Subject: Re: [PATCH 6/9] user namespaces: convert all capable checks in kernel/sys.c
Posted by akpm on Fri, 18 Feb 2011 23:59:28 GMT
View Forum Message <> Reply to Message

On Thu, 17 Feb 2011 15:03:42 +0000
"Serge E. Hallyn" <serge@hallyn.com> 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.
>
>
> ...
>
> --- a/kernel/sys.c
> +++ b/kernel/sys.c
> @@ -118,17 +118,29 @@ EXPORT_SYMBOL(cad_pid);
>
>  void (*pm_power_off_prepare)(void);
>
> +/* called with rcu_read_lock, creds are safe */
> +static inline int set_one_prio_perm(struct task_struct *p)
> +{
> + const struct cred *cred = current_cred(), *pcred = __task_cred(p);
> +
> + if (pcred->user->user_ns == cred->user->user_ns &&
> +     (pcred->uid  == cred->euid ||
> +      pcred->euid == cred->euid))
> +  return 1;
> + if (ns_capable(pcred->user->user_ns, CAP_SYS_NICE))
> +  return 1;
> + return 0;
> +}

uninline.  Document return value?

>
> ...
>

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