Subject: Re: [RFC] [PATCH -mm] oom_kill: remove uid==0 checks
Posted by serue on Fri, 21 Dec 2007 14:46:06 GMT
View Forum Message <> Reply to Message

Quoting Andrew Morton (akpm@linux-foundation.org):
> On Wed, 12 Dec 2007 15:06:17 -0800
> Andrew Morgan <morgan@kernel.org> wrote:
>
> > Serge E. Hallyn wrote:
> > > Andrew, I've cc:d you here bc in doing this patch I noticed that your
> > > 64-bit capabilities patch switched this code from an explicit check
> > > of cap_t(p->cap_effective) to using __capable(). That means that
> > > now being glossed over by the oom killer means PF_SUPERPRIV will
> > > be set. Is that intentional?
> >
> > Yes, I switched the check because the old one didn't work with the new
> > capability representation.
> >
> > However, I had not thought this aspect of this replacement through. At
> > the time, it seemed obvious but in this case it actually depends on
> > whether you think using privilege (PF_SUPERPRIV) means "benefited from
> > privilege", or "successfully completed a privileged operation".
> >
> > I suspect, in this case, the correct thing to do is add the equivalent of:
> >
> > #define CAPABLE_PROBE_ONLY(a,b) (!security_capable(a,b))
> >
> > and use that in the code in question. That is, return to the old
> > behavior in a way that will not break if we ever need to add more bits.

Oh, I'm sorry - Andrew Morgan, I somehow read that email to say you were
going to post such a patch, and let it fall off my todo list. Should I
go ahead and post a patch or do you have one ready?

> I'm struggling to understand whether the above was an ack, a nack or a
> quack.
>
> > Thanks for finding this.
>
> >From that I'll assume ack ;)

It actually wasn't an ack of my patch. But I'm not sure where to look
for that.

thanks,
-serge

_____
Containers mailing list