Subject: Re: [RFC] [PATCH -mm] oom_kill: remove uid==0 checks
Posted by Andrew Morgan on Wed, 12 Dec 2007 23:06:17 GMT
View Forum Message <> Reply to Message

-----BEGIN PGP SIGNED MESSAGE-----
Hash: SHA1

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.

Thanks for finding this.

Cheers

Andrew

```
>
> Signed-off-by: Serge Hallyn <serue@us.ibm.com>
> ---
>  mm/oom_kill.c |    2 +-
>  1 files changed, 1 insertions(+), 1 deletions(-)
>
> diff --git a/mm/oom_kill.c b/mm/oom_kill.c
> index 016127e..9fd8d5d 100644
> --- a/mm/oom_kill.c
> +++ b/mm/oom_kill.c
> @@ -128,7 +128,7 @@ unsigned long badness(struct task_struct *p, unsigned long uptime,
>    * Superuser processes are usually more important, so we make it
>    * less likely that we kill those.
>    */
> - if (__capable(p, CAP_SYS_ADMIN) || p->uid == 0 || p->euid == 0)
> + if (__capable(p, CAP_SYS_ADMIN) || __capable(p, CAP_SYS_RESOURCE))
>    points /= 4;
>
>   /*
```

-----BEGIN PGP SIGNATURE-----
Version: GnuPG v1.4.7 (Darwin)
Comment: Using GnuPG with Mozilla - http://enigmail.mozdev.org

iD8DBQFHYGln+bHCR3gb8jsRAgNwAKDQED4YNy479LKfDL1fhVGWMK22eACgjPMh
JcFgzPsvIQkoatjvJ1vtHQ8=
=50l1
-----END PGP SIGNATURE-----