Subject: Re: [PATCH 0/3] capabilities: per-process capbset
Posted by serue on Tue, 02 Oct 2007 03:24:03 GMT
View Forum Message <> Reply to Message

Quoting James Morris (jmorris@namei.org):
> On Mon, 1 Oct 2007, Serge E. Hallyn wrote:
>
> > Here is a new per-process capability bounding set patchset
> > which I expect to send to linux-kernel soon. It makes
> > the capbset per-process. A process can only permanently
> > remove bits from it's bounding set, not add them. To
> > remove bits, CAP_SYS_ADMIN is currently needed. Maybe
> > that's not the best choice, but some privilege should
> > probably be required.
>
> I'm not clear on why privilege would required for a process to remove
> capability bits from its set. (Sure, if running setuid).

I don't know what you mean by "Sure, if running setuid."

> Doesn't that just make it more difficult to write safe applications ?

My concern is that an unprivileged user could find a setuid root
application that will partially run but fail unsafely if it can get,
say, CAP_SETUID but not CAP_CHOWN, and is fooled in completing part of
it's privileged operation, then quitting or dying due to lack additional
privs, leaving the operation in an exploitable state.

No concrete examples offhand, but it feels like by not requiring
privilege we're inviting things similar to the old sendmail capabilities
bug. (I don't think that exact flow is possible, but something
similar) After all the setuid root applications assume that if
they ran as root, then they have full capabilities.

-serge

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