mail-archive.com/busybox@busybox.net/msg07013.htm> I

There are four top level categories of filesystem: Block backed, ram backed, pipe backed (network and fuse and so on), and synthetic (sysfs, procfs, devtmpfs...). And that's not counting bind mounts (which are internal to the VFS and not really a filesystem), and loopback devices (which are sort of the _opposite_ of a filesystem)...

> I suppose I could hack lxc-restart to do it. But I'm sort of hoping we
> can get something less hacked and more true to the 'real' upstream
> code.

Which upstream code?

> So do you know of anyone who's been working on re-creation of mounts
> in the kernel? If not, what have you been doing, hand-scripting
> all container creation, checkpoint, and restart?

I express interest in this topic.

Rob

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