
Subject: Re: [PATCH 6/9] user namespaces: convert all capable checks in kernel/sys.c

Posted by [Daniel Lezcano](#) on Sat, 19 Feb 2011 17:52:46 GMT

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

On 02/17/2011 04:03 PM, Serge E. Hallyn wrote:

> This allows setuid/setgid in containers. It also fixes some
> corner cases where kernel logic foregoes capability checks when
> uids are equivalent. The latter will need to be done throughout
> the whole kernel.
>
> Changelog:
> Jan 11: Use nsown_capable() as suggested by Bastian Blank.
> Jan 11: Fix logic errors in uid checks pointed out by Bastian.
> Feb 15: allow prlimit to current (was regression in previous version)
>
> Signed-off-by: Serge E. Hallyn<serge.hallyn@canonical.com>

Acked-by: Daniel Lezcano <daniel.lezcano@free.fr>

>
> - if (!ns_capable(current->nsproxy->uts_ns->user_ns, CAP_SYS_ADMIN))
> + if (!ns_capable(current->nsproxy->uts_ns->user_ns, CAP_SYS_ADMIN)) {
> + printk(KERN_NOTICE "%s: did not have CAP_SYS_ADMIN\n", __func__);
> return -EPERM;
> + }
> + printk(KERN_NOTICE "%s: did have CAP_SYS_ADMIN\n", __func__);

A couple of printk left here.

Containers mailing list

Containers@lists.linux-foundation.org

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