

---

Subject: Re: [RFC] [PATCH -mm] oom\_kill: remove uid==0 checks

Posted by [akpm](#) on Fri, 21 Dec 2007 00:34:42 GMT

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

---

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.

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 ;)

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---