Subject: Re: [PATCH 4/9] allow killing tasks in your own or child userns
Posted by akpm on Fri, 18 Feb 2011 23:59:21 GMT
View Forum Message <> Reply to Message

On Thu, 17 Feb 2011 15:03:25 +0000
"Serge E. Hallyn" <serge@hallyn.com> wrote:

> /*
> + * called with RCU read lock from check_kill_permission()
> + */
> +static inline int kill_ok_by_cred(struct task_struct *t)
> +{
> + const struct cred *cred = current_cred();
> + const struct cred *tcred = __task_cred(t);
> +
> + if (cred->user->user_ns == tcred->user->user_ns &&
> + (cred->euid == tcred->suid ||
> + cred->euid == tcred->uid ||
> + cred->uid == tcred->suid ||
> + cred->uid == tcred->uid))
> + return 1;
> +
> + if (ns_capable(tcred->user->user_ns, CAP_KILL))
> + return 1;
> +
> + return 0;
> +}

The compiler will inline this for us.

> +/*
> * Bad permissions for sending the signal
> * - the caller must hold the RCU read lock
> */
> static int check_kill_permission(int sig, struct siginfo *info,
> struct task_struct *t)
> {
> - const struct cred *cred, *tcred;
> struct pid *sid;
> int error;
>
> @@ -656,14 +676,8 @@ static int check_kill_permission(int sig, struct siginfo *info,
> if (error)
> return error;
>
> - cred = current_cred();
> - tcred = __task_cred(t);
> if (!same_thread_group(current, t) &&

> -   (cred->euid ^ tcred->suid) &&
> -   (cred->euid ^ tcred->uid) &&
> -   (cred->uid ^ tcred->suid) &&
> -   (cred->uid ^ tcred->uid) &&
> -   !capable(CAP_KILL)) {
> +   !kill_ok_by_cred(t)) {
>   switch (sig) {
>   case SIGCONT:
>   sid = task_session(t);

_____

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