Subject: Re: [patch 1/2] [RFC] Simple tamper-proof device filesystem.
Posted by serue on Tue, 18 Dec 2007 01:55:57 GMT
View Forum Message <> Reply to Message

Quoting Serge E. Hallyn (serue@us.ibm.com):
> Quoting Tetsuo Handa (penguin-kernel@i-love.sakura.ne.jp):
> > Hello.
> >
> > Serge E. Hallyn wrote:
> > > CAP_MKNOD will be removed from its capability
> > I think it is not enough because the root can rename/unlink device files
> > (mv /dev/sda1 /dev/tmp; mv /dev/sda2 /dev/sda1; mv /dev/tmp /dev/sda2).
>
> Sure but that doesn't bother us :)
>
> The admin in the container has his own /dev directory and can do what he
> likes with the devices he's allowed to have. He just shouldn't have
> access to others. If he wants to rename /dev/sda1 to /dev/sda5 that's
> his choice.
>
> > > To use your approach, i guess we would have to use selinux (or tomoyo)
> > > to enforce that devices may only be created under /dev?
> > Everyone can use this filesystem alone.
>
> Sure but it is worthless alone.
>
> No?

Oh, no, I'm sorry - I was thinking in terms of my requirements again.
But your requirements are to ensure that an application accessing a
device at a well-known location get what it expect.

So then the main quesiton is still the one I think Al had asked - what
keeps a rogue CAP_SYS_MOUNT process from doing
mount --bind /dev/hda1 /dev/null ?

thanks,
-serge

> What will keep the container admin from doing 'mknod /root/hda1 b 3 1'?
>
> > But use with MAC (or whatever access control mechanisms that prevent
> > attackers from unmounting/overlaying this filesystem) is recomennded.
>
> -serge

_____
Containers mailing list
Containers@lists.linux-foundation.org